# SCIENTIFIC REPORT 2016 - 2018

Max Planck Institute for Intelligent Systems

Stuttgart & Tübingen

MAX-PLANCK-GESELLSCHAFT

MAX-PLANCK-GESELLSCHAFT

# Max Planck Institute for Intelligent Systems

# SCIENTIFIC REPORT 2016 – 2018

March 2019

This report was prepared on the occasion of the Scientific Advisory Board Meeting to be held on April 1 – 3, 2019.

## Imprint

| | |
|---|---|
| Address: | Max Planck Institute for Intelligent Systems<br><br>Stuttgart Location:<br>Heisenbergstraße 3<br>70569 Stuttgart, Germany<br><br>Tübingen Location:<br>Max-Planck-Ring 4<br>72076 Tübingen, Germany<br><br>https://is.mpg.de |
| Managing Director: | Prof. Dr. Bernhard Schölkopf |
| Deputy Managing Director: | Dr. Katherine J. Kuchenbecker |
| Scientific Coordination: | Dr. Matthias Tröndle |

# Contents

# PREFACE

This report presents research done at the Max Planck Institute for Intelligent Systems from January 2016 to December 2018. It is our third report since the founding of the institute in 2011.

This status report is organized as follows: we begin with an overview of the institute, including its organizational structure (Chapter 1). The central part of the scientific report consists of chapters on the research conducted by the institute's departments (Chapters 2 to 5) and its independent research groups (Chapters 6 to 18), as well as the work of the institute's central scientific facilities (Chapter 19). For entities founded after January 2016, the respective report sections cover work done from the date of the establishment of the department, group, or facility.

Bernhard Schölkopf, Managing Director, March 2019

### Electronic Version of the Report:

Citations in the PDF document are linked to the corresponding bibliography entry. To see a tooltip with the citation, hover over the right end of the citation number with your mouse pointer, using a PDF viewer that supports tooltips (e.g., Acrobat), as shown below.

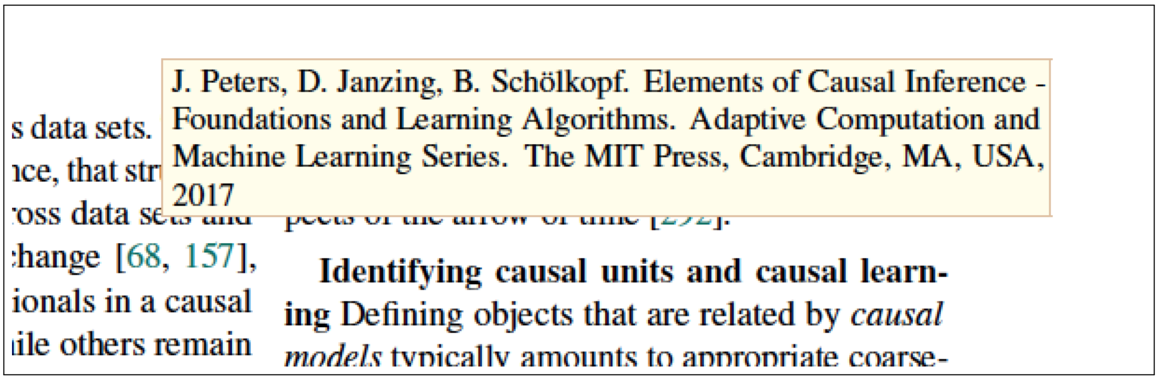


In Adobe Acrobat, the presence of tooltips can cause the following message to appear at the top of the document. Click on the symbol on the left in order to collapse the message and the green bar.

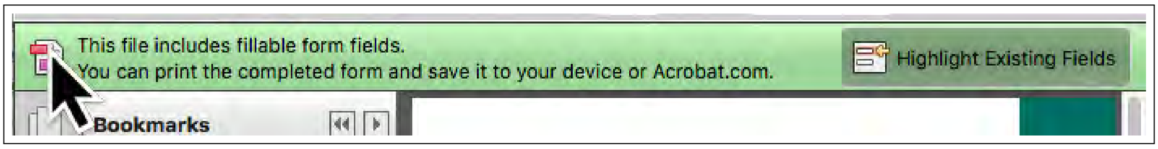

# 1 OVERVIEW

## 1.1 Scientific Concept of the Institute

The Max Planck Institute for Intelligent Systems (MPI-IS) strives to understand the principles of perception, action, and learning in autonomous systems that successfully interact with complex environments. In addition to gaining scientific understanding about natural intelligent systems, the institute's researchers further aim to use such insights to design future artificially intelligent systems that can benefit humanity. With campuses in Tübingen and Stuttgart, we combine theory, software, and hardware expertise in a single interdisciplinary center to enable the pursuit of creative and impactful research on a wide range of connected topics within the thriving research field of intelligent systems.

Our Tübingen site focuses more on *computational* aspects of intelligence, with departments in the broad fields of computer vision and machine learning, plus independent research groups that bring in perspectives from theory, algorithms, and robotics. The institute's Stuttgart site concentrates on *physical* realizations of intelligent systems with departments in the broad fields of mobile micro-robots and haptics, with independent research groups that emphasize smaller scales, biological inspiration, and control. Of course, many intelligent systems that we study involve both computational and physical aspects, providing ample opportunities for cross-site inspiration and collaboration.

The concept of the institute is inspired by the observation that biological systems have developed sophisticated abilities through interaction, evolution, and learning that enable them to act successfully in complex environments, and that a similar approach – although poorly understood at present – can be taken to synthesize effective autonomous systems. Regardless of their origin, form, or scale, we refer to such systems as *intelligent systems*. Understanding the computational and physical intelligence that underlies nature's systems is a major challenge: we work both to adequately describe such behavior and to use its key characteristics to invent synthetic, biohybrid, and human-assistive systems.

The field of artificial intelligence (AI) has long sought to create intelligent synthetic systems with capabilities similar to those of biological systems; the approach has traditionally involved explicit rules encoded in a computer program. This approach has succeeded for some well-defined tasks such as playing chess, but it has thus far dramatically failed for many tasks that are easily solved by biological systems (e.g., vision or movement control). These challenging tasks are often characterized by high-dimensional inputs and/or state spaces and complex stochastic nonlinear environments, in which case explicit models and simple behavioral rules are hard to design – humans have learned them over many years, and there is no reason for evolution and learning to produce solutions that are easy to comprehend post hoc.

A paradigm shift has taken place over the last few decades in major subfields of artificial intelligence, such as information retrieval and computer vision. Rule-based systems have been all but superseded by approaches that use machine learning. Consequently, machine learning methods have found their way into most fields of research and development and are now widely used in many areas that are highly relevant to human society.

However, researchers in fields that involve hardware systems (such as robots) still commonly try to handcraft solutions to challenging problems. In other words, these areas are in a state comparable to computer vision some time ago. Hand-engineered solutions are typically brittle and often work only in highly structured environments (e.g., factory assembly lines); they usually cannot generalize to novel situations and unstructured environments. We believe that machine learning has the potential to enable a performance jump in hardware-related intelligent systems similar to the jump that it facilitated in visual object detection and recognition.

Importantly, learning-based approaches are not limited to traditional engineering domains where systems have been designed to allow for classical control techniques: in principle, a learning-based approach can also be pursued in domains that are not amenable to standard engineering methods (e.g., bio-hybrid systems), as long as the basic mechanisms of perception, action, and adaptivity can be incorporated into such systems. Creating these new intelligent systems will likely also spur the invention of learning methods that can transcend the assumptions of problems in a single domain. The Max Planck Society (in German the Max-Planck-Gesellschaft, or MPG) founded the MPI for Intelligent Systems in 2011 to create a community of strong scientists who can drive the development of this vision and study the fundamental principles of perception, action, and learning in both computational and physical realizations of intelligent systems.

A major organizing principle for the institute is the concept of a perception-action-learning loop. Biological systems function by perceiving relevant aspects of the world and their own state, processing these aspects, and continuously deciding what actions to take based upon the results. While the computations involved in the processing of perceptual data and the generation of actions – and in particular the closed-loop properties of perception-action-learning systems – are far from understood in biological systems, they appear to involve aspects of learning and adaptation, as well as processes of inference, in the face of uncertain and even hidden information. Synthetic autonomous systems interacting with the world fundamentally face the same challenges: based on noisy measurements taken in a complex world, they need to perform actions that let them carry out critical tasks, such as maintaining their structural integrity, navigating a dynamic environment, manipulating an object, or harvesting energy.

As in biology, researchers at the MPI for Intelligent Systems investigate systems at length scales as small as nanometers. Small-scale mobile robots represent some of the most basic synthetic or bio-hybrid perception-action-learning systems one can hope to build in the next decade or two. By inventing such systems, we hope to gain insight into the intelligence that the small-scale biological world exhibits, including how the physical environment influences the perception-action cycle. Our scientists also study and learn from the animal world, including insects, fish, birds, and mammals. Inherent limitations in on-board sensing, perception, computation, power, and actuation capabilities necessitate advanced physical design, materials, interactions, computing, and controls. At the length scale of meters, we also take inspiration from the archetypal example of all intelligent systems: humans. In contrast to microsystems, human-like systems exploit significantly more memory and processing power in their perception-action loops, leading to potentially different mechanisms of perception, action, and learning (and also providing opportunities for technology to augment human capabilities). However, we believe that *the methods of learning and control in perception-action loops share essential aspects across platforms and scales*.

In summary, the Max Planck Institute for Intelligent Systems aims to establish the scientific foundations of *artificial* intelligence – perception, action and learning – through interdisciplinary research aimed at understanding and synthesizing intelligent behavior. We focus on constructing, modeling, and analyzing synthetic virtual, physical, biological, and bio-hybrid perception-action-learning systems using learning and self-organization, both for handling sensory complexity and for building structural and physical complexity. Methods of learning and inference are crucial for perceiving and acting in both the cyber world and the physical world. Our institute uniquely brings together expertise in theory, software, and hardware to address the fascinating and important research field of intelligent systems. The potential of this approach for basic science rests in the hope that even systems that are too complex to be analyzed in their entirety may be reducible to organizing principles that can be studied, comprehended, and applied in new contexts.

## 1.2 Organization

**Departments** The institute currently has four departments dedicated to intelligent systems research:

**Empirical Inference**
Prof. Dr. Bernhard Schölkopf (Tübingen)

**Haptic Intelligence**
Dr. Katherine J. Kuchenbecker (Stuttgart)

**Perceiving Systems**
Dr. Michael J. Black (Tübingen)

**Physical Intelligence**
Dr. Metin Sitti (Stuttgart)

The next four chapters of this report outline the research of each of these departments.

The Stuttgart site presently has two additional departments that were part of the former Max Planck Institute for Metals Research: the **Modern Magnetic Systems** department is headed by Prof. Dr. Gisela Schütz, and the **Theory of Inhomogeneous Condensed Matter** department is headed by Prof. Dr. Siegfried Dietrich. The **Phase Transformation, Thermodynamics and Kinetics** department of the former Max Planck Institute for Metals Research was closed upon the retirement of its director, Prof. Dr. Eric J. Mittemeijer, in October 2016. Given their distinct research foci, these three departments are not covered in this report.

Dr. Stefan Schaal, the former director of the **Autonomous Motion** department in Tübingen, left the Max Planck Society on April 30, 2018; he no longer has any formal affiliation with our institute. Michael J. Black has headed this department since May 2018 on a temporary basis. Given their scientific relevance to intelligent systems, the scientific publications of the Autonomous Motion department are listed at the end of this report.

In its final state, we anticipate that the Max Planck Institute for Intelligent Systems will have four departments in each site, for a total of eight, plus a large number of independent research groups and central scientific facilities.

**Research Groups** In addition to the departments, the institute hosts eleven independent research groups in the new research direction: one permanent Max Planck Research Group (PMPRG), four fixed-term Max Planck Research Groups (MPRG), five fixed-term Cyber Valley Max Planck Research Groups (CVRG), and one fixed-term ERC starting grant group (ERC). Like departments, each group has a home campus but interacts with researchers at both sites. Our independent research groups are as follows:

**Autonomous Learning** (MPRG)
Dr. Georg Martius (Tübingen)

**Autonomous Vision** (MPRG)
Prof. Dr. Andreas Geiger (Tübingen)

**Dynamic Locomotion** (MPRG)
Dr. Alexander Spröwitz (Stuttgart)

**Embodied Vision** (CVRG)
Dr. Jörg Stückler (Tübingen)

**Intelligent Control Systems** (CVRG)
Dr. Sebastian Trimpe (Stuttgart)

**Locomotion Biorobotics and Morphological Intelligence** (CVRG)
Dr. Ardian Jusufi

**Micro, Nano, and Molecular Systems** (PMPRG)
Prof. Dr. Peer Fischer (Stuttgart)

**Movement Generation and Control** (ERC)
Prof. Dr. Ludovic Righetti (Tübingen)

**Physics for Inference and Optimization** (CVRG)
Dr. Caterina De Bacco (Tübingen)

**Probabilistic Numerics** (MPRG)
Prof. Dr. Philipp Hennig (Tübingen; ended in October 2018)

**Rationality Enhancement** (CVRG)
Dr. Falk Lieder (Tübingen)

The institute hosts the following two Max Planck Fellows, who are excellent local university professors that were each recently appointed by the Max Planck Society to lead a small group at the institute:

**Physical Reasoning and Manipulation**
Prof. Dr. Marc Toussaint (Stuttgart)

**Statistical Learning Theory**
Prof. Dr. Ulrike von Luxburg (Tübingen)

The leaders of all thirteen of these groups outline their research in this report. In addition, the Stuttgart site of our institute hosts the ERC starting grant group of Prof. Dr. Laura Na Liu; because its research is not closely related to intelligent systems, the **Smart Nanoplasmonics** group is not covered in this report.

**Central Scientific Facilities** The MPI for Intelligent Systems also operates several central scientific facilities (CSFs, called ZWEs in German). Led by one or more scientists holding a Ph.D. and staffed with engineers and technicians, each CSF provides a particular type of scientific and technical support to our departments and groups. Such centralized support is a hallmark of the Max Planck Society and greatly expands the scope and quality of the research our scientists can pursue. The Materials CSF in Stuttgart was created by merging CSFs of the former MPI for Metals Research, while the rest of the Stuttgart CSFs and all of those in Tübingen were custom designed for our institute. Chapter 19 describes the work of our current CSFs.

**Management** The four directors in the new direction plus the two remaining directors from the former MPI for Metals Research constitute the Board of Directors of the MPI for Intelligent Systems. This board is headed by the institute's overall Managing Director (currently Bernhard Schölkopf), with support from the deputy Managing Director (currently Katherine J. Kuchenbecker); these two directors hail from our two campuses and each lead their sub-institute's Board of Directors, which handle decisions relevant to the respective local institute site.

The central administration that we share with the MPI for Solid State Research in Stuttgart, the local Stuttgart and Tübingen managing offices, and the Scientific Coordination Office (SCO) support the directors in all administrative and scientific matters. This latter entity, SCO, was newly established in 2018 to bridge and coordinate the scientific and externally facing activities for the entire institute and both sub-institutes. SCO's responsibilities include public relations, website, event management, reporting and documentation, grants, and policy. The Cyber Valley initiative, the International Max Planck Research School for Intelligent Systems (IMPRS-IS), and the Max Planck ETH Center for Learning Systems (CLS) are all closely associated with the Scientific Coordination Office.

## 1.3 Scientific Advisory Board

To ensure superior scientific quality, each Max Planck Institute is regularly evaluated by a Scientific Advisory Board (SAB). The President of the Max Planck Society appoints experts from leading international research institutions and universities to serve on the SAB.

The Scientific Advisory Board of the MPI for Intelligent Systems currently consists of the following eight members:

- Prof. Dr. Dario Floreano (École polytechnique fédérale de Lausanne, Lausanne, Switzerland)
- Prof. Dr. William T. Freeman (Massachusetts Institute of Technology, Cambridge, USA)
- Prof. Dr. Zoubin Ghahramani (University of Cambridge, United Kingdom)
- Prof. Dr. Danica Kragic (Royal Institute of Technology, Stockholm, Sweden)
- Prof. Dr. Barbara Mazzolai (Istituto Italiano di Tecnologia, Pontedera, Italy)
- Prof. Dr. Massimiliano Pontil (University College London, United Kingdom)
- Prof. Dr. Helge Ritter (Universität Bielefeld, Germany)
- Prof. Dr. Yair Weiss (Hebrew University of Jerusalem, Israel)

The following individuals served on our Scientific Advisory Board for the indicated time spans:

- Prof. Dr. Andrew Blake (Samsung AI Research and University of Cambridge, Cambridge, United Kingdom), 2012 – 2018
- Prof. Dr. Vijay Kumar (University of Pennsylvania, Philadelphia, USA), 2016 – 2018
- Prof. Dr. Josef A. Käs (Universität Leipzig, Germany), 2012 – 2016
- Prof. Dr. Itamar Willner (Hebrew University of Jerusalem, Israel), 2008 – 2012

The Scientific Advisory Board meets at the institute every three years to perform the main part of their evaluation. An extended evaluation takes place every six years, including additional rapporteurs who participate in several Scientific Advisory Board evaluations across the Max Planck Society.

Since the institute's founding in 2011, the Scientific Advisory Board of the Max Planck Institute for Intelligent Systems has had three scheduled meetings, as follows:

- January 12 – 14, 2013
- April 18 – 20, 2016 (extended evaluation)
- April 1 – 3, 2019

## 1.4 Board of Trustees

The Board of Trustees is an assembly of influential representatives from politics, industry, science, and the media. Its objective is to connect the institute to the public, particularly in the local region around Stuttgart and Tübingen. The Board of Trustees of the MPI for Intelligent Systems currently includes the following members:

## 1.5 New Department: Haptic Intelligence

The institute's newest director, Katherine J. Kuchenbecker, signed her contract to join the Max Planck Society on May 30, 2016, just shortly after the last meeting of the SAB. At the time, she was a tenured Associate Professor at the University of Pennsylvania (Penn), which is an Ivy League school in Philadelphia, USA. She started part-time Max Planck employment on June 15, 2016, and used her partial 2016 budget to buy a refurbished da Vinci robot to pursue research on teleoperation interfaces and robotic surgery. She visited the institute approximately every month for the rest of 2016 to plan her move, including onboarding, hiring, space, and the start of the International Max Planck Research School for Intelligent Systems (IMPRS-IS), for which she is spokesperson.

In January 2017 Katherine started full-time Max Planck employment and moved to Stuttgart with her husband and their cat. As shown in Fig. 1.1, her Haptic Intelligence (HI) department began with only three other people: a technician and an engineer transferred from another department, plus a newly hired assistant. For the first half of 2017, Katherine was still supervising a small research group at Penn, including four Ph.D. students and one masters thesis student. This previous group quickly transformed into her new MPI-IS department, united by a weekly group meeting held by video conference.

In addition to authoring a Springer encyclopedia entry on "Haptics and Haptic Interfaces," Katherine spent 2017 working to recruit great postdocs, Ph.D. students, engineers, and technicians. Her department's first new postdoc started at MPI-IS on March 1, and her first new Ph.D. student (who had just finished her masters thesis with Katherine at Penn) joined the Max Planck ETH Center for Learning Systems on July 1. Previous and current members of her Penn team, all of her MPI team members, and several of her future MPI employees met together in Munich in June 2017 at the IEEE World Haptics Conference. She and her remaining three Ph.D. students closed her main lab at Penn in August 2017.

Many new HI team members arrived in fall

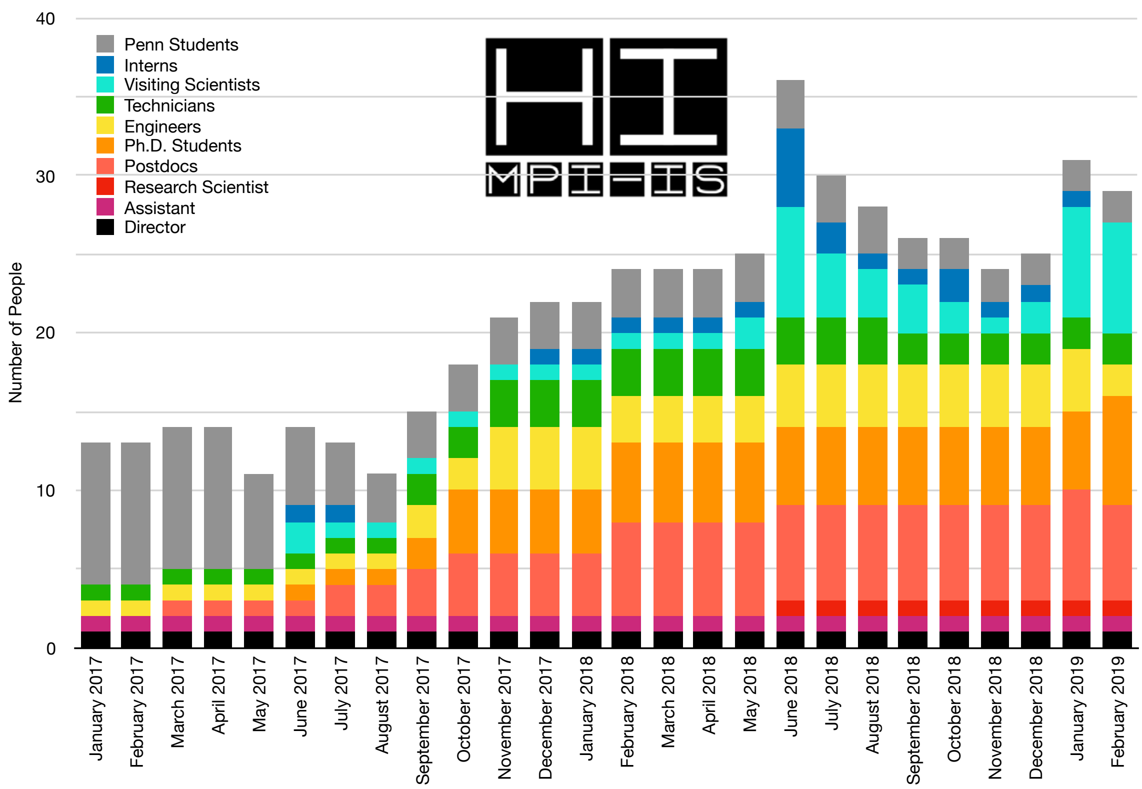


Figure 1.1: Monthly composition of the Haptic Intelligence Department since its inception.

2017, including the department's first three IMPRS-IS Ph.D. students. More department members joined slowly after that, with many interns and visiting scientists in the summer of 2018, plus the department's first research scientist, Ad Spiers, in June of 2018.

The HI department had possession of all of their offices (located on the fifth floor of the institute's main building in Stuttgart) from the beginning, but they started with only a few labs on the third floor, most of which are shared with the new Robotics Central Scientific Facility. They took over their main labs on the fifth floor in February 2018 and started using them immediately, as no renovations were needed. They are still in the process of planning the renovations to two labs on the third floor and several offices on the fifth floor, which will be reconfigured to provide a climate-controlled room for human-subjects research, a human motion-capture lab, and a larger meeting room optimized for video conferencing.

The start of 2018 brought a big push for the IEEE Haptics Symposium, which Katherine co-chaired both scientifically and organizationally. This biannual, four-day, 375-person conference took place in March 2018 in San Francisco, USA, and was attended by six other department members. The HI team presented three work-in-progress (WIP) papers, and Katherine presented a cross-cutting challenge. In June most department members attended EuroHaptics 2018 in Pisa, Italy, where the HI team presented two demos and one work-in-progress paper. Four HI team members traveled to Songdo, South Korea, in November for AsiaHaptics 2018, to present one workshop and two hands-on demonstrations. Katherine also delivered a keynote address at AsiaHaptics 2018, making a total of four keynotes in 2018: the earlier ones were at IEEE Virtual Reality (VR), Robotics: Science and Systems (RSS), and Hand, Brain, and Technology, all top venues in their respective fields.

As described in Chapter 5, the department's research centers on understanding haptic interaction while simultaneously inventing and optimizing helpful human-computer, human-machine, and human-robot systems that take advantage of the unique capabilities of the sense of touch. This interdisciplinary focus connects well with the physical, perceptual, and learning research taking place in the institute's three other departments, as well as with that of many of the institute's independent research groups.

The first two years of the HI department saw many victories. Katherine's fifth Ph.D. student Naomi Fitter graduated from Penn in May of 2017, completed a postdoc at USC, and started as an Assistant Professor at Oregon State University in the USA in January 2019. Alex Burka, another Penn Ph.D. student, graduated in August 2018 and quickly took a job at a local robotics startup. Katherine's first MPI masters thesis student, Maria Paola Forte, defended her thesis with the highest marks in July 2018, stayed on as an engineer, and is now a finalist for both the IMPRS-IS and CLS Ph.D. programs. The department's first postdoc, Gunhyuk Park, started a faculty position at GIST in South Korea in February 2019; two of the department's other postdocs are presently interviewing for faculty positions in North America, and another is interviewing for a permanent scientist position in France. Many team members were also successful in winning funding for their research, including a Humboldt Research Award for visiting Professor Brent Gillespie from the University of Michigan.

Katherine is particularly proud that the HI department is gender balanced, with equal numbers of men and women at all levels. The team is highly international and highly interdisciplinary, with a strong positive spirit. Department members are working hard and are beginning to see the fruits of their labor, such as the acceptance of a CHI 2019 paper and two ICRA 2019 papers. Many conference paper and journal article submissions are under review and already planned for 2019. Katherine will share new results in her keynotes at both RehabWeek 2019 (six combined conferences) in June in Toronto, Canada, and IROS 2019 in November in Macau, China. Thus, it seems the new MPI-IS Haptic Intelligence department is off to a good start!

## 1.6 New Institute Building in Tübingen


Figure 1.2: The new building of the MPI for Intelligent Systems in Tübingen.

The founding of the Max Planck Institute for Intelligent Systems was intimately connected with the establishment of a new institute site in Tübingen to leverage synergies with the local research institutions existing there, especially the University of Tübingen, the MPI for Biological Cybernetics, the MPI for Developmental Biology, and the Friedrich Mischer Laboratory. These connections complement the physical proximity and research collaborations between our institute site in Stuttgart and both the MPI for Solid State Research and the University of Tübingen. Germany's state of Baden-Württemberg supported the construction of the new institute building (Fig. 1.2) in Tübingen with funding of more than 41 million euros.

An architectural competition took place in 2012, and the winning team – ArGe Architekten, Waldkirch, Germany – started to plan the new building. After the architects and the Max Planck Society completed the planning process, the public and the neighboring residents were informed about the project, and construction began on the Max Planck Campus in Tübingen in September 2014. The laying of the foundation stone was publicly celebrated on April 27, 2015, after the basement of the institute building took shape.


Figure 1.3: The building inauguration ceremony featured participation by Martin Stratmann (left), who is the President of the Max Planck Society, the robot Apollo (center), and Winfried Kretschmann (right).

With the inauguration of this new building, our institute has established itself as the fourth major research institution on the Max Planck campus in Tübingen. The building's opening ceremony was held on July 12, 2017, and fittingly involved teamwork between humans and machines (Fig. 1.3). The robot Apollo held a red ribbon while the Minister-President of Baden-Württemberg, Winfried Kretschmann, and the President of the Max Planck Society, Martin Stratmann, cut the symbolic cordon to officially open the new building of the Max Planck Institute for Intelligent Systems in Tübingen.

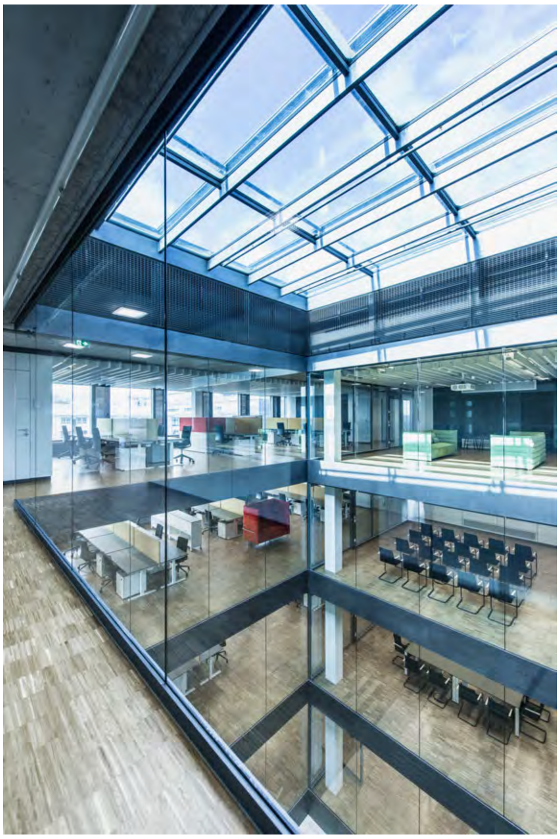

Figure 1.4: A view into the atrium of our new institute building.

The building's location is characterized by the hilly topography of Tübingen and a panorama of the Swabian Alb to the south and east. Communication zones enclose an atrium that spans five levels and is naturally exposed via a square skylight (Fig. 1.4). The upper floors are occupied by offices and administrative areas, with one department on each floor, while the lower level mainly contains laboratories. External balconies, shared tea kitchens, and meeting rooms with glass walls incorporate the viewing of the institute's beautiful surroundings into everyday working life.

Understanding and engineering complex technical systems are central topics of MPI-IS. Self-learning systems often control a large number of parameters in a way that optimizes a given objective function. The close connection between the new building project and the research orientation of the institute offered the opportunity to install a scientific energy-monitoring system, with the aim of recording the energy usage, performance, external influences, and user behavior. By equipping our building in this way, we enabled research projects that aim to minimize energy consumption and operating costs by means of intelligent controls.

### Timeline

- 2011: Commitment of the state of Baden-Württemberg to provide special financing of over €41 million
- 2012: Architectural competition
- May 2014: Building permit
- September 2014: Start of construction
- April 27, 2015: Laying of the foundation stone
- October 27, 2015: Roof-topping ceremony
- April 2017: First users move in
- July 12, 2017: Opening ceremony

### Data about the building

- 5, 6, or 7 floors tall in different locations
- 14,307 $m^2$ gross floor area
- 62,150 $m^3$ gross volume
- 5,900 $m^2$ of office, lab, and seminar space
- Office space for more than 250 employees
- Approximately €46.1 million total construction costs, of which the state of Baden-Württemberg funded €41 million

CyberValley

## 1.7 Cyber Valley

The Stuttgart-Tübingen region is at the heart of the German state of Baden-Württemberg. Not only is the region one of the engines of the German economy, but also a European powerhouse in terms of innovation. Metropolitan Stuttgart is home to some of the largest companies in the automotive and technology sectors, as well as to thousands of successful small and medium-sized enterprises. The region is well known for its high productivity, commitment to quality, and inventive spirit. Simultaneously, Tübingen is a major hub for education, research, and medicine.

The academic and industrial environment in the Stuttgart-Tübingen region provides ideal support for interdisciplinary research on intelligent systems. The region is already among the leaders in AI research: researchers from Stuttgart and Tübingen occupy the top position in scientific publications in Germany. In the field of machine learning, the region even occupies the top position in Europe and is among the top 10 locations worldwide.

Our institute recently joined forces with key players from industry, science, and politics to establish the **Cyber Valley** initiative, a regional cluster in the field of artificial intelligence. With this initiative, the Max Planck Society, the state of Baden-Württemberg, the Universities of Stuttgart and Tübingen, and companies from the IT and industrial sectors are working together to promote a start-up culture and significantly strengthen local research activities in the field of intelligent systems.

Cyber Valley is one of the largest research cooperations in Europe in the field of artificial intelligence, with partners from politics, science, business, and society. Cyber Valley strengthens research and education in the fields of machine learning, computer vision and robotics, as well as connections between these scientific disciplines. By promoting the exchange between science and industry, and by fostering start-ups, this initiative gives strong support to technology transfer in this important research field.

### 1.7.1 Motivation and Objectives

As core elements of Cyber Valley, new research groups have been established at the Max Planck Institute for Intelligent Systems (MPI-IS), the University of Stuttgart, and the University of Tübingen. The Cyber Valley partners and supporting foundations fund these groups. Over the next few years, several professorships will be established at the universities as additional core elements of Cyber Valley. Some of these new professorships will be financed by endowments. The establishment of a new doctoral program was another key element of the project: the International Max Planck Research School for Intelligent Systems is described in the next section.

With a new model of cooperation between science and industry, Cyber Valley aims to create a stimulating ecosystem for technology transfer in the field of artificial intelligence. After all, when it comes to the development of intelligent systems, the path from basic research to commercialization is often very short. Ideas and startups that originate in the research environment drive this development, and promoting these startups calls for collaboration between science and industry.

Since it was founded at the end of 2016, Cyber Valley has become a trademark for a large research initiative that spans the fields of machine learning, computer vision, and robotics. The research network was initiated by MPI-IS. The other founding partners were the University of Stuttgart, the University of Tübingen, the State of Baden-Württemberg, and the following seven industrial partners: Amazon, the BMW Group, IAV GmbH, Daimler AG, Porsche AG, Robert Bosch GmbH and ZF Friedrichshafen AG. Together, these partners all aim to strengthen the region by increasing research activities, creating an ecosystem for start-ups and technology transfer, and increasing the visibility of Baden-Württemberg as a global hotspot for AI research and development.

### 1.7.2 Brief History

Figure 1.5: Cyber Valley inauguration ceremony on December 15, 2016.

Together with the other project participants, Winfried Kretschmann (Minister-President of Baden-Württemberg), Theresia Bauer (Minister of Science in Baden-Württemberg), and Martin Stratmann (President of the Max Planck Society) launched the Cyber Valley initiative on Thursday, December 15, 2016, at Stuttgart's Neues Schloss (see Fig. 1.5).

In the two years since the initial letter of intent was signed, Cyber Valley has reached several important milestones. In 2017 and 2018, the focus was on hiring and onboarding research group leaders, as well as on expanding and consolidating academic partnerships and programs. In addition, as a result of the growing public interest in AI and its development, public and media relations have become an important part of the Cyber Valley managing office's work.

**New Cyber Valley research groups** In September 2017, all Cyber Valley partners took part in a joint symposium at MPI-IS. The symposium aimed to hire as many as ten new Cyber Valley research group leaders for MPI-IS and the Universities of Stuttgart and Tübingen. Twenty finalists presented their research and met with several Cyber Valley representatives.

The institute's hiring effort was very successful, with all five selected candidates accepting offers to become Cyber Valley group leaders. All of these outstanding early-career scientists took up their new positions between February and July 2018. Two are based mainly in Stuttgart, and the other three are in Tübingen. The University of Tübingen successfully hired three people to fill its two Cyber Valley group leader positions, one of which spends the other 50 percent of her time as a postdoctoral fellow at MPI-IS. While the University of Stuttgart did not hire any group leaders from the initial hiring symposium, it is currently in the process of hiring two Cyber Valley group leaders following another hiring symposium in January 2019.

**New Cyber Valley professorships** A total of ten new professorships will be part of the Cyber Valley initiative, in addition to a general expansion of activities in computer science, machine learning, and robotics at both universities. Two professors at the University of Tübingen were appointed in 2018. Further appointment procedures are currently underway at both universities, and most of the ten planned professorships are likely to be filled in 2019.

### 1.7.3 Organization

**Financial contributions** In total, all Cyber Valley partners – the Max Planck Society, the Universities of Stuttgart and Tübingen, several foundations, the private sector partners, and the State of Baden-Württemberg – are investing 165 million euros in the location as a first step towards establishing an internationally competitive AI hotspot.

The State of Baden-Württemberg, the Max Planck Society, and the Universities of Stuttgart and Tübingen are investing in new research buildings in Stuttgart and Tübingen, ten new university chairs, ten new research groups, a new graduate school, and central facilities to support Cyber Valley. From 2018 to 2022, the industrial partners are supporting the research groups at MPI-IS and the two universities with a total of 7.5 million euros. They will also finance two endowed professorships: Robert Bosch GmbH is funding a Chair for Machine Learning in Tübingen, and Daimler AG is funding a Chair for Digital Entrepreneurship in Stuttgart.

**Cyber Valley Research Fund** Most contributions from the industry partners are going towards the Cyber Valley Research Fund, for which only Cyber Valley research group leaders can apply. The fund is used to finance independent research projects, over which industry has no influence. A joint commission of all partners, in which representatives of science form the majority, selects topics based on scientific excellence. While the participating companies are allowed to propose topics, the scientists themselves are completely free to decide which projects they wish to pursue.

Those researchers or institutions that have discovered and invented something new retain intellectual property. Groundbreaking scientific innovations that could be potentially be turned into attractive products or services should primarily be used to set up new companies. If the inventors are not interested in starting a company themselves, the companies participating in Cyber Valley can apply to use an invention financed by the research fund, provided that they pay a customary license fee.

**Organization and management** The Cyber Valley research cooperation is governed by a cooperation contract, which specifies several boards that are in charge of decisions regarding the cooperation:

**Pleanary Assembly** All core partners are represented in the plenary assembly, in which each of the three main groups - i) the Max Planck Society, ii) the state of Baden-Württemberg and its universities, iii) and the industry partners - each cast one third of the votes.

**Executive Board** The Cyber Valley spokesperson and his or her two deputies form the Executive Board. The three members are representatives of the three main groups of the cooperation.

**Research Fund Board** Decisions about the distribution of funds from the Cyber Valley Research Fund are made by the research fund board, which comprises an equal number of members from academia and industry. In the event of a tie, the academic side decides.

The Cyber Valley managing office is located at the Max Planck Institute for Intelligent Systems and provides services for the entire cooperation. In terms of science and education, economic development, digitization, and international relations, Cyber Valley is increasingly important for the development of the state of Baden-Württemberg. The Cyber Valley managing office will thus continue to focus its efforts on meeting these demands. As the relationships with our two academic partners have grown stronger, the Cyber Valley management team has been expanded to include members from both the Universities of Stuttgart and Tübingen.

**Public relations** The heightened interest in AI-related topics has led to a major increase in the number of media requests for information and interviews. Cyber Valley has been featured in all major German print media, both in the science and politics/economics sections. Furthermore, regional radio and television broadcasters have reported on Cyber Valley and its progress many times. Cyber Valley has thus become the most visible research initiative of its kind in Germany.

Beyond this heightened media attention, events and visits have been important in conveying the goals and research of Cyber Valley to the public and informing policymakers of the importance of investments in this field. Cyber Valley has organized visits for public organizations, and for representatives of political parties at the local, regional, and state levels, as well as at the national and European levels (see Fig. 1.6). The initiative was also present with information booths at the MPI-IS open house days at both sites and at the new.New start-up festival. The Cyber Valley team also co-hosted the “AI and Society” public symposium on Max Planck Day in Tübingen, and participated in AI-CON at Bosch in Renningen, to name only two examples.

Figure 1.6: Visit of the President of Austria Alexander Van der Bellen (center) to Cyber Valley on November 29, 2018, upon invitation of Baden-Württemberg's Minister-President Winfried Kretschmann (left).

### 1.7.4 Further Development

**Start-up support** The Cyber Valley initiative aims to help researchers realize their ideas by setting up their own companies. Today, innovations in the field of digitization often take place outside the traditional technology transfer and are spearheaded by young, highly innovative start-ups. In this way, know-how not only remains in the region, but also new jobs and companies are created that can make a significant contribution to securing Germany's economic strength. While we plan to roll out our formal start-up support program in 2019, several actions to promote entrepreneurial activities and to support startups have already been implemented.

**Infrastructure** The state of Baden-Württemberg, the Max Planck Society, and the universities will create additional infrastructure by constructing two new buildings; one will be located in Tübingen, and the other will be located in Stuttgart. Moreover, we will establish new central scientific facilities dedicated to supporting the Cyber Valley cooperation.

**Towards an AI ecosystem** In addition to the core Cyber Valley research cooperation, further large-scale initiatives in AI have been successfully established by the partners in the region, such as new clusters of excellence at the universities and a new nationally funded competence center on machine learning and AI, which are all described in an upcoming section of this report. Cyber Valley has also helped attract world-leading researchers and new industry research activities. To further pursue the vision of Cyber Valley, we therefore aim to strengthen a regional ecosystem for education, basic research, cooperation, and innovation in AI through new partnerships and activities in the coming years.

**AI network in Baden-Württemberg** Beyond the Stuttgart-Tübingen region, we aim at building up a strong research network in Baden-Württemberg. Recently, ten new junior professorships in AI have been established at the universities of Freiburg, Heidelberg, Hohenheim, Konstanz, Mannheim and at the Karlsruhe Institute of Technology to strengthen competence throughout the entire state in synergy with Cyber Valley.

**A hub for European AI** The federal government has recently announced a German AI development strategy. The four newly founded competence centers for machine learning and AI funded by the Federal Ministry of Edcuation and Research, one of them hosted in Tübingen, are to play a key role in the government's strategy. Joining forces only within Germany will not be sufficient in order to be internationally competitive with the tremendous investments in China and the US. We therefore envision Cyber Valley to serve as a leading hub in building a prominent European network for modern AI research. We strongly believe that Europe needs to organize and take a worldwide leadership role as the economic and societal influence of AI increases.

## 1.8 International Max Planck Research School for Intelligent Systems

The International Max Planck Research School for Intelligent Systems (IMPRS-IS) is our institute's main doctoral program and a key element of the Cyber Valley initiative, which was described in the previous section. IMPRS-IS was founded in early 2017 to bring a new generation of young scientists and engineers into our highly multi-disciplinary research environment, enabling them to learn to tackle the fundamental challenges of intelligent systems while simultaneously earning a Ph.D. under the mentorship of our directors, our group leaders, and top faculty in aligned research areas at our neighboring universities.

IMPRS-IS brings together the MPI for Intelligent Systems, the University of Tübingen, and the University of Stuttgart to form a highly visible and unique graduate school of internationally renowned faculty working at the leading edges our field. At a high level, this program aims to advance human knowledge about intelligent systems and directly support our institute's research mission by recruiting, educating, and supporting outstanding doctoral students. The participating faculty, students, and staff are committed to pushing this graduate training program forward, making it among the best worldwide. After only two years since the founding of the IMPRS-IS, we believe we are well on our way to achieving this ambitious goal.

### 1.8.1 Motivation for Establishing IMPRS-IS

Our society is changing rapidly due to the immense progress made in the research fields of artificial intelligence and robotics. Researchers, politicians, business leaders, and philosophers – anyone who tries to assess the future – agree that these new technologies are already changing the way we live, work, learn, and travel. Research in this field is progressing rapidly, but it lags far behind the natural sciences and is only beginning to gain momentum. Whatever the future may look like, it is important that we in Europe confidently rely on our academic and social values and traditions to play a leading role in the development of the intelligent systems of the future. We should not leave such research solely to scientists in North America or China.

The founding of our institute and the creation of the Cyber Valley initiative represent a major strategic investment in basic research designed to enable Germany to play a leading role in the field of intelligent systems. Tomorrow's learning methods will be invented today, most likely by curious, creative, hard-working young people; this is where the IMPRS-IS comes in. One needs an interdisciplinary education in order to understand intelligent systems. Such an educational opportunity did not exist in Germany in a structured form until the IMPRS-IS came into being. In particular, Max Planck Institutes on their own are research institutions and cannot confer degrees of any kind; although doctoral research can be performed within many types of institutions, only universities have the right to grant a Ph.D. in Germany.

Unsurprisingly, intelligent systems research is being pursued with vigor worldwide; there is thus fierce competition for the best minds in the scientific community at all levels, including doctoral students. Consequently, the first mission of the IMPRS-IS is to help our faculty members attract a diverse and excellent set of doctoral students from around the world. In particular, we look to recruit candidates with a strong academic background and a master's degree in engineering, computer science, cognitive science, mathematics, control theory, neuroscience, materials science, physics, or related fields. Students who join this doctoral program benefit from near-perfect research conditions: they can conduct basic research freely, within a uniquely strong scientific community that is supported by the open and democratic society of southern Germany.

### 1.8.2 Brief History

International Max Planck Research Schools have been an integral part of the Max Planck Society's approach to doctoral research since the year 2000, granting talented students from both Germany and abroad the opportunity to earn a doctorate under excellent conditions. Common characteristics of these graduate programs are the close cooperation between an MPI and one or more local universities, a curriculum of research seminars and soft-skill workshops, and a strong academic community that brings together students, faculty, and other researchers.

Our IMPRS is one of 67 such schools. Its creation was enabled by a central grant from the Max Planck Society to MPI-IS with an initial duration of six years, as well as significant funding from the state of Baden-Württemberg for doctoral education at both the University of Tübingen and the University of Stuttgart. The successful grant proposal was led by Michael J. Black in summer of 2016 with support from Katherine J. Kuchenbecker and our scientific coordinator, Matthias Tröndle. To give Michael time to manage the larger Cyber Valley initiative, Katherine took on the role of IMPRS-IS spokesperson when she joined the institute in January 2017. The school's coordinator, Dr. Leila Masri, joined the institute in April 2017 and immediately took over management of the first recruitment round, which was well under way.

The International Max Planck Research School for Intelligent Systems enrolled its first class of 28 scholars in fall of 2017; as shown in Fig. 1.7, these students come from 11 different countries and have backgrounds ranging from mathematics, computer science, and machine learning to physics, engineering, and robotics. The class was selected in a multi-stage process from the 362 total applications that were submitted online. After completeness checks, review by multiple scientists, and a round of video-chat interviews, the 48 strongest applicants were invited to the first IMPRS-IS interview symposium, which took place at the two sites of the MPI for Intelligent Systems in June 2017. There, each applicant presented his or her past research projects and future research interests, took part in multiple interviews with selected IMPRS-IS faculty, and got to know our institute and our surrounding research community. Faculty each evaluate a set of candidates, but final admission decisions are made strictly by the seven-person executive board of the IMPRS-IS; this board includes the school spokesperson, a deputy spokesperson from each of our partner universities, another director from each site of MPI-IS, an independent group leader, and a faculty member elected as the school's equal opportunity officer.

Figure 1.7: The 28 members of the first class of IMPRS-IS scholars, along with a handful of faculty, staff, and associated students at the school's first bootcamp.

Word about the new IMPRS-IS spread quickly, and we held our second round of Ph.D. student recruiting just seven months later, following a similar selection process. The second generation of IMPRS-IS scholars included 28 individuals who interviewed on site in January 2018 and joined the school around April 2018, making a total of 56 scholars. To maintain the school's quality and focus, the executive board decided that even students who already hold a Ph.D. contract with one of the IMPRS-IS faculty members must formally apply online, take part in all parts of the interview symposium, and be admitted by a vote of the IMPRS-IS executive board. This process is the only way in which scholars can be admitted to the school.

The standard MPI-IS employment offer is for each scholar to receive a three-year-long contract equivalent to a 65% E13 position in the TVöD scheme for public employment in Germany; this pay level is higher than the minimum of 50% permitted within MPG because recruiting in the field of intelligent systems is competitive. Importantly, the source from which a particular doctoral student is funded has no bearing on his or her ability to become an IMPRS-IS scholar. A small number of scholars at each institution are co-funded by the school, with these funding allocations decided by the executive board on the basis of criteria that are aligned with our school's values and communicated clearly to all faculty.

The third hiring round of the IMPRS-IS closed at the beginning of February 2019, reaching a new height in popularity. After receiving a total of 675 applications, we selected 103 candidates to participate in our third interview symposium in January 2019, with each applicant giving a 15-minute talk and taking part in between two and ten 45-minute interviews. Approximately 70 offers for funded doctoral positions are in the process of being made for this third round. It is hard to believe believe how popular our school has become!

We will soon have about 130 Ph.D. scholars in the IMPRS-IS, which will make it the largest of all 67 such schools. In just two years, IMPRS-IS has far exceeded our initial vision of recruiting 100 Ph.D. scholars over the course of the initial six-year grant proposal. These high enrollment numbers are testimony to the success of our faculty members at obtaining third-party funding and their knack for devising attractive research projects that captivate qualified applicants.

We are very proud of our scholars: each one is highly talented and motivated and is working to contribute to cutting-edge research that spans the interconnected fields of computational cognitive science, computer graphics, computer vision, control systems, human-computer interaction, machine learning, micro- and nano-robotics, perceptual inference, and robotics.

To help contribute to the management of the school, the entire population of scholars elects a set of four student representatives on an annual basis, with one representative for each site of our institute as well as one per university. One of these representatives is then internally selected to represent all of the scholars at meetings of the IMPRS-IS executive board.

Post-masters and doctoral students who are doing research with an IMPRS-IS faculty member but who are not IMPRS-IS scholars can become associated students of the school. Application requires only a current curriculum vitae and the endorsement of the supervising faculty member. Associated students take part in the activities of the school when space is available.

### 1.8.3 Objectives

The IMPRS-IS offers an interdisciplinary doctoral research program with opportunities to study all key areas of intelligent systems, ranging from computer science to control theory, from mechanical engineering to neuroscience, and from mathematics to materials science.

In addition to ongoing support from their primary advisor, each scholar benefits from regular interactions with their Thesis Advisory Committee (TAC). This cross-institutional panel is formed near the start of the Ph.D. and includes three or more IMPRS-IS faculty members (including the scholar's main advisor) from at least two different institutions. One TAC member must be part of the faculty at the university where the student wants to earn his or her doctorate. The student drives the process of proposing and securing consent from the faculty members they believe would add the most to their research.

The TAC reviews and provides suggestions on the Ph.D. project, monitors the progress of the scholar, and advises her or him on further studies and research. The TAC also approves a thesis as ready for submission. When it can be foreseen that a Ph.D. project will not be successfully finished, the TAC can recommend to discontinue the project; this decision must be confirmed by the IMPRS-IS executive board.

The scholar meets with his or her TAC at least once per year and documents these meetings by written minutes with an executive summary confidentially communicated to the IMPRS-IS coordinator. These meetings preferentially take place in months 6-9, 15-18 and 30-33 after the beginning of a Ph.D. project. On request of the student or one of the TAC members, the TAC can also meet more frequently than once per year.

Beyond the scholar's own research, the IMPRS-IS program includes the following courses and events:

**Courses** Over 40 specific scientific courses ranging from "Advanced Artificial Neural Networks" to "Theory of Machine Learning", which are taught by IMPRS-IS faculty, associated faculty, and instructors.

**Cross-institutional exchanges** Every department and research group within the broader IMPRS-IS community participates in cross-site bridging events where interested members of our community can visit each other and learn from lab tours. In addition to the scientific stimulation of such experiences, these occasions help scholars build their professional networks and foster a strong community feeling in an informal environment.

**Bootcamps** The entire IMPRS-IS community participates in an intensive one-week off-site bootcamp each year. Both faculty and scholars present their latest findings through talks and poster sessions. Additionally, interdisciplinary tutorials, soft-skill seminars offering tailored career development, design competitions, and social activities are organized to widen the perspective of the scholars and increase their efficiency and effectiveness as Ph.D. students.

**Annual chats** The school's coordinator (Leila Masri) meets for 30 to 45 minutes with each IMPRS-IS scholar on an annual basis to hear how their research is going and help them strategize about any challenges they may be facing. Providing such coaching helps our school both support individual scholars as well as hear about and handle any larger issues that may affect our community. The coordinator also meets personally with each new faculty member to explain the details of the school and checks in regularly with all faculty, particularly during recruiting season.

**Networking and social activities** The IMPRS-IS organizing team and the school's elected student representatives cultivate a friendly environment by organizing frequent social activities to help students advance academically and personally, all the while driving artificial intelligence to the forefront of excellent research and setting the groundwork for future research leaders in this field.

**Support programs for women in science** IMPRS-IS seeks to increase the number of women in areas where they are underrepresented, so we explicitly encourage women to apply. IMPRS-IS has also established a program that fully funds internships for female masters students doing research with one of our MPI-IS faculty for a duration of up to six months. In addition, IMPRS-IS is a key founder and cooperator of the Athena Group, a bottom-up organization that provides a mentoring program for our young female scientists and supports them in their career development; all female IMPRS-IS scholars are strongly encouraged to join this mentoring program. In collaboration with the Athena Group, the school has implemented a series of talks wherein female scientists present a lecture on their research at one of our institute sites and then also talk about their personal career experiences with a small group of individuals in a more intimate setting.

### 1.8.4 Organization

The current faculty of the IMPRS-IS includes world-renowned researchers drawn from the MPI for Intelligent Systems, the University of Tübingen, and the University of Stuttgart. All are established leaders in intelligent systems and have the right to advise scholars within the IMPRS-IS.

The following seven faculty form the **Executive Board** of the IMPRS-IS:

**Prof. Dr. Frank Allgöwer** Professor at the University of Stuttgart and Deputy Spokesperson

**Dr. Caterina De Bacco** Cyber Valley Research Group Leader at MPI-IS and Equal Opportunity Officer of the Executive Board

**Dr. Michael J. Black** Director at MPI-IS

**Dr. Katherine J. Kuchenbecker** Director at MPI-IS and IMPRS-IS Spokesperson

**Prof. Dr. Hendrik Lensch** Professor at the University of Tübingen and Deputy Spokesperson

**Dr. Metin Sitti** Director at MPI-IS

**Dr. Alexander Spröwitz** Max Planck Research Group Leader at MPI-IS

### 1.8.5 Faculty

Any director, professor, or independent research group leader associated with the MPI for Intelligent Systems, the University of Stuttgart, or the University of Tübingen may apply to become a member of the faculty of the IMPRS for Intelligent Systems. After securing the support of his or her institution's spokesperson or deputy spokesperson, the prospective faculty member submits a curriculum vitae and a motivational letter. The school's executive board decides on each nomination in an in-person meeting based on the employment status, research track record, relevance of the research direction, and training record of the nominee.

IMPRS-IS presently has 34 faculty, including all 4 MPI-IS directors, 11 MPI-IS group leaders, 9 University of Tübingen professors, 3 University of Tübingen group leaders, and 7 University of Stuttgart professors. Both of the institute's Max Planck Fellows are included in these numbers, as are all of the current Cyber Valley group leaders. After taking up their positions, the two Cyber Valley group leaders presently being recruited to join the University of Stuttgart are also expected to join the IMPRS-IS faculty through the standard nomination procedure.

In addition to these faculty members, IMPRS-IS currently has six associated faculty, who may co-advise scholars with members of our faculty and also help teach at the boot camp. These individuals are Ph.D.-holding scientists in the three participating institutions, including leaders of central scientific facilities, non-independent group leaders, and university faculty working at the fringes of intelligent systems. Finally, interested postdocs and research scientists working with our faculty and associated faculty can also contribute to the educational mission of the school and become IMPRS-IS instructors upon approval of the executive board. For full lists of our faculty, associated faculty, and instructors, please see https://imprs.is.mpg.de/people.

### 1.8.6 Future Plans

The future of the IMPRS-IS looks highly promising; we intend to keep operating it for the lifetime of our institute. The vast increase in applications over our three recruitment rounds underlines the timeliness and importance of interdisciplinary research in intelligent systems. Simultaneously, the high quality and diversity of the students enrolling in this doctoral program show that we are already competitive with top university programs worldwide. Both the MPI for Intelligent Systems and our two partner universities are committed to finding ways to continue growing our school to meet this demand without compromising quality, focus, or personal interactions with our scholars and faculty.

The school's founding in 2017 constituted a key element of the Cyber Valley initiative, confirming that the Stuttgart/Tübingen region is a major AI and robotics research hub in Europe and the world. Every day, the International Max Planck Research School for Intelligent Systems further strengthens this ecosystem in southern Germany, bringing together faculty and students at our three institutions to foster research collaborations, interdisciplinary scientific insights, and friendships that will last a lifetime.

Max Planck ETH Center for Learning Systems

## 1.9 Max Planck ETH Center for Learning Systems


The Max Planck ETH Center for Learning Systems (CLS) is an international and inter-institutional endeavor, and also a widely cross-disciplinary effort bringing together Europe's leading basic research institutions to tackle one of the most prominent scientific challenges of our time. Its core element is an elite cross-institutional co-advised doctoral program that enrolls between five and ten new students per year. Many of the biggest scientific opportunities and most impactful innovations around machine intelligence and learning systems require expertise from multiple disciplines. At ETH Zürich this collaborative opportunity has been mirrored by the organizational challenge not only to bring together researchers from MPI-IS and ETH, but also to coordinate researchers within ETH who are associated with different academic departments.

CLS is based on a bottom-up participation-based approach and prioritizes inter-institutional scientific projects conducted by Ph.D students doing research at both institutions. This grass-roots approach has helped tremendously to build a vivid CLS community, something that has been further supported by annual off-sites and a multitude of other events such as targeted workshops and summer schools. Only a few years after its founding in 2015, today CLS has established a network and productive exchange of Ph.D. students and researchers. There has been a growing identification of members and fellows with the center on the inside and an increased interest in CLS positions from the outside, resulting in a candidate pool that stands out, even relative to the selective programs at both institutions.

### 1.9.1 Motivation for Establishing CLS

One of the most significant challenges of science today is to cope with the complexity that arises in biology, medicine, engineering, economics, sociology and many other areas of high societal relevance. Humans and animals are constantly managing multiple sensory inputs and have become experts in using the integrated knowledge to their advantage. Using this approach of natural learning as a basis, artificial intelligence systems have been built by computer scientists. Machine learning approaches are able to process large and complex information, while adjusting and adapting their behavior to environmental influences. Examples of artificial learning systems are robots that can adapt their behavior to their environments, or machine-learning-based software systems that are able to make predictions based on big data sets. Natural and artificial learning systems are both susceptible to highly unreliable, stochastic factors. Both the natural and engineering sciences (with their complementary scientific methods of analysis and synthesis) explore such learning systems by interacting with them, by modeling them, and by explicit construction or reconstruction. However, to date a general understanding and comprehensive approach within learning systems, especially with respect to the integration of natural and artificial features aiming to expand the analysis, is still elusive.

The excellent engineering competences of the faculty and research team members at the Eidgenössische Technische Hochschule (ETH) Zürich in Switzerland ideally complement the competences in natural sciences and computer science at the Max Planck Institute for Intelligent Systems, Tübingen/Stuttgart (MPI-IS) in Germany.

## 1.9.2 Brief History

### Research Network on Learning Systems

Since the institute's founding in 2011, MPI-IS researchers have built up strong connections with scientists at ETH, leading to the initiation of the "Research Network on Learning Systems" (RNLS) founded by MPI-IS and ETH in 2013. The goals of RNLS were to achieve a fundamental understanding of perception, learning and adaptation in complex systems, by providing a platform for exchange in research and education. Under the umbrella of RNLS, several events – including a Summer School with more than 100 participants – were initiated to bring together leading researchers and junior scientists from MPI-IS and ETH.

### Max Planck ETH Center for Learning Systems

The efforts started under RNLS received additional support, and we realized that we needed to form a new entity in order to boost our shared research capacities and pool scientific know-how in the emerging field of learning systems in Europe. MPI-IS and ETH thus successfully proposed the establishment of the "Max Planck ETH Center for Learning Systems," also known as CLS, which had an initial run time of five years and can be extended for another five years after the completion of a successful evaluation in 2019. CLS was officially inaugurated by the presidents of ETH Zürich and the Max Planck Society on November 30, 2015. Many internationally recognized researchers in this field attended the inauguration ceremony and have been involved in the long-term development of CLS. The center has become an internationally recognized academic program aiming to provide cross-disciplinary doctoral training and networking to highly qualified individuals. As a joint program of two leading institutes in this field of research, CLS can draw from a variety of highly recognized experts in the field of artificial intelligence.

The Max Planck ETH Center for Learning Systems is the first joint center between the renowned German Max Planck Society and this world-famous Swiss research institution. Ph.D. Students at both institutes benefit from joint supervision through ETH professors and MPI-IS directors and group leaders. Students and young scientists also participate in topic-specific summer schools and workshops, and they have the possibility of taking an internship in industry. By offering these unique opportunities, CLS ensures that young academics in the field of learning systems at ETH and MPI-IS have the chance to engage in outstanding collaborations and exchanges early in their careers. These connections build a strong positive basis for their professional careers after the completion of their Ph.D.

## 1.9.3 Objectives

The main objective of CLS is to foster cross-institutional exchange for the larger benefit of driving artificial intelligence to the forefront of excellent research and setting the groundwork for our future research leaders of this field. This goal is achieved by:

1. Agreement on Ph.D. research topics by a supervisory team consisting of an ETH professor and an MPI director or independent research group leader
2. Training through specialized workshops
3. Networking through science retreats

4. Short- and long-term cross-institutional exchanges

The Max Planck ETH Center for Learning Systems addresses cross-disciplinary research questions in the design and analysis of natural and man-made learning systems. A particular focus of the center is the development of a unified view of learning systems found in natural sciences and engineering disciplines. On the one hand, scientists are investigating learning systems in the natural sciences, such as nervous systems and adaptive systems, including immune systems or metabolic dynamics in living cells. On the other hand, research in engineered learning systems, such as robots capable of adapting their dynamic behavior to that of the environment or self-organizing systems like the Internet, belongs to the cutting-edge science projects of our time. Both artificial and natural systems share the property of having many degrees of freedom while having to combine computations with strong dynamical fluctuations. Rather than perceiving stochastic theory as a menace, scientists are in the process of discovering beneficial sides of stochastic dynamics, ranging from stochastic resonance in sensing and signal processing with increased sensitivity, to randomized algorithms with an expected fast run time. A generally accepted unifying viewpoint in the design and analysis of complex stochastic systems has not yet been developed; an alliance of ETH and MPI-IS with their outstanding engineering competences and excellent knowledge on natural learning systems is situated in a unique position to achieve substantial progress on this research front. To this end, the center hosts activities in research fields that build our understanding of complex systems via the pursuit of the following six central goals:

**G1** Machine Learning and Empirical Inference of Complex Systems

**G2** Perception-Action-Cycle for Autonomous Systems

**G3** Robust Model-Based Control for Intelligent Behavior

**G4** Robust Perception in Complex Environments

**G5** Design, Fabrication, and Control of Synthetic, Bio-Inspired, and Bio-Hybrid Micro/Nanoscale Robotic Systems

**G6** Neurotechnology and Emergent Intelligence in Nervous Systems

ETH and MPI-IS have strong research groups with international visibility in all of these areas, and pooling these competences has created a salient network of expertise in learning systems. Our center has already attracted some of the best students and researchers who have joined CLS to further share, develop and foster their own skills within the community. Several events such as retreats and workshops have been initiated successfully in order to bring together researchers from MPI-IS and ETH to tackle these challenges.

### 1.9.4 Organization

The Max Planck ETH Center for Learning Systems has three locations: the MPI for Intelligent Systems in Tübingen, the MPI for Intelligent Systems in Stuttgart, and the ETH Zürich. The center is equally financially supported by ETH Zürich (50%) and by the Max Planck Society and the Max Planck Institute for Intelligent Systems (50%) with a total budget of 5 million EUR for 2015-2020.

#### Involved Institutes and Departments

Learning systems research and teaching activities are pursued in several departments of ETH Zürich, each with its specific application and research context. The Department of Computer Science supports significant activities in machine learning and computational intelligence, while the Department of Information Technology and Electrical Engineering provides expertise on information theory and practice of control. The Department of Mechanical and Process Engineering focuses on autonomous robots, control of robotic systems and on micro- and nanotechnology for robotics. The Department of Health Sciences and Technology adds competences in physical human-robot interaction and medical applications. The Institute for Neuroinformatics is devoted to studying the brain and its learning capabilities both experimentally and by simulation.

The Max Planck ETH Center for Learning Systems is officially connected to the ETH Department of Computer Science, but it is jointly operated by the engineering departments of Computer Science, Mechanical and Process Engineering, Information Technology and Electrical Engineering. Furthermore, professors of Health Sciences and Technology, Mathematics, Biosystems Science and Engineering and Civil, Environmental and Geomatic Engineering are actively involved in the center activities.

On the Max Planck side, the center is affiliated with both sites of the Max Planck Institute for Intelligent Systems. All departments and almost all research groups are participating, giving CLS connections into all learning-related facets of our institute's research.

## Steering Board

The Max Planck ETH Center for Learning Systems is headed by two co-directors, one from each site. These are currently:

**Prof. Dr. Thomas Hofmann** Professor for Data Analytics, Department of Computer Science, ETH Zürich

**Prof. Dr. Bernhard Schölkopf** Director, Empirical Inference Department, Max Planck Institute for Intelligent Systems

The steering board of the center supports and advises the co-directors with respect to the strategic development of the center, including hiring. It is composed of:

**Prof. Dr. Luc van Gool** Professor for Computer Vision, Department of Computer Science, ETH Zürich

**Prof. Dr. Andreas Geiger** Group Leader at the Max Planck Institute for Intelligent Systems

**Prof. Dr. Bradley Nelson** Professor of Robotics and Intelligent Systems, ETH Zürich

**Dr. Metin Sitti** Director, Physical Intelligence Department, Max Planck Institute for Intelligent Systems

From 2015 to 2018, Professor Philipp Hennig (then an independent research group leader at MPI-IS in Tübingen) greatly helped the development of CLS by being a member of the steering board. His place was taken over by Professor Andreas Geiger.

The co-directors and the steering board are supported in programmatic and organizational matters by the executive office of the center, headed by the scientific coordinator. The development and activities of CLS are coordinated in close interaction with the board and the faculty of the center. The scientific coordinator organizes scientific events such as meetings, workshops and educational courses, while also acting as a key contact person for all members and applicants and communicating with the public on behalf of the program.

## Faculty and Students

CLS researchers participate in the center activities either as members, associated members, fellows, or associated fellows.

**Members** form the senior core faculty of CLS. These are full professors or directors at ETH and MPI-IS, respectively.

**Associated members** are formed by associated/assistant professors at ETH and independent group leaders at MPI. Moreover, full professors in related areas can also join CLS as associated members.

Currently, 48 members and associated members are engaged at CLS, including 32 professors from ETH Zürich and 16 directors and group leaders from MPI-IS. CLS currently has 10 alumni members. For a full list, see http://learning-systems.org/members.

**Fellows** on the Ph.D. and postdoc level are at the heart of CLS. Ph.D. fellows are co-supervised by ETH and MPI researchers, and they spend at least one year at each institution. The center supports 50% of the salary for fellows, and Ph.D. fellows earn their doctorates from ETH.

**Associated fellows** are students and postdoctoral researchers who are fully engaged in CLS activities, and who can draw on the benefits of the CLS program. In particular, associated doctoral fellows from MPI can graduate from ETH.

The Center for Learning Systems establishes a common pool of qualified applicants that can benefit from both institutions and produce publications in high impact conferences and journals. Doctoral students co-supervised by scientists in both institutions have the opportunity to gain experience from working in two distinct research laboratories, with two mentors who necessarily have different insights and capabilities. Most doctoral students need to complete a postdoctoral research fellowship in order to gain the new perspectives and interdisciplinary insights that naturally stem from being co-advised across institutions. Joint events organized by CLS guarantee sustainable and long-term collaborations between scientists from different fields. All in all, this combination attracts highly motivated and talented scientists of all relevant levels, backgrounds and disciplines. CLS gives these future leaders a deep and broad scientific network, challenges them to contribute to top-tier international research, and exposes them to different top-performing research environments.

Since 2015, 95 Ph.D. students and postdocs joined the CLS as fellows or associated fellows. Currently, 18 full Ph.D. fellows, 36 associated Ph.D. fellows, and nine associated postdoctoral fellows are actively involved in the center. Postdoctoral fellows were appointed at the beginning of CLS from existing MPI-IS and ETH postdocs with the goal of kick-starting CLS; specifically, these postdoctoral fellows benefit from collaboration with the other institute and the opportunity to help guide the research of CLS Ph.D. students. Partly as a result of this effort, CLS developed its own momentum and the focus shifted to emphasize Ph.D. students, so no further postdoc fellows were appointed. CLS currently has 32 alumni fellows.

So far, CLS has published four calls for doctoral students who are being jointly advised by one supervisor from ETH and one from MPI-IS, while working on projects that span research at both institutes. Although CLS was established only in 2015, its calls for applications have received tremendous attention, and so far more than 1000 students have applied to be admitted to CLS in total across the four recruiting rounds.

The selection of CLS Ph.D. students is highly competitive, with an acceptance rate of about 2.5%. In 2016, four Ph.D. students took up their studies, while the calls of 2017 and 2018 resulted in 5 and 8 new doctoral students, respectively. The statistics for the 2016, 2017, 2018, and ongoing 2019 application rounds can be found in Table 1.1.

Overall, 5 full and 11 associated doctoral fellows have successfully completed their Ph.D. within CLS after 3.5 to 4 years (sometimes having transferred to CLS after starting their doctorate). 14 postdoctoral fellows have completed their research while being supported by CLS. Currently, 20% of all CLS doctoral and postdoctoral fellows are female, and CLS aims to increase gender equality by establishing dedicated support for female researchers.

The structured Ph.D. education and joint recruiting of CLS have especially helped the junior faculty at both institutions get access to strong students. The increased student quality in turn has helped us retain and build attractiveness also when recruiting top group leaders and faculty. This recruitment cycle helps us especially in cases that are particularly challenging, such as international recruitment and recruitment of female top candidates.

| **Description** | **2016** | **2017** | **2018** | **2019** |
|---|---|---|---|---|
| Number of Applications | 209 (100%) | 228 (100%) | 252 (100%) | 386 (100%) |
| Invited to Panel Interview | 15 (7.2%) | 20 (8.8%) | 20 (7.9%) | 22 (5.9%) |
| Offer after Panel Interview | 6 (2.9%) | 6 (2.6%) | 10 (4.0%) | – |
| Accepted Offer | 4 (1.9%) | 5 (2.2%) | 8 (3.2%) | – |

Table 1.1: Summary statistics of all CLS Ph.D. application calls and procedures. The 2019 recruitment round is underway, with up to about 10 offers to be made.

### 1.9.5 Impact

The creation of the Max Planck ETH Center for Learning Systems in 2015 constituted a key development for the visibility of the Zürich region and the Stuttgart/Tübingen region as two major AI hubs in Europe. After the establishment of the Center for Learning Systems, the foundation of the Cyber Valley initiative in Stuttgart/Tübingen as well as the International Max Planck Research School for Intelligent Systems further strengthened this regional AI ecosystem. In bringing together two major European research institutions and regions, today, CLS is in the ideal position to form the backbone for a European research network on AI and robotics.

134 papers have been published since 2015 through the direct support of the CLS program. In particular, more than 20 joint papers co-authored by MPI-IS and ETH researchers were published in areas not covered by either institution independently. Our doctoral and postdoctoral researchers within the program state their CLS affiliation in publications and at conferences. Moreover, fellows also recognize CLS as a funding body in the acknowledgments section of their publications. Many more papers have been published that were indirectly supported through CLS. The center has thus allowed a large number of young scientists to be part of two institutions and to benefit from cross-disciplinary collaborations, attend relevant workshops, and build their own strong network on the foundation of this cross-institutional academic program.

### 1.9.6 Further Development

CLS combines Europe's leading computer science department, several other outstanding academic departments, and the strongest basic research institute in modern AI. We have created a joint Ph.D. program that is internationally competitive, providing students with the chance to work with some of the most influential AI faculty worldwide. In the next phase of CLS, we want to jointly build a lighthouse for machine learning (ML) and modern artificial intelligence (AI) in Europe. To this end, we will

1. deepen the existing links,
2. reach out to a select group of outstanding researchers in ML and AI throughout Europe,
3. build partnerships worldwide
4. strengthen the possibilities for industry co-sponsorship while retaining full academic independence, and
5. seek synergies with the Cyber Valley initiative, which was launched since the inception of CLS.

The above vision is ambitious and would be unrealistic without the foundations that we have laid during the first CLS funding period. We believe we stand a good chance of **building the strongest European center for modern AI research**. As the economic and societal influence of AI increases, so will likely the investments, and CLS will influence Europe's further development in this field by proving that it is possible to be internationally competitive despite significantly larger investments in China and the US.

## 1.10 Collaborations

In addition to the major collaborations described in the three previous sections (**Cyber Valley**, the **International Max Planck Research School for Intelligent Systems**, and the **Max Planck ETH Center for Learning Systems**), the institute has established several additional major collaborations worldwide, which we highlight in this section.

**Cambridge Tübingen Ph.D. program in Machine Learning** The MPI-IS Department of Empirical Inference, headed by Bernhard Schölkopf, and the Machine Learning Group of the University of Cambridge are among the world-leading centers of research in machine learning. Launched in 2014, the Cambridge Tübingen Ph.D. fellowship in Machine Learning is a collaborative program in which scientists at both institutions jointly supervise a small group of elite doctoral students. Each year, about two to three Ph.D. fellows are jointly selected via a symposium, and 10 Ph.D. students are currently enrolled in the program. Students spend at least one year at each institution, and they benefit from the excellent research environment that both institutions provide. Successful graduates will receive their doctoral degrees from the University of Cambridge.

**Carnegie Mellon University / MPI-IS Collaborative Ph.D. program** The Max Planck Institute for Intelligent Systems and Carnegie Mellon University (CMU) in Pittsburgh, USA, have initiated a joint Ph.D. program, aiming to strengthen the ties between these two leading research institutions in the field of intelligent systems. Starting with the signing of the cooperation contract in spring of 2018, some of the world's most promising Ph.D. students interested in robotics now have the opportunity to benefit from the strong and unique research environments provided by both CMU and MPI-IS.

Students who get accepted to participate in this joint Ph.D. program will receive top education and training in robotics at CMU, which is one of the world's leading universities in this field. The students begin their research with a professor at CMU for 1.5 to 2 years and continue their research under the supervision of a director or group leader at the MPI-IS for at least another 2 years. Successful graduates will eventually receive their Ph.D. degree from CMU.

The program was kicked off in a ceremonial meeting on December 15, 2017, in Pittsburgh, USA. Roughly 2 to 3 new Ph.D. students are estimated to be admitted to this program each year over the next 5 years. The CMU/MPI-IS collaborative Ph.D. program was initiated by Metin Sitti, director of the Physical Intelligence Department. Before joining our institute, he spent twelve years as a professor at the Department of Mechanical Engineering and the Robotics Institute at Carnegie Mellon University.

**National Competence Center on AI and Machine Learning** Since October 2018, the Max Planck Institute for Intelligent Systems and the University of Tübingen have jointly hosted a competence center for artificial intelligence and machine learning. This center is one of four locations where Germany's Ministry of Education and Research is pooling scientific projects in artificial intelligence. The Tübingen AI Center will provide research groups at the university and at MPI-IS with a place to develop learning systems. The German Federal Ministry of Education and Research will sponsor the center with approximately 6.6 million euros for an initial runtime of four years. The center is expected to play a key role in the German government's artificial intelligence development strategy.

Researchers at the center mainly work on new concepts and principles to create more robust learning systems. Learning algorithms must be able to deal with external and unexpected influences. At the same time, their reactions must be predictable and transparent. The center will also raise the question of and investigate possible misuses of AI technologies. One junior research group will look at ways of protecting sensitive data and finding better solutions. Part of the work at the Tübingen AI center will involve benchmark competitions aimed at defining important scientific problems. Research groups will compete to find the best solutions. This new center is another part of the major Cyber Valley initiative in the Stuttgart-Tübingen region.

**Clusters of Excellence at the Universities of Stuttgart and Tübingen** The German national Excellence Strategy recently decided to fund a large number of Clusters of Excellence. The 57 winners were chosen from 88 full proposals and 195 original applications. The Max Planck Institute for Intelligent Systems was involved in four applications, all of which were selected for funding as Clusters of Excellence. In Stuttgart, MPI-IS scientists were directly involved in both of the successful proposals at the University of Stuttgart. In Tübingen, MPI-IS is involved in two of the University of Tübingen's three approved clusters. This quadruple success demonstrates the strength of the Stuttgart/Tübingen region in artificial intelligence and robotics, as well as the synergy between the MPI for Intelligent Systems and its two neighboring universities. Somewhat based on the Max Planck Society's strategic investment in this topic in this region, this success will further strengthen the institute's cooperation with the University of Stuttgart and the University of Tübingen. Just seven years after its founding, the MPI for Intelligent Systems is already helping fuel the research excellence of the region.

The Stuttgart Cluster of Excellence, "Integrative Computational Design and Construction for Architecture" (IntCDC) highlights strong research areas at the University of Stuttgart that are complementary to the expertise of MPI-IS researchers and have the potential to benefit humanity. Importantly, the University of Stuttgart recently realigned its research profile under the vision 'Intelligent Systems for a Sustainable Society', thereby creating intensive synergies with research at both sites of the Max Planck Institute for Intelligent Systems. Institute scientists are contributing deep expertise in robotics to the IntCDC cluster, focusing on multi-robot systems and haptic human-robot interaction. There are also close links between MPI-IS and the second successful cluster at the University of Stuttgart, "Data-integrated Simulation Sciences" (SimTech).

The Tübingen Cluster of Excellence "Machine Learning in Science" is particularly closely linked to the MPI-IS. The cluster's spokesperson, Professor Ulrike von Luxburg, is a Professor of Computer Science at the University of Tübingen and a Max Planck Fellow at the Tübingen site of the MPI for Intelligent Systems. Von Luxburg was the first doctoral student in the Department of Empirical Inference, with which Bernhard Schölkopf established the research field of machine learning in Tübingen in 2001. The cluster's second spokesperson, Professor Philipp Berens, is also an alumnus of Schölkopf's department and has close ties to the MPI-IS. The institute's participation in the other Tübingen cluster "Image-Guided and Functionally Instructed Tumor Therapies", is based on a long-term cooperation in the application of machine learning methods in medical imaging.

Up until 2018, the Max Planck Institute for Intelligent Systems was already involved in the previous version of the SimTech Cluster of Excellence at the University of Stuttgart, as well as the now-closed Cluster of Excellence on integrative neuroscience at the University of Tübingen.

**Max Planck Fellows and Joint Groups** The Max Planck Fellow program of the Max Planck Society aims to promote cooperation between outstanding university professors and Max Planck researchers. By strengthening ties with universities, the program fosters cutting-edge research on both sides. During their initial five-year term, which can be extended, Max Planck Fellows typically set up a small group at the host institute, funded by an annual budget of 100,000 euros that is provided jointly by the hosting institute and the Max Planck Society. The institute currently hosts one Max Planck Fellow group on each campus, headed by Professor Marc Toussaint (University of Stuttgart) and Professor Ulrike von Luxburg (University of Tübingen).

**Tübingen Research Campus** The University of Tübingen and the other research institutes in Tübingen have joined forces to create the Tübingen Research Campus (TRC). Together, these partners work on common research focus areas, keeping science made in Tübingen at the forefront of academic endeavor. They cooperate on specific research projects, through shared facilities, in nurturing emerging academics, and in welcoming scientists from all over the world to Tübingen.

**Max Planck Fraunhofer cooperation** Within the framework of the national Pact for Research and Innovation, the Max Planck Society and the Fraunhofer Society intend to continue and intensify their cooperation across research areas and disciplines. With Fraunhofer's focus on application, this collaboration is of particular interest to the Max Planck Society. The aim of such a venture is to bring to application the knowledge resulting from collaborative efforts, thereby making a direct contribution to the development and impact of new technologies. The institute is currently participating in two projects under the Max Planck Fraunhofer cooperation scheme:

**Theory and practice for reduced learning machines** Cooperation between the Empirical Inference department of our institute and the Fraunhofer Institute for Telecommunications - Heinrich Hertz Institute in Berlin.

**Acoustic holograms** Cooperation between the Micro, Nano, and Molecular Systems group and the Empirical Inference department of our institute and the Fraunhofer Institute for Biomedical Engineering in St. Ingbert and Sulzbach.

**Industry Collaborations** Many companies are partnering with and supporting research at the MPI for Intelligent Systems. The Cyber Valley core partners Amazon, BMW, Bosch, IAV, Porsche, and ZF Friedrichshafen are supporting the Cyber Valley cooperation with at least 7.5 million euros in the period from 2018 to 2022. The majority of this funding is made available to the Cyber Valley research groups via a competitive procedure within the framework of a research fund. Amazon also supports the Max Planck Institute for Intelligent Systems with Amazon Research Awards (ARA) worth 420,000 euros per year. These research funds are used to finance research by doctoral and postdoctoral students at our institute. Intel founded the "Network on Intelligent Systems (NIS)" in 2017 with MPI-IS as a key partner. NIS supports leading academic research and collaboration in robotics, computer vision, motor control, and machine learning. NVIDIA supports AI research at the MPI-IS with financial contributions and state-of-the-art hardware through an NVIDIA AI Lab (NVAIL) grant. In addition, researchers at MPI-IS collaborate with researchers from other international companies such as Google, Microsoft, Facebook, and others. In all cases, cooperation with an industrial partner is subject to clear rules to protect research freedom and intellectual property rights.

**Collaboration within the Max Planck Society** The Max Planck Institute for Intelligent Systems has close collaborations, both structural and scientific, with several other Max Planck Institutes, including the following highlights:

**MPI for Solid State Research, Stuttgart** In addition to a joint administration and a joint building, we share a large set of scientific-technical services with the MPI for Solid State Research, including workshops, laboratories, and a library.

**MPI for Biological Cybernetics, Tübingen** In addition to close collaboration on campus matters, our institute has traditionally also shared close scientific links with the MPI for Biological Cybernetics, where Bernhard Schölkopf began his career as a Max Planck director. With its recent scientific re-orientation and the appointment of Peter Dayan as a director, the MPI for Biological Cybernetics will most likely build even stronger scientific links with our institute in the coming years.

**MPI for Developmental Biology, Tübingen** In addition to close collaboration on campus matters, scientific collaborations exist in the field of genomics and bioinformatics, in particular with the departments of Ruth Ley and Detlef Weigel.

**CS@max planck** The Max Planck Institutes working in the broad area of computer science joined forces in 2018 by founding the Max Planck Graduate Center for Computer and Information Science. CS@max planck is a highly selective multi-site doctoral program that grants admitted students full financial support to pursue doctoral research in the broad area of computer and information science; its faculty come from the participating Max Planck Institutes and some of the best German universities. The other two participating

Max Planck Institutes are the MPI for Software Systems in Kaiserslautern and Saarbrücken and the MPI for Informatics in Saarbrücken. The first round of CS@max planck doctoral recruiting is underway, with ten admitted students scheduled to meet possible doctoral advisors and host labs in the near future.

**BiGmax** This initiative of the Max Planck Society and MaxNet centers on Big-Data-Driven Materials Science (abbreviated BiGmax). Eleven institutes of the Max Planck Society, among them the MPI for Intelligent Systems, combine their know-how in data-driven material science. The aim is a better use of the possibilities associated with analyzing large amounts of data. The first funding started in 2017 and will run for five years.

**Professorships** The close collaboration between researchers at the Max Planck Institute for Intelligent Systems and universities worldwide is also reflected in professorships held by the current directors and independent research group leaders of the institute.

**Michael J. Black** (director) is an Honorary Professor at the University of Tübingen in the department of Computer Science. He is also an Adjunct Professor (Research) at Brown University (Providence, USA) in the department of Computer Science. From April 2014 to April 2016 he was a Visiting Professor at ETH Zürich in the Electrical Engineering department.

**Peer Fischer** (permanent group leader) is a Professor in the Faculty of Chemistry at the University of Stuttgart.

**Andreas Geiger** (group leader) is a Professor of Computer Science at the University of Tübingen. He was a Visiting Professor at ETH Zürich from 2016 to 2018.

**Katherine J. Kuchenbecker** (director) is an Adjunct Professor at the University of Pennsylvania (Philadelphia, USA) in the department of Mechanical Engineering and Applied Mechanics.

**Ludovic Righetti** (group leader) is an Associate Professor in the Electrical and Computer Engineering Department and in the Mechanical and Aerospace Engineering Department at the Tandon School of Engineering of New York University (New York, USA).

**Bernhard Schölkopf** (director) is an Honorary Professor at the Technical University Berlin in the department of Computer Science. He is also an Honorary Professor at the University of Tübingen in the department of Mathematics and Physics. From 2010 to 2012 he held an extraordinary professorship at Stellenbosch University (Stellenbosch, South Africa) in the department of Mathematical Science. Since January 1, 2019, he is an Affiliated Professor at ETH Zürich in the Computer Science department.

**Metin Sitti** (director) has been an honorary professor within SimTech at the University of Stuttgart since 2017. He is an Adjunct Professor at Carnegie Mellon University (Pittsburgh, USA) in the department of Mechanical Engineering. Since October 2018, he has also been a Professor in both the Medical School and the Engineering School at Koç University (Istanbul, Turkey).

# 2 EMPIRICAL INFERENCE

## 2.1 Research Overview

The problems studied in the department can be subsumed under the heading of *empirical inference*, i.e., inference performed on the basis of empirical data. This includes statistical learning, but also the inference of causal structures from statistical data, leading to models that provide insight into the underlying mechanisms, and make predictions about the effect of interventions. Likewise, the type of empirical data can vary, ranging from biological measurements (e.g., in neuroscience) to astronomical observations. We are conducting theoretical, algorithmic, and experimental studies to try and understand the problem of empirical inference.

The department was started around statistical learning theory and kernel methods. It has since broadened its set of inference tools to include a stronger component of Bayesian methods, including graphical models with a strong focus on issues of causality. In terms of the inference tasks being studied, we have moved towards tasks that go beyond the relatively well-studied problem of supervised learning, such as semi-supervised learning or transfer learning. Finally, we have continuously striven to analyze challenging datasets from biology, astronomy, and other domains, leading to the inclusion of several application areas in our portfolio.

The most competitive publication venues in empirical inference are NeurIPS (Neural Information Processing Systems), ICML (International Conference on Machine Learning), UAI (Uncertainty in Artificial Intelligence), and for theoretical work, COLT (Conference on Learning Theory). The presence at these conferences makes us one of the top international machine learning labs. In addition, we sometimes submit our work to the leading application oriented conferences in neighboring fields including computer vision (ICCV, ECCV, CVPR) and data mining (KDD, ICDM, SDM), as well as to specialized journals.

Our work has earned us a number of awards, including best paper prizes at the major conferences in the field (NeurIPS, ICML, UAI, COLT, ALT, CVPR, ECCV, ISMB, IROS, KDD). Recent awards include IEEE SMC 2016, ECML-PKDD 2016, a honorable mention at ICML 2017, and the test-of-time award for Olivier Bousquet at NeurIPS 2018, received for work he started while he was still member of our department (at MLSS 2003 in Tübingen with Leon Bottou).

Theoretical studies, algorithms, and applications often go hand in hand. For instance, it may be the case that someone working on a specific application will develop a customized algorithm that turns out to be of independent theoretical interest. Such serendipity is a desired side effect caused by interaction across groups and research areas, for instance during our frequent departmental talks. It concerns cross-fertilization of methodology (e.g., kernel independence measures used in causal inference), the transfer of algorithmic developments or theoretical insights to application domains (e.g., causal inference in neuroscience), or the combination of different application areas (e.g., using methods of computational photography for magnetic resonance imaging).

The linear organization of the text does not permit an adequate representation of all these connections. Below, we have opted for an organization of the material that devotes individual sections to our main application areas (computational imaging, robot learning, and neuroscience), and that comprises four methodological sections, on learning algorithms, causal inference, probabilistic inference, and statistical learning theory. We begin with the latter.

**Statistical Learning Theory** A machine learning algorithm is given training data and tries to learn a model that is well-suited to describe the data and that can be used to make predictions. The goal of statistical learning theory is to assess to which extent such algorithms can be successful in principle. The general approach is to assume that the training data have been generated by an unknown random source, and to develop mathematical tools to analyze the performance of a learning algorithm in statistical terms: for example, by bounding prediction errors ("generalization bounds") or by analyzing large sample behavior and convergence of algorithms on random input ("consistency").

The department has made various contributions to this area, especially in areas of machine learning where statistical learning theory is less well developed. These include settings like active and transfer learning, privacy preserving machine learning as well as unsupervised generative modeling. Our goal is to contribute statistical foundations to these exciting new challenges where pioneering work can be done.

Active learning exploits structure and information in unlabeled data to reduce label supervision by requesting labels only for a small set of points from a large pool. Developing a novel active query procedure that takes in an unlabeled data and constructs a compressed version of the underlying labeled sample, we showed that active learning can provide label savings even in non-parametric learning settings. A formal analysis of compressing a data sample so as to encode a set of functions consistent with (or of minimal error on) the data was then conducted in [295].

In recent years the kernel mean embedding (KME) of distributions started to play an important role in various machine learning tasks, including independence testing, density estimation, and many more. Inspired by the James-Stein estimator, in [164] we introduced a new type of KME estimators called kernel mean shrinkage estimators (KMSEs) and showed that it can converge faster than the empirical KME estimator. We have studied the optimality of KME estimators in the minimax sense in [128] and shown that the convergence rate for the KME, and many other methods published in the literature, is optimal and can not be improved. The advances and characterizations for kernel mean embeddings also play a role for privacy preserving machine learning.

Recently significant progress has been made in the field of *unsupervised (deep) generative modeling* with generative adversarial networks (GANs), variational autoencoders (VAEs) and other deep neural network based architectures, significantly improving the state of the art in the quality of samples, especially in the domain of natural images. Traditionally the training objectives in VAEs and GANs have been based on f-divergences. We showed, starting from Kantorovich's primal formulation of the optimal transport problem, that this can be equivalently written in terms of probabilistic encoders, which are constrained to match the latent posterior and prior distributions.[1] This leads to a new training procedure of latent variable models, so called Wasserstain auto-eoncoders (WAEs) as described in [213]. While WAEs share many of the nice properties of VAEs, the generated samples often exhibit better quality and leads to interesting properties of the learned latent representations as described in [178] and [177]. Another theoretical study of generative modeling led us to propose the *AdaGAN*, a boosting approach to greedily build mixtures of generative models (e.g., GANs or VAEs) by solving, at each step, an optimization problem that results in the best additional model to reduce the discrepancy between the current mixture model and the target [261].

**Learning Algorithms** Learning algorithms based on kernel methods have enjoyed considerable success in a wide range of supervised learning tasks such as regression and classification. One reason for the popularity of these approaches is that they solve difficult nonparametric problems by mapping data points into high dimensional spaces of features, specifically reproducing kernel Hilbert spaces (RKHSes), in which linear algorithms can be brought to bear, leading to solutions taking the form of kernel expansions [125].

Based on foundational work from our department, kernel methods underlie methods determining the goodness of fit of a model and more recently of differentially private learning. In [190], we address the problem of measuring the relative goodness of fit of two models using kernel mean embeddings. Given two candidate models, and a set of target observations, the goal is to produce a set of so-called informative features, which indicate the regions in the data domain where one model fits better than the other. The task is formulated as a statistical test whose runtime complexity is linear in the sample size. Privacy-preserving machine learning algorithms aim to come up with database release mechanisms that allow third-parties to construct consistent estimators of population statistics while ensuring that the privacy of each individual contributing to the database is protected. In [219], we develop privacy-preserving algorithms based on the kernel mean embedding, allowing us to release a database while guaranteeing the privacy of the database records.

Optimization lies at the heart of most machine learning algorithms and our department aims to understand the convergence property of coordinate descent as well as Frank-Wolfe optimization algorithms under different sampling schemes and constraints. In [252], we provide a theoretical understanding of greedy coordinate descent for smooth functions. Similarly, in [236] we propose an adaptive recursive sampling scheme based on the min-max optimal solution of the variance reduction problem to achieve faster convergence for coordinate descent, which can also be applied to stochastic gradient descent. A connection between matching pursuit and Frank-Wolfe optimization is explored in [208]

**Causal inference** The detection and use of statistical dependences form the core of statistics and machine learning. In recent years, machine learning methods have enabled us to perform rather accurate prediction, often based on large training sets, for complex nonlinear problems that not long ago would have appeared completely random. However, in many situations we would actually prefer a causal model to a purely predictive one; i.e., a model that might tell us that a specific variable (say, whether or not a person smokes) is not just statistically associated with a disease, but causal for the disease.

Pearl's graphical approach to causal modeling generalizes Reichenbach's common cause principle and characterizes the observable statistical (conditional) independences that a causal structure should entail. Many causal inference methods build on these independences to infer causal graphs from data. This "graphical models" approach to causal inference has several weaknesses that we try to address in our work: it only can infer causal graphs up to Markov equivalence, it does not address the hardness of conditional independence testing, and it usually does not worry about the complexity of the underlying functional regularities that generate statistical dependences in the first place. Our work in this field is characterized by the following three

[1] O. Bousquet, S. Gelly, I. Tolstikhin, C. J. Simon-Gabriel, B. Schölkopf. From Optimal Transport to Generative Modeling: the VEGAN cookbook. *CoRR* **abs/1705.07642**, 2017.

aspects:

1. We often work in terms of structural equation models (SEMs) or functional causal models (FCMs), i.e., we do not take statistical dependences as primary, but rather study mechanistic models which give rise to such dependences. In FCMs, each variable is modeled as a deterministic function of its direct causes and some noise variable $N$, e.g., $Y = f(X, Z, N)$; all noise variables are assumed to be jointly independent. FCMs do not only allow us to model observational distributions; one can also use them in order to model what happens under interventions (e.g., gene knockouts or randomized studies).

Under suitable model assumptions like additive independent noise, causal knowledge and the framework of SCMs admits novel techniques of noise removal via so-called half-sibling regression [153, 166], or revealed previously unknown aspects of the arrow of time [151, 292].

2. Viewed from an FCM perspective, the crucial assumption of the graphical approach to causality is statistical independence of all noise terms. Intuitively, it is clear that as the noises propagate through the graph, they pick up dependences due to the graph structure, hence the assumption of initial independence of the noise terms allows us to tease out properties of that structure. We believe, however, that much can be gained by considering a more *general independence assumption* related to notions of invariance and autonomy of causal mechanisms. Here, the idea is that causal mechanism are autonomous entities of the world that (in the generic case) do not depend on or influence each other, and changing (or intervening on) one of them often leaves the remaining ones invariant.

In the context of the classical pattern recognition task with handwritten digits, [203] shows that learning causal models that contain independent mechanisms helps in transferring information across substantially different data sets. Theoretical work in [216] shows that the independence of causal mechanisms can be formalized via group symmetry.

3. This leads to the third characteristic aspect of our work on causality. Wherever possible, we attempt to establish connections to machine learning, and indeed we believe that some of the hardest problems of machine learning (such as those of domain adaptation and transfer) are best addressed using causal thinking. To this end, one may assume, for instance, that structural equations remain constant across data sets and only the noise distributions change [68], or that some of the causal conditionals in a causal Bayesian network change, while others remain constant [96] or that they change independently [245], which results in new approaches to domain adaptation [300].

Our lab has played a major role in putting causal inference on the agenda of the machine learning community, including a recent award-winning textbook [68], and we expect that causal inference will have practical implications for many inference problems (e.g., in astronomy [153] and neuroscience [162]), as well as increasingly in machine learning methods, including deep learning [203] and reinforcement learning [191]. We also expect that it will play a major role in societal aspects of AI, including fairness [237] and interpretability/accountability. Causality touches statistics, econometrics, and philosophy, and it constitutes one of the most exciting field for conceptual basic research in machine learning today. We expect that going forward, causality will play a major role in taking *representation learning* to the next level, moving beyond the mere representation of statistical dependence structures towards models that support intervention and planning (and thus Konrad Lorenz' notion of *thinking* as *acting in an imagined space*).

**Probabilistic Inference** The probabilistic formulation of inference is one of the main research streams within machine learning. One of our main themes in this field has been nonparametric inference on function spaces using Gaussian process models [187, 283, 284]. The approaches developed at the department allow finding the best kernel bandwidth hyperparameter efficiently and are especially well-suited for online learning.

A crucial bottleneck in Bayesian models is the marginalization of latent variables. This can be computationaly demanding, so approximate inference routines reducing computational complexity are a major research theme. In [188, 215], we study the convergence properties of this approach from a modern optimization viewpoint, establishing connections to the classic Frank-Wolfe algorithm. The analyses yield novel theoretical insights regarding the sufficient conditions for convergence, explicit rates, and algorithmic simplifications.

Members of the department have also been working on aspects of probabilistic programming and studied the problem of representing the distribution of $f(X)$ for a random variable $X$ and a function $f$ [288]. We use kernel mean embedding methods to construct consistent estimators of the mean embedding of $f(X)$. The method is applicable to arbitrary data types on which suitable kernels can be defined. It thus allows us to generalize (to the probabilistic case) functions in computer programming which are originally only defined on deterministic data types.

As algorithmic decision making systems are becoming ubiquitously trained from historical data collected from, as well as implemented in a wide variety of online as well as offline services; there is a growing concern that these automated decisions can lead to a lack of fairness, i.e., their outcomes can disproportionately hurt (or, benefit) particular groups of people sharing one or more sensitive attributes (e.g., race, gender). As a consequence, there has been an increase in research on computational (un)fairness. The contributions of the department on in this domain are twofold. First, we have focused on proposing new definitions and metrics of fairness, as well as on designing automatic decision systems that incorporate a fairness definition in their training step to avoid discrimination towards particular groups of people sharing certain sensitive attributes, while providing clear mechanisms to trade-off fairness and accuracy [79, 223]. Second, we have done pioneering work in connecting fairness to causality, showing that whether an algorithmic decision is fair or not should really take into account the underlying causal graph rather than just the observational distribution [237].

**Computational Imaging** Handheld video cameras are now commonplace and available in every smartphone, and thus images and videos are recorded in unprecented amounts. Our research focus is on digital image restoration that aims at computationally enhancing the quality of images and recovering the most likely original image by undoing the adverse effects of image degradation such as noise and blur.

To recover a high-resolution image from a single low-resolution input, we proposed [263] a novel method for automated texture synthesis in combination with a perceptual loss focusing on creating realistic textures rather than optimizing for a pixel-accurate reproduction of ground truth images during training. By using feed-forward fully convolutional neural networks in an adversarial training setting, we achieve a significant boost in image quality even at high magnification ratios.

For recovering an image from corrupted measurements, e.g., due to unwanted camera shake during recording, or moving objects in the scene, we developed methods for propogating information between multiple consecutive blurry observations to help restore the desired sharp image or video. In a number of works [198, 247, 248], we have developed efficient recurrent network architectures to deblur frames taking temporal information into account, which can efficiently handle both ego and object motion for arbitrary spatial and temporal input sizes.

**Robot Learning** Research in robotics and artificial intelligence has lead to the development of complex robots ranging from anthropomorphic arms to complete humanoids. In order to be meaningfully applied in human-inhabited environments, robots need to possess a variety of physical abilities and skills. However, programming such skills is a labor- and time-intensive task which often leads to brittle solutions and requires a large amount of expert knowledge. In particular, it often involves transforming intuitive concepts of motions and actions into formal mathematical descriptions and algorithms.

While kinematic optimization allows for efficient representation and online generation of hitting trajectories, e.g., in robot table tennis, learning to track such dynamic movements with inaccurate models remains an open problem. To achieve accurate tracking for such tasks in a stable and efficient way, we have proposed a series of novel adaptive Iterative Learning Control (ILC) algorithms that are implemented efficiently and enable caution during learning [72]. Moreover, we have built a muscular robot system in order to study the problem of how to accurately control musculoskeletal robots by learning control. Muscular systems are hard to control by classical methods, but offer beneficial properties to achieve human-comparable performance in complex tasks [56].

In real robot experiments on a Barrett WAM, we have studied the properties of optimal trajectory generation in robot table tennis strikes [100], and how to robustly learn such primitives from multiple demonstrations as well as adapt them to new goals [71]. We have also recently demonstrated how a table tennis serve can be captured and successfully be reproduced [49, 70].

Robot table tennis has a number of components that are representative of tasks encountered by natural intelligent systems, including perception and action, as well as various aspects of social interaction (opponent modeling, competition, collaboration). We have recently completed the move of our robotics laboratory to the new building, now operating a two-robot setup. In the long term, this will enable us to study a rich set of problems, including cooperative game play.

**Machine Learning in Neuroscience** The neurosciences present some of the steepest challenges to machine learning. Nearly always there is a very high-dimensional input structure — particularly relative to the number of exemplars, since each data point is usually gathered at a high cost. To avoid overfitting, inference must thus make considerable use of domain knowledge. Relevant regularities are often subtle, the rest being made up of noise that may be of much larger magnitude (often composed largely of the manifestations of other neurophysiological processes, besides the ones of interest). In finding generalizable solutions, one usually has to contend with a high degree of variability, both between individuals and across time, leading to problems of covariate shift and non-stationarity.

One specific neuroscientific application area in which we have a long-standing interest is that of brain-computer interfacing, or BCI (see page 39). BCIs hold promise in restoring communication for completely locked-in stage (CLIS) patients in late stages of amyotrophic lateral sclerosis (ALS). Despite more than two decades of research, however, late-stage ALS patients remain incapable of operating BCIs, arguably because such systems currently rely on brain processes that are impaired as a result of disease progression. In a series of studies, we have investigated how ALS affects neural- and cognitive processes [85, 126].

Building upon these novel insights, we have developed and validated a new type of cognitive BCI for late-stage ALS patients [149, 308]. To translate this system from laboratory- into home-use, we have pioneered a transfer learning approach for BCIs [127, 159]. These advances now enable us to build high-performance, cognitive BCIs with off-the-shelf hardware, thereby (in ongoing work) rendering BCI systems available for large-scale application outside of laboratory environments.

We have also designed machine learning techniques to assist with the interpretation of experimental brain data (see page 49). Unsupervised learning tools based on non-negative Matrix Factorization were designed to automatically identify activity patterns among large populations of recorded neurons [93]. Finding causal relationships between neural processes is also of particular interest to neuroscientists, but hard to address in living neural systems due to ethical and practical concerns. Combining machine learning with detailed computational biophysical network models provided theoretical insights into biological learning by identifying key neural circuits underlying the reliable replay of memorized events [95].

We conclude this section with a short summary of our contributions to cognitive science and vision research, performed in collaboration with Felix Wichmann, professor of computational neuroscience at the local University Tübingen and until recently also part-time member of our department. In one line of work, we developed methods to gain more information from psychophysical data, linking traditional methods with machine learning approaches [145, 303]. We investigate reliable supra-threshold psychophysical paradigms which are more intuitive and require less training and are thus more likely to yield reliable crowd-sourcing data [120, 129]. A second focus was the development of a predictive image-based model of spatial vision. We integrated the large psychophysical literature on simple detection and discrimination experiments and proposed a model based on maximum-likelihood decoding of a population of model neurons predicting the most important spatial vision data sets simultaneously, using a single set of parameters [122]. Third, we have explored similarities and differences of DNNs and the human visual system [121, 173].

## 2.2 Selected Research Projects

## Astronomy

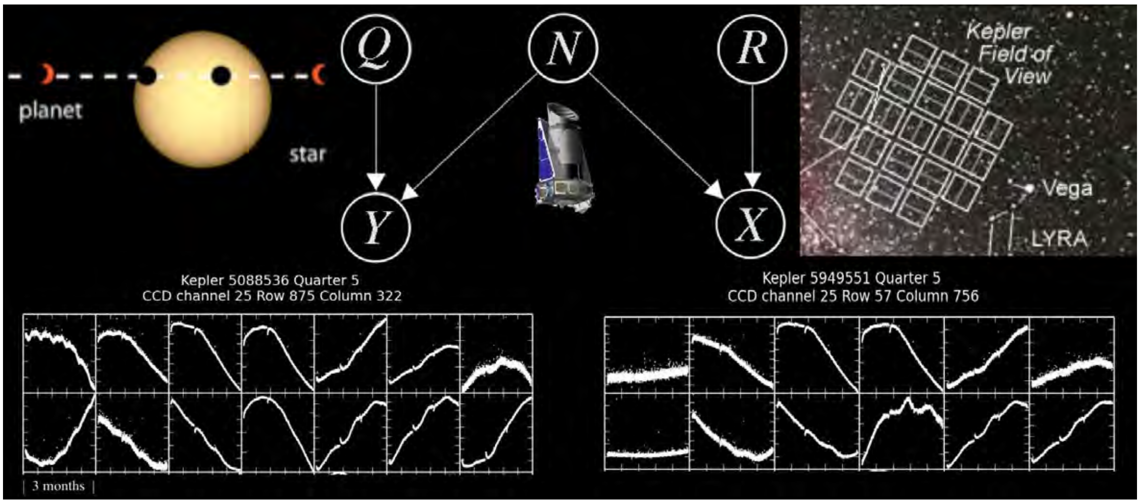


Figure 2.1: We observe a set of pixel light curves $Y$ (all belonging to the same star). They are affected by astronomical signal $Q$ and instrument $N$. Recovering $Q$ from $Y$ is impossible since we do not know $N$. However, at the same time we are measuring other stars $X$, which have no causal link to $Q$, using the same instrument (making them half-siblings). Subject to suitable assumptions, regression of $Y$ on $X$ allows recovery of $Q$ almost surely.

Satellite missions and ground-based observations generate datasets of unprecedented quality and size. New fields such as exoplanet science and gravitational wave physics have emerged, impossible without large-scale data analysis.

**Exoplanet search** The Kepler space telescope, launched in 2009 and retired in 2018, monitored the brightness of 150,000 stars to find exoplanets causing decreases of brightness by temporary occlusions. The measurements are corrupted by systematic noise due to the telescope. However, since the stars can be assumed to be causally independent of each other (being light years apart) as well as of the instrument noise, we can denoise the signal of a single star by removing all information that can be explained by the measurements of the other stars. Subject to an additivity assumption, we provide theoretical guarantees regarding the quality of the reconstruction of our *half-sibling regression* method [166]. We used half-sibling regression to develop a practical method to denoise pixel light curves [153] and an exoplanet search pipeline which discovered 21 subsequently confirmed exoplanets.

More recently, we have focused on detecting single exoplanet transits using supervised learning. Detecting transits is hard already when periodically reoccurring transit events are available, and harder still when only single events are observed. In this setting, we need to resort to a strong data-driven model based on a large set of representative transits [150]. This is of current interest as NASA's TESS mission is starting to release data, and many of the most interesting planets will only have a single transit.

**Gravitational wave detection** The detection of gravitational waves from a binary black hole merger in 2015 was a milestone in modern physics. However, despite the unparalleled sensitivity of the LIGO detectors, data analysis remains a challenge. We have developed a dilated, fully convolutional neural net to be applied directly on the time series strain data to identify simulated GW signals from black hole mergers in real, non-Gaussian background measurements from the LIGO detectors. The system efficiently runs on strain data of arbitrary length from any number of detectors in real time [235]. It has the potential to develop into a complementary trigger generator in the existing LIGO search pipeline. To explore this, several department members are currently members of the LSC (Ligo Scientific Collaboration).

More information: https://ei.is.mpg.de/project/astronomy

## Brain-Computer Interfaces

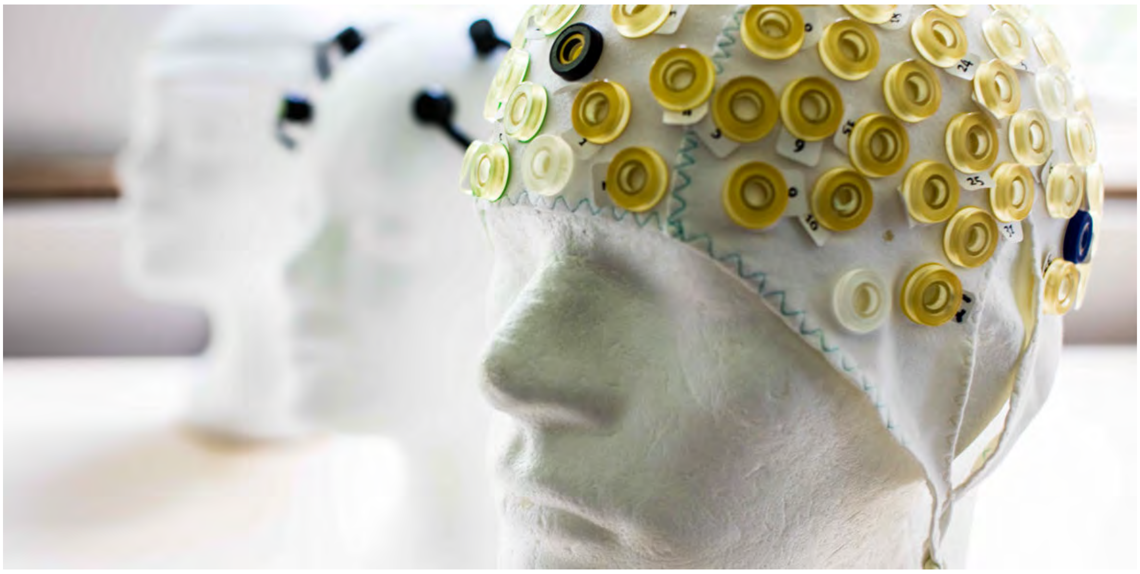

Brain-computer interfaces (BCIs) hold promise in restoring communication for completely locked-in stage (CLIS) patients in late stages of amyotrophic lateral sclerosis (ALS). Despite more than two decades of research, however, late-stage ALS patients remain incapable of operating BCIs, arguably because such systems currently rely on brain processes that are impaired as a result of disease progression. To establish communication with CLIS-ALS patients, it is thus crucial to, first, understand how ALS affects neurophysiological as well as cognitive processes and, second, develop BCI systems that target brain processes which remain functional into late stages of the disease.

In a series of studies, we have investigated how ALS affects neural and cognitive processes. In particular, we could show that the neural correlates of self-referential processing are already impaired in the early stages of ALS, i.e., neural deficits emerge before cognitive deficits manifest in overt behavior [126]. Following up on this work, we recorded electrophysiological signals in two CLIS-ALS patients and demonstrated that the locked-in stage is accompanied by a major slow-down of the brain's dominant rhythm [85]. The cognitive implications of this slow-down are yet to be understood.

Building upon these novel insights, we have developed and validated a new type of cognitive BCI for late-stage ALS patients. This system takes into account neurophysiological changes and targets brain processes that are likely to remain functional into late disease stages [308]. We successfully validated this system in a long-term study with two ALS patients, who achieved stable communication for more than one year [149].

To translate this system from laboratory- into home-use, we have pioneered a transfer learning approach for BCIs that drastically reduces the amount of training data while maintaining decoding accuracy [159] and developed a novel brain-decoding feature, based on task-induced frequency modulations of canonical brain rhythms, that is particularly useful for low-channel setups [127]. These advances should enable us to build high-performance cognitive BCIs with off-the-shelf hardware, eventually leading to BCI systems available for large-scale application outside of laboratory environments.

More information: https://ei.is.mpg.de/project/brain-computer-interface

## Causal Inference

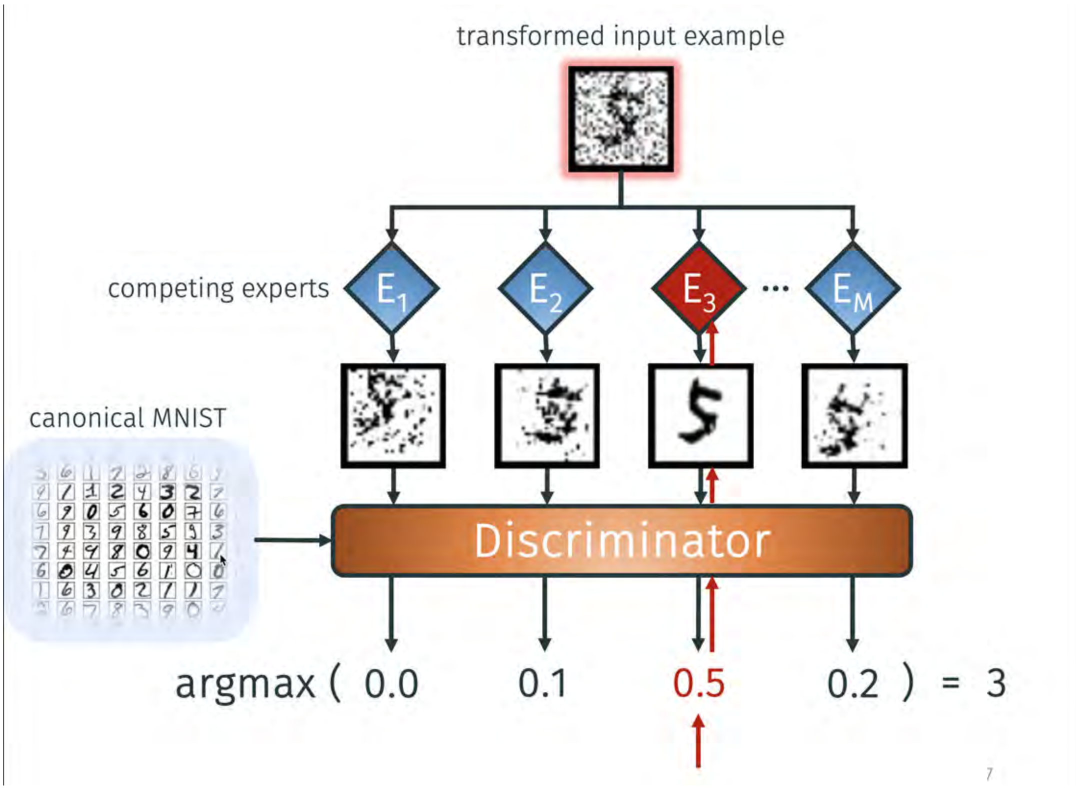


Figure 2.2: Through a competition of experts, inverse causal mechanisms are learned as independent modules in an image task. This results in generalisable and broadly applicable network modules, which can subsequently be combined and applied to new domains.

In the terminology of a book we recently published [68], the term causal inference comprises both *causal reasoning* and *causal discovery,* two somewhat inverse scenarios: While the former employs causal modelds for inferring about the expected observations (often, about their statistical properties), the latter is concerned with inferring causal models from empirical data. Both parts crucially depend on assumptions on the statistical properties entailed by hypothetical causal structures. During the past decade various assumptions have been proposed and assayed that go beyond traditional assumptions like the causal Markov condition and causal faithfulness. This has implications for both scenarios: it improves identifiability of causal structure, and it also entails additional statistical predictions if the causal structure is known.

Traditional causal discovery assumes that the units are connected by a causal directed acyclic graph a priori (mostly as random variables). In contrast, real-world observations are not necessarily a priori structured into those units (e.g. objects in images). The task of identifying reasonable units that admit *causal* models is challenging for both human and machine intelligence, but it aligns with the general goal of modern machine learning to learn *meaningful* representations for data, where 'meaningful' can for instance mean *interpretable*, *robust,* or *transferable.* The general idea that causal structure is a concept that often remains invariant across changing background conditions (first discussed by Herb Simon, and discussed in detail in our book) can be utilized for both causal reasoning and causal discovery. *Identifying causal units* will therefore be a third subsection below.

**Causal reasoning** Subject to sufficiently specific model assumptions (here: additive independent noise), causal knowledge admits novel techniques of noise removal such as our method of half-sibling regression published in PNAS [166], applied to astronomic data from NASA's Kepler space telescope [153].

Apart from entailing statistical properties for a *fixed* distribution, causal models also suggest how distributions *change* across data sets. To this end, one may assume, for instance, that structural equations remain constant across data sets and only the noise distributions change [68, 157], that some of the causal conditionals in a causal Bayesian network change, while others remain constant [96], or that they change *independently* [245], which results in new approaches to domain adaptation [300].

**Causal discovery** The toy problem of telling cause from effect in bivariate distributions, which we have earlier shown to be insightful also for more general causal inference problems, has been further explored [217, 307]. The performance of a broad variety of new approaches has been extensively studied in a long JMLR paper [165], suggesting that classification of cause and effect is indeed possible above chance level. New results for the multi-variable setting deal with, for instance, the problem of learning structural equation models in the presence of selection bias [297] and the idea of employing generalized score functions [209]. In [123], we introduce a kernel-based statistical test for joint independence of random variables which is a key component of multi-variate additive noise based causal inference.

Apart from progress on those 'classical' causal inference problems the domain of causal inference has been extended in several directions. Causal discovery for rare events has been further developed in terms of interacting Hawkes processes [298]. To study causal signals in images, the CVPR paper [258] infers whether the presence of an object on an image is the cause or the effect of the presence of another one, using additive noise based cause-effect inference.

In a study connecting principles of causal inference and foundations of physics [151], we relate asymmetries between cause and effect to asymmetries between past and future, deriving the thermodynamic arrow of time from the basic assumption of (algorithmic) independence of causal mechanisms. Within machine learning and time series modeling, new causal inference methods have revealed previously unknown aspects of the arrow of time [292].

**Identifying causal units and causal learning** Defining objects that are related by *causal models* typically amounts to appropriate coarse-graining of more detailed models of the world (e.g., physical models). Subject to appropriate conditions, causal models like structural equation models can arise from coarse-graining of 'microscopic' models including microscopic structural equation models [249], ordinary differential equations [220], temporally aggregated time series [250], or temporal abstractions of recurrent dynamical models [191]. Although every causal models in economics, medicine, or psychology uses variables that are abstractions of more elementary concepts, it is challenging to state general conditions under which coarse-grained variables admit causal models with well-defined interventions. Our work [249] provides some sufficient conditions.

Theoretical work in [216] shows that the independence of causal mechanisms can be formalized via group symmetry. The plausibility of a scene inferred from an image, for instance, can be formally assessed by testing whether some 'contrast function' attains values that are *typical* among those reached by symmetry transformations. This way, statistical independence (with permutations as corresponding symmetry group) is just a special case of a more general notion of independence.

In the context of a classical image recognition task, our recent ICML paper [203] shows that learning causal models that contain *invariant* mechanisms helps in transferring information across substantially different data sets (see figure). We plan to further pursue this direction, with the goal of moving from statistical representation learning towards causal world representations that should be more robust and support notions of intervention and planning.

More information: https://ei.is.mpg.de/project/causal-inference

## Computational Imaging

Figure 2.3: Super-resolution aims at estimating a high-resolution image from a single low-resolution input (left). Traditional methods tend to produce over-smoothed images that lack high frequency textures and do not look natural (second picture). One focus of our research was the development of algorithms that are able to create realistic textures (third picture) rather than a pixel-accurate reproduction of ground truth (right picture). This boost in perceived image quality is achieved using neural networks in an adversarial training setting [263].

Handheld video cameras now being available in every smartphone, images and videos have become ubiquitous. The amount of visual content on the internet has been ever increasing and digital images and videos have become the main carrier of information over the last few decades.

In our computational imaging group we are interested in a range of signal and image processing problems both in computational photography and scientific imaging. Our focus is on digital image restoration that aims at computationally enhancing the quality of images and recovering probable original images by undoing the adverse effects of image degradation such as noise and blur. Advances in convolutional neural networks have revolutionized computer vision and the field of digital image restoration has been no exception to this rule.

An important problem in digital image restoration is super-resolution, aiming at recovering a high-resolution image from low-resolution input. Traditionally, the performance of algorithms for this task is measured using pixel-wise reconstruction measures such as peak signal-to-noise ratio (PSNR) which have been shown to correlate poorly with the human perception of image quality. In addition, super-resolution is ill-posed and usually multiple plausible high-resolution "explanations" are possible for a given input image. Algorithms minimizing these metrics thus tend to produce over-smoothed images that lack high-frequency textures and do not look natural despite yielding high PSNR values. In a 2017 CVPR paper, we proposed a novel method for automated texture synthesis in combination with a perceptual loss focusing on creating realistic textures rather than optimizing for a pixel accurate reproduction of ground truth images [192, 263]. By using feed-forward fully convolutional neural networks in an adversarial training setting, we achieve a significant boost in image quality even at high magnification ratios (see figure above). Enforcing time consistency leads to novel approaches for the challenging problem of video super-resolution [218] and video prediction [259].

Another focus area has been the problem of blind deblurring. Images often exhibit blur due to unwanted camera shake or moving objects in the scene. Removing the blur is hard as neither the sharp image nor the motion blur kernel is known. Propagating information between multiple consecutive blurry observations can help restore the desired sharp image or video. We have developed efficient recurrent network architectures to deblur frames taking into account temporal information, which can efficiently handle both ego and object motion for arbitrary spatial and temporal input sizes [198, 247, 248, 266].

In addition to work applying image processing to MR images [101, 239], we have recently also tackled the problem of MTF (modulation transfer function) estimation directly from natural images [210], thus avoiding the need to use expensive equipment to characterize the quality of photographic equipment.

More information: https://ei.is.mpg.de/project/computational-imaging

## Fairness

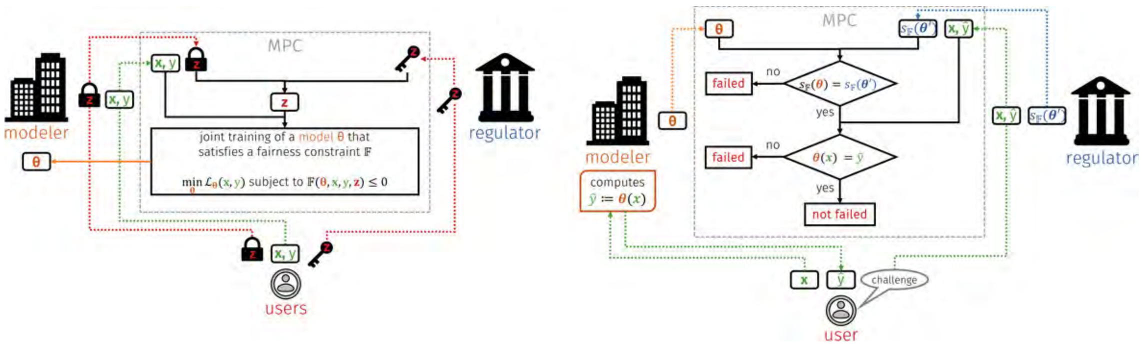


Figure 2.4: Protocols based on secure multi-party computation to learn (left) and verify (right) a fair model using only encrypted sensitive attributes; for details, see [206].

Algorithmic decision making processes are increasingly becoming automated and data-driven in both online (e.g., spam filtering, product personalization) as well as offline (e.g., pretrial risk assessment, mortgage approvals) settings. However, as automated data analysis supplements and even replaces human supervision in decision making, there are growing concerns from civil organizations, governments, and researchers about potential unfairness of these algorithmic decision systems towards people from certain demographic groups (e.g., gender or ethnic groups).

To assay the set of problems and alleviate these concerns, a number of recent studies in the emerging field of ethical machine learning have proposed and analyzed mechanisms to ensure that algorithmic decision systems do not lead to unfair outcomes, or perpetuate historic biases and harmful stereotypes. Such work has the potential to ensure that AI systems are used in a way that is compatible with human values, and it can also help us better understand some of our own biases as we take decisions driven by data and prior knowledge.

**Definitions, Metrics, and Mechanisms** Different forms of legally-problematic discrimination are commonly divided into several categories, the first one being *disparate treatment* (or *direct discrimination*), which occurs if individuals are treated differently according to their sensitive attributes (with all others equal). To avoid disparate treatment, one should not inquire about individuals' sensitive data ("fairness by unawareness"). While this has some intuitive appeal, a significant concern is that sensitive attributes may often be accurately predicted ("reconstructed") from non-sensitive attributes.

Hence, the second form, *disparate impact* (or *indirect discrimination*), occurs when the outcomes of decisions disproportionately benefit or hurt individuals from subgroups with particular sensitive attribute settings. Much recent work in fair learning has focused on approaches to avoiding various notions of disparate impact. Specifically, *demographic parity* demands that the proportion of people in each sensitive group receiving the favorable outcome must be equal. Building on similar ideas, *equality of opportunity* or *disparate mistreatment* demand that among all people who deserve to receive the favorable outcome (for instance, people who would pay back a loan, or people who do not go on to reoffend if released from prison on parole), the fractions of people actually receiving it are equal across sensitive groups. Most of the existing criteria are observational: They depend only on the joint distribution of predictor, protected attribute, features, and outcome, and are formulated as potentially approximate conditional independencies.

In our work, we have worked on defining, measuring, and efficiently mitigating the (un)fairness of a decision-making process regarding people from different sensitive groups. Specifically, we proposed a novel preference-based notion of (un)fairness, which is inspired by the fair division and envy-freeness literature in economics and game theory [223]. This definition provides a more flexible alternative to previous notions that are mostly based on parity of distributions of outcomes. It focuses on whether any group of users would collectively prefer its treatment regardless of the (dis)parity as compared to the other groups, when they are given the choice between various sets of decision treatments.

Moreover, we contributed to algorithmic solutions to mitigate unfairness by developing flexible constraint-based frameworks to enable the design of fair margin-based classifiers [79]. The main technical innovation of our framework is a general and intuitive measure of decision boundary unfairness, which serves as a tractable proxy to several of the most popular computational definitions of unfairness from the literature, such as disparate impact and mistreatment, or preferred fairness. We can thus reduce the design of fair margin-based classifiers to adding tractable constraints on their decision boundaries.

Avoiding both disparate impact and disparate mistreatment is a major challenge. First, to avoid disparate mistreatment, the modeler often needs access to sensitive attributes. However, actively taking sensitive attributes into account introduces disparate treatment, an apparent contradiction. Second, individuals are unlikely to want to entrust sensitive attributes to modelers in all application domains. Finally, legal barriers – for example EU's General Data Protection Regulation (GDPR) – may limit collection and processing of sensitive personal data. We introduce cryptographic methods to resolve these tensions in the intersection of privacy, accountability, and fairness. By encrypting sensitive attributes, we show in our recent ICML paper how a fair model may be learned, checked, or have its outputs verified and held to account, without users revealing their sensitive attributes, cf. title figure [206].

Given the local expertise on causality, we early on realized a major limitation of fairness criteria based solely on the distribution of the observational data [237]. The way humans reason about fairness in decision-making processes crucially hinges on the causal relations underlying the observed decisions. It is conceptually insightful as well as practically relevant to incorporate the causal pathways into fair training procedures. Most fairness notions are based solely on the joint distribution of all variables at play and can thus not distinguish between different causal structures (leading to the same observation distribution) that can have vastly different intuitive interpretations of what is fair. Within the causal framework, we propose novel fairness criteria as well as training methods to achieve fair classifiers that come with theoretical guarantees under certain regularity assumptions.

In particular, we distinguish two context-dependent scenarios. In the *skeptic viewpoint*, we assume that any causal influence from the sensitive attribute on the outcome amounts to harmful discrimination unless it is mitigated via *resolving variables*, which we deem fair to use in our decision as measured. In the *benevolent viewpoint*, we specifically identify *proxy variables* (causal descendents) of the sensitive attribute that must not influence the decision (e.g., the name of a job applicant), but allow for causal pathways from the sensitive attribute to the outcome that do not go through proxies.

**Fairness in Human Decision Making** As described above, there has been a flurry of work on developing computational mechanisms to make sure that the machine learning methods that fuel algorithmic decision making are fair. In contrast, there is a lack of machine learning methods to ensure accuracy and fairness in human decision making, which is still prevalent in a wide range of critical applications such as, e.g., jail-or-release decisions by judges, or accept-or-reject decisions in academia. In this context, each decision is taken by an expert who is typically chosen uniformly at random from a pool of experts. In our recent NeurIPS paper [189], we showed that a random assignment might result in undesirable results in terms of both accuracy and fairness, and propose an algorithm to perform an assignment between decisions and experts which allows optimizing the accuracy and fairness of the overall decision-making process.

More information: https://ei.is.mpg.de/project/fairness

## Learning Algorithms

A *learning algorithm* is the backbone of machine learning that distinguishes it from traditional computer programming by allowing data-driven model building. In the past years, we have developed learning algorithms using a number and tools and for diverse application domains, as outlined below.

**Learning with Kernels** Kernel methods offer a mathematically elegant arsenal to help tackle several problems that arise in machine learning ranging from probabilistic inference to deep learning. Recently, a subfield of kernel methods known as *Hilbert space embedding of distributions* has gained increasing popularity [125], thanks to foundational work done in our department during the last 10+ years. For a probability distribution $\mathbb{P}$ over a measurable space $\mathcal{X}$, the kernel mean embedding of $\mathbb{P}$ can is defined as the mapping $\mu : \mathbb{P} \mapsto \int k(x,\cdot)\,\mathrm{d}\mathbb{P}(x)$ where $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is a positive definite kernel function. Its applications include, but are not limited to, comparing real-world distributions on the basis of samples, differentially private learning, and determining goodness-of-fit of a model.

Our department has an ongoing history of contributions to the state-of-the-art in this area. In [219], we develop privacy-preserving algorithms based on the kernel mean embedding that allow one to release a database while guaranteeing the privacy of each record in the database.

In applications such as probabilistic programming, transforming a base random variable $X$ with a function $f$ forms a basic building block to manipulate a probabilistic model. It then becomes necessary to characterize the distribution of $f(X)$. In [288], we show that for any continuous function $f$, consistent estimators of the mean embedding of a random variable $X$ lead to consistent estimators of the mean embedding of $f(X)$. For Matèrn kernels and sufficiently smooth functions, we also provide rates of convergence.

In [190], we address the problem of measuring the relative goodness of fit of two models using kernel mean embeddings. Given two candidate models, and a set of target observations, the goal is to produce a set of interpretable examples (so-called informative features) which indicate the regions in the data domain where one model fits better than the other. The task is formulated as a statistical test whose runtime complexity is linear in the sample size.

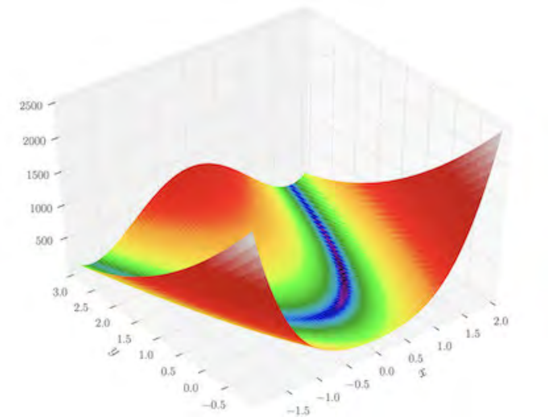

Figure 2.5: The Rosenbrock function, a non-convex function which serves as a test-bed for optimization algorithms (image credit: Wikipedia)

**Optimization for Machine Learning** Optimization lies at the heart of most machine learning algorithms. Characteristics of modern large-scale machine learning problems include: high-dimensional, noisy, and uncertain data; huge volumes of batch or streaming data; intractable models, low accuracy, and reliance on distributed computation or stochastic approximations. Optimizing under these settings with approaches such as coordinate descent and the Frank-Wolfe algorithm has shown promosing results in recent years. The high-level goal of research in optimization in our department is to understand the convergence property of coordinate descent as well as Frank-Wolfe optimization algorithms under different sampling schemes and constraints.

It is well known that greedy coordinate descent (CD) converges faster in practice than the randomized version, however the properties of greedy CD were less well understood. In [252], we provide a theoretical understanding of greedy coordinate descent for smooth functions. We also propose an approximate greedy CD approach which is computationally cheap and always provably better than the randomized version. Similarly, in [236] we propose an adaptive recursive sampling scheme based on the min-max optimal solution of the variance reduction problem to achieve faster convergence for CD. The proposed approach can also be applied to stochastic gradient descent.

Matching pursuit (MP), Frank-Wolfe (FW), and coordinate descent do have a similar structure of the optimization problem. A connection between MP and coordinate descent is explored in [208]. We also prove the rate for accelerated matching pursuit, which was not known previously. The MP algorithm for optimization over conic hulls is proposed in [234]. In [207], an easy-to-implement conditional gradient method is proposed for a composite minimization problem, which converges at the rate of $O(1/\sqrt{k})$. In a different line of work, we propose a Frank-Wolfe based approach to boost variational inference [215], which enables us to analyze the convergence of the proposed framework under suitable assumptions.

**Extreme Classification** Extreme multi-label classification refers to supervised multi-label learning involving hundreds of thousands or even millions of labels. It has been shown that machine learning problems arising in tasks such as recommendation, ranking, and web-advertising can be reduced to the framework of extreme classification. It had been long conjectured that a binary-relevance-based one-vs-rest scheme is not statistically and computationally tenable for such scenarios. Surprisingly, we have been able to show in our recent work [265], that a Hamming loss minimizing one-vs-rest paradigm is key to getting good prediction performance, as well as to efficient training (by enabling parallel training). DiSMEC [265], when published in 2016, surpassed the contemporary state-of-the-art methods by up to 10% points on various datasets consisting of up to a million labels. Since then, it has been a top-performing benchmark method in this domain for over two years now. The concurrent training coupled with model pruning paradigms in DiSMEC have motivated algorithms by Microsoft research which have been used in the Bing Search engine for dynamic search advertising and related searches.

**Neural Networks** Research interest in *deep neural networks*, especially in the generative adversarial network (GAN) approach [176, 199, 214, 261, 263], has increased substantially in recent years. In [214], we propose a simple module to improve a GAN by preprocessing samples with a network that initially makes the task of the discriminator harder (akin to a data smoothing), thus simplifying the generator's task. This leads to a tempered learning process for both generator and discriminator. In a number of experiments, the proposed method can improve quality, stability and/or convergence speed across a range of different GAN architectures (DCGAN, LSGAN, WGAN-GP). In [261], we propose the AdaGAN, a boosting style meta-algorithm which can be combined with various modern generative models (including GANs and VAEs) to improve their quality. We provide an optimal closed form solution for performing greedy updates to approximate an unknown distribution with a sequentially built mixture in any given f-divergence. The paper establishes a fruitful connection between learning theory and neural network research and has already attracted a large amount of follow-up work.

The work [262] develops a deep neural network that can learn to write programs from a corpus of program induction problems. The approach leads to an order of magnitude speedup over strong baselines and an approach based on a recurrent neural network (RNN).

Ideas from causality are beginning to influence our work on machine learning, and the notion of *independent causal mechanism* has been adopted in several areas including semi-supervised learning, domain adaptation, and transfer learning. From a deep learning perspective, we investigated whether a set of independent mechanisms can be recovered using deep neural networks [203]. We proposed an algorithm that enables a set of experts (i.e., deep neural networks) to recover independent (inverse) mechanisms from a data set that has undergone unlabelled transformations. Using a competitive training procedure, the experts specialize to different mechanisms. Not only can the mechanisms be learned successfully, but the system also generalizes to transformed data in other domains.

More information: https://ei.is.mpg.de/project/learning-algorithms

## Learning for Control

Figure 2.6: Muscle-based robotic arm serving as a testbed for Learning for Control. While it offers unique possibilities in terms of high accelerations, extreme speeds, and variable stiffness actuation, classical control methods are unable to unlock these abilities.

Control of complex plants or systems, especially robots actuated by pneumatic artificial muscles, is a challenging task due to nonlinearities, hysteresis effects, massive actuator delay and unobservable dependencies such as temperature. Such plants and robots require much more from the control than classical methods can deliver. Therefore, we aim to develop novel methods for learning control that can deal with high-speed dynamics and muscular actuation.

Highly dynamic tasks that require large accelerations and precise tracking usually rely on accurate models and/or high gain feedback. While kinematic optimization allows for efficient representation and online generation of hitting trajectories, learning to track such dynamic movements with inaccurate models remains an open problem. To achieve accurate tracking for such tasks in a stable and efficient way, we have proposed a series of novel adaptive Iterative Learning Control (ILC) algorithms that are implemented efficiently and enable caution during learning [72].

Muscular systems offer many beneficial properties to achieve human-comparable performance in uncertain and fast-changing tasks [56]. For example, muscles are backdrivable and provide variable stiffness while offering high forces to reach high accelerations. Nevertheless, these advantages come at a high price as such robots defy classical approaches for control. We have built a muscular robot system to study how to accurately control musculoskeletal robots by learning control. We have shown how probabilistic forward dynamics models can be employed to control complex musculoskeletal robots on an antagonistic pair of pneumatic artificial muscles using only one-step-ahead predictions of the forward model and incorporating model uncertainty.

In addition, we have continued to work on reinforcement learning problems, at the intersection of control and machine learning. We have extended several approaches in reinforcement learning for continuous control (NAF, Q-Prop, IPG, TDM) to handle function approximations with significantly improved sample efficiency [195, 212, 238, 243, 280]. In [241], we have shown that our approach scales to learning a door opening task. Aside from fundamental algorithmic problems such as sample efficiency and stability, we also proposed algorithms that enable learning on real-world robots with less human interventions during learning. In [211], we propose the Leave No Trace (LNT) algorithm that significantly reduced the number of hard resets required during learning, paving a path toward autonomous, reset-free learning in real environments. Lastly, we made a contribution to the field of hierarchical reinforcement learning with the HIRO algorithm [170], a scalable off-policy HRL algorithm with substantially improved sample efficiency on difficult continuous control benchmarks over previous methods.

More information: https://ei.is.mpg.de/project/learning-4-control

## Learning on Social Networks

Social media and online social networking sites contain many opinionated, inaccurate or false facts that are often refuted over time. Spread of misinformation may have confused and misled voters in the last U.S. presidential election or the Brexit referendum. To overcome this problem, online platforms deploy evaluation mechanisms for their users to further curate information within these platforms. For example, users can remove inaccurate contents from *Wikipedia*, mark a correct answer as verified in *Stack Overflow* and flag a story as misinformation in *Facebook* and *Twitter*.

We developed [222] a unified computational framework that leverages the temporal traces left by the aforementioned examples of noisy evaluations to estimate robust, unbiased and interpretable notions of information reliability and source trustworthiness. The key idea is that unreliable contents are often removed quickly while reliable contents remain on platforms such as *Wikipedia* for a long time. Similarly, information contributed by sources which systematically spread misinformation are often removed quickly on a wide range of entries, while contents contributed by thrustworthy sources remain on the platform for a long time. By applying our framework to *Wikipedia* data, we are able to answer questions such as whether bbc.co.uk provides more reliable information compared to newyorktimes.com in *Wikipedia* entries related to the U.S politics, and at which point in time a particular *Wikipedia* entry, such as *Barack H. Obama*, was unreliable due to ongoing controversies.

Next, we focused on developing a machine learning method to detect and reduce the spread of harmful misinformation in online social networking sites through the power of the crowd and fact-checking [169]. Given limited reviewing resources, the main question is how to prioritize growing amounts of questionable contents. Some cases of misinformation are not identified until a large number of users have been already exposed to it. However, many cases may only have a limited impact on people and unnecessarily consume reviewing resources for fact-checking. To address these challenges, we developed a robust methodology with provable guarantees to minimize the impact of potentially harmful contents on a large number of people. Results of applying this algorithm on datasets from *Twitter* and *Weibo* suggest that our method can identify cases of misinformation earlier than alternative methods and also uses fact-checking resources more efficiently.

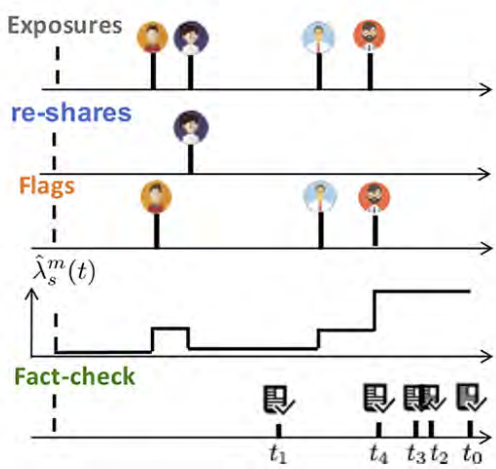


Figure 2.7: Fact checking of contents on social networks using observed exposure, reshare and flag events. The rate of fact checking, $\hat{\lambda}_s^m(t)$, is updated after every observed event.

More recently, we applied some of the above techniques to human learning. In this project, due to appear in PNAS [83], we develop a methodology to optimize spaced repetition algorithms, a widely used procedure to memorize new information, e.g., vocabulary when learning a foreign language in online learning platforms such as *Duolingo*. The promise of online platforms such as *Duolingo* is that automated fine-grained monitoring and a greater degree of control will result in more effective spaced repetition algorithms. However, the algorithms used in these services tend to be simple rule-based heuristics. We introduce a principled and data-driven method to characterize spaced repetition algorithms. Results from applying this framework to a dataset from *Duolingo* indicate that this method help learners memorize more effectively and efficiently.

More information: https://ei.is.mpg.de/project/learning-on-social-networks

## Neuroscience

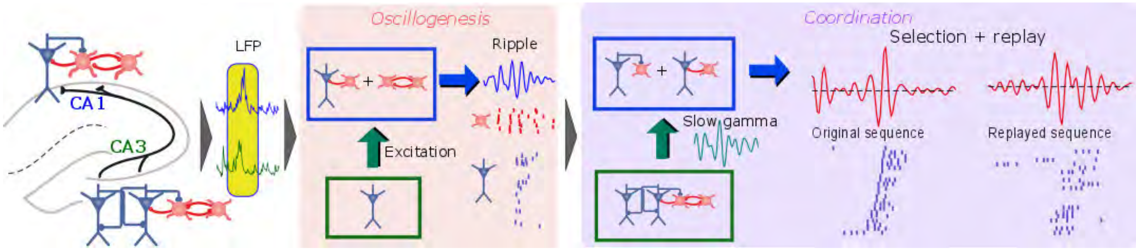


Figure 2.8: The mechanisms involved in replay of memory traces in the hippocampus. Left: connectivity of excitatory (blue) and inhibitory (red) cells in regions CA3 and CA1. Center: key circuits involved in generating synchronized neuronal activity. Right: circuits involved in generating replays of sequential activity of a sparse population encoding a memory trace. See [95].

Biological neural networks are characterized by a dense connectivity at multiple scales that makes it challenging to assess their organization and function, and generates complex and high dimensional dynamics. We develop machine learning and statistical tools to study the functional organization of these systems.

One illustration of the challenges brought by modern neuroscience is the diversity of functions and information encoded by neurons in higher level associative regions such as the prefrontal cortex (PFC) of non-human primates. In order to characterize the dominant activity patterns among large populations of recorded neurons in this region without prior assumptions, we exploit unsupervised learning tools such as Non-negative Matrix Factorization. This led to uncovering that a large proportion of cells in PFC exhibit general task monitoring activities, independently from the population representing the conscious content of visual percepts [93]. This finding is significant for the field, as such task monitoring activities represent a potential confounder when studying the neural basis of consciousness investigated in PFC.

Beyond characterizing neuronal responses in various contexts, a fundamental question is how networks self-organize in order to coordinate information processing among their units, thereby giving rise to reliable internal representations. One fascinating example is the organization of *episodic memory*, where evidence has shown that recollection of a particular event is associated with the precisely coordinated *replay* of sequences of spikes in hippocampal cell assemblies (see Figure). The mechanisms allowing the coordination of this phenomenon are still largely unknown and require accounting for detailed biophysical properties through computational modeling. Due to their intractable dynamics and high dimensionality, biologically realistic computational models do not provide immediate mechanistic knowledge of the underlying mechanisms of neural dynamics, but require sophisticated data analysis techniques to be interpreted. In order to investigate the replay of *memory traces* in the hippocampus of primates, we designed a detailed computational model able to reproduce a broad range of in-vivo experimental observations [95]. We established the key neural circuits underlying the reliable replay of memorized events with the help of supervised machine learning and interventions on the model. This led in particular to the discovery of the pivotal role played by *feedback inhibition* loops between excitatory and inhibitory neurons (see Figure), shaping the sparse sequential activity of neurons encoding a memory trace during replay.

Beyond these successful examples, our work is directed towards exploiting causality principles [216] to establish and exploit a synergy of machine learning with experimental and computational modeling approaches, and ultimately bridge the gap between the complexity of brain phenomena and their underlying biological mechanisms.

More information: https://ei.is.mpg.de/project/neuroscience

## Privacy

A high percentage of modern data is collected from devices enabling human-machine interactions. For instance, robot vacuum cleaners often have a built-in camera to navigate where to clean, but the camera can also capture scenes of people in the room. Similarly, user-facing artificial intelligence agents such as virtual assistants collect, store, and stream potentially privacy-sensitive data. These data might be used to fit machine learning models to solve certain statistical tasks. Many recent papers indicate that fitted machine learning models can expose sensitive information from the dataset they were trained on.

For developing privacy-preserving machine learning methods, the first question we need to answer is how to define privacy and how to measure it. While there are various privacy notions developed in the field, *differential privacy* is currently the gold standard for privacy, due to its rigorous privacy guarantees. A randomized algorithm $\mathcal{M}(\mathbf{X})$ that takes a dataset $\mathbf{X}$ and outputs a quantity (scalar or vector) is said to be $(\epsilon, \delta)$-differentially private if the probabilities of the algorithm outputting the same quantity $\mathcal{S}$ given two datasets $\mathbf{X}$ and $\mathbf{X}'$ are related as

$$Pr(\mathcal{M}(\mathbf{X}) \in \mathcal{S}) \leq \exp(\epsilon) Pr(\mathcal{M}(\mathbf{X}') \in \mathcal{S}) + \delta \tag{2.1}$$

for all measurable subsets $\mathcal{S}$ of the range of $\mathcal{M}$ and for all neighbouring datasets $\mathbf{X}, \mathbf{X}'$ that differ by a single entry (either by excluding that entry or replacing it with a new entry). Here, an entry usually corresponds to a single individual's private data. Intuitively, the definition states that the probability of any event does not change much when a single individual's data is modified, thereby limiting the amount of information that the algorithm reveals about any one individual.

Equipped with such a notion of privacy, the development of differentially private data analysis methods requires first to decide which quantity we want to guard. Depending on whether we want to share data, statistics, or model parameters, each will require developing different techniques. Below, we look into two different types of sharing, namely, data sharing and parameter sharing.

**Model sharing** is probably the most popular way to achieve privacy in the current differential privacy literature. Model sharing often focuses on guarding model parameters by adding noise to them before releasing them. Generally, there are two ways of achieving private models. First, one could add noise to the objective function ("objective perturbation") such that the resulting estimates guarantee some level of privacy. Second, one could add noise to the output of an optimization routine ("output perturbation"). Often, objective perturbation techniques end up adding less noise than output perturbation techniques. However, to be able to analyze the relationship between the amount of noise that needs to be added to an objective function and the guaranteed privacy level of a corresponding estimate after optimizing the perturbed objective, it is unavoidable to make some stringent assumptions (e.g., strong convexity). Many learning problems in machine learning have non-convex objective functions, which limits the usefulness of the existing objective perturbation techniques. We are currently working on addressing the problem with non-convex objective functions, especially in the context of deep learning [171].

**Data sharing** requires adding noise to the dataset itself before releasing it. Most of the existing data sharing frameworks are designed for particular data types (e.g., count data, low-dimensional data) or particular purposes (e.g., decision tree algorithms) only. We thus need algorithms to add noise in a way that is truthful to the statistical properties of the raw data while being independent of downstream tasks, for a better utility in various statistical analyses. At ICML 2018, we have published a kernel method for this task, which relies upon kernel mean embedding on datasets in reproducing kernel Hilbert spaces, which (for suitable kernels) retains all information about a distribution [219].

More information: https://ei.is.mpg.de/project/privacy

## Probabilistic Inference

The probabilistic formulation of inference-conditioning probability measures encoding prior assumptions by multiplying with a likelihood describing the data given the generative process (refer to Figure 2.9)-remains one of the main research streams within machine learning, and therefore, of the EI Department, where we address the main aspects of probabilistic inference.

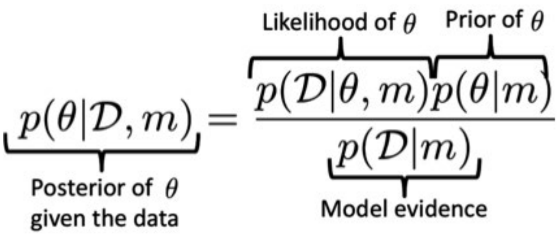


Figure 2.9: Illustration of probabilistic inference.

**Nonparametric inference on function spaces** Gaussian Process models allow for nonparametric inference in function spaces. They have a strong inductive bias specified by the covariance (kernel) function between observations, which allows for data-efficient learning. One of our main themes in this field has been nonparametric inference on function spaces using Gaussian process models, with its main practical challenge for inference being the cubic complexity in terms of the number of training points. In our work carried out on this topic, we provided a more thorough theoretical and practical analysis of two classes of approximations [283], highlighting both the theoretical merits but also the optimisation problems associated with variational inference. In another project dealing with Gaussian Processes, we have explored models that can account for invariances between data points [187]. To this end, we construct a covariance kernel that explicitly takes these invariances into account by integrating over the orbits of general symmetry transformations. We provide a tractable inference scheme that generalises recent advances in Convolutional Gaussian Processes.

In the work on the Mondrian kernel [284], we provide a new algorithm for approximating the Laplace kernel in any kernel method application (including in Gaussian processes) using random features. Attractive properties of the algorithm include that it allows finding the best kernel bandwidth hyperparameter efficiently and it is well-suited for online learning.

**Variational Inference** Variational inference is a popular technique to approximate a possibly intractable Bayesian posterior with a more tractable one. Recently, boosting variational inference has been proposed as a new paradigm to approximate the posterior using a mixture of densities by greedily adding components to the mixture. However, its theoretical properties were unclear. In our research [215], we study the convergence properties of this approach from a modern optimization viewpoint by establishing connections to the classic Frank-Wolfe algorithm. Our analysis yields novel theoretical insights regarding the sufficient conditions for convergence, explicit rates, and algorithmic simplifications.

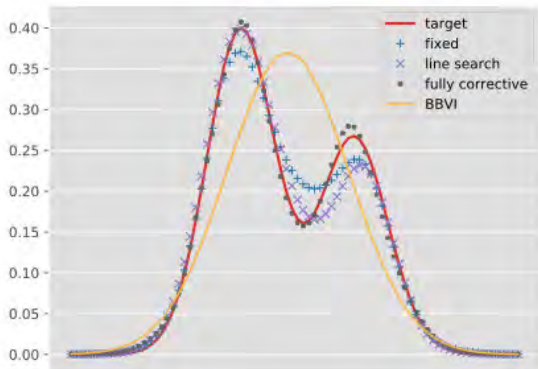


Figure 2.10: Comparison between black box variational inference (BBVI) and three variants of our boosting BBVI method on a mixture of Gaussians [188].

Unfortunately, the above work imposes stringent assumptions that require significant effort for practitioners. Specifically, they require a custom implementation of the greedy step (called the LMO) for every probabilistic model with respect to an unnatural variational family of truncated distributions. In a more recent work [188], we fix these issues with novel theoretical and algorithmic insights. On the theoretical side, we show that boosting variational inference (VI) satisfies a relaxed smoothness assumption which is sufficient for the convergence of the functional Frank-Wolfe (FW) algorithm. Furthermore, we rephrase the LMO problem and propose to maximize the Residual ELBO (RELBO) which replaces the standard ELBO optimization in VI.

These theoretical enhancements allow for black box implementation of the boosting subroutine. As a result, the proposed boosting black box variational inference algorithm can be readily implemented in any probabilistic programming framework based on variational inference. As shown in Figure 2.10, it leads to richer posterior approximations that standard black box variational inference (BBVI).

**Inference on discrete graphical models** Another open challenge is to perform efficient inference in discrete graphical models, where estimating normalising constants and sampling is often difficult. In this context, in [240], we show that the Gumbel trick – known to convert either the partition-function-estimation problem (an integration problem), or the sampling problem, into an optimisation problem – is just one method out of an entire family and that other methods of the family can sometimes work better. For partition function estimation this means that other members of the family can lead to estimators of the partition function that have lower variance, or lower mean-squared-error, than the estimator obtained from the standard Gumbel trick. This paper received a Best Paper Honourable Mention at ICML 2017.

In a more applied piece of work [78], a Bayesian inference algorithm on slice sampling and particle Gibbs with ancestor sampling is developed, to efficiently deal with the combinatorial number of states in the infinite factorial finite state machine model, which is here used to address the problem of joint channel parameter and data estimation in a multiuser communication channel in which the number of transmitters is unknown.

**Probabilistic programming** Recent probabilistic programming languages aim to enable data scientists to express sophisticated probabilistic models appropriate for her data as programs, without needing to worry about the inference step, since they come with the implementation of multiple algorithms for performing posterior inference for models and data sets expressed in these languages. In this sense, probabilistic programming languages will help pave the way to statistical data science and AI.

In [196], we present an architectural design of a library for Bayesian modeling and inference in modern functional programming languages. The novel aspect of the approach is modular implementations of existing state-of-the-art inference algorithms. The design relies on three inherently functional features: higher-order functions, inductive data-types, and support for either type-classes or an expressive module system. We provide a performant Haskell implementation of this architecture, demonstrating that high-level and modular probabilistic programming can be added as a library in sufficiently expressive languages.

Although in probabilistic programming languages, sophisticated inference algorithms are often explained in terms of the composition of smaller parts, neither their theoretical justification nor their implementation reflects this modularity. In [197], it is shown how to conceptualise and analyse such inference algorithms as manipulating intermediate representations of probabilistic programs using higher-order functions and inductive types, and their denotational semantics.

A theoretical study relevant to probabilistic programming [288] deals with the problem of representing the distribution of $f(X)$ for a random variable $X$ and a function $f$. We use kernel mean embedding methods to construct consistent estimators of the mean embedding of $f(X)$. The method is applicable to arbitrary data types on which suitable kernels can be defined. It thus allows us to generalize (to the probabilistic case) functions in computer programming which are originally only defined on deterministic data types.

More information: https://ei.is.mpg.de/project/probabilistic-inference

## Psychology

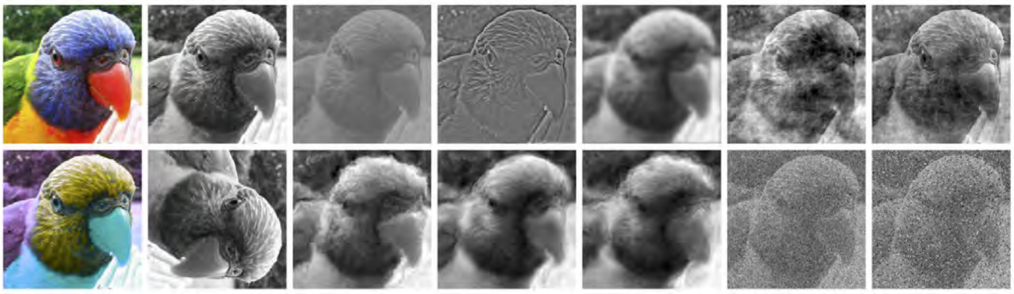

Figure 2.11: Human-DNN-robustness comparison: stimulus images. From left to right, image manipulations in the top row are: undistorted, greyscale, low contrast, high-pass, low-pass (blurring), phase noise, power equalisation. Bottom row: opponent colour, rotation, Eidolon I, II and III, additive uniform noise, salt-and-pepper noise. Typically human observers are more robust when the signal-to-noise ratio decreases. Additionally, we find progressively diverging patterns of classification errors between humans and DNNs with weaker signals.

This project investigates human perception, combining psychophysical experiments and computational modeling. We have five research foci:

First, we develop methods to gain more information from psychophysical data, linking traditional methods with machine learning approaches [145, 303]. We characterized serial dependencies in behavioral responses and introduced methods to correct for them. Our software package to perform Bayesian inference for the psychometric function for non-stationary data is freely available and widely used [146]. In addition, we investigate reliable supra-threshold psychophysical paradigms which are more intuitive and require less training and are thus more likely to yield reliable crowd-sourcing data [120, 129].

Our second focus is the development of an image-based model of spatial vision. We integrated the large psychophysical literature on detection and discrimination experiments and proposed a model based on maximum-likelihood decoding of a population of model neurons predicting the most important spatial vision data sets simultaneously, using a single set of parameters [122]. Remarkably, the model generalises well to data it was not trained upon (natural images), produces a highly sparse code, and is useful in applications (in perceptual visual quality metrics). In the future, we plan to assay whether the model could be useful as a preprocessing step for deep neural networks (DNNs).

Third, we are exploring similarities and differences of DNNs and the human visual system. One line of work uses DNNs as generative models for texture synthesis, casting doubt on the popular notion that the peripheral visual system's internal representations are texture-like [121]. In a second line of work, we compare the robustness of humans and current convolutional DNNs on object recognition under various types of image degradation, finding the human visual system (still) to be much more robust [173]. These differences cannot be overcome by training on distorted images (i.e., data augmentation): While DNNs cope well with the exact distortion they were trained on, they still show a strong generalisation failure towards unseen distortions.

Fourth, we investigate perception of causality, focusing on the arrow-of-time that we previously studied theoretically [292]. Preliminary data suggest that human observers can discriminite forward and backward played movies of autoregressive (AR) motion with non-Gaussian additive independent noise, i.e., they appear sensitive to the subtle temporal dependencies of the residuals of the AR-motion. Currently we are testing how long the motion sequences have to be for successful arrow-of-time discrimination and how well observers can generalise from one type of non-Gaussian noise to another.

Another project pertaining to the wider field of psychology is described in the project report on social networks, dealing with the optimization of spaced repetition algorithms for learning [83].

More information: https://ei.is.mpg.de/project/psychology

## Robotics

Figure 2.12: In the past, we have studied a single robot playing table tennis against a human or ball gun. In 2018, we migrated to a new lab and a setup with two robots. We aim to enable the unique high-speed Barrett WAM robot arms PING and PONG to jointly play table tennis in a collaborative game using methods from robotic skill learning.

Creating autonomous robots that can learn to assist humans in situations of daily life is a fascinating challenge for machine learning. While this aim has been a long-standing vision of artificial intelligence, we have yet to create robots that can learn to accomplish many different tasks triggered by environmental context or higher-level instruction. Our goal is the investigation of the ingredients for such a general approach to motor skill learning, in order to get closer towards human-like performance in robotics. We thus focus on the solution of basic problems in robotics by developing domain-appropriate machine-learning methods.

While many machine learning methods work in theory, in simplified simulations and text-book control plants, it is essential to study real robot systems to understand the learning of high-performance motor skills. We focus on the problem of learning robot table tennis as our "Drosophila" (i.e., a model system used in experiments to gain insight that we hope will be generalizable) of robot learning. This task has a number of components that are representative of tasks encountered by natural intelligent systems, including perception and action, as well as various aspects of social interaction (opponent modeling, competition, collaboration).

In our studies, we focus on using machine learning approaches for improved tracking, imitation of demonstrated behavior, and self-improvement by robot reinforcement learning with a strong focus on high-speed skill learning.

Learning approaches have to generalize a complex hitting behavior from relatively few demonstrated trajectories, which neither cover all ball trajectories nor all desired hitting directions. Therefore, past approaches that only modeled a deterministic mean behavior without capturing the variability of the movement have been fairly limited. Recent work on capturing trajectory distributions using probabilistic movement representations [71] opens new possibilities for robot table tennis. We have presented several methods to adapt probabilistic movement primitives, e.g., for adapting hitting movements learned in joint space to have a desired end effector position, velocity and orientation [47], as well as to find the initial time and duration of the movement primitive in order to intercept a moving object like the table tennis ball [71]. The resulting methods rely on straightforward operations from probability theory and provide a more principled approach to solve some of the challenges of robot table tennis compared to previous approaches. This also enabled us to learn several other tasks.

We have worked on various other questions of robot motion control with the context of robot table tennis in real robot experiments on the Barrett WAM. Among these approaches, we have studied the properties of optimal trajectory generation in robot table tennis strikes [49], learning striking controllers [72]. We have also recently demonstrated how a table tennis serve can be captured and successfully reproduced [70].

More information: https://ei.is.mpg.de/project/robotics

## Statistical Learning Theory

Machine learning algorithms are designed to generalize from past observations across different problem settings. The goal of learning theory is to analyze statistical and computational properties of learning algorithms and to provide guarantees on their performance. To do so, it poses these tasks in a rigorous mathematical framework and deals with them under various assumptions on the data-generating process.

In [295] we initiate a formal analysis of **compressing a data sample** so as to encode a set of functions consistent with (or of minimal error on) the data. We propose several formal requirements (exact versus approximate recovery and worst case versus statistical data generation) for such compression and identify parameters of function classes that characterize the resulting compression sizes. In [291] we provide a novel analysis for a **life-long learning** setup where performance guarantees are required for every encountered task. Such a setup had previously been analyzed only under rather restrictive assumptions on the data-generating process. Our study generalizes a natural lifelong learning (where at every step, if possible, a predictor is created as an ensemble over previously learned ones using only little data) and identifies conditions of task relatedness that render such a scheme data efficient. In [290] we show that **active learning** can provide label savings in non-parametric learning settings. Previously this had mostly been done in parametric learning of a classifier from a fixed class of bounded capacity. We develop a novel active query procedure that takes in unlabeled data and constructs a compressed version of the underlying labeled sample while automatically adapting a number of label queries. We then show that this procedure maintains performance guarantees of nearest neighbor classification.

In recent years the **kernel mean embedding (KME)** of distributions started to play an important role in various machine learning tasks, including independence testing, density estimation, implicit generative models, and more. Given a reproducing kernel and its corresponding reproducing kernel Hilbert space (RKHS), KME maps a distribution $P$ over the input domain to the element $\mu_P := \int k(X,\cdot)dP(X)$ in the RKHS. An important step in many KME-based learning methods is to estimate the distribution embedding $\mu_P$ using observations $x_1, \ldots, x_n$ sampled from $P$. Inspired by the James-Stein estimator, in [164] we introduced a new type of KME estimators called **kernel mean shrinkage estimators (KMSEs)** and proved that it can converge faster than the empirical KME estimator $\hat{\mu}_P := \sum_{i=1}^n k(x_i,\cdot)/n$. This improvement is due to the bias-variance tradeoff: the shrinkage estimator reduces variance substantially at the expense of a small bias. We also empirically showed that KMSE is particularly useful when the sample size $n$ is small compared to the input space dimensionality.

We have studied the optimality of KME estimators in the **minimax sense**. In [128] we show that the rate $O(n^{-1/2})$ achieved by $\hat{\mu}_P$, KMSE, and many other methods published in the literature is optimal and cannot be improved. This holds for any continuous translation-invariant kernel and for various classes of distributions, including both discrete and smooth distributions with infinitely differentiable densities. In [287] we also study the minimax optimal estimation of the **maximum mean discrepancy (MMD)** between two probability distributions, which is defined as the RKHS distance between their KMEs: $\mathrm{MMD}(P,Q) := \|\mu_P - \mu_Q\|$. We show that for any radial universal kernel the rate $O(n^{-1/2} + m^{-1/2})$ achieved by existing estimators is minimax optimal.

The properties of MMD are known to depend on the underlying kernel and have been linked to three fundamental concepts: **universal, characteristic, and strictly positive definite kernels**. In [103] we show that these concepts are essentially equivalent and give the first complete characterization of those kernels whose associated MMD metrizes the weak convergence of probability measures. Finally, we show that KME can be extended to Schwartz-distributions and analyze properties of these distribution embeddings.

While MMDs are known to metrize convergence in distribution, the underlying conditions are too stringent when one only aims to metrize convergence to a fixed distribution, which is the case for instance in goodness-of-fit tests. To address this, we derive necessary and sufficient conditions for MMD to **metrize tight convergence** to a fixed target distribution.[2] We use our characterizations to analyze the convergence properties of the targeted kernel Stein discrepancies (KSDs) commonly employed in the goodness-of-fit testing. The results validate the use of KSDs for a broader set of targets, kernels, and approximating distributions.

The problem of **estimating a distribution of functions of random variables** plays an important role in the field of probabilistic programming, where it can be used to generalize functional operations to distributions over data types. In [288] we proposed a non-parametric way to estimate the distribution of $f(X)$ for any continuous function $f$ of a random variable $X$. The proposed KME based estimators are proven to be asymptotically consistent. We provide finite-sample guarantees under stronger assumptions.

Motivated by recent advances in **privacy-preserving machine learning** and building upon the results of [288], we have proposed a theoretical framework for a novel database release mechanism that allows third-parties to construct consistent estimators of population statistics while ensuring that the privacy of each individual contributing to the database is protected [219]. Our framework is based on newly introduced differentially private and consistent estimators of KMEs, of interest in their own right.

Visually impressive progress in machine learning has been made in the field of **unsupervised generative modeling** with generative adversarial networks (GANs), variational autoencoders (VAEs) and other deep neural network based architectures, significantly improving the state of the art in the quality of generated samples, especially in the domain of natural images.

In [261] we study the **training of mixtures of generative models** from a theoretical perspective. We find a globally optimal closed form solution for performing greedy updates while approximating an unknown distribution with mixtures in any given f-divergence. We then derive a boosting style meta-algorithm which can be combined with many modern generative models (including GANs and VAEs).

While training objectives in VAEs and GANs are based on f-divergences, it has been recently shown that other divergences, in particular, **optimal transport distances**, may be better suited to the needs of generative modeling. Starting from Kantorovich's primal formulation of the optimal transport problem, we show that it can be equivalently written in terms of probabilistic encoders, which are constrained to match the latent posterior and prior distributions.[3] We then apply this result to train latent variable generative models in [213]. When relaxed, the constrained optimization problem leads to a new **regularized autoencoder algorithm** which we call Wasserstein auto-encoders (WAEs). WAEs share many of the properties of VAEs (stable training, nice latent manifold structure) while generating samples of better quality, as measured by the Frechet Inception score across multiple datasets.

In [177, 178] we focus on **properties of the latent representations** learned by WAEs and draw several interesting conclusions based on various experiments. First, we show that there are fundamental problems when training WAEs with deterministic encoders when the intrinsic dimensionality of the data is different from the latent space dimensionality. Second, we point out that training WAEs with probabilistic encoders is a challenging problem, and propose a heuristic approach with promising results on several datasets.

Many deep neural network based architectures have been proven vulnerable to so-called **adversarial attacks**. In the case of natural image classifiers, carefully chosen but imperceptible image perturbations can lead to drastically changing predictions. We showed that adversarial vulnerability increases with the gradients of the training objective when viewed as a function of the inputs.[4] For most current network architectures, we prove that the $\ell_1$-norm of these gradients grows as the square root of the input size. These nets therefore become increasingly vulnerable with growing image size.

More information: https://ei.is.mpg.de/project/statistical-learning-theory

[2] C. J. Simon-Gabriel, L. Mackey. Targeted Convergence Characteristics of Maximum Mean Discrepancies and Kernel Stein Discrepancies. In *PhD thesis: Distribution-Dissimlarities in Machine Learning (University of Tübingen)*, 2018.

[3] O. Bousquet, S. Gelly, I. Tolstikhin, C. J. Simon-Gabriel, B. Schölkopf. From Optimal Transport to Generative Modeling: the VEGAN cookbook. *CoRR* **abs/1705.07642**, 2017.

[4] C. J. Simon-Gabriel, Y. Ollivier, B. Schölkopf, L. Bottou, D. Lopez-Paz. Adversarial Vulnerability of Neural Networks Increases with Input Dimension. *CoRR* **abs/1802.01421**, 2018.

## 2.3 Equipment

**Computing Infrastructure.** The desktop computing environment of the Department of Empirical Inference is based on Intel PCs, and currently uses the centrally managed operating systems Ubuntu Linux, Microsoft Windows, and Mac OS X. Silent PCs with no moving parts are used to provide the best possible work environment.

For data storage, the department currently offers 80TB of online file system storage on two file servers, used for shared data as well as personal home directories. Backups are done through a tape library system, as well as a disk-based snapshotting system for the most important data. Off-site storage from RZ Garching is offered for long-term archival of scientific data.

For numerical experiments, the group has access to a high performance computing cluster, shared with the departments of the MPI for Intelligent Systems. The cluster nodes and their supporting infrastructure are maintained by a central scientific facility. The Empirical Inference department is using around 40

Figure 2.13: Some of the nodes of the computer cluster.

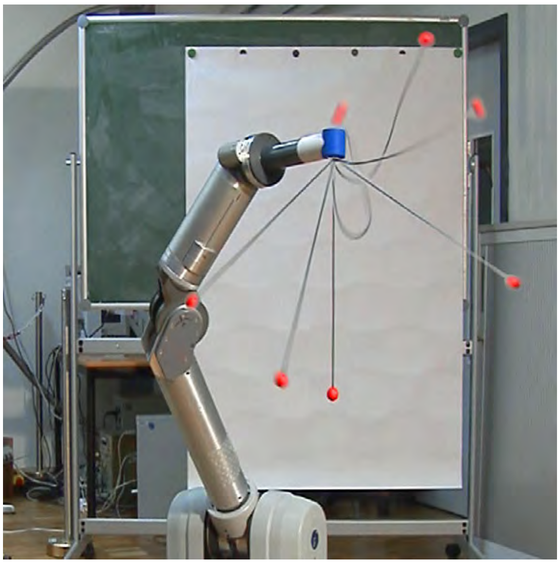

Figure 2.14: The Barrett WAM robot arm

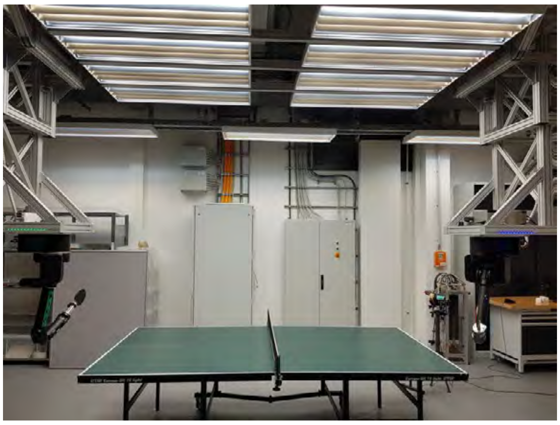

Figure 2.15: Two Barrett WAM robot arms playing cooperatively or adversarily

**Robot Learning Lab: High-Speed Robot Arms.** The Robot Learning Lab of the Empirical Inference Department focuses on finding task-appropriate machine learning methods for acquiring and refining motor skills. Current objectives include learning high-speed compliant control, learning simple skills and composite tasks such as table tennis. These goals require a maintenance-intensive set-up.

For fast as well as compliant control, we have two unique, custom-made high-voltage version of Barrett WAM robot arms specifically designed for high-speed control while having seven degrees of freedom. Torque-level access to the robot, back drivability, and little backlash enables the use of this platform for learning control experiments and as a haptic input device during imitation learning or robot task assistance. The two Barrett WAM robots oppose each other at a table tennis setup, enabling both robots to potentially play cooperatively as well as adversarily. The robot control makes use of a custom four-camera high-speed (200 Hz) vision setup. Strong LED lights with adjustable intensity and color ease tracking of objects at such high frame rates. All software is based on a real-time Linux operating system and on the robot programming framework SL.

A tendon-driven pneumatic artificial muscle robot arm was developed in-house in order to study antagonistically actuated joints, allowing for light-weight segments, incorporating strong muscles, and using co-contraction for compliant control. We aim at approaching performance as observed in humans by designing such a system and enabling safer applications of learning control approaches whilst extending the variety of possible trajectories. We aim to use this robot to show that learning can be particularly beneficial when classical methods fail.

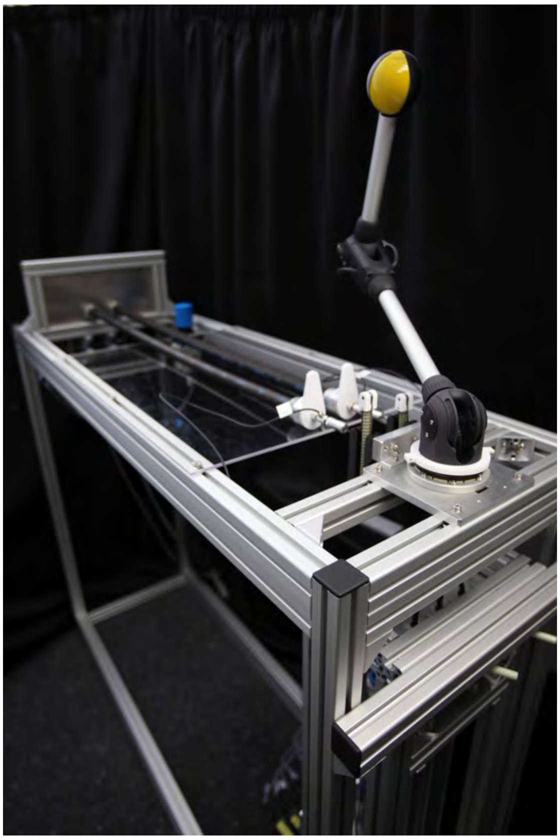

Figure 2.16: Lightweight robot actuated by eight pneumatic muscles. Strong muscles enable to generate high accelerations while the antagonistic actuation enables this robot to explore in high-speed regimes.

**Brain-Imaging Equipment.** The Brain-Computer Interfaces (BCI) Lab of the Empirical Inference department divides its experimental interests into two main areas: high-density, clinical studies, and large-scale neuroscientific recordings. For the clinical studies, we have an mrShield electromagnetically shielded cabin by CFV and a 128-channel BrainAmp system with a BrainCap to conduct high-fidelity, low-noise recordings within the institute. For studies that require high quality data recorded outside of the laboratory, the QuickAmp 136 (Figure 2.17) is a high input-impedance amplifier capable of measuring and recording up to 128 channels of electroencephalographic (EEG) data, as well as eye (EOG), muscle (EMG), and other physiological signals. We have further acquired the Armeo Power, a robotic arm exoskeleton manufactured by Hocoma AG (Switzerland) for rehabilitation. This exoskeleton enables us to both support and perturb movements of the upper extremities during motor learning, which enables us to study the neural basis of motor learning in healthy subjects and patient populations with movement disorders. Lastly, we have recently also purchased an HTC Vive VR headset with additional sensors in order to explore the possibilities of rehabilitation and brain-computer interfacing in virtual environments.

Figure 2.17: The BCI lab of the empirical inference department

For large-scale neuroscientific recordings, we have acquired a stock of commercial recording headsets. This compromises 34 Muse headbands (InteraXon, Canada), two Emotiv EEGs (Emotiv, USA), and one Dreem headband (Dreem, France). While these devices are not sufficient for high-fidelity reconstruction of neural signals, they are key to scaling the basic research efforts in our laboratory, as they can be used without expert assistance. When paired with a computer or iPad, it is possible to have patients conduct certain studies fully at home, with researcher involvement limited to experimental design and data analysis. In addition, the Dreem headset is optimized for recordings during sleep, which allows for long-term investigations of circadian rhythms and their effect on BCI usability.

**Computational Imaging.** In computational imaging, it is desirable to develop deblurring and denoising algorithms that work not just on synthetic but also on real data. To be able to take such photographs for varying setups and with controlled distortions, we use several SLR and CCD cameras (including Canon 5DS R, Sony A7R Mk. II, SBIG STX16803) with various lenses. For quality assessment and quantitative comparisons, we use an image quality analysis system (iQ Analyer by Image Engineering). Additionally, we have developed a precise test-panel that allows us to efficiently measure the properties of imaging optics (such as MTF) in a single shot (see Figure 2.18).

Figure 2.18: Test panel with over 4000 point sources (see zoom-in) for precise point-spread-function mapping.

To generate images and image sequences with controlled camera shake, we also use a Stewart platform (hexapod robot, see Figure 2.19) that allows a high accuracy for repeated movements.

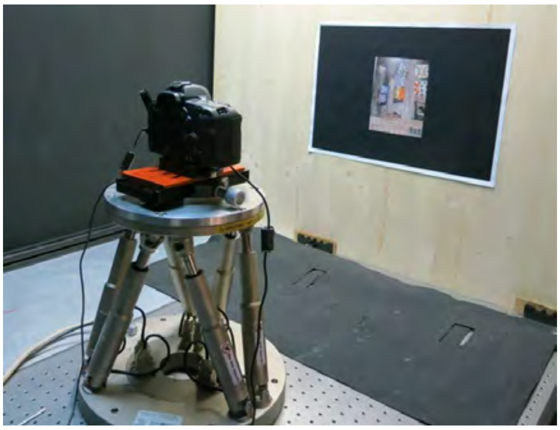

Figure 2.19: Camera-shake simulator: hexapod robot performing reproducible trajectories.

In addition to usual photographic lenses, we also make use of telescopes (including a 12" Astro Physics Riccardi-Honders) that allow us to work on real-world astronomical image sequences affected by turbulence. Alongside the telescope, we employ an array of up to five telephoto lenses (see Figure 2.20) to simultaneously acquire multiple images for joint super-resolution algorithms.

Figure 2.20: Telescope array for simultaneous multi-image acquisition (right).

More information: https://ei.is.mpg.de/pages/equipment

## 2.4 Awards & Honors

### 2018

**Olivier Bousquet**, Leon Bottou: Test of Time Award, 2018 Neural Information Processing Systems conference, for their 2007 paper on "The Tradeoffs of Large Scale Learning" reporting work originally started in 2003 when Olivier Bousquet was a member of our lab.

**Bernhard Schölkopf** receives the Landesforschungspreis 2018 for Basic Research by the State of Baden-Württemberg.

**Dominik Janzing, Jonas Peters, Bernhard Schölkopf**: Causality in Statistics Education Award (American Statistical Association).

**Bernhard Schölkopf**: Gottfried Wilhelm Leibniz-Preis 2018, Deutsche Forschungsgemeinsschaft (DFG).

### 2017

**Bernhard Schölkopf**: ACM Fellow for contributions to the theory and practice of machine learning.

**Matej Balog**, Nilesh Tripuraneni, Zoubin Ghahramani, Adrian Weller: Honorable Mentions at ICML 2017 for their paper: "Lost Relatives of the Gumbel Trick".

### 2016

**Moritz Grosse-Wentrup**, **Vinay Jayaram**: IEEE Brain Initiative Best Paper Award at the IEEE SMC 6th Workshop on Brain-Machine Interface Systems, at the Systems, Man, and Cybernetics Annual Conference in Budapest.

Daniel, C.; van Hoof, H.; Neumann, G..; **Peters, Jan**: Best Student Paper Award of ECML-PKDD 2016 for their paper: "Probabilistic Inference for Determining Options in Reinforcement Learning".

**Jonas Peters**: elected New Member of the German "Young Academy".

**Bernhard Schölkopf**: has been elected member of the Leopoldina (German National Academy of Science).

## 2.5 Director profile: Bernhard Schölkopf

Bernhard Schölkopf studied Physics, Mathematics and Philosophy in Tübingen and London. In 1994 he joined Bell Labs to work on a Ph.D. with Vladimir Vapnik. Following researcher positions at GMD, Microsoft Research, and a biotech startup, Schölkopf started his lab at the Max Planck Institute for Biological Cybernetics (Tübingen) in 2002. In 2011, he became a founding director of the Max Planck Institute for Intelligent Systems.

Bernhard Schölkopf has been program chair of NIPS and COLT and is currently co-editor-in-chief of the flagship journal in machine learning (JMLR). He has been elected to the boards of the NIPS foundation and of the International Machine Learning Society. With Alex Smola, he initiated the *Machine Learning Summer Schools* series in 2002, which has meanwhile been organized, by various teams, 35 times. Many of his past students and postdocs have gone into academia (around 30 tenured or tenure-track positions) as well as to R&D labs (around 25 tenured), and he is one of the most highly cited researchers in Computer Science worldwide.[5]

### Current Appointments

| | |
|---|---|
| 2019–present | Affiliated Professor, ETH Zürich |
| 2018–present | Managing Director, MPI for Intelligent Systems, Stuttgart & Tübingen |
| 2017–present | Distinguished Amazon Scholar |
| 2015–present | Co-Director of the Max Planck ETH Center for Learning Systems |
| 2015–present | Member of the Max Planck Campus Triumvirate, Tübingen |
| 2011–present | Director, Max Planck Institute for Intelligent Systems |
| 2010–present | Honorary professor, University of Tübingen, Faculty of Science |
| 2002–present | Honorary professor, Technical University Berlin, Dept. of Computer Science |

### Awards & Honors (2009 – 2018)

| | |
|---|---|
| 2018 | Hector Science Award |
| 2018 | Landesforschungspreis Baden-Württemberg |
| 2018 | Gottfried Wilhelm Leibniz Prize |
| 2017 | Fellow of the Association for Computing Machinery (ACM) |
| 2016 | Member of the German National Academy of Science (Leopoldina) |
| 2015 | Overseas Visiting Scholarship, St. John's College, Cambridge, UK |

[5] http://web.cs.ucla.edu/~palsberg/h-number.html

| | |
|---|---|
| 2014 | Royal Society Milner Award, London, UK |
| 2013 | XXVIIIth Courant Lectures, New York University |
| 2012 | Academy Prize, Berlin-Brandenburg Academy of Sciences and Humanities |
| 2011 | Posner keynote lecturer at the NIPS Conference, Granada, Spain |
| 2011 | Max Planck Research Award, Max Planck Society |
| 2011 | Annual BCI Research Award (with Moritz Grosse-Wentrup) |
| 2010 | Inclusion in the list of ISI Highly Cited Researchers |

### Current Memberships

Machine Learning in Science, Cluster of Excellence, University of Tübingen; Image-Guided and Functionally Instructed Tumor Therapies, Cluster of Excellence, University of Tübingen; Data-integrated Simulation Science, Cluster of Excellence, Stuttgart University; Tübingen AI Center, Competence Center for Machine Learning; Bernstein Center for Computational Neuroscience (BCCN, Tübingen); German Association for Pattern Recognition (DAGM); Deutsche Mathematiker Vereinigung (DMV); Association for Computing Machinery (ACM); Institute of Electrical and Electronics Engineers (IEEE, Senior Member); European Academy of Sciences and Arts.

### Committees, Service, Board Memberships

Co-founder of DALI – Data, Learning, and Inference (2015 –); Core Committee Member, MPI for Biological Cybernetics (2015 –); The Future of AI – A New York University Symposium on Science, Technology, Reason and Ethics (2016); ACM Heidelberg Laureate Forum Committee (2014–), Section Panel for Mathematics in Science and Technology, International Congress of Mathematicians (ICM) 2014; General Chair of the International Conference on Artificial Intelligence and Statistics (AISTATS) 2012; Initial Training Network for Machine Learning for Personalized Medicine (MLPM); PASCAL/PASCAL2 EU Network of Excellence; Forum Scientiarum at the University of Tübingen; Snowbird Learning workshop; Machine Learning Summer Schools (MLSS); International Machine Learning Society (IMLS); Neural Information Processing Systems Foundation (NIPS).

### Review Panels

Member of the pool of experts in the process to establish the Alan Turing Institute; Centre for Doctoral Training (CDT) in Data Science, Edinburgh, UK; Italian Institute of Technology (IIT); Computer Science Department, École Normale Supérieure, Paris; Neural Computation and Adaptive Perception Program of the Canadian Institute of Advanced Research (Chair); Gatsby Computational Neuroscience Unit, Quinquennial Review 2010; Machine Learning Program, NICTA (Sydney).

### Editorial Board

Journal of Machine Learning Research (JMLR, Co-editor-in-chief); International Journal of Computer Vision (2004 – 2010); Foundations and Trends in Machine Learning (since 2007); SIAM Journal on Imaging Sciences (2007 – 2012); Series *ACM Books*, Association for Computing Machinery.

### Invited Talks 2016 – 2018 (Selection)

International Conference on Learning Representations (ICLR), Vancouver (2018); SIAM International Conference on Data Mining (SDM), San Diego (2018); Asian Conference on Machine Learning (ACML), Seoul (2017); International Conference on Machine Learning (ICML), Sydney (2017); Conditional Independence Structures and Extremes, München (2016); Machine Learning for Signal Processing (MLSP); Vietri sul Mare (2016); ARES, Salzburg (2016);

# 3 PERCEIVING SYSTEMS

## 3.1 Research Overview

Computer vision is often treated as a problem of pattern recognition, 3D reconstruction, or image processing. While these all play supporting roles, our view is that the goal of computer vision is to *infer what is not in the picture* – to recognize the unseen. This is different from the Aristotelian view that the goal of vision is "to know what is where by looking." We see vision as the process of inferring the causes and motivations behind the images that we observe; that is, we want to infer the story behind the picture.

The most interesting stories involve people. Consequently, our research focuses on understanding humans and their actions in the world. We aim to recover human behavior in detail, including human-human interactions, and human interactions with the environment.

Humans interact with each other and manipulate the world through their bodies, faces, hands and speech. If computers are to understand humans and our behavior, then they are going to have to understand much more about us than they currently do. For example, they need to recognize when we are picking up something heavy and might need help. They need to understand when we are distracted. They need to understand that changes in our behavior may signal medical or psychological problems.

To address this, we are developing the datasets, tools, models, and algorithms to recover human movement in unconstrained scenes at a level not previously possible. From single images or videos, we estimate full 3D body pose, including the motion of the face and the pose of the hands. We also recover the 3D structure of the world, its motion, and the objects in it so that human movement can be placed *in context.* We are not just interested in pose but also what the person is in contact with, what they are holding, where they are looking, who they are interacting with, and what they may do next.

This is quite different from previous work in which the human body is treated in isolation, removed from the world around it, and 3D scene analysis happens on static scenes without humans. We think the interesting research problems involve analyzing human behavior when people are present in, and interacting with, the 3D world. By building 3D models of people and how they move, we are able to place them in context and reason about the goals behind their behavior and the physical contraints on this behavior.

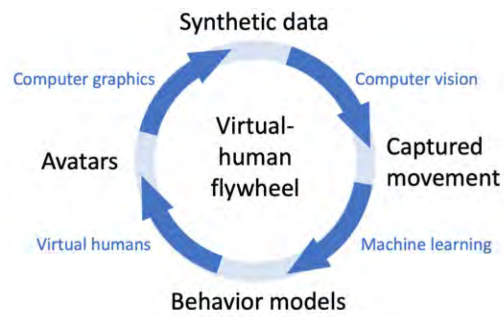


Figure 3.1: The virtual-human flywheel. By building better models of humans, we can simulate better training data, which leads to better vision algorithms, that help us gather more data about how humans behave, which helps us build better models of humans. This touches the core competencies of Perceiving Systems: Computer Vision, Machine Learning, Computer Graphics, and Virtual Humans.

To advance this agenda, Perceiving Systems combines computer vision with machine learning and computer graphics. For example, our 3D graphics models of the body enable us to generate training data for machine learning methods, which improve our computer vision algorithms [362, 378]. These improved algorithms give us better data from images and video with which to improve our graphics models, leading to a virtuous cycle as illustrated in Fig. 3.1.

This cycle is producing better and better virtual humans. We see the virtual human as more than a useful artifact. We see it as a testbed for evaluating our models of human behavior. If we can simulate a virtual human in a virtual world behaving in ways that are indistinguishable from a real human, then we assert that we have captured something essential about what it means to be human. This forces us to go beyond capturing human movement and to model the causes of that movement.

Over the history of Perceiving Systems, we have built the foundational technology for this effort. Specifically, we learn realistic 3D models of human body shape and pose deformation from thousands of detailed 3D scans. We have built many models, but SMPL[1] has become the de facto standard for research on human pose. To go beyond SMPL, we have learned a face model (FLAME) using a novel dataset of 4D facial sequences [338]. FLAME captures realistic 3D head shape, jaw articulation, eye movement, blinking, and facial expressions. Similarly, we developed MANO, a 3D hand model learned from around 2000 hand scans of different people in many poses [337].

In the last three years we have shown how to fit SMPL to image data and how to train deep networks end-to-end to extract full-body shape and pose from single images or video. This includes the following methods, which provide foundational tools for capturing and analyzing human motion in natural settings

- SMPLify [393],
- Unite the People [377],
- Human Mesh Recovery [366],
- Neural Body Fitting [356].

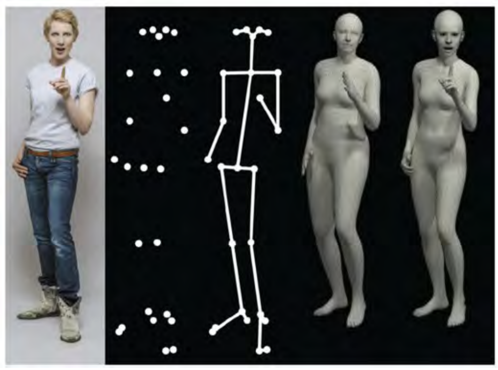

Figure 3.2: Bodies are not a collection of joints and bodies are not a skeleton. Bodies have shape, can move, can express emotion, and can interact with the world. Hence virtual bodies need faces and hands and the ability to move and use them. Our new SMPL-HF model (right) has the expressiveness needed to model human interactions with the world and between people.

### 3.1.1 Expressive Bodies

SMPL, alone, however is not enough to understand human behavior. Thus, we have combined SMPL, FLAME and MANO into a single model,

[1] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, M. J. Black. SMPL: A Skinned Multi-Person Linear Model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)* **34** (6): 248:1–248:16, Oct. 2015

SMPL-HF, that goes beyond previous work to represent 3D body shape, pose deformations, facial expression, and hand pose in a unified model. Figure 3.2 illustrates how SMPL-HF is more expressive than previous representations of the body that are commonly used today. With SMPL-HF, for the first time, we can estimate information about the body together with hand-object interaction, gestures, and facial expression. We have developed algorithms to estimate the parameters of SMPL-HF from a single image and are working on extending these methods to video and RGB-D. This is the first step towards expressive motion capture in complex scenes that will underpin our future research on human behavior.

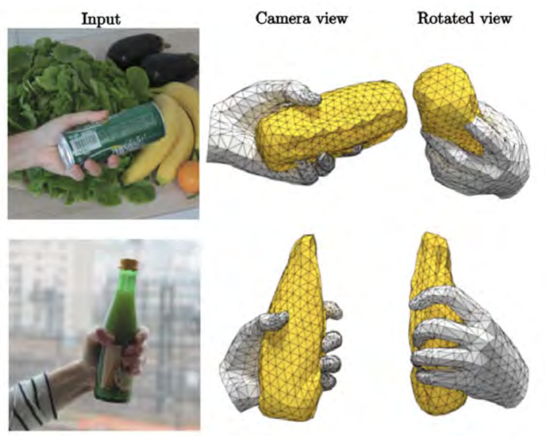


Figure 3.3: Hand pose and object shape are often estimated separately. These methods do not generalize to hand-object interaction. We jointly estimate both the 3D object and the 3D hand from a single image. This joint estimation enables us to penalize interpenetration and encourage contact, leading to more stable grasps.

A key step in this direction is to analyze hand-object interaction in images and video. To that end, in collaboration with colleagues at INRIA, we have generated a novel synthetic training set of hands interacting with 3D objects; this is a good example of the virtual-human flywheel in action. Using this data, we train a novel neural network to estimate both 3D hand pose and 3D object shape. We observe that these two processes are synergistic and that estimating them together produces better results because occlusion and contact can be modeled. Specifically we train the model in such a way that we penalize hand-object interpenetration and encourage contact when parts of the hand are close to an object surface. Despite training only on synthetic data, we obtain realistic 3D hand pose/shape and object shape as seen in Fig. 3.3. We show that, by incorporating constraints about hand-object interaction during training, we achieve more stable grasps than when training separate hand and object networks.

Figure 3.4: Without 3D supervision, RingNet learns a mapping from the pixels of a single image to the 3D facial parameters of the FLAME model. Top: Images are from the CelebA dataset. Bottom: estimated shape, pose and expression.

Like hands, faces and facial expressions are critical to understanding human behavior. The estimation of 3D face shape from a single image must be robust to variations in lighting, head pose, expression, facial hair, makeup, and occlusions. Robustness requires a large training set of in-the-wild images, which by construction, lack ground truth 3D shape. Consequently, to train a network without any 2D-to-3D supervision, we developed *RingNet*, which learns to compute 3D face shape from a single image (Fig. 3.4). Our key observation is that an individual's face shape is constant across images, regardless of expression, pose, lighting, etc. RingNet uses a novel loss that encourages the face shape to be similar when the identity is the same and different for different people. We achieve invariance to expression by representing the face using our FLAME face model. Once trained, our method takes a single image and outputs the parameters of FLAME, which can support the analysis of human behavior.

The above models capture the surface shape of the body and how it varies across people and with pose. To generalize to settings that we have never seen, we learn a physics-based model of soft-tissue. We extend SMPL from a triangulated mesh model to a volumetric tetrahedral mesh. From 4D scans, we infer the material properties and thickness of the fat under the skin. We can then simulate soft-tissue dynamics and compression using finite-element methods [342].

Since humans wear a wide variety of clothing, we also develop models of clothing and how clothing drapes on the body. Specifically, ClothCap exploits 4D clothing scans to learn how clothing deforms with pose [341]. This enables us to retarget clothing from one person to people of other shapes.

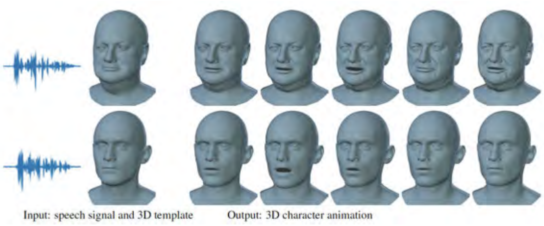


Figure 3.5: We learn a model (VOCA) that relates 3D facial deformations to speech signals. Given an arbitrary speech signal and a static 3D face mesh as input (left), VOCA outputs a realistic 3D character animation (right).

The field has focused on capturing the kinematics of the human body or the parameters of facial muscle activations. These have proven useful for many applications and can be used to animate models of the body. But, we are intersted in more. We want more semantic controls of human activity that relate goals and kinematics in context. As a first step in this direction we have captured an extensive database of 4D scans of people talking. We record both the facial shape and the speech and then train a neural network to relate the two. Using this method (VOCA), we can then animate any 3D face shape saying anything in any language and we can do so in a variety of speaking styles (Fig. 3.5). This is a step towards non-kinematic character control and opens up many avenues for the joint analysis of speech (and associated semantics, sentiment, emotion, etc.) and human behavior.

### 3.1.2 Bodies in Scenes

Humans and animals live in, and interact with, the 3D world around them. To understand humans then, we must understand the surfaces that support them and the objects with which they interact. To that end, we develop methods to estimate the structure and motion of the world from a single image, video, or multiple images. We approach this by combining unsupervised learning with physical knowlege of the world.

For moving scenes, we compute the optical flow representing the projection of the 3D motion field into the image. In doing so, we exploit the geometric structure of the problem to simplify it. If the scene is rigid and only the camera moves, then the optical flow is completely described by the depth of the scene and the camera motion. Real scenes, however, contain rigid structure and independently moving objects. To deal with this, we segment the scene into regions corresponding to the different types of motion [372]. To do so we exploit different constraints that are both geometric and semantic.

In the latter case, we know that certain objects like animals and cars can move independently while others, like buildings, cannot. Additionally objects like the road are typically planar and hence their motion is simply modeled. Thus, a semantic segmentation of the scene, provides information about what motions may be present where [400]. We argue that segmentation and motion estimation go hand in hand and we have explored methods to do both in a coupled way.

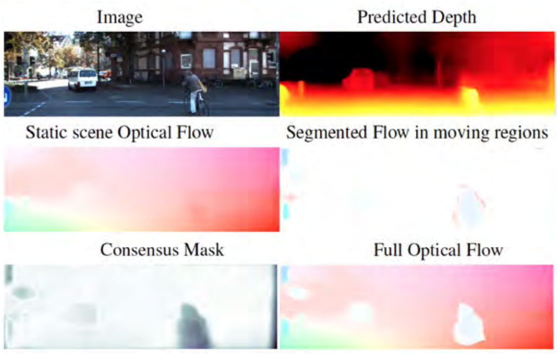


Figure 3.6: Unsupervised learning of depth, camera motion, motion segmentation (mask), and optical flow. Depth and camera motion gives rise to the flow in rigid regions. The networks learn to segment the non-rigid regions and compute flow in these. Everything is learned without supervision using a novel “collaborative competition” method.

Our most recent work focuses on the unsupervised learning of motion, scene depth, camera motion, and segmentation (Fig. 3.6). Training exploits *competitive collaboration* in which different neural networks vie to explain the motion in the scene [333]. To make this work, the physical constraints about motion in rigid scenes are critical. This geometric information allows us to learn depth from a single image without supervision. Motion makes this possible because scene structure is constant over time, allowing motion over time to inform the network about scene depth.

This work demonstrates that classical physical and geometric constraints that we already know about the world and its motion are not only compatible with deep learning but can play a critical role in enabling unsupervised learning. In a sense, the physics of the world provides a form of pre-existing supervision.

Our ongoing work combines people, scenes, and unsupervised learning to derive coherent explanations of the 3D world. We posit that the joint estimation of people and scenes will improve both.

### 3.1.3 Beyond Mocap

To understand human behavior, we must record that behavior across long periods of time, record interactions between people, and capture how people interact with their environment. Motion capture in lab environments is highly accurate but, necessarily, restricted in terms of realism and the length of sequences that can be captured. Vision-based methods for monocular mocap are designed to work in arbitrary scenes but still lack accuracy. No current method enables accurate capture of human motion in natural settings over long periods of time. Consequently, we are developing a range of motion capture technologies that can break free of the laboratory and capture more realistic behavior. As a first step, we are developing capture technologies that use IMUs, hand-held cameras, or swarms of drones to enable 3D human motion estimation in natural settings like outdoors, around town, or at home (Fig. 3.7).

Figure 3.7: To take mocap out of the lab we use a small number of IMUs mounted on the body. Top: We train a neural network to regress full body pose from 6 IMUs. Bottom: We combine IMUs with a hand-held camera to obtain 3D poses in natural videos.

The above methods, models, and datasets support our long-term goal of teaching computers to see us and, understand our behavior, and to mimic this behavior in virtual human avatars.

### 3.1.4 Impact and Outreach

While our focus is basic research on computer vision, graphics and machine learning, we want to have an impact beyond these academic disciplines. Consequently, we pursue collaborations that allow our work to have a broader impact.

Figure 3.8: Our body models are now used for problems in psychology and medicine. Our SMPL body model captures realistic body shape and can easily run in MR/VR applications. We have recently developed simple methods where clinicians can create avatars using only the controllers of an HTC Vive.

For example we develop applications in medicine and psychology in collaboration with medical colleagues. We have collaborative efforts to relate the distribution of adipose tissue in the body to the risk of diabetes and cardiovascular disease [346]. We have developed a 3D model of infants and use it to track their movement to aid the early diagnosis of cerebral palsy [364]. Our 3D body model has also played an important role in understanding how women who suffer from anorexia nervosa see their body and the bodies of others [334–336]. We continue to collaborate with psychologists and doctors on a range of related topics and, through these collaborations, have developed tools like the Virtual Caliper (Fig. 3.8), which makes it easy for practitioners to create realistic 3D body avatars and animate them in virtual reality (VR) or mixed reality (MR).

More information: https://ps.is.mpg.de/field/medicine-and-psychology

Figure 3.9: From a few snapshots of an animal, we reconstruct the detailed 3D shape. This can then be animated or 3D printed.

We are also collaborating with researchers on animal conservation. Specifically we work with Wildbook on a project to protect the Grevy Zebras in Africa. We have developed technology that takes a few images of an animal and creates realistic 3D models (Fig. 3.9). With Wildbook, we are working on the first methods to analyze herd shape from photos and on developing drone-based surveying methods.

More information: https://ps.is.mpg.de/project/capturing-animal-shape

We are also active in patenting, technology licensing, and startups. We have spun off two companies that are using our 3D body model technology. One of these, **Body Labs Inc.**, was acquired by Amazon in 2017. The second, **Meshcapade GmbH**, started in 2018 and provides services and software for processing 3D scans and motion capture data.

We also make code and data available open source or for license. For example, over the last three years, our open source differentiable renderer (OpenDR) had a significant impact and has kick-started research on this topic. Additionally our SMPL body model is now in wide use in both academia and industry for representing 3D body shape. Our code for estimating SMPL from images (SMPLify) has helped drive the field to solve this challenging problem.

Finally we are responsible for, or contribute to, widely used datasets and evaluation benchmarks that help push the state of the art and provide a platform for industry to understand what works, how well, and why. We have played central roles in many influential datasets and evaluations in the field including Middlebury Flow, Sintel, KITTI, HumanEva, FAUST, JHMDB, and others. Over the last three years we have released several major datasets related to faces, hands, 3D bodies, clothing, animals, optical flow, and IMUs.

Datasets and code released in the last three years include: SMPL[1]; SMPLify [393]; FLAME [338] (model, fitting code, and registered meshes); CoMA [360] (code, model and data); MANO [337] (hand scans and model); Dynamic FAUST [381] (precise 4D scans in correspondence); 3D poses in the wild (3DPW) [358]; SURREAL [378] (synthetic humans for training deep networks); SlowFlow [380] (optical flow in real scenes); Unite the People [377] (training set for 3D human pose from images); BUFF [388] (body shape under clothing); SMAL [383] and SMALR [367] models of 3D animals; SMIL [364] infant body model; DIP [330] (code and data for 3D human pose from IMUs); AirCap [332, 339] (design and software for aerial motion capture).

More information: https://ps.is.mpg.de/code

### 3.1.5 About Us

The Perceiving Systems department was founded in January 2011 and today has about 50 members from all over the world. This includes support staff, technicians, students, guests, and scientists at various career stages. We have about 60 alumni including eight graduated Ph.D. students. Many of our alumni have gone on to academic positions, founded companies, or joined major research laboratories.

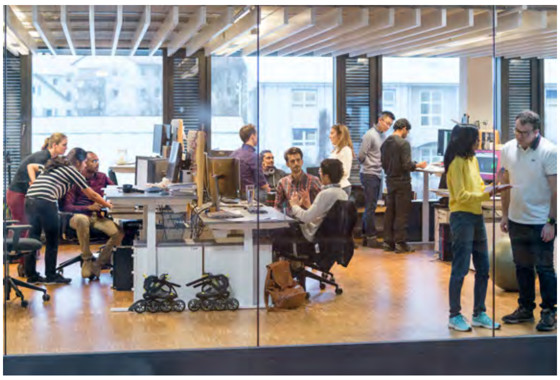

Figure 3.10: Perceiving Systems is highly international, diverse, and collaborative.

We take diversity seriously and, between 2016 and 2018, 30% of our papers had female first authors and 50% had at least one female author.

The department hosts two group leaders, Siyu Tang and Aamir Ahmad, who lead groups focused on “Holistic Vision” and “Robot Perception" respectively. Group leaders receive department funding and raise external funds to support their research. Details about these groups can be found on our website: https://ps.is.mpg.de/field/robot-perception-group and https://ps.is.mpg.de/field/holistic-vision-group

We have regular sabbatical and long-term visitors. Between 2016 and 2018 we hosted Cordelia Schmid (INRIA) as a Humboldt Professor, Andrew Blake (Samsung and FiveAI), and Hedvig Kjellström (KTH) as a sabbatical visitor. We also have a highly active visitor program and lecture series. We have had over 175 invited speakers, including many of the leaders in the field. A full list is here https://ps.is.tuebingen.mpg.de/talks

In what follows, we present a sampling of our research projects over the last three years. Our website provides information about many other projects as well as greater detail: https://ps.is.tue.mpg.de

A broader view of the department activities, including more of the social life, can be found on our Facebook page: https://www.facebook.com/PerceivingSystems/

### 3.1.6 Facilities

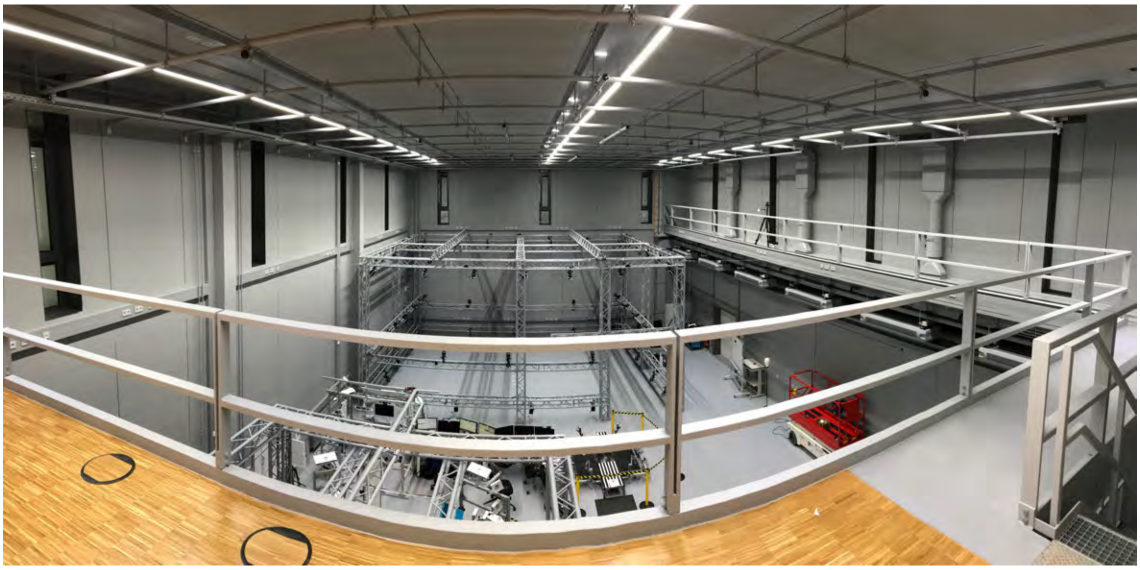

Two years ago, Perceiving Systems moved into the new building and we finally had laboratory space for our research on human shape and motion. In our Capture Hall we run a large and complex range of equipment and experiments in approximately 830 square meters of space. During 2016–2018 we captured $2,334,643$ 3D scans broken down into $938,017$ full body scans, $1,082,752$ face scans, $236,225$ hand scans, and $77,649$ foot scans. We know of no scanning effort of this size elsewhere in the world. Keeping this running is a professional staff of two human subjects coordinators, two scanner technicians, and three software engineers who support the custom software that processes all this data.

Figure 3.11: The 4D body scanner captures 3D meshes of the full body at 60 fps.

The centerpiece of this facility is our 4D body scanner (Fig. 3.11) made by 3dMD, which was the first of its kind that could capture the full range of human movement at 60 fps. At each time instant the system captures a 3D point cloud with about $150k$ points. We capture people both in minmal clothing and in normal street clothes. The scan data is then processed so that our 3D body model is aligned to the data using our own methods to produce detailed meshes that are in correspondence across people and poses.

Figure 3.12: The 4D face and hand scanner captures detailed facial or hand shape at 60 fps together with high-quality texture.

The full body scanner has limited resolution in the face region and, given the importance of the face for communication we purchased a dedicated 3dMD system that can capture the full 3D head in detail at 60 fps together with high-quality texture maps (Fig. 3.12). We also reconfigure this system to perform hand scanning but plan (in collaboration with the Haptic Intelligence department) to install a dedicated hand system as well to study hand-object interaction.

Figure 3.13: The 4D foot scanner captures a full 3D view of the foot during contact. The system images the bottom of the foot through a glass plate so that the deformation of the shape is captured. It is able to capture 4D sequences, revealing how feet interact with surfaces.

Feet receive much less attention than faces and hands but are literally the foundation of human movement. Many problems with the capture and animation of human avatars can be traced to the feet, which are typically approximated as rigid shapes. In fact, the feet are complex and highly deformable. If one wants to model the physical interactions between the body and the world, one needs a detailed foot model. Consequently, we have been capturing feet interacting with the ground, both in shoes and barefoot using a new 4D foot scanner (Fig. 3.13) created by 3dMD. We plan to add a detailed foot model to our 3D human model in the near future.

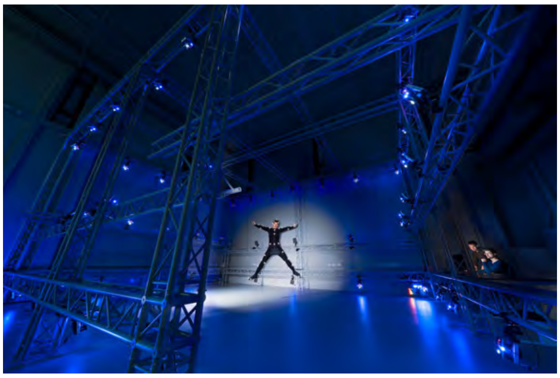

Figure 3.14: We recently added a 54-camera Vicon motion capture system to capture multiple people and complex motions.

While our 4D systems give unprecedented details of the human body in motion, the capture volume of these systems is very limited and this limits the kinds of motions we can study. In particular, we can only capture information about one person at a time. To study human interaction and more complex behaviors, we have installed a Vicon marker-based motion capture system with 54 high-resolution Vantage V16 cameras (Fig. 3.14). This system enables us to capture the motion of multiple people interacting, including the full body together with the face and hands. The high density of cameras minimizes problems with occlusion and the high resolution of the cameras means that we can resolve very small markers on the face and hands.

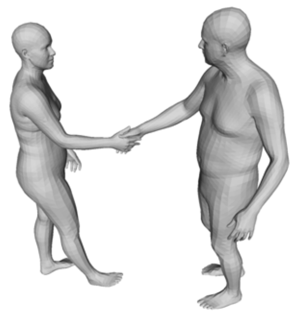

Figure 3.15: Capturing interaction. We are extending all of our capture and modeling methods to deal with human-human and human-object interaction. Here MoSh is extended to extract multiple people interacting and in contact. This is the beginning of a major effort to accurately capture and model contact and behavior.

[2] M. M. Loper, N. Mahmood, M. J. Black. MoSh: Motion and Shape Capture from Sparse Markers. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* **33** (6): 220:1–220:13, Nov. 2014

While such a system gives sparse marker data, our MoSh technology[2] converts this into a full 3D representation of the actors and their movements. We are extending MoSh to capture hands and faces and to automatically fit multiple people (Fig. 3.15). At the same time we are designing new capture scenarios to study human communication, human-human contact, and human-object interaction.

Figure 3.16: As few as six inertial measurement units worn on the body provide enough information to infer human pose. This makes motion capture practical in many scenarios.

Going beyond the confines of the capture hall, we work extensively with IMU-based motion capture (Fig. 3.16). This includes full-body IMU capture as well as hand pose capture. While IMUs enable us to capture people in natural settings, they have many drawbacks. We have developed several methods to make them more practical by training a method to estimate body pose from only 6 IMUs and combining IMUs with a hand-held video data to eliminate drift.

Figure 3.17: We have developed custom octo-copters with on board processing and cameras. These form the basis of our flying-motion-capture system, which aims to capture human and animal motion outdoors without any worn sensors.

IMUs also require the subject to cooperate and wear sensors that could affect their movement. This becomes impractical when we want to capture animal movement in the wild. To address this, we are developing a flying motion capture system based on micro-aerial vehicles that fly autonomously, track the subject in real-time, and then estimate the 3D pose of the body over time (Fig. 3.17). Our goal is to make this practical and accurate enough to be used for animal behavior analysis in the wild.

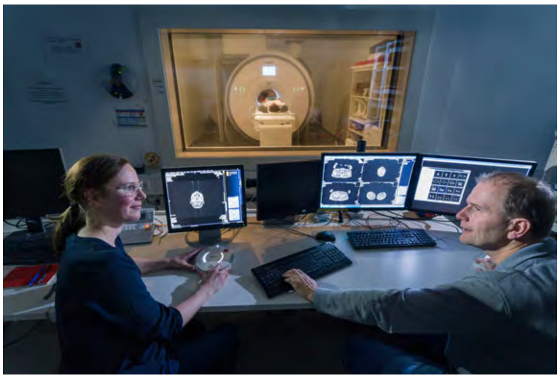

Figure 3.18: To look inside the body, we also scan people in an MRI scanner operated by the MPI for Biological Cybernetics.

Finally, the above work focuses on the outside of the body and its movement but we also capture the inside of the body using MRI scans in humans (Fig. 3.18) and CT scans in animals (rodents). We capture full-body MRI scans of human subjects lying down and 3D surface scans of them lying on a glass table. This enables us to relate the two datasets. Subjects are also scanned standing up so that we can learn how to relate such scans to the compressed shape of people lying down. Together with collaborators at the University Hospital Tübingen, we use this data to analyze the distribution of body fat and to relate this to body shape. In the case of rodents, we again collaborate with the University of Tübingen to model the relationship between the outer surface of the animal and the internal structure. The goal is to predict the motion of the bones and internal organs from only external observations.

## 3.2 Selected Research Projects

## Expressive Body Models

*Michael Black, Dimitris Tzionas, Timo Bolkart, Ahmed Osman, Vassilis Choutas, Georgios Pavlakos, Nima Ghorbani*

Figure 3.19: We learn a new 3D model of the human body called SMPL+HF that jointly models the human body, face and hands. We fit the female SMPL+HF model to single RGB images and show that it captures a rich variety of natural and expressive 3D human poses, gestures and facial expressions.

Bodies in computer vision have often been an afterthought. Human pose is often represented by 10-12 body joints in 2D or 3D. This is inspired by Johannson's moving light displays, which showed that some human actions can be recognized from the motion of the major joints of the body. But the joints do not capture everything. The skeletal structure of the body is also a popular representation but is only approximate and is never actually observed in images.

In our work we have focused on 3D body shape, represented as a triangulated mesh. Shape gives us more information about a person related to their health, age, fitness, and clothing size. But shape is also useful because our body surface is critical to our physical interactions with the world. We cannot interpenetrate objects and they cannot interpenetrate us.

It has taken a few years for the field to catch on to this idea but now our SMPL[1] body model is widely used in research and industry. It is simple, efficient, posable, and compatible with most graphics packages. It is also differentiable and easy to integrate into optimization or deep learning methods.

While popular, SMPL has drawbacks. Pose deformations are non-local, the face does not move, the hands are rigid, there is no clothing and no hair. We are addressing these issues in ongoing work (see the theme on Clothing [341, 388] and projects on Faces [338] and Hands [337]). Our latest work is putting bodies, faces and hands together in a simple model that can be fit to data or animated. Like all our body models, we train this from scans of people to capture the realism and statistics of the population [381].

Such models provide the foundation for our analysis of human movement, emotion, and behavior.

More information: https://ps.is.mpg.de/project/expressive-body-models

## Faces and Expressions

*Michael Black, Timo Bolkart, Anurag Ranjan, Soubhik Sanyal, Tianye Li, Javier Romero, Cassidy Laidlaw*

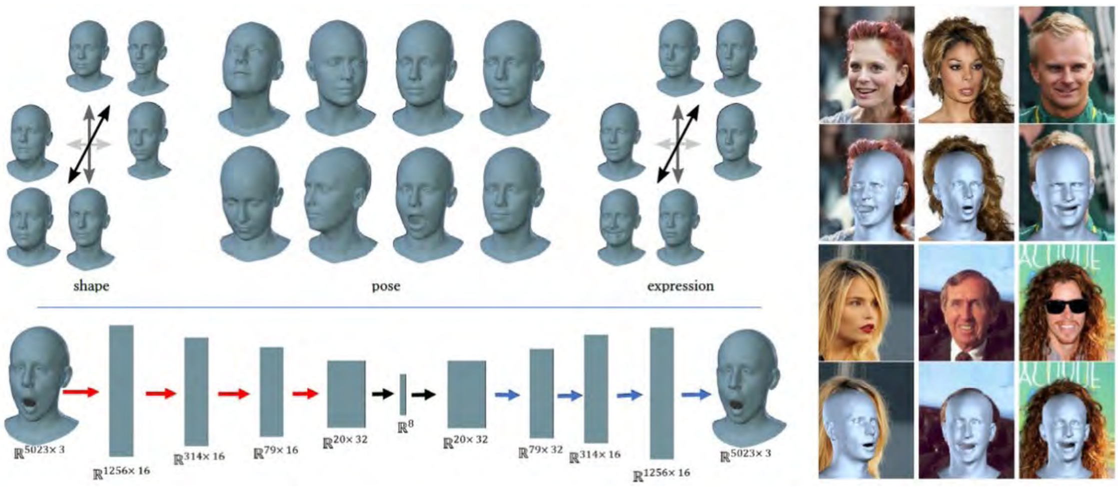


Figure 3.20: 3D Face Analysis: Top left: FLAME [338] captures face shape, pose, and expression with a linear model. Bottom left. CoMA models non-linear deformations using a novel mesh auto-encoder [360]. Right: Learning 3D shape, pose and expression from 2D images without 3D supervision.

Faces, their shape, and their motion are essential to communication. Consequently, we want a model of the face that can capture the full range of face shapes and expressions. Such a model should be realistic, easy to animate, easy to fit to data, and should support inference about human emotion and speech. Additionally we need the tools to estimate faces, their shape, pose, expression, gaze, and movement from images.

To that end, we trained a 3D face model called FLAME [338] from 4D scans. Because it is learned from large-scale, expressive, data of real people, it is more realistic than previous models. FLAME uses a linear shape space trained from 3800 scans of human heads and combines this with an articulated jaw, neck, eyeballs, pose-dependent corrective blendshapes, and additional global expression blendshapes. The pose and expression dependent articulations are learned from 4D face sequences to which we accurately register a template mesh. In total the model is trained from over 33,000 scans.

While expressive, it is difficult to capture the non-linear deformations of extreme expressions with FLAME's low-D linear subspace. While neural networks would be a natural choice for representing such deformations in a low-D latent space, existing convolutional neural networks do not generalize to 3D meshes in a straightforward way. To address this, we introduce a versatile encoder-decoder framework for meshes using spectral convolutions on a mesh surface [360]. Additionally, we introduce mesh up- and down-sampling operations that enable a hierarchical mesh representation that captures non-linear variations in shape at multiple scales. Our CoMA mesh convolution algorithm is generic and now widely used.

To capture, model, and understand facial expressions, we need to estimate the parameters of our face models from images and videos. Training a neural network to regress model parameters from image pixels is difficult because we lack paired training data of images and the true 3D face. To address this we learn this mapping using only 2D image features. The key is to leverage multiple images of a person with a novel loss that encourages the face shape to be similar when the identity is the same and different for different people. FLAME enables the network to factor out changes in expression so that it can exploit this shape constancy.

More information: https://ps.is.mpg.de/project/human-face-analysis

## Hands in Action

*Dimitris Tzionas, Javier Romero, Michael Black, Gul Varol, Cordelia Schmid, Yana Hasson, Igor Kalevatykh, Ivan Laptev*

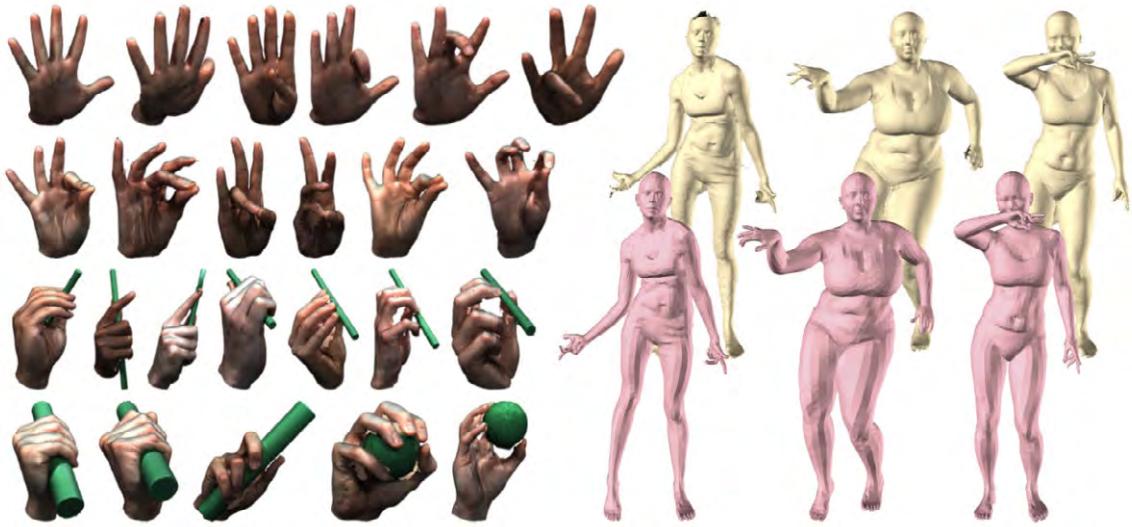

Figure 3.21: We collect 3D scans of human hands (left) from multiple people and model pose and shape variation across people and poses by learning a statistical model of the human hand, called MANO. We then combine the MANO hand model with our SMPL body model to build a holistic model called SMPL+H. The figure (right) shows example 3D scans (white) from our 4D sequences and corresponding fits of SMPL+H to these scans (pink). SMPL+H is able to capture natural motions even under challenging conditions, such as severe missing data due to fast motion, occlusion, finger-webbing, or noise.

Hands are important to humans for signaling and communication, as well as for interacting with the physical world. Capturing the motion of hands is a very challenging computer vision problem that is also highly relevant for other areas like computer graphics, human-computer interfaces, and robotics.

We focus on building an accurate and realistic model of the human hand [337] that captures the pose and shape variation across a human population. For this we collect many examples of human hands with our 3D scanner, following a systematic grasp taxonomy [348]. We then combine the hand model with our SMPL body model to build a seamless model of the body together with hands, called SMPL+H. This allows us to naturally capture the motion of people with expressive body and hand motion using our 4D scanner.

A strong hand model can be used to regularize fitting to noisy input data to reconstruct hands and/or objects [413]. We focus on hands that interact with other hands or known objects [349], using either a single RGB-D camera or multiple synchronized RGB cameras. Interaction cues can also reveal information that helps to reconstruct unknown properties of the object, like the kinematic skeleton [392].

Our current work focuses on estimating hands performing tasks from a single image or video. We use our hand model to generate synthetic training data of hand-object interaction and use deep learning to reconstruct hand-object configurations jointly from a single RGB image. By estimating the 3D hand and object shape together, we are able to reason about interactions such as proximity, contact, grasp stability, and forces while preventing interpenetration.

In collaboration with the Haptic Intelligence department we are extending our hand capture and modeling to account for the soft tissue deformation of the hand during contact and manipulation [556, 557]. This is critical for realistic physical reasoning about grasp.

More information: https://ps.is.mpg.de/project/hands-in-action

## Clothing Capture and Animation

*Gerard Pons-Moll, Sergi Pujades, Christoph Lassner, Peter Vincent Gehler, Michael Black, Sonny Hu, Chao Zhang*

Figure 3.22: (Left, top to bottom) Image-based generative model of people in clothing (ClothNet) [376]; estimating 3D human body shape under clothing (BUFF) [388]; 4D clothing capture for garment modeling, retargeting, and virtual try-on (ClothCap) [341]. (Right) Virtual try-on: a) Scan of a subject wearing clothes; b) image with a new body shape; c) estimated new body shape and pose; d) new body wearing the captured clothes.

While our detailed models of the body capture important aspects of body shape, they are missing something important – clothing. Most people appear in images in clothing and to analyze this we seek models of the body and clothing. Modeling clothing is hard, however, because of the variety of garments, varied topology of clothing, varied appearance, and the complex physical properties of cloth.

Standard methods for clothing 3D bodies rely on 2D patterns and physics simulation. These require expert knowlege, do not work for all types of clothing, and can be difficult to apply to arbitrary body shapes and poses. Consequently, we take a data-driven approach to learn the shape, movement, and appearance of clothing.

ClothNet [376] is a conditional generative model that is directly learned from images of people in clothing. Given a body silhouette, the model produces different people with similar pose and shape in different clothing styles by using a variational autoencoder, followed by an image-to-image translation network that generates the texture of the outfit.

To dress people in 3D, the minimally-clothed body shape is needed. To estimate this from clothed bodies, our BUFF method [388] estimates body shape under clothing from a sequence of 3D scans. In a scan sequence, different poses will make the clothing tight on body in different regions. All frames in a sequence are brought into an unposed canonical space and fused into a single point cloud. We optimize for the body shape to robustly fit the inside of this fused point cloud. This produces a remarkably accurate personalized body shape.

Given the underlying body shape, we can then model how clothing deviates from the body by capturing 3D and 4D scans of clothed people. ClothCap [341] is a pipeline that captures dynamic clothing on humans from 4D scans, segments the clothing from the body, segments it into pieces, and models how the clothing deviates from the body. ClothCap can then retarget the captured clothing to different body shapes paving the way towards virtual clothing try-on.

More information: https://ps.is.mpg.de/project/clothing

## Physics of Body Shape and Motion

*Michael Black, Sergi Pujades, Meekyoung Kim, Gerard Pons-Moll, Javier Romero, Ludovic Righetti*

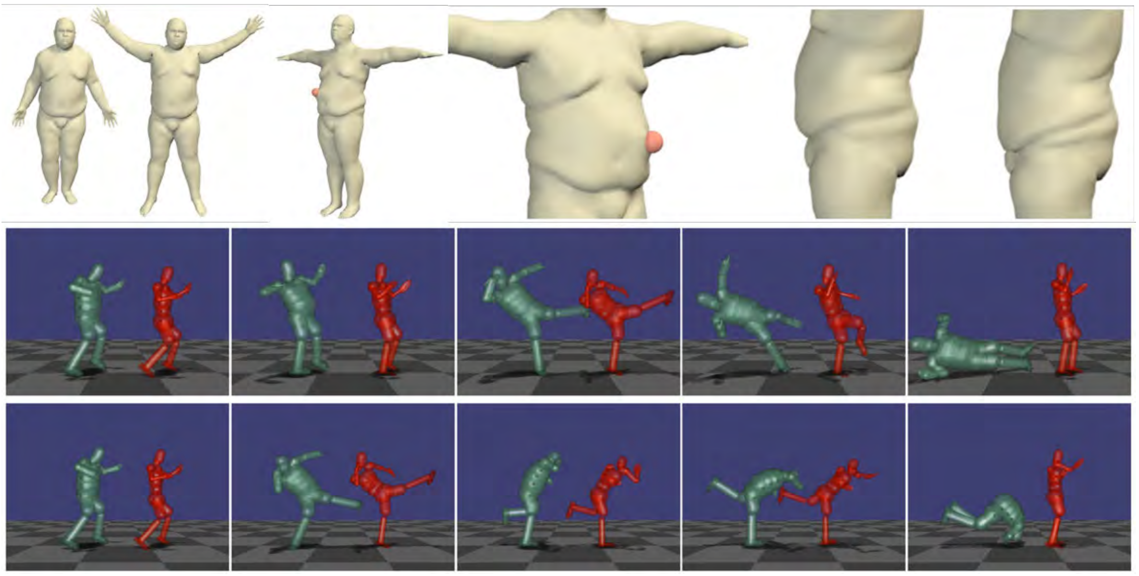

Figure 3.23: Top: We create avatars reproducing the soft tissue motions seen on real humans. We show how such avatars generalize to external forces (applied with the red sphere), and how they deform if gravity is 7 times higher (far right). Bottom: The red avatar performs a motion. A green heavier avatar tries to reproduce the same motion, struggling to raise the leg and loosing balance afterwards.

Humans live in a real world governed by the laws of physics; that is, we apply and exploit forces, such as gravity, in our daily interactions with the world. In this project we allow virtual humans to interact with a virtual world subject to the laws of physics. How would one's body shape deform in case of a collision with an object? How would our walk pattern look like if we weighed a few kilos more?

To model soft-tissue dynamics, we learn a layered volumetric body model from data [342]. To enable this we extend the triangulated mesh of the SMPL body model with a volumetric tetrahedral model called VSMPL. VSMPL contains an inner "rigid" layer and an outer soft-tissue layer. Given 4D sequences of people in motion, we learn the physical properties (Young's modulus) of the outer tetrahedra. We do this such that, when simulated in motion using a finite element method, the surface motion of the VSMPL model resembles the observations. The learned model is a realistic full-body avatar that generalizes to novel motions and external forces.

We also address the problem of retargeting the captured motion of one person onto a different person with a different body shape and physical properties (e.g. taller, heavier, thinner) such that the new morphology is taken into account [331]. We obtain visually plausible simulations using a simplified representation of human body shape that we animate using physically-based retargeting. We develop a novel spacetime optimization approach that learns and robustly adapts physical controllers to new bodies and constraints. The method automatically adapts the motion to a subject with the novel target body shape, respecting the physical properties, and producing an appropriate movement. This makes it easy to create a varied set of motions from a single mocap sequence by simply varying the characters.

More information: https://ps.is.mpg.de/project/physics-shape-motion

## Capturing Animal Shape

*Silvia Zuffi, Angjoo Kanazawa, Michael Black*

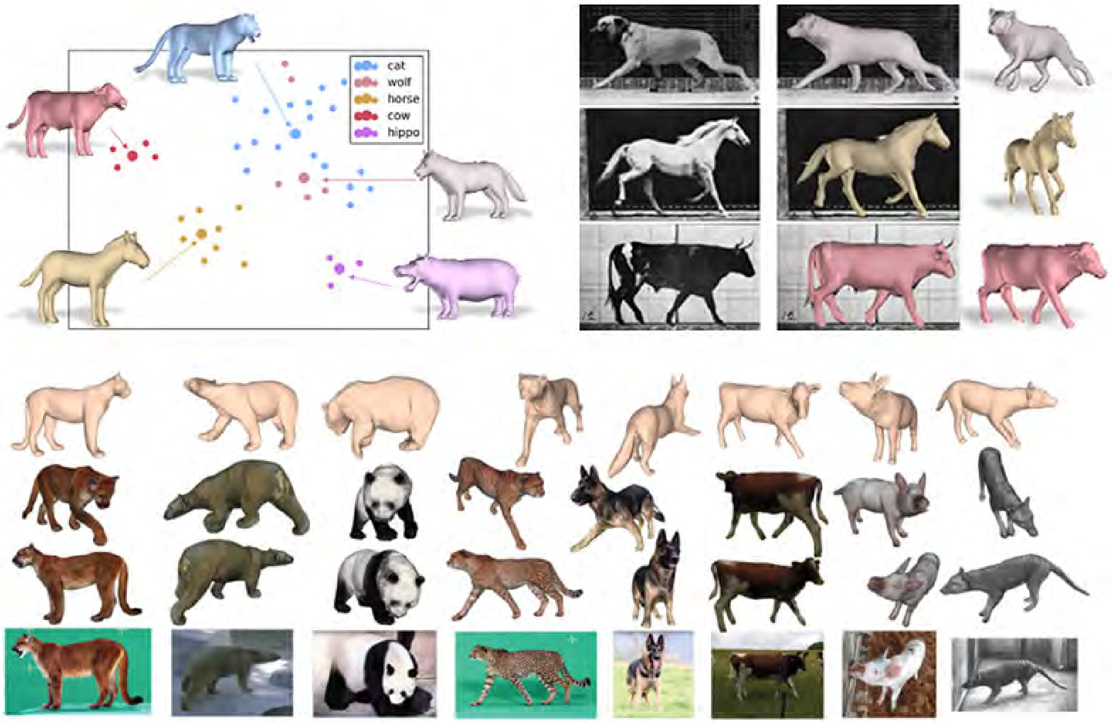


Figure 3.24: We learn an articulated, 3D, statistical shape model of animals (SMAL) that can represent quadrupeds of different species using very little training data (top). We fit SMAL to a set of uncalibrated images estimating pose, shape and vertex displacements and recover 3D textured meshes for a wide range of species (bottom).

In the past 15 years impressive advances have been made in capturing, modeling and tracking the human body. Animals have received much less attention, despite many applications in biomechanics, biology, neuroscience, robotics, and entertainment. The main reason for the lack of 3D animal models is that the methods for modeling the human body cannot be easily applied to animals: animals are not cooperative, cannot be brought to the lab in large numbers, and current scanners cannot be taken into the wild. Additionally they vary significantly in shape and even in the type of body parts they have.

In this project we develop methods to learn 3D articulated statistical shape models that can represent a wide variety of species in the animal kingdom, allowing intra- and inter-species analysis of 3D shape and the automatic and non-invasive assessment of animal shape from images.

From scans of toy animals, we learn the SMAL (Skinned Multi Animal Linear) model [383], a 3D articulated statistical shape model able to represent animal shapes for different species: big cats, dogs, cows, horses, zebras, and hippos. To capture animals outside the SMAL space, we developed SMALR (SMAL with Refinement) [383]. SMALR estimates a detailed 3D textured mesh using a small set of uncalibrated, non-simultaneous images of the animal.

Today animal motion is mostly captured indoors for domestic species with marker-based systems. To address this we are exploiting our 3D articulated animal shape models to develop a markerless motion capture system that will capture the shape and articulated motion of wild animals in their natural environment.

More information: https://ps.is.mpg.de/project/capturing-animal-shape

## 3D Body Shape and Pose from Images

*Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Vincent Gehler, Javier Romero, Michael Black, Gerard Pons-Moll, Martin Kiefel, Yinghao Huang, Ijaz Akhter*

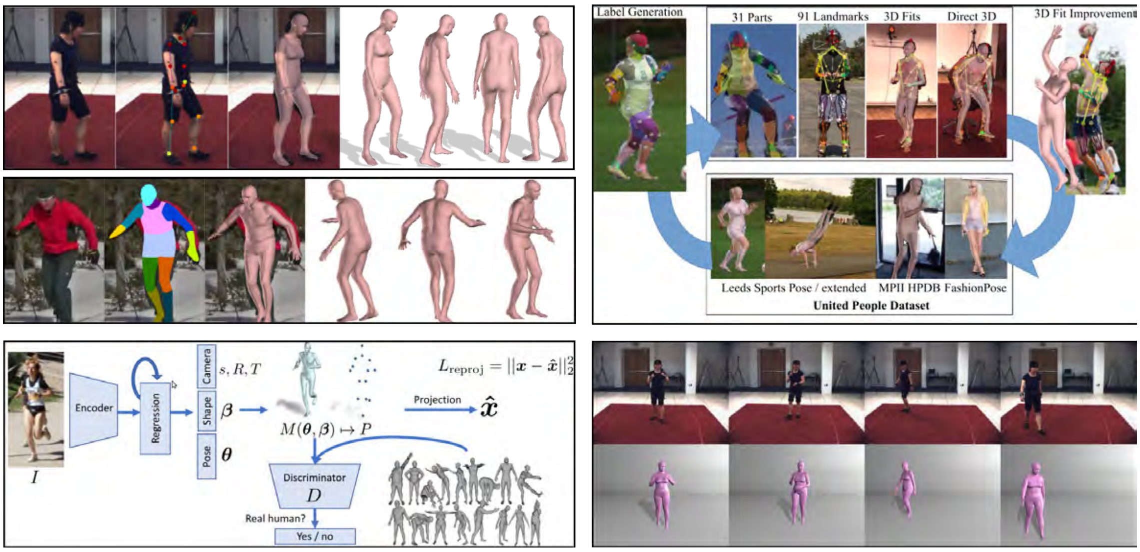


Figure 3.25: Overview of the methods and sample results. Clockwise from upper left: SMPLify [393], Unite the People [377], Multi-view SMPLify (MuVS) [369], Human Mesh Recovery (HMR) [366], Neural Body Fitting (NBF) [356]

Much of the field has focused on estimating 2D joints, 3D joints, or the skeleton of the body. We focus on estimating the full 3D shape and pose. This is crucial for reasoning about interactions. Having the ability to do so from RGB images enables markerless motion capture and provides the foundation for human behavior analysis. We explore two strategies: classical top-down model fitting and feed-forward regression.

SMPLify [393] combines bottom-up 2D feature detection with top-down 3D model fitting. The shape and 3D pose of a person are estimated by minimizing the error between the projected 3D joints of the SMPL model and 2D detected landmarks. Unite the People [377] adds a new loss term, creates a pseudo ground-truth dataset and trains discriminative models for detailed 2D landmark detection and 3D pose estimation. The whole process is repeated multiple times to refine the results and increase the quantity and quality of available data. In [369] the optimization pipeline of SMPLify is extended to handle multi-view imagery and video. Using temporal information helps resolve left/right ambiguities while giving better estimates of global orientation and body shape.

Human Mesh Recovery [366] learns to regresses the shape and the 3D pose directly from a single RGB image, by minimizing the reprojection error of 3D SMPL keypoints during training. This is not sufficient though, so we add an adversarial loss that forces the model to produce SMPL parameters that the discriminator is unable to distinguish from real ones drawn from a database of 3D human meshes. An advantage of this approach is that it can be trained without any expensive paired 2D-to-3D data.

In Neural Body Fitting [356] the shape and 3D pose parameters of SMPL are regressed from body part segmentations given by an intermediate network. Since the whole pipeline is differentiable, different types of supervision can be used, depending on the available information. Extensive experiments show that the body part segmentation is a good intermediate representation for lifting to 3D, as well as that competitive performance can be achieved with limited paired 2D-to-3D data.

More information: https://ps.is.mpg.de/project/3d-pose-and-shape-from-images

## Groups and Crowds

*Siyu Tang, Michael Black, Peter Vincent Gehler*

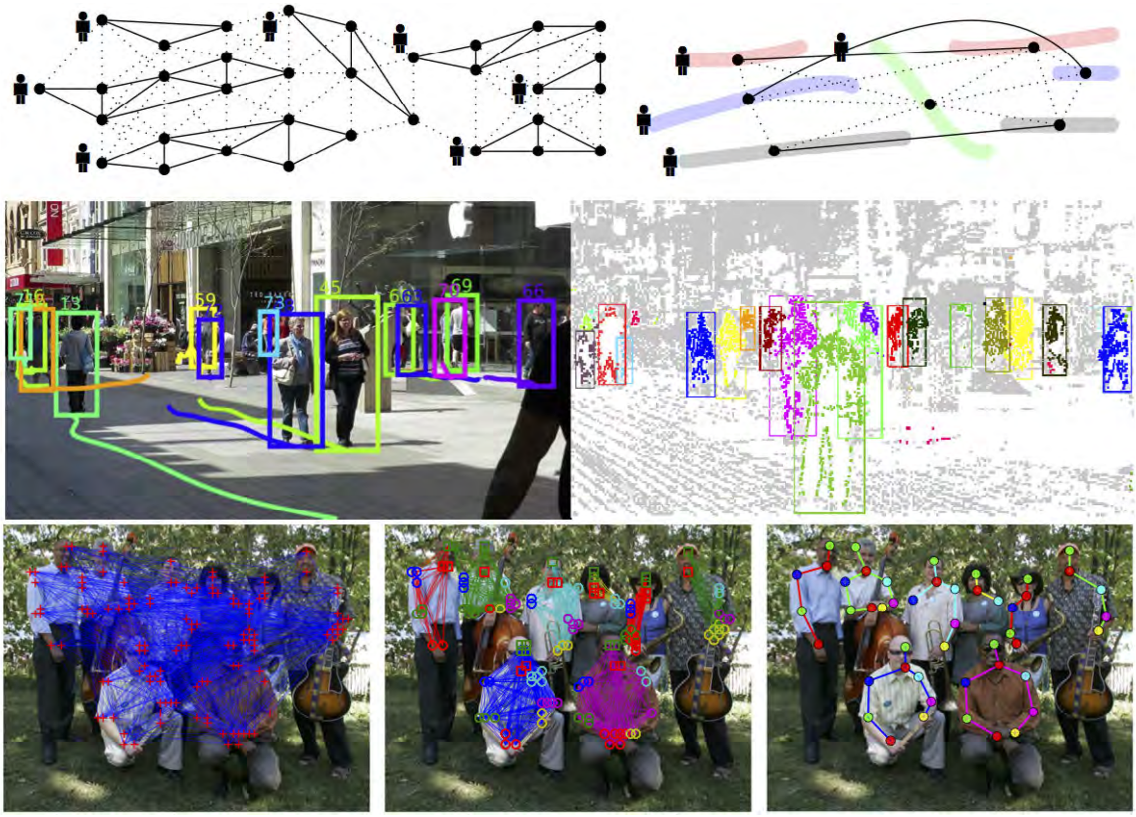

Figure 3.26: Top row illustrates the hierarchical correlation clustering formulation for multi-person tracking [353]. A dotted line indicates that the edge is a cut. The detection graph is partitioned into 7 components, indicating 7 people (top left), which are associated by the global clustering, resulting in 4 persons (top right). Middle row shows qualitative results of tracking and segmentation on the MOT16 benchmark. The solid line under each bounding box indicates the lifetime of the track. Bottom row illustrates the Deepcut model [398] for multi-person pose estimation. Initial detections (bottom left) and pairwise terms between all detections are jointly clustered and each part is labeled corresponding to its part class. Bottom right shows the predicted pose sticks.

People are often a central element of visual scenes. It has been a long-standing goal in computer vision to develop computational models that enable machines to detect crowds of people, analyze their motion and poses, infer their actions and reason about the consequences. Our research addresses a wide range of challenges in visual understanding of people in real-world crowded scenes. These include multi-person tracking [353] [371], multi-person pose estimation [398], segmentation [328] and person re-identification [365].

For multi-target tracking, our work [353] proposed to link, cluster and track targets jointly across space and time. We defined a novel mathematical abstraction for tracking in the form of a minimum cost multicut problem. In order to avoid that distinct but similar looking targets are assigned to the same track, we formulated tracking as a minimum cost lifted multicut problem [371].

Our work [365] presented a novel method to re-identify people in different images, where a second-pooling method is utilized to fuse the feature maps from the pose and the appearance estimator. The method significantly advanced the state-of-the-art on many challenging public benchmarks.

This work forms a foundation for our ongoing work on estimating detailed 3D motions of people in crowded scenes.

More information: https://ps.is.mpg.de/project/groups-and-crowds

## AirCap: 3D Motion Capture

*Aamir Ahmad, Eric Price, Nitin Saini, Guilherme Lawless, Roman Ludwig, Igor Martinovic, Michael Black*

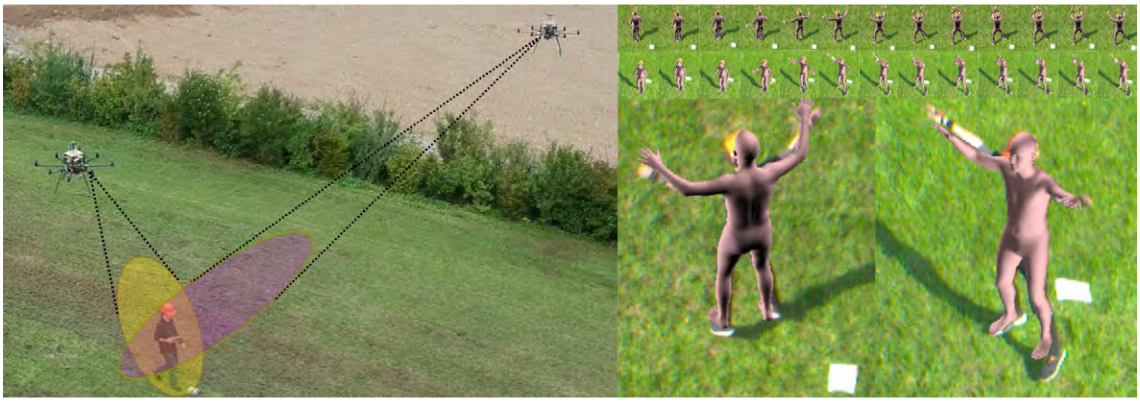

Figure 3.27: Two of our self-designed aerial robots cooperatively detecting and tracking a person on-board in real time (left). Cropped region of interests (ROIs) of images from both MAVs (right). The SMPL mesh with shape and pose estimated using our method is overlaid on all the images in a motion sequence.

Our goal is markless, unconstrained, human and animal motion capture outdoors. To that end, we are developing a flying mocap system using a team of aerial vehicles (MAVs) with only on-board, monocular RGB cameras. To realize such an outdoor motion capture system we need to address research challenges in both control and perception. In a separate ongoing project we solve the control-related challenges, with perception problem in the loop.

The perception functionality of AirCap is split into two phases, namely, i) online data acquisition, and ii) offline pose and shape estimation.

During the online data acquisition phase, the MAVs detect and track the 3D position of a subject while following them. To this end, they perform online and on-board detection using a deep neural network (DNN)-based detector. DNNs often fail at detecting small-scale objects or those that are far away from the camera, which are typical in scenarios with aerial robots. In our solution [332], the mutual world knowledge about the tracked person is jointly acquired by our multi-MAV system during cooperative person tracking. Leveraging this, our method actively selects the relevant region of interest (ROI) in images from each MAV that supplies the highest information content. Our method not only reduces the information loss incurred by down-sampling the high-res images, but also increases the chance of the tracked person being completely in the field of view (FOV) of all MAVs. The data acquired in the online data acquisition phase consists of images captured by all MAVs (see, for example, the left image above) and their estimated camera extrinsic and intrinsic parameters.

In the second phase, which is offline, human pose and shape as a function of time are estimated using only the acquired RGB images and the MAV's self-localization (the camera extrinsics). Using state-of-the-art methods like VNect and HMR, one obtains only a noisy 3D estimate of the human pose. Our approach is to exploit multiple noisy 2D body joint detectors and noisy camera pose information. We then optimize for body shape, body pose, and camera extrinsics by fitting the SMPL body model to the 2D observations. This approach uses a strong body model to take low-level uncertainty into account and results in the first fully autonomous flying mocap system.

More information: https://ps.is.mpg.de/project/aircap

## AirCap: Perception-Based Control

*Aamir Ahmad, Rahul Tallamraju, Eric Price, Roman Ludwig, Michael Black*

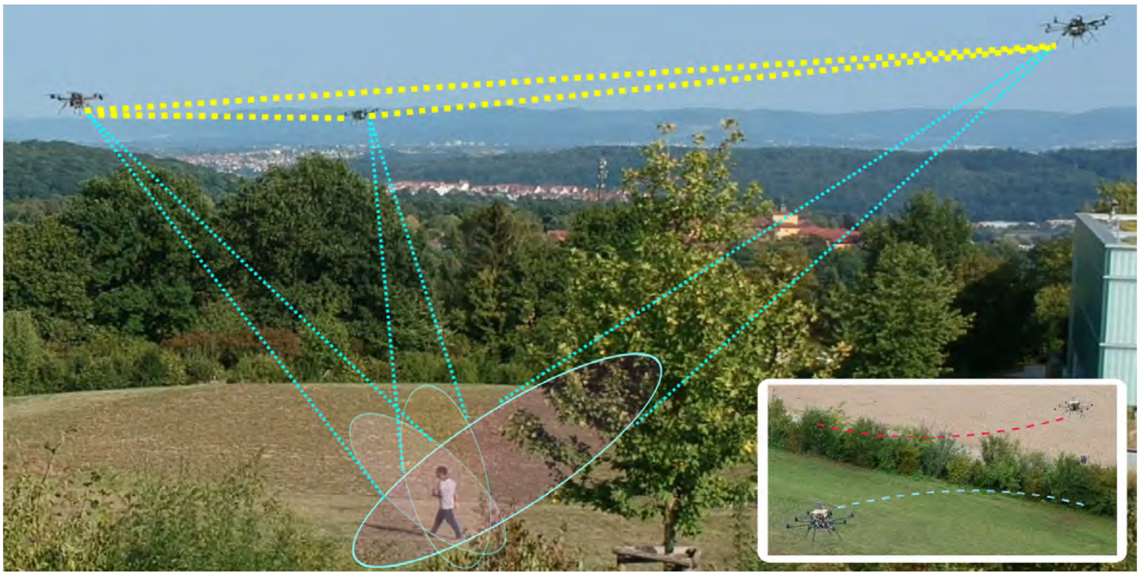

Figure 3.28: Perception-driven formation control of aerial robots tracking a person. The jointly estimated uncertainty in the person's 3D position estimate is minimized while avoiding inter-robot collisions.

Autonomous MoCap systems, like AirCap, rely on robots with on-board cameras that can localize and navigate autonomously. More importantly, these robots must detect, track and follow the subject (human or animal) in real time. Thus, a key component of such a system is motion planning and control of multiple robots that ensures optimal perception of the subject while obeying other constraints, e.g., inter-robot and static obstacle collision avoidance.

Our approach to this formation control problem is based on model predictive control (MPC). An important challenge is to handle collision avoidance as the constraint itself is non-convex and leads to local minima that are not easily identifiable. A possible approach is to treat it as a separate planning module that modifies the MPC-generated optimization trajectory using potential fields. This leads to sub-optimal trajectories and field local minima. In our work [363] we provide a holistic solution to this problem. Instead of directly using repulsive potential field functions to avoid obstacles, we replace them by their exact value at every iteration of the MPC and treat them as external input forces in the system dynamics. Thus, the problem remains convex at every time step. As long as a feasible solution exists for the optimization, obstacle avoidance is guaranteed. Even though field local minima issues remain, they become easier to identify and resolve. To this end, we propose and validate multiple strategies.

In ongoing work we address the complete problem of perception-driven formation control of multiple aerial robots for tracking a human using multiple aerial vehicles. For this, a decentralized convex MPC is developed that generates collision free formation motion plans while minimizing the jointly estimated uncertainty in the tracked person's position estimate. This estimation is performed using a cooperative approach[3] similar to the one developed in our recent work [339]. We validated the real-time efficacy of the proposed algorithm through several field experiments (see image above) with 3 self-designed octocopters and simulation experiments in a realistic outdoor environmental setting with up to 16 robots.

More information: https://ps.is.mpg.de/project/autonomous-mocap

[3] A. Ahmad, E. Ruff, H. Bülthoff. Dynamic baseline stereo vision-based cooperative target tracking. In *19th International Conference on Information Fusion*, pages 1728–1734, 2016.

## IMU-based Human Motion Capture Systems

*Gerard Pons-Moll, Michael Black, Yinghao Huang, Timo von Marcard, Roberto Henschel, Bodo Rosenhahn, Manuel Kaufmann, Emre Aksan, Otmar Hilliges*

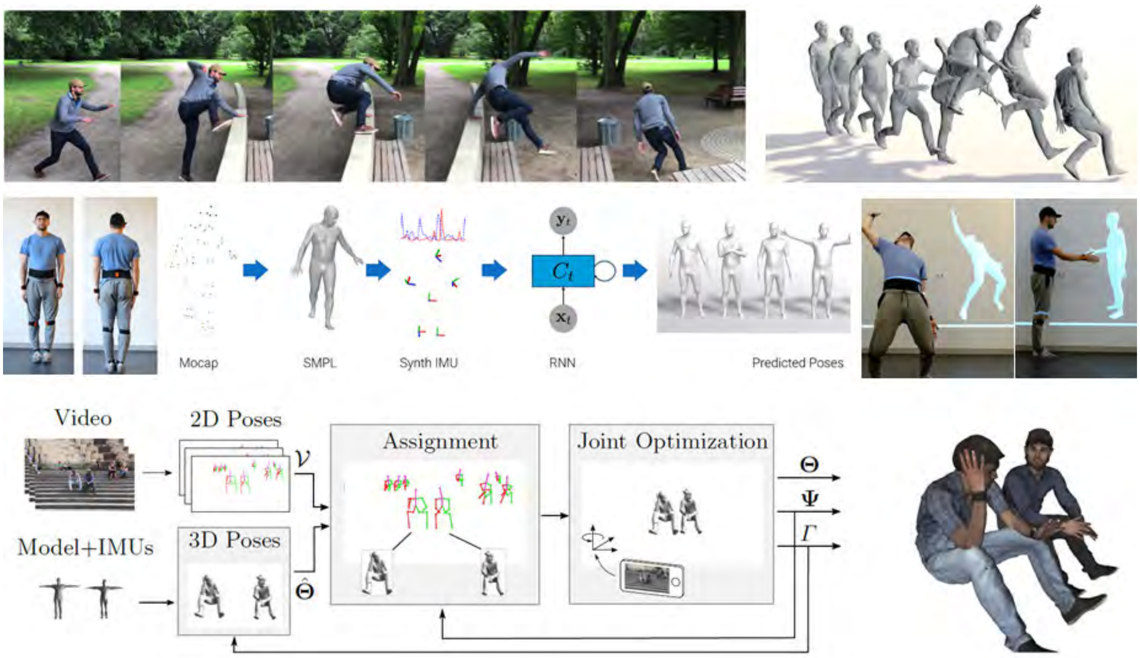


Figure 3.29: Overview. Top row: Unconstrained human motion capture using SIP. Mid row: In DIP, we synthesize an IMU dataset, and leverage that to train an RNN regressor, improving SIP both in accuracy and runtime. Bottom row: Using VIP, we combine videos with sparse IMUs to collect 3DPW, a new dataset of accurate 3D human poses in natural scenes, containing variations in person identity, activity and clothing.

Marker-based optical motion capture (mocap) systems are intrusive and restrict motions to controlled laboratory spaces. Therefore, simple daily activities like biking, or having coffee with friends cannot be recorded with such systems. To address this, and record human motion in everyday natural situations, we develop novel systems based on Inertial Measurement Units (IMUs), that can track the human pose without cameras, making them more suitable for outdoor recordings.

Existing commercial IMU systems require a considerable number of sensors, worn on the body or attached to a suit. These are cumbersome and expensive. To make full-body IMU capture more practical, we developed Sparse Inertial Poser (SIP) [344], which recovers the full 3D human pose from orientation and acceleration measrued by only 6 IMUs attached to the wrists, lower legs, waist and head. This setup is a minimally intrusive solution to capture human activities.

SIP gives an offline, non-intrusive, mocap system that can be used in unconstrained settings of daily life, but the method does not run in real time. In Deep Inertial Poser (DIP) [330], we go beyond the accuracy of SIP and further make it real time. To this end, we synthesize a large IMU dataset from motion capture data and leverage that to learn a deep recurrent regressor that produces SMPL pose parameters in real time from 6 IMU sensor recordings.

While portable, IMU systems are prone to drift. To address this we combine IMUs with a moving camera and current 2D pose-detection methods. Our VIP system [358] solves for the body movements that match the IMU data and project into the image to match 2D joints. Using VIP, we collected the 3DPW dataset, that includes videos of humans in challenging scenes with accurate 3D parameters that will provide the means to quantitatively evaluate monocular methods in difficult scenes.

More information: https://ps.is.mpg.de/project/imu-mocap

## Modeling Human Movement

*Julieta Martinez, Judith Bütepage, Javier Romero, Hedvig Kjellström, Michael Black, Ludovic Righetti, Partha Ghosh, Sergey Prokudin*

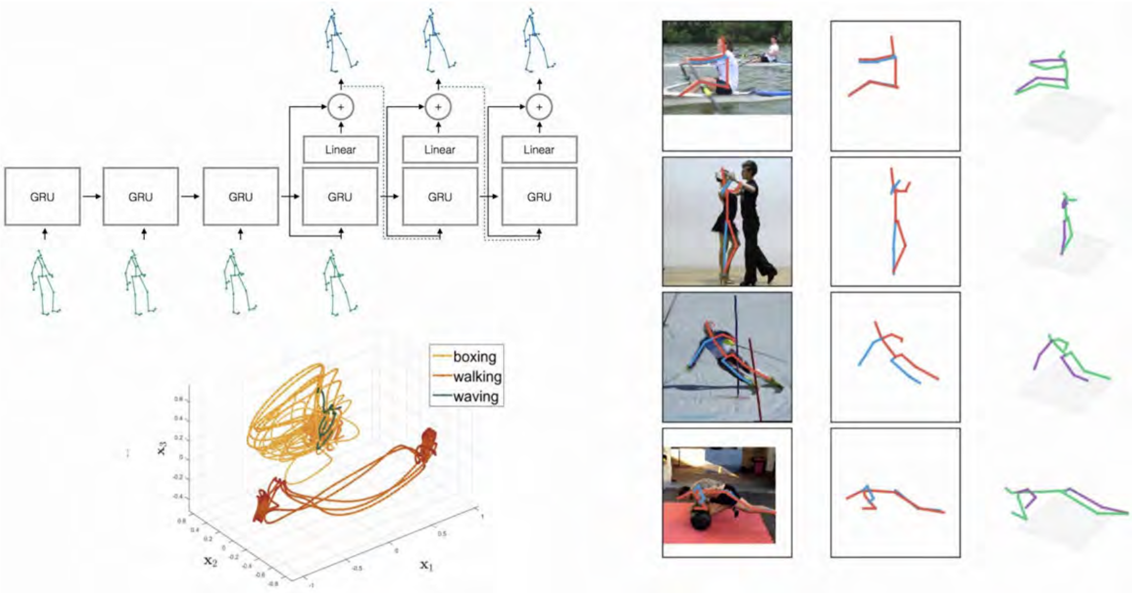


Figure 3.30: Clockwise from upper left: Recurrent neural network for motion prediction [373]; predicting 3D human pose from 2D images [374]; visualization of learned representations of motion sequences [379].

An expressive model of human motion is essential for action classification, motion prediction and synthesis. To that end, we are exploring several deep network architectures to predict human movement.

Current methods for motion prediction typically do not work for a wide range of actions and suffer from "regression to the mean". We show that, surprisingly, state-of-the-art performance can be achieved by a simple baseline that does not model motion at all. We investigate this and propose three changes to the standard RNN models typically used for human motion, which result in a simple and scalable RNN architecture that obtains state-of-the-art performance on human motion prediction [373].

We have also shown that a simple encoder/decoder architecture that takes a set of past poses and predicts a set of future poses works well and is simpler than RNN models. By forcing the encoding through a bottleneck, the approach learns features of human movement that are useful for action recognition. Our feed-forward networks outperform recurrent approaches for short- and long-term predictions and generalize to novel subjects and actions [379].

We have worked on several methods to estimate 3D pose from 2D joints. We show that this can actually be solved with a very simple network that outperforms previous, more complex, methods by a substantial margin. This suggests that "lifting" from 2D to 3D is not the really hard problem but, rather, that extracting the relevant information from the 2D image is the key [374].

Neural networks, however, may not generalize to scenarios that they have never seen – imagine someone floating in zero gravity. Hence we also explore physics-based controllers of human movement [331]. We envision a future that combines the best of both approaches with learned models of behavior combined with physical constraints coming from environmental interaction.

More information: https://ps.is.mpg.de/project/modeling-human-movement

## Action and Behavior

Siyu Tang, Yan Zhang

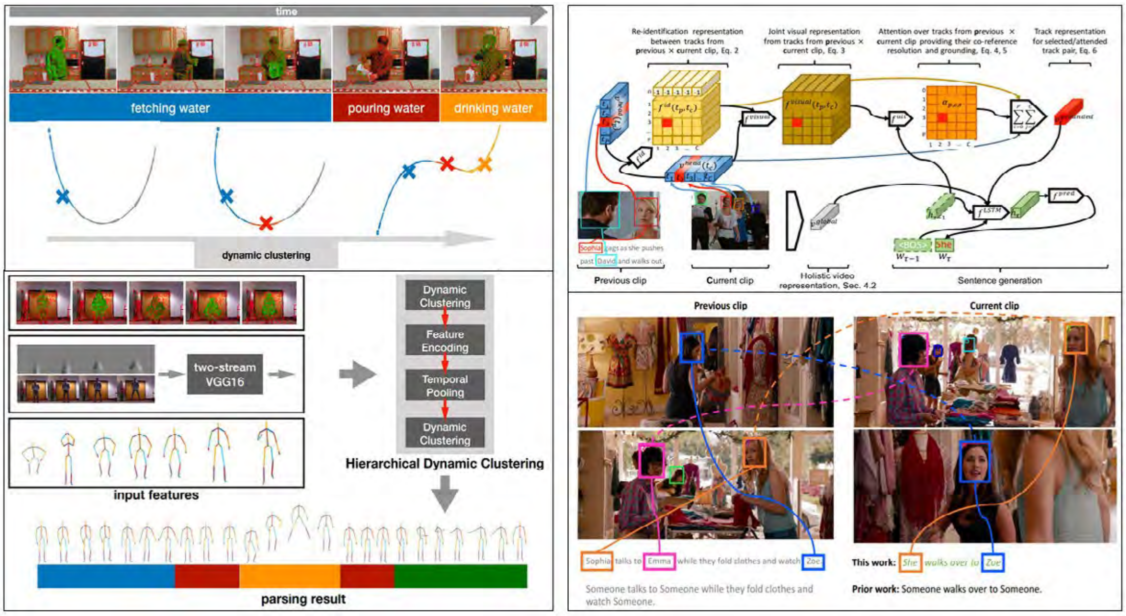


Figure 3.31: Left: A dynamic clustering method for low-level action understanding enables unsupervised human motion parsing [357]. Right: Our method and a result of generated descriptions with grounded and co-referenced people, linking scripts and people in a movie [384].

Human behavior can be described at multiple levels. At the lowest level, we observe the 3D pose of the body over time. Poses can be organized into primitives that capture coordinated activity of different body parts. These further form more complex “actions” or “behaviors”. Finally, underlying all of the above are the goals, motives, and emotions of the person; that is, the *cause* of the movement. Our ultimate goal is to extract this high-level causal information from video.

**Low-level understanding.** Humans can readily differentiate biological motion from non-biological motion. They can do this from sparse visual cues like moving dots and without any explicit supervision. In this spirit, we perform behavior analysis at a low-level using a novel dynamic clustering algorithm that groups actions in an online fashion [357]. As a building block, dynamic clustering is employed in a computational pipeline, where low-level visual cues are aggregated to high-level action patterns via temporal pooling. Our experiments show that this hierarchical dynamic clustering scheme is reliable, generic for diverse input features and fast.

**High-level understanding.** Here we relate low-level behavior to high-level concepts by identifying individual actors and synthesizing natural-language descriptions of their actions and interactions. We do so using weak supervision provided by scripts associated with the video. As a first attempt, we generate descriptions with grounded and co-referenced people [384]. Specifically, we first learn to localize characters by relating their visual appearance to mentions in the descriptions via a semi-supervised approach. We then provide this (noisy) supervision to a description model, which greatly improves its performance. Our proposed description model improves over prior work w.r.t. generated description quality and additionally provides grounding and local co-reference resolution.

Ongoing work leverages richer models of the human body and its motion as well as richer models of the scene and the objects in it.

More information: https://ps.is.mpg.de/project/action-and-behavior

## Learning Optical Flow

*Michael Black, Andreas Geiger, Anurag Ranjan, Jonas Wulff, Deqing Sun, Varun Jampani, Laura Sevilla, Joel Janai, Fatma Güney*

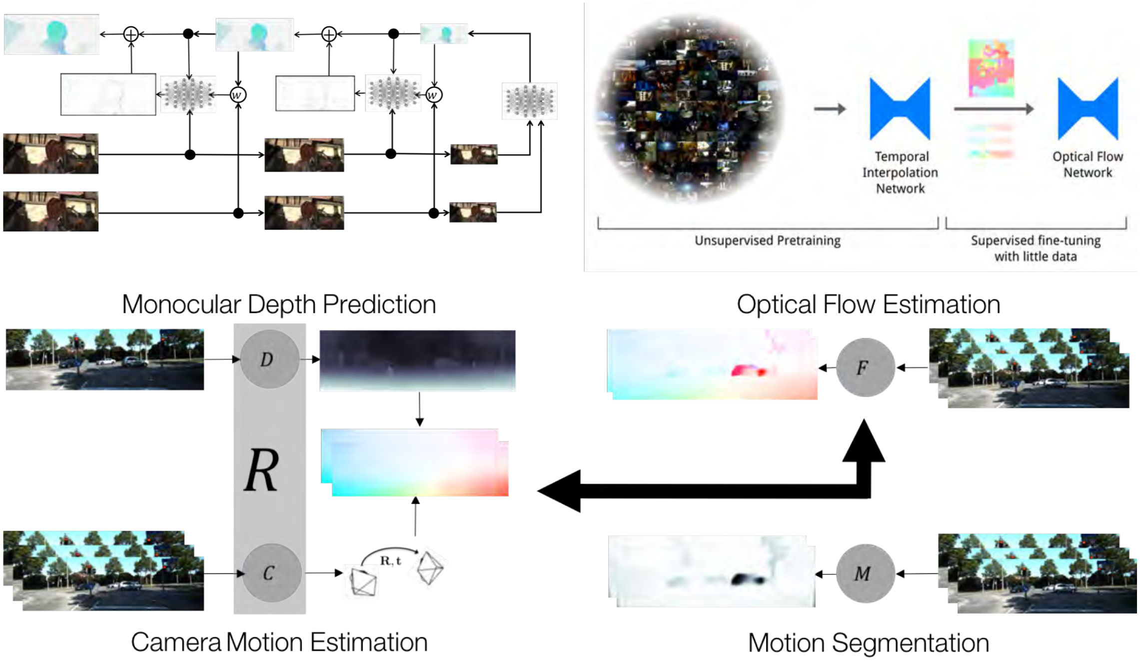


Figure 3.32: Top Left: A Spatial Pyramid Network used for optical flow estimation. Top Right: Optical flow estimates improve when the network is pretrained for temporal interpolation. Bottom: A fully unsupervised approach to solving depth, camera motion, optical flow and motion segmentation.

We view optical flow as the projection of the 3D motion field into the image plane. Until recently, optical flow algorithms were designed by hand and incorporated various heuristics. Deep learning methods provide an opportunity to move away from hand-crafted models but have several limitations. The key one is that they require significant amounts of training data and there are no sensors that give ground truth optical flow for real image sequences.

To deal with large image motions in a compact network, we developed the Spatial Pyramid Networks (SpyNet) [375], which computes optical flow by combining a classical coarse-to-fine flow approach with deep learning. At each level of a spatial pyramid, the deep network computes and update to the current flow estimate. SpyNet is 96% smaller than FlowNet, is very fast, and can be trained end-to-end, making it easy to incorporate into other networks for tasks like action recognition [354].

Synthetic data, used for training most deep flow methods is currently far from realistic. Consequently, we train our IPFlow [355] method on a temporal frame interpolation task using a movie database such that it is encouraged to learn about image motion in complex scenes. We show that this netwok can then be easily fine tuned to compute flow using a small amount of ground truth data.

We go further to address the problem of unsupervised learning. To make this feasible, we build in known geometric information about optical flow in rigid scenes. We introduce the Competitive Collaboration framework [333] and use it to train four different networks that estimate monocular depth, camera pose, optical flow and non-rigid motion segmentation. All of these models compete and collaborate to explain the image sequence. This produces the most accurate unsupervised flow results to date.

Additionally, occlusion boundaries give important information about scene structure and we have worked on learning to detect these [397]. Furthermore, we model occlusions and multiple frames in a video sequence for unsupervised learning of optical flow [359].

More information: https://ps.is.mpg.de/project/learning-optical-flow

## Optical Flow and Human Action

*Anurag Ranjan, Laura Sevilla, Javier Romero, Yiyi Liao, Fatma Güney, Varun Jampani, Andreas Geiger, Michael Black*

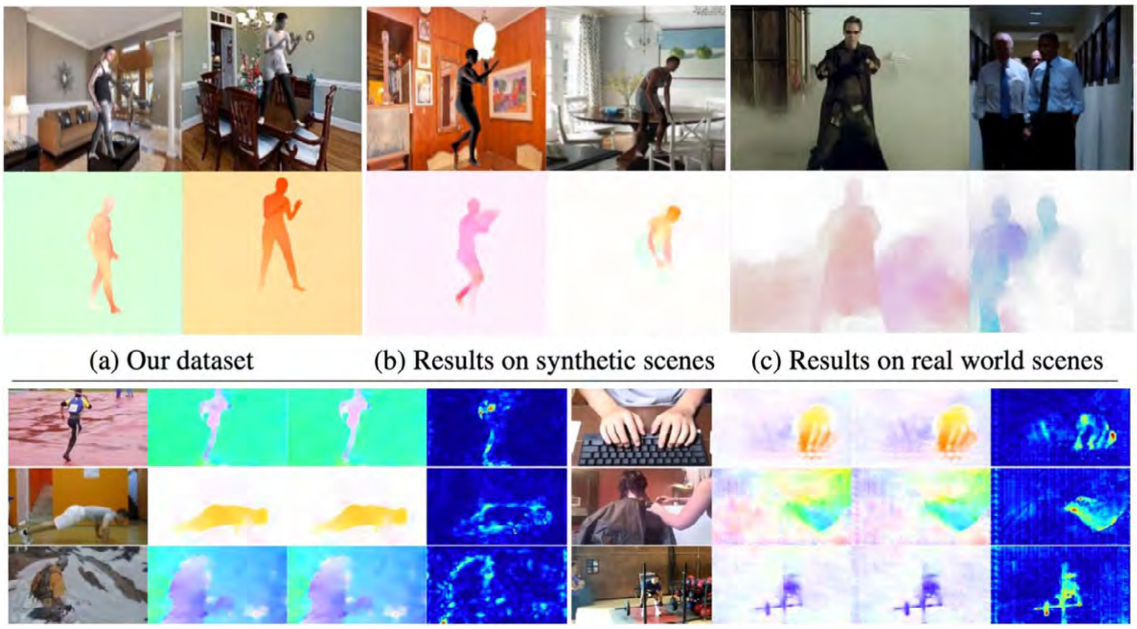


Figure 3.33: Top: We learn human flow from synthetically generated flow fields and find that this generalizes to real videos of human movement. Bottom: We fine tune an optical flow algorithm to produce flow that improves action recognition. Left columns: SpyNet. Right columns: FlowNet. In each set, left to right: first image in sequence, original flow, flow when trained on action recognition, differences in the flow are focused on the human action.

Understanding human action requires modeling and understanding human movement. While we mostly focus on 3D human movement, what is directly observable in videos is the 2D optical flow. Previous work has shown that flow is useful for action recognition and, consequently, we explore how to better estimate human flow and improve action recognition.

Specifically, we trained a neural network to compute human optical flow [362]. To enable this we created a new synthetic training database of image sequences with ground truth human flow. For this we use the 3D SMPL body model and motion capture data to synthesize realistic flow fields; this effectively extends the SURREAL dataset [378]. We then train a convolutional neural network (SpyNet [375] with some modifications) to estimate human flow from pairs of images. The new network is more accurate than a wide range of top methods on held-out test data and generalizes well to real image sequences. When combined with a person detector/tracker, the approach provides a full solution to the problem of 2D human flow estimation.

Most of the top performing action recognition methods use optical flow as a "black box" input. In [354], we take a deeper look at the combination of flow and action recognition, and investigate why optical flow is helpful, what makes a flow method good for action recognition, and how we can make it better. Specifically, we fine tune two neural-network flow methods end-to-end on the UCF101 action recognition dataset. Based on these experiments, we make the following five observations: 1) optical flow is useful for action recognition because it is invariant to appearance, 2) optical flow methods are optimized to minimize end-point-error (EPE), but the EPE of current methods is not well correlated with action recognition performance, 3) flow accuracy at boundaries and for small displacements is most correlated with action recognition performance, 4) training optical flow to minimize classification error instead of EPE improves recognition performance, and 5) optical flow learned for the task of action recognition mostly differs from traditional optical flow inside and at the boundary of the human body.

More information: https://ps.is.mpg.de/project/learning-optical-flow

## Image Segmentation and Semantics

*Raghudeep Gadde, Varun Jampani, Peter Vincent Gehler, Martin Kiefel, Daniel Kappler, Jun Xie*

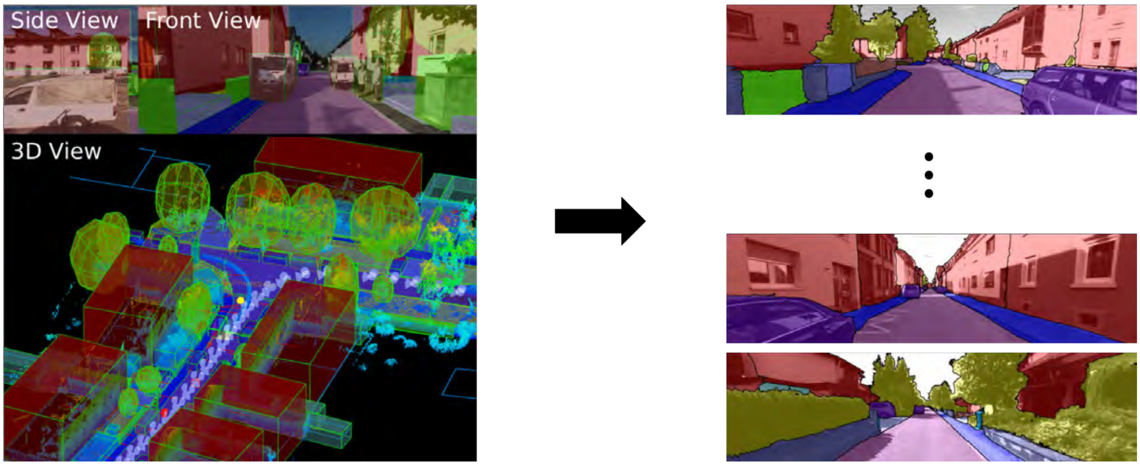


Figure 3.34: 3D to 2D label transfer.

Semantic segmentation is a fundamental problem of computer vision that requires answering what is where in a given image, video or 3D point cloud. The best performing recent techniques require human annotations to obtain ground truth used to train deep neural networks. Such annotation is costly and time consuming to obtain. Consequently, in this project, we address the following two questions:

- How to acquire accurate training data with minimal human cost [394]?
- How to build fast and efficient models for test time inference leveraging the collected data [343], [53]?

In [394], we developed a scalable technique to generate pixelwise annotations for images. For a given 3D reconstructed scene, we annotate static elements in a rough manner and transfer annotations into the image domain using a novel label propagation technique leveraging geometric constraints. We leverage our method to obtain 2D labels for a novel suburban video dataset that we have collected, resulting in 400k semantic and instance image annotations.

In [343], [53] we introduced fast and efficient techniques for semantic segmentation to propagate information using well established Auto-Context and Bilateral filter techniques.

Bilateral filters have wide spread use due to their edge-preserving properties. We generalize the approach to derive a gradient descent algorithm so the filter parameters can be learned from data [399]. This allows us to learn high dimensional linear filters that operate in sparsely populated feature spaces. We build on the permutohedral lattice construction for efficient filtering.

We further introduce a new "bilateral inception" module [53] that can be inserted in existing CNN architectures and performs bilateral filtering, at multiple feature-scales, between superpixels in an image. The feature spaces for bilateral filtering and other parameters of the module are learned end-to-end using standard backpropagation techniques. The bilateral inception module addresses two issues that arise with general CNN segmentation architectures. First, this module propagates information between (super) pixels while respecting image edges, thus using the structured information of the problem for improved results. Second, the layer recovers a full resolution segmentation result from the lower resolution solution of a CNN.

More information: https://ps.is.mpg.de/project/image-segmentation-and-semantics

## Multi-View Stereo

*Osman Ulusoy, Andreas Geiger, Michael Black*

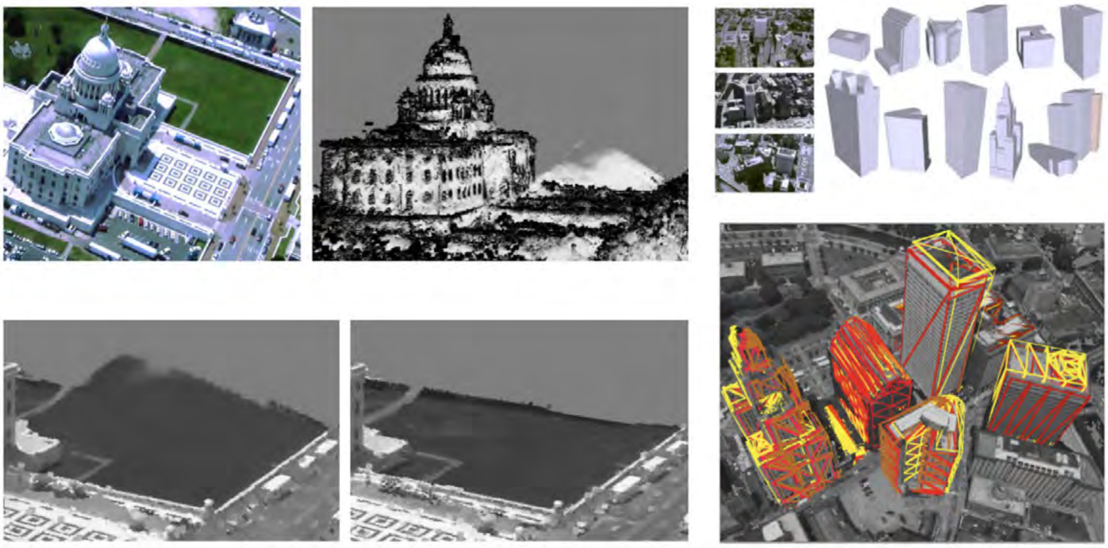

Figure 3.35: In the left: Patches, Planes and Probabilities [395] proposed a planarity prior that regularizes over large distances and helps reconstruct the correct surface. In the Right: Semantic Multi-view Stereo [382] jointly reconstructs a dense 3D model of the entire scene and solves for the existence and pose of each object model.

Dense 3D reconstruction from RGB images is a highly ill-posed problem due to occlusions, textureless or reflective surfaces, varied scene geometry, and spatial discontinuities. We propose algorithms that bring in various types of geometric information that imposes long-range, or semantic, knowlege to address these ambiguities

Our work on Patches, Planes and Probabilities [395] proposed a novel Markov random field model based on ray potentials and a non-local structured prior for volumetric multi-view 3D reconstruction. It was inspired by the planar nature of many elements in man-made environments, i.e., 3D range images of generic scenes can be approximated by piecewise smooth regions with discontinuities at object boundaries. The prior encourages planarity within image segments and regularizes over large voxel neighborhoods. The method was able to resolve reconstruction ambiguities of textureless and partially reflective surfaces and achieved state of-the-art results in reconstruction accuracy for highly challenging aerial datasets.

In our work on Semantic Multi-view Stereo [382], we address ambiguities in 3D reconstruction by presenting a probabilistic approach that integrates object-level shape priors with image-based 3D reconstruction. Our method can infer not only a dense 3D reconstruction of the scene but the existence and precise 3D pose of the objects in it as well. Thus our method not only yields an accurate mapping of the environment but also a semantic understanding in terms of the objects in the environment. The proposed prior allows for powerful regularization that can resolve large ambiguities common in 3D reconstruction. For instance, our shape prior can help reconstruct the back-side of an object even though it is occluded in the images.

More information: https://ps.is.mpg.de/project/multi-view-stereo

## Scene Models for Optical Flow

*Michael Black, Jonas Wulff, Anurag Ranjan, Laura Sevilla, Fatma Güney, Varun Jampani, Andreas Geiger, Deqing Sun*

Figure 3.36: Reasoning about the structure of the scene improves optical flow estimation. Semantic segmentation helps to impose meaningful motion priors based on object identity (left). By segmenting the scene into a static background and moving objects an algorithm can use strong geometric constraints in the background region, simplifying the flow problem (right).

Historically, optical flow methods make generic, spatially homogeneous, assumptions about the spatial structure of the 2D image motion. In reality, optical flow varies across an image depending on object class. Simply put, different objects move differently. For rigid objects, the motion is related to the 3D object shape and relative motion. For articulated and non-rigid objects, the motion may be highly stereotyped. Consequently, we should be able to leverage knowledge about objects in the scene, their semantic category, and their geometry, to better estimate optical flow.

We proposed a method for semantic optical flow (SOF) [400] estimation that exploits recent advances in static semantic scene segmentation to segment the image into objects of different types. We define different models of image motion in these regions depending on the type of object. For example, we model the motion on roads with homographies, vegetation with spatially smooth flow, and independently moving objects like cars and planes with affine motion plus deviations. We then pose the flow estimation problem using a novel formulation of localized layers, which addresses limitations of traditional layered models for dealing with complex scene motion. At time of publication, SOF achieved the lowest error of any monocular method in the KITTI-2015 flow benchmark and produces qualitatively better flow and segmentation than recent top methods on a wide range of natural videos.

Furthermore, the optical flow of natural scenes is a combination of the motion of the observer and the independent motion of objects. Existing algorithms typically focus on either recovering motion and structure under the assumption of a purely static world or optical flow for general unconstrained scenes. We combine these approaches in an optical flow algorithm that estimates an explicit segmentation of moving objects using appearance and physical constraints. In static regions, we take advantage of strong constraints to jointly estimate the camera motion and the 3D structure of the scene over multiple frames. This allows us to also regularize the structure instead of the motion. Our formulation uses a Plane+Parallax framework, which works even under small baselines, and reduces the motion estimation to a one-dimensional search problem, resulting in more accurate estimation. In moving regions the flow is treated as unconstrained, and computed with an existing optical flow method. The resulting Mostly-Rigid Flow (MR-Flow) method [372] achieved state-of-the-art results on both the MPISintel and KITTI-2015 benchmarks.

These methods are optimization-based methods that tend to be slow. Furthermore, we manually define constraints, which are often strong simplifications of the real world. To overcome this, we present the Collaborative Competition framework [333], which reasons about the whole scene in a joint, data-driven fashion, and is able to learn to compute the segmentation and the geometry of the scene, and the motion of objects and the background, without explicit supervision.

More information: https://ps.is.mpg.de/project/scene-models-for-optical-flow

## Video Segmentation

*Varun Jampani, Raghudeep Gadde, Yi-Hsuan Tsai, Michael Black, Peter Vincent Gehler*

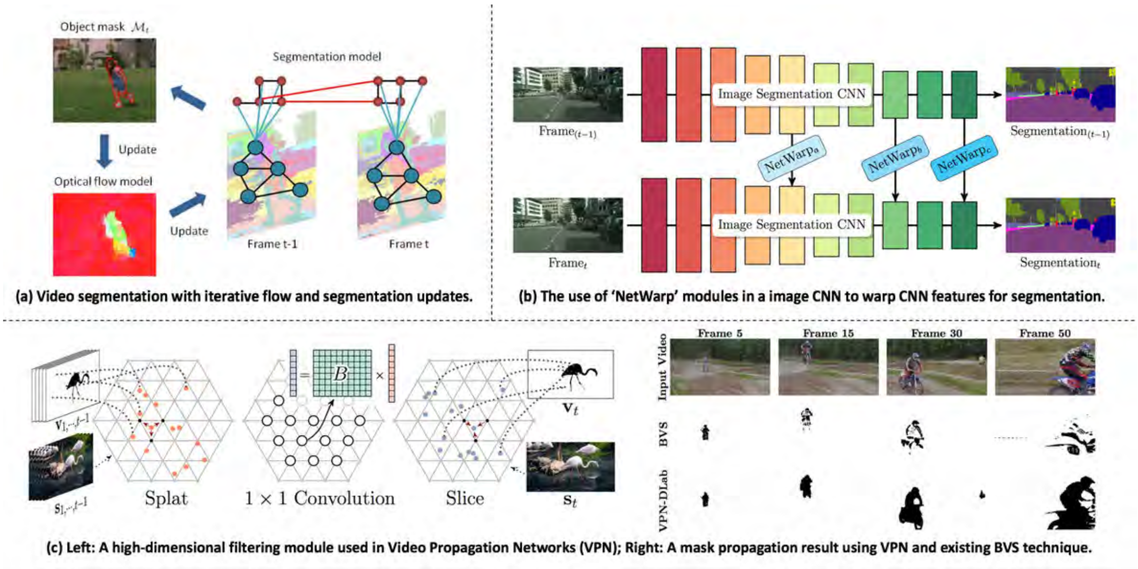


Figure 3.37: Illustration of different video segmentation and propagation techniques: (a) Object Flow [396]. (b) NetWarp [370]. (c) Video Propagation Networks [387].

Videos provide a much richer scene information compared to still images. Despite this, most existing techniques for video segmentation are dominated by per-frame techniques. Video segmentation is a challenging problem due to fast moving objects, deforming shapes and cluttered backgrounds. At Perceiving Systems, we study the use of motion information or pixel correlation that is present across video frames to overcome some of these challenges and obtain better video segmentations.

In [396], we propose an efficient algorithm that considers video segmentation and optical flow estimation simultaneously. We formulate a principled, multiscale, spatio-temporal objective function that uses optical flow to propagate information between frames. For optical flow estimation, we compute the flow independently in the segmented regions and recompose the results. We call the process "object flow" and demonstrate the effectiveness of jointly optimizing optical flow and video segmentation using an iterative scheme.

We also propose one of the first deep neural networks that can be used for general information propagation across video frames. In [387], we project video pixels into a six dimensional XYRGBT space and learn a deep network in this high-dimensional space thereby learning the efficient long-range information propagation across several video frames. Experiments on video object segmentation, video color propagation and semantic video segmentation demonstrate the generality and the effectiveness of our video propagation network.

More recently, we propose a fast and lightweight neural network module called "NetWarp" [370] that can learn to warp intermediate deep feature representations across video frames for better semantic segmentation. Introducing these NetWarp modules in already trained networks and then fine-tuning results in consistent improvements in segmentation accuracy.

More information: https://ps.is.mpg.de/project/video-segmentation

## Learning from Synthetic Data

*Javier Romero, Anurag Ranjan, Michael Black, Jonas Wulff, David Hoffmann, Dimitris Tzionas, Siyu Tang, Naureen Mahmood, Gül Varol, Xavier Martin, Igor Kalevatykh, Ivan Laptev, Cordelia Schmid*

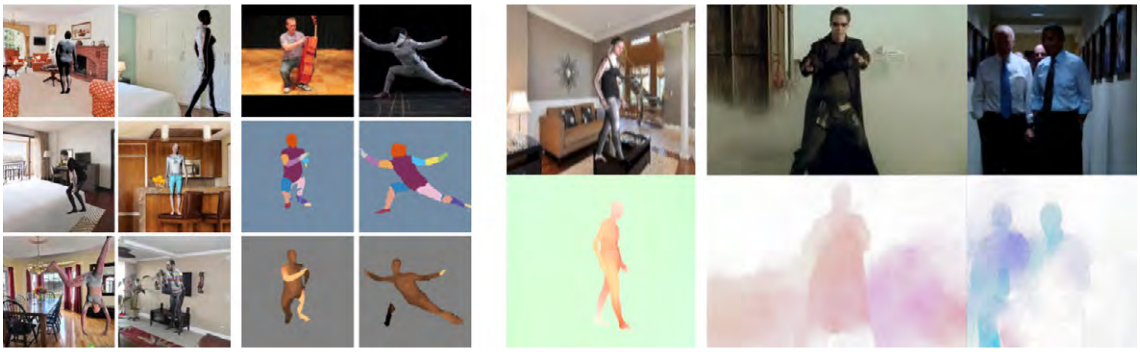

Figure 3.38: Left side: First two columns: SURREAL training data uses SMPL and MoSh to render textured humans in natural poses against random image backgrounds. Right two columns: Training on SURREAL enables the estimation of part segments, depth maps, and more from real images. Right side: Human Flow training data extends SURREAL to moving sequences, giving ground truth optical flow of people in movement (left column). Right two columns: a flow method trained on synthetic data generalizes to human flow in real sequences.

Deep learning has brought rapid progress for many computer vision problems but current methods require large training datasets with annotated ground truth. Human annotators tend to be reasonably efficient for tasks like sparse 2D joint estimation, however annotation for other tasks like dense optical flow estimation or 3D pose estimation is intractable.

To make progress on these tasks, we exploit our 3D body models to generate synthetic training data. In early work, we showed that synthetic data was useful for evaluating optical flow (Middlebury[4] and Sintel[5]). Progress in computer graphics has enabled rendering of synthetic scenes and people and, while not completely realistic, the trends are clear – the quality of such data will steadily improve. Synthetic rendering is appealing for creating training datasets, as it is easily scalable and automatically generates ground truth for a wide variety of problems such as 3D human joints, part segmentations, 3D pose, depth maps, optical flow, body shape, etc.

We focus on learning from synthetic data, using as realistic data as possible about humans, like their motion, body shapes, body textures and backgrounds. We create the SURREAL dataset (Synthetic hUmans foR REAL tasks) and learn deep models for depth estimation and body part segmentation for humans [378]. While not fully realistic, we show that pre-training on this data is valuable and reduces the amount of labeled real data that is needed.

We further create the Human Optical-Flow dataset [362] for learning optical flow of humans in motion. This uses motion capture sequences, processed by MoSh[2] , to produce realistic human optical flow.

Our current work focuses on extending synthetic rendering and inference to multiple people in a single image, for tasks like optical flow, 2D and 3D pose estimation. We further focus on rendering and reconstructing hand-object interactions with realistic hand shapes and poses, object shapes, textures, as well as realistic hand-object grasps. We then plan to extend synthetic data generation to more complex and realistic scenes to reduce the domain gap between real and synthetic data.

More information: https://ps.is.mpg.de/project/learning-from-synthetic-data

[4]S. Baker, D. Scharstein, J. Lewis, S. Roth, M. Black, et al. A Database and Evaluation Methodology for Optical Flow. In *Int. Conf. on Computer Vision, ICCV*, pages 1–8, 2007.

[5]D. J. Butler, J. Wulff, G. B. Stanley, M. J. Black. A naturalistic open source movie for optical flow evaluation. In *European Conf. on Computer Vision (ECCV)*. Part IV, LNCS 7577, pages 611–625, Oct. 2012.

## Psychology and Body Shape

*Simone Mölbert, Anne Thaler, Betty Mohler, Stephan Streuber, Javier Romero, Naureen Mahmood, Sergi Pujades, Alejandra Quiros-Ramirez, Silvia Zuffi, Joachim Tesch, Michael Black*

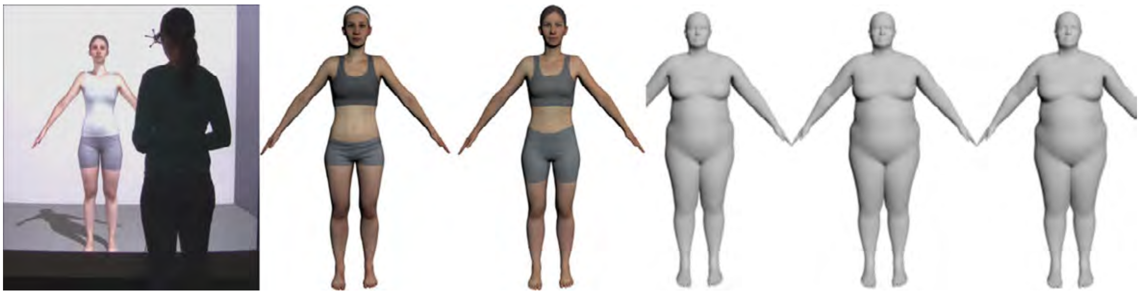

Figure 3.39: Illustration of the virtual reality mirror setup with weight manipulated individual avatars, the manipulation of avatar identity and our figure rating scale (partial view).

Body representation is an essential part of a person's self-concept and also shapes how we see the world. A disturbed body representation also plays a role in clinical conditions such as eating disorders and stroke. So far, a major hurdle for research was the lack of ecologically valid body stimuli. In this project, we cooperate with partners from the Max-Planck-Institute for Biological Cybernetics and the University Hospital Tübingen to develop ecologically valid methods for the assessment of body representation.

For example, we explore how body shape is perceived by people with anorexia nervosa (AN). It was thought that AN patients might perceive their bodies in a distorted way. Using virtual reality and body scans of AN patients we explored this by varying the shape of personal and other avatars to test their perception. We found that anorexics perceive body shape veridically but prefer unhealthy weights.

Based on a 3D body scan and a statistical body model learned from the CAESAR dataset, we generate individual avatars of the participants that can be distorted in terms of weight. Through texture manipulations, we are able to vary the identity of the displayed person. As a major improvement to the existing artist-generated figural drawing scales, we also created a biometric figure rating scale [336] and a desktop tool. In different projects, we assessed >100 participants from the general population as well as >30 women with anorexia nervosa.

Our results in [334] show that in the general population, the accuracy of own body size estimation is predicted by personal BMI, such that participants with lower BMI underestimated their body size and participants with higher BMI overestimated their body size. Critically, these biases suggest that people tend to perceive their weight in an exaggerated way, while there was no hint of a general denial in underweight or overweight persons. The same underestimation bias also occurred in women with anorexia nervosa. Further, we consistently observed that women with anorexia nervosa favored a much thinner body as ideal weight than healthy women. This observation has major clinical implications, because it questions the common idea that misperception of body dimensions may be a maintaining mechanism of this eating disorder. Rather, it suggests that treatment should support patients in accepting a healthy body weight for their own.

Through this work, we have developed a range of body shape modeling, animation, and VR technologies that can be clinically deployed. To study body shape perception we created a virtual reality mirror scenario [335] and our *virtual caliper* allows subjects to create a realistic 3D human avatar using only the controllers of a VR game system. We also studied how we see body shape and describe it with language [345]. We had subjects rate 3D bodies along many dimensions and then built a statistical model relating words and shape. With this "Body Talk" system, we could recover 3D shape from the descriptions of people [347]. This opened up research on body shape and subjective judgements.

More information: https://ps.is.mpg.de/project/anorexia-and-body-shape

## Medical Diagnosis

*Michael Black, Sergi Pujades, Javier Romero, Nikolas Hesse*

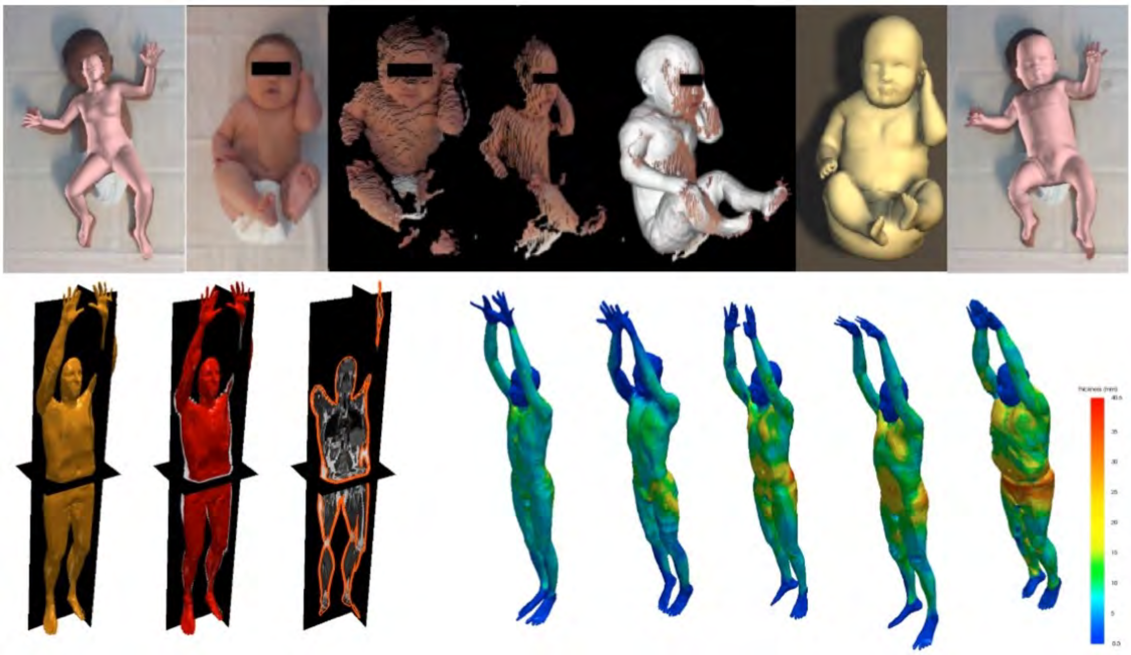

Figure 3.40: Top: Existing adult body models are not suitable for infants, as body proportions differ. We learn an infant body model from RGB-D sequences and use this to recover the shape and pose of freely moving infants. Bottom: We use an adult body model to guide the segmentation of the subcutaneous adipose tissue in a full body MRI (left). The thickness of the adipose tissue is illustrated on the bodies (right).

Body shape and movement are related to human health. For example, our shape tells us about our body fat and our movement tells us something about the health of our motor system. Using our 3D models of body shape we analyze movement and shape to create non-invasive and deployable methods of analyzing human health.

For example, if Cerebral Palsy (CP) is detected early, there are effective therapies to minimize the impact in later life. CP can be diagnosed in infants based on their spontaneous, undirected movements. Unfortunately, this requires expert training that is not widely available. If we can automatically track infant movement, we may automate the early detection of CP. The vision community has made great progress on 3D tracking of adults. Infants have a very different body shape from adults (see figure), which makes it difficult to directly extend prior work to infants. To address this, we learn a model of infant body shape [364] and use it to track 3D movement in RGB-D sequences. Previous models of 3D humans[1] were learned from thousands of high quality 3D scans, which is not practical with infants. Consequently we developed a novel method that learns infant body shape directly from low quality, incomplete, RGB-D scan sequences and deployed this in hospitals where we scanned over 30 infants.

Another example involves the distribution of adipose tissue in the body. Not all fat is the same. Visceral adipose tissue (around the organs) is highly correlated with diabetes and cardiovascular disease. In contrast, sub-cutaneous adipose tissue (fat under the skin) is relatively benign. Today an analysis of this fat distribution requires an MRI scan to reveal where fat is stored. We are developing methods to estimate this fat distribution purely from the surface shape of the body. To that end, we fit our 3D body models to full-body MRI scans [346] to model both the external surface and the subcutaneous fat layer. We are collecting a dataset of matched MRI data and 3D surface shape and our ongoing work is focused on predicting what is inside solely from the surface.

More information: https://ps.is.mpg.de/project/medical-diagnosis

## 3.3 Awards & Honors

### 2018

**Michael J. Black**, Alumni Research Award, Department of Computer Science, University of British Columbia, 2018

**Siyu Tang**, DAGM MVTec 2018 Dissertation Award at the German Conference on Pattern Recognition (GCPR) for her thesis “People Detection and Tracking in Crowded Scenes”.

**Christoph Lassner** and **Gerard Pons-Moll**, Best Student Paper at 3DV 2018 for the paper “Neural Body Fitting: Unifying Deep Learning and Model-Based Human Pose and Shape Estimation.”

**Varun Jampani, Christoph Lassner, Juergen Gall, Yi-Hsuan Tsai, Julietta Martinez, Fatma Güney, Lars Mascheder, Gernot Riegler, Deqing Sun**, Outstanding Reviewer Awards (active members and alumni), CVPR 2018.

**Alejandra Quiros-Ramirez, Betty Mohler, Michael Black, Simone Mölbert** and collaborators, Best Poster Award, Deutsche Gesellschaft für Essstörungen (DGESS), 2018, for the paper “Körper Sprache: Sprachliche Repräsentation von Körpern bei Patientinnen und Patienten mit Essstörungen.

### 2017

**Andreas Geiger** and collaborators, Best Student Paper at 3DV for the paper “Sparsity Invariant CNNs.”

**Siyu Tang**, Early Career Research Grant from the University of Tuebingen to start a research group.

**Siyu Tang**, winner of the CVPR 2017 Multi-Object Tracking Challenge.

**Varun Jampani, Osman Ulusoy, and Silvia Zuffi**, Outstanding Reviewer Award, CVPR 2017.

**Gerard Pons-Moll and Michael Black** and collaborators, Best Paper Award, Eurographics 2017, for the paper “Sparse Inertial Poser: Automatic 3D Human Pose Estimation from Sparse IMUs.”

### 2016

**Federica Bogo, Javier Romero, Matthew Loper, and Michael Black** Dataset Award at the Eurographics Symposium on Geometry Processing 2016. The award encourages and recognizes the importance of the distribution of high-quality datasets on which geometry processing algorithms are tested.

### Faculty Appointments

**Laura Sevilla** appointed Lecturer at the University of Edinburgh, Scotland.

**Sergi Pujades** appointed Associate Professor at Université Grenoble Alpes, France.

**Stefan Streuber** appointed Assistant Professor at the University of Konstanz, Germany.

**Gerard Pons-Moll** appointed independent group leader at the MPI for Informatics, Saarbrücken.

## 3.4 Director profile: Michael J. Black

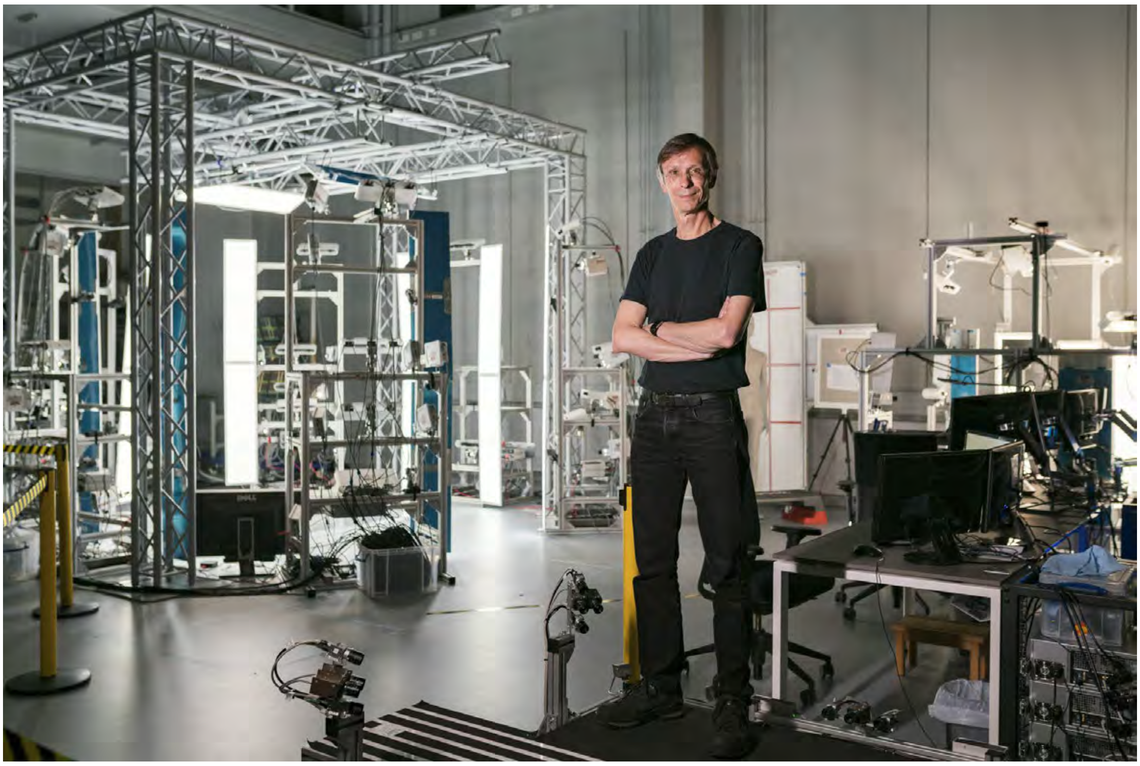

Michael J. Black received his B.Sc. in Honours Computer Science from the University of British Columbia (1985), his M.S. in Computer Science from Stanford University (1989), and his Ph.D. in Computer Science from Yale University (1992). As a graduate student he performed research at the NASA Ames Research Center, Aerospace Human Factors Research Division. After one year as an assistant professor at the University of Toronto, he joined the Xerox Palo Alto Research Center in 1993 as a member of research staff. He went on to managed the Image Understanding Area and found the Digital Video Analysis Area. In 2000 he joined the faculty of Brown University in the Department of Computer Science as an Associate Professor with tenure. He was promoted to Full Professor in 2004.

In 2011 he joined the Max Planck Society as a Scientific Member and one of the founding directors of the Max Planck Institute for Intelligent Systems in Tübingen, Germany. Since 2017 he is also an Distinguished Amazon Scholar.

Dr. Black's research spans computer vision, computer graphics, and machine learning. He is most known for his work on optical flow, robust statistical methods, human motion capture and analysis, 3D body shape modeling, neural prosthetics, and motor-cortical decoding.

Dr. Black is a foreign member of the Royal Swedish Academy of Sciences. He is a recipient of the 2010 Koenderink Prize for Fundamental Contributions in Computer Vision and the 2013 Helmholtz Prize for work that has stood the test of time. His work has won several paper awards including the IEEE Computer Society Outstanding Paper Award (CVPR'91) and Honorable Mention for the Marr Prize in 1999 and 2005. His early work on optical flow has been widely used in Hollywood films including for the Academy-Award-winning effects in "What Dreams May Come" and "The Matrix Reloaded." He has contributed to several influential datasets including the Middlebury Flow dataset, HumanEva, and the Sintel dataset. Black has coauthored over 200 peer-reviewed scientific publications.

Dr. Black was a co-founder and member of the board of directors of Body Labs Inc., which commercialized his team's research on 3D human body shape. Body Labs was acquired by Amazon in 2017.

## Dr. Michael J. Black

### Appointments (2016-2018)

| | |
|---|---|
| 01/2011 – present | Director at the Max Planck Institute for Intelligent Systems |
| 03/2018 – 11/2018 | Managing Director of the MPI for Intelligent Systems, Stuttgart and Tübingen |
| 09/2017 – present | Distinguished Amazon Scholar |
| 05/2012 – present | Honorary Professor, Department for Computer Science, University of Tübingen |
| 04/2014 – 04/2016 | Visiting Professor, Dept. of Inf. Tech. and Electrical Eng., ETH Zurich |
| 01/2011 – present | Adjunct Professor, Dept. of Computer Science, Brown University |

### Awards & Honors (Selected)

| | |
|---|---|
| 2018 | Alumni Research Award, Dept. Computer Science, Univ. British Columbia |
| 2017 | Best Paper Award, Eurographics 2017 |
| 2015 | Elected foreign member of the Royal Swedish Academy of Sciences |
| 2013 | Helmholtz Prize for work that has stood the test of time |
| 2010 | Koenderink Prize for Fundamental Contributions in Computer Vision |
| 1999 & 2005 | Marr Prize, Honorable Mention, Int. Conf. on Computer Vision, ICCV |
| 1991 | IEEE Computer Society, Outstanding Paper Award, CVPR |

### Selected Organization and Community Service (2016-2018)

| | |
|---|---|
| 2017 | Co-organizer, Scenes from Video (SfV) Workshop, III, Lago di Garda, Italy |
| 2016 | Chair, Scientific Advisory Board, Computer Science Department, *École Normale Supérieure*, Paris |
| 2016 | SIGGRAPH Course, Co-organizer, "Learning human body shapes in motion," Anaheim, CA, 2016 |
| 2016 | PAMI Young Investigator Award Committee |

### Selected Memberships (2016–2018)

Royal Swedish Academy of Science, since 2015
European Association for Computer Graphics, since 2017
Intel Network on Intelligent Systems (NIS), since 2017
Association of Computing Machinery (ACM), member since 2014
MPI-ETH Center for Learning Systems, Member since 2015
Werner Reichardt Center for Integrative Neuroscience (CIN), Tübingen University, member since 2011
Institute for Electrical and Electronics Engineers (IEEE): Senior Member since 2008

### Startup Activity and Board Memberships (2016 – 2018)

Meshcapade GmbH, Tübingen, angel investor, Nov. 2018
Body Labs Inc., New York, NY, Co-founder, Member of the Board, 2013 – 2017

### Selected Keynote, Conference, and Public Talks (2016-2018)

"Estimating Human Motion: Past, Present, and Future." *40 Years DAGM - Invited Talks*, GCPR 2018, Stuttgart, Oct. 2018.
"The Digital Body: Capturing, Modelling and Animating Realistic 3D Humans," *Public Lecture Series on 'What Beings are We?'*, Institute for Art and Architecture, IKA, Vienna, Austria, May 2018.
"Building digital humans by scanning real ones," Keynote, *13th European Conference on Visual Media Production (CVMP)* London, Dec. 12–13, 2016.
"Human body shape modeling: A tutorial," *Invited Tutorial: European Conference on Computer Vision and the ACM Multimedia Conference*, Amsterdam, Oct. 2016.
"The future of generative models: A case study of human bodies in motion," *Int. Computer Vision Summer School, ICVSS*, Sicily, July 2016.
"On building digital humans," *Shape Analysis and Learning by Geometry and Machine*, Inst. for Pure and Applied Mathematics (IPAM), UCLA, Feb. 2016.

# 4 PHYSICAL INTELLIGENCE

## 4.1 Research Overview

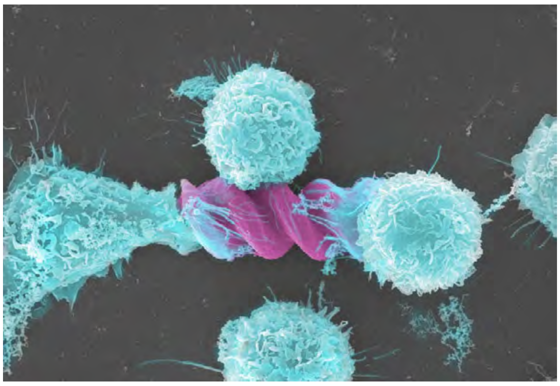

The Physical Intelligence Department started its research activities in fall 2014 at the Stuttgart site of our institute. Our department aims to understand the principles of design, fabrication, control, and learning of single and large numbers of small scale-mobile robots made of smart and soft materials as our physical intelligence platforms. The intelligence of such robots would dominantly come from their physical design, materials, interactions, and control, more than, or in addition to, their computational perception, learning, and control. Such physical intelligence methods are indispensable at the small scale, especially since small-scale mobile robots are inherently limited in onboard computation, actuation, powering, perception, and control capabilities. Our department has three main research thrust fields to realize new physical intelligence platforms to understand them in detail: Advanced materials, mobile millirobots, and mobile microrobots.

## 4.2 Research Fields



### 4.2.1 Advanced Materials


To realize new physical intelligence platforms at the small scale for operation in complex and dynamic environments with their inherent limitations in all functionalities, novel smart, (multi)functional, highly integrated, safe, and advanced materials are indispensable for improved performance and adaptive behavior.

**Bio-inspired reversible and switchable adhesives** Our group has pioneered novel dry, reversible, fibrillar adhesives inspired by gecko's foot-hairs. Although, considerable advances have been made, significant scientific and engineering challenges remain. Therefore, we have been recently expanding our studies to include the self-cleaning mechanics of insect. There is no clear understanding on how insects and geckos can clean their footpads when they are soiled with particles. Our team has studied the self-cleaning mechanism of insects and proposed a robust mechanism for artificial adhesives or grippers.

In robotics and precision manufacturing, adhesion-controlled grasping of complex 3D surfaces is very challenging and current systems are limited by a trade-off between 3D surface conformability and high adhesion strength. Inspired by the natural adhesive systems, we developed elastomeric microfibrillar surfaces and intergraded these into an adhesion-based soft gripping system, enabling improved load sharing and handling of complex 3D and deformable objects.

Although, remarkable advances have been made in reversible and switchable adhesives, conformal attachment and robust adhesion to rough surfaces or soft skin remains a grand challenge. Therefore, we are exploring different approaches to overcome these challenges. In the first approach, the fascinating wrinkling phenomena of the highly bendable thin Gallium oxide skin and the phase changing ability of Gallium near room temperature was studied to enable reversible adhesion. We explored the mechanical behavior of the highly bendable native oxide skin on a Ga droplet, it′s wrinkling phenomenon, and its implications on the adhesion energy necessary to separate the interface. Moreover, the phase change ability of Gallium from liquid to solid phases was studied and demonstrated reversible adhesion performance on rough and wet surfaces. In the second approach, we investigated composite microfibrillar adhesive films, which provide superior conformation and adhesion to the hierarchical topography of the skin. The soft and stretchable skin-adhesive films are composed of microfibers decorated with conformal and mushroom-shaped tips, which offer direct crosslinking of the tips on the skin surface and greatly enhance the skin adhesion through their excellent shape conformation. As applications of these reversible adhesives, we are designing, fabricating, and demonstrating novel soft inflatable, adhesion-based grippers using microfibrillar reversible adhesives, and novel transfer printing methods under dry and wet conditions employing Gallium and composite microfibrillar adhesives for biocompatible highly sensitive wearable devices.

**Soft sensors and actuators** In recent years, new emerging functional materials, such as, conductive polymers, dielectric elastomers, liquid-crystal elastomers, nanomaterials, and magnetic composites, have considerably improved the performance of strain sensors and soft actuators. Considerable attention is being paid to wearables owing to their seamless integration with the human body, allowing continuous monitoring of important vital signs. Despite remarkable advances, their attachment to skin and the lack of sensitivity remains a grand challenge. Therefore, we developed ultrasensitive wearable strain sensors by composite structures of nanomaterials and polymers, which can be attached onto the clothing or even directly mounted on the human skin, enabling accurate monitoring of vital signals such as respiratory and heart rate.

Next, soft actuator materials are being actively pursued owing to their importance in soft robotics, artificial muscles and biomimetic devices. However, most of the current systems are stiff and rigid, exhibit small deformations and require high voltages. In this study, we introduce a new soft actuation method based on a combination of inflated hyperelastic membranes and dielectric elastomer actuators, which does not require either pumps, compressors, or valves.

Finally, we have added sensing function to our multiresponsive actuators composed of paper and polypropylene film. We have shown that the combination of functional materials overcomes the selfsensing limitation of current soft actuators. We are planning to use our highly sensitive sensors and soft actuators for remote healthcare applications, human interactive electronic devices, and emerging soft untethered robotics capable of multimodal locomotion.

**Ferrofluid-containing liquid infused porous surfaces (FLIPS)** Developing adaptive materials with changing geometries in response to external stimuli provides fundamental insights into the links between the physical forces involved and the resultant morphologies and creates a foundation for technologically relevant dynamic systems. Recently, gels, polymers, composites and slippery surfaces have been explored as reconfigurable surfaces to control interfacial properties. However, these concepts exhibit a limited range of topographical changes and thus a restricted scope of function. Here we introduce a hierarchical magnetoresponsive composite surface, made by a ferrofluid-containing liquid-infused porous surfaces. We demonstrate various topographical reconfigurations at multiple length scales and a broad range of associated emergent behaviours. We envision that FLIPS could be used as part of integrated control systems for the manipulation and transport of matter, thermal management, microfluidics and fouling-release materials.

More information: https://pi.is.mpg.de/field/advanced-materials

### 4.2.2 Mobile Microrobots

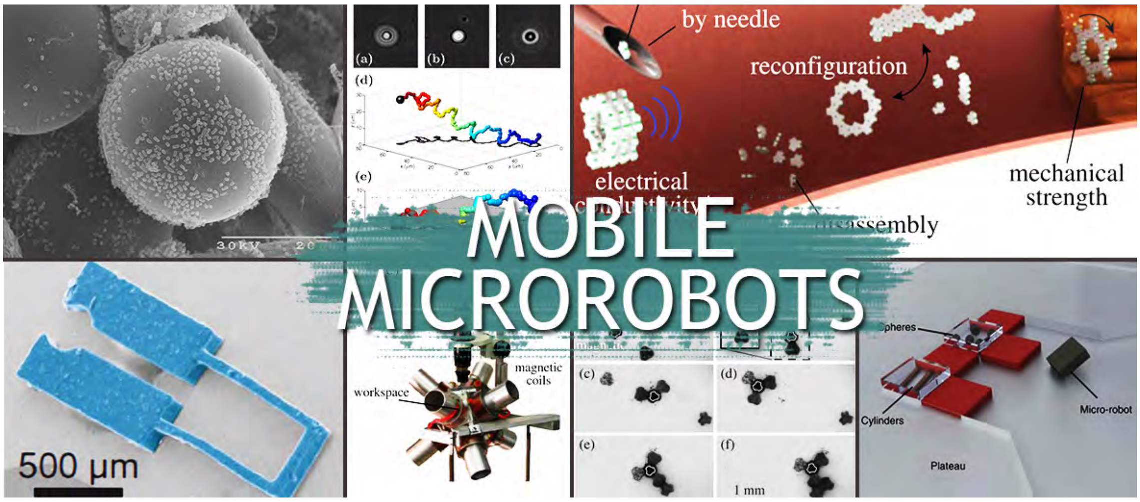


Mobile robotic devices smaller than one millimeter could enable new directions in healthcare, biotechnology and manufacturing. Our focus in healthcare is on medical microrobots that can leverage minimally invasive interventional and targeted delivery strategies with high precision and repeatability by navigating and performing in hard-to-reach and delicate inner body sites, such as, the central nervous system, the circulatory system and the eye. Active navigation of highly concentrated therapeutic and diagnostic agents to the site of action could represent a state-of-the-art application of microrobots, considering the limited delivery and distribution efficiencies offered by the systemic routes and local diffusion. By this means, it is possible to minimize the effects of systemic toxicity and increase the overall efficacy of single-dose administration. Autonomous release of multiple types of payloads with programmable kinetics based on the environmental sensing of local cues, *e.g.*, disease markers, in the living milieu could pave the way for microrobotic therapy and diagnosis in the form of an orderly executed, programmable operation. A conventional robot responds to the changes in its environment by means of its on-board sensors and computational capabilities. Achieving such capabilities at the smaller dimensions, where such computational capabilities do not exist; however, remains a major research question. In nature, microorganisms, such as slime molds and bacteria have evolved to use physical intelligence as the main route of making decisions in complex and evolving conditions. Accordingly, programmed physical and chemical properties of materials can enable a robust design route for making microrobotic systems with the capabilities of motion, sensing, and functioning in dynamic interaction with their local environment. Such goals require integrated design and engineering strategies, where powering, actuation, control, environmental sensing and medical functionality need to be considered altogether.

**Biohybrid microrobots** Unicellular motile organisms, *e.g.* flagellated bacteria and algae, are physically integrated with a microrobot that can be used as on-board microactuators and microsensors that harness the biochemical energy in the microenvironment to power the mobility of the microrobot. We exploit the microorganisms that can achieve high propulsion speeds (tens of their body lengths per second), thereby providing high thrust power to the biohybrid microrobots. Integrated sensing and motility of microorganisms further enable steering using cues in the environment chemical and pH gradients.

**Synthetic microrobots** Fabrication of synthetic microrobots presents unique challenges concerning design, fabrication process, and encoding operational capabilities. Conventional microfabrication techniques usually provide relatively simple geometric structures, such as tubes, spheres, and surfaces, with limited design flexibility and function. Integration of computer-aided design to microfabrication technologies has been a significant advancement to realize sophisticated 3D designs that could not be conceivable with the alternative microfabrication methods. To this end, we explore the potential of additive manufacturing processes for making microrobots. We put a special emphasis on making new materials with tailorable local 3D chemical properties that would enable on-board programmable functionalities, and hence novel design opportunities for diverse types of microrobots.

**Reconfigurable and programmable self-assembly and collective behavior of magnetic microrobotic swarms** Collectives of microrobots working cooperatively could significantly enhance their functionalities in many applications such as micro-manufacturing and biomedical engineering. We are working on programmable and reconfigurable self-assembly and collective behavior of microrobots that could enable parallel and distributed massive tasks eventually in the living environment.

More information: https://pi.is.mpg.de/field/mobile-microrobots

### 4.2.3 Mobile Millirobots


The department currently has two focus areas in the field of mobile millirobots. The first one is on soft mobile millirobots which would have high impact in healthcare. This includes soft mobile millirobots and medical pill-size soft capsule robots. The second focus is on design and control of self-contained robots with capabilities of dealing with complex terrains. These robots are bio-inspired by organisms such as jumping-gliding vampire bat and locust.

**Soft mobile millirobots** Soft systems, actuators, and robots are becoming increasingly popular for dynamic task environments where safe interaction with unknown or unpredictable objects is required. Our department puts special attention on magnetic soft materials due to their flexibility to be programmed and compatibility for medical applications. Not only have we built up a complete theoretical and computational system to design and fabricate the miniature magnetic soft devices, with such knowledge and inspiration from small-scale animals, we have also developed miniaturized soft mobile robots with superior locomotion capabilities. These robots can transit reversibly between different liquid and solid terrains, as well as switch between locomotive modes. Therefore, they are very suitable to operate in highly unstructured environments such as those inside the human body. They can additionally execute pick-and-place and cargo-release tasks. With further optimization, these robots could open great potentials in biomedical applications.

**Medical pill-size soft capsule robots** As another effort to realize minimally invasive medical applications, we have been designing, fabricating, and controlling novel soft capsule robots to operate inside the GI tract. Building upon the established pill-camera technology, these untethered soft capsule robots with embedded magnetic materials, a microcamera, and therapeutic microtools could actively image the GI tract, deliver drugs actively or passively, conduct a biopsy, and conduct an elastography.

Precise 6-DOF localization of such medical milli/microrobots inside the human body, their safe design and operation, 3D mapping inside the human body using such active imaging systems, and preclinical animal experiments of such robots are current research foci.

**Bio-inspired millirobots with multilocomotion** Long flying times, safe flapping designs, agility, and low flying speeds make flapping robots an excellent agent in rescue and mapping scenarios. Getting off the ground takes, however, considerably more energy than staying in the air, and bio-inspired jumping mechanisms can change this. Inspired by a common vampire bat (*Desmodus Rotundus*), we have developed a very lightweight robot (Multimo-bat) that incorporate jump-off mechanisms, gliding mechanisms, and control to guide the robot within a large operating space. These robots will be able to consecutively jump off, glide or flap, and land.

The robotic system we developed also allows us to directly test scientific questions initially raised by biologists. Specifically, we investigate the desert locust (*Schistocerca gregaria*). We report the slipping behavior, dynamic attachment, passive mechanics, and interplay between the spines and adhesive pads, studied through both biological and robotic experiments, which contribute to the locust's ability to jump robustly from diverse surfaces. We found slipping to be surface-dependent and common, yet the morphological intelligence of the feet produces a significant chance to reengage the surface. Additionally, a discovered noncontact-type jump, further studied robotically, broadens the applicability of the morphological adaptations to both static and dynamic attachment. Our results demonstrate the potential contribution of morphological intelligence in solving complex dynamic locomotion problems.

More information: https://pi.is.mpg.de/field/mobile-millirobots

## 4.3 Selected Research Projects

## Bio-inspired microfibers for skin adhesion and signal amplification of wearable sensors

*Dirk Drotlef, Morteza Amjadi, Metin Sitti*

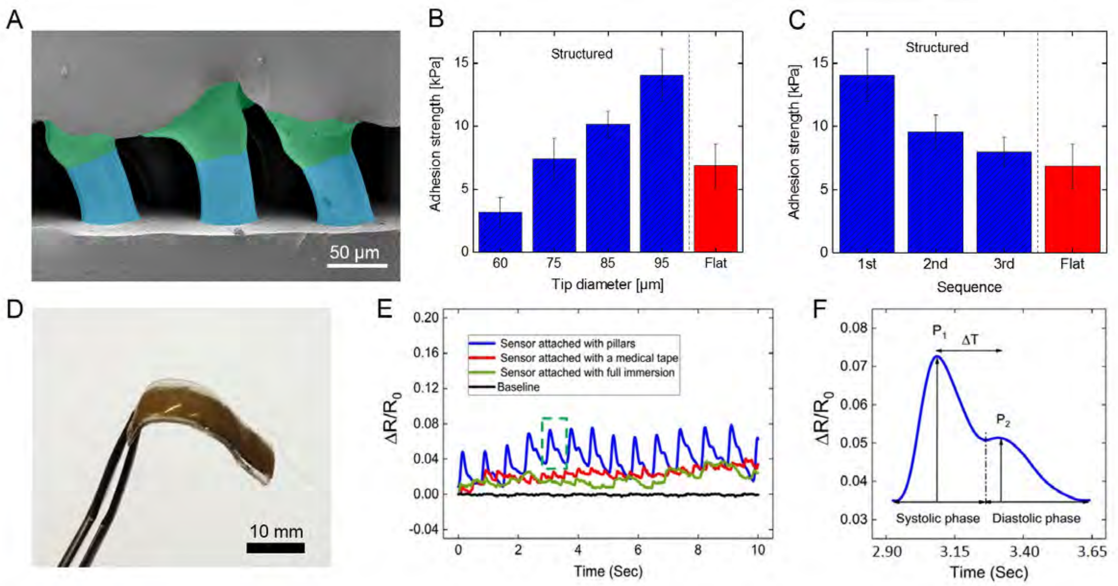


Figure 4.1: (a) Scanning electron microscope image of the skin adhesive film attached to a skin replica. (b) Adhesion of adhesive films with different vinyl siloxane tip diameters and (c) adhesion after multiple inking and attachment process. (d) Photograph of a highly flexible strain sensor on an adhesive film. (e) Output signal of the strain sensor mounted onto the wrist. (f) Response of the strain sensor during one cardiac cycle.

Considerable attention is being paid to wearable medical systems owing to their seamless integration with the human body, allowing continuous monitoring of important vital signs, such as respiratory and heart rate and greatly assists early diagnosis of diseases and subsequent therapy. Despite remarkable advances have been made in wearable medical devices, their conformal attachment to the rough and soft texture of the skin remains a grand challenge.

Herein, we propose a facile method for superior conformation and adhesion of bioinspired composite microfibers to the hierarchical topography of the skin (Figure 4.1(a)). The soft and stretchable skin-adhesive films are composed of poly(dimethylsiloxane) microfibers decorated with conformal and mushroom-shaped vinylsiloxane tips. Direct crosslinking of the viscous tips on the skin surface can greatly enhance the skin adhesion through their excellent shape conformation. High adhesion strength of 18 kPa is achieved after optimizing the pattern geometries and processing parameters (Figure 4.1(b)). The reusability of our adhesive films was tested by multiple inking and attachment process (Figure 4.1(c)).

We demonstrated our skin-adhesive films as a wearable device application by integrating them with flexible strain sensors. Silver nanoparticle thin film based strain sensors were fabricated on the top of the micropatterned PDMS films (Figure 4.1(d)). The ultrahigh gage factor of the strain sensor (ca. 767) together with strong bonding to the skin enables recording of the bloodflow pressure with improved signal-to-noise ratio of 59.7. Figure 4.1(e) depicts the response of the skin-adhesive sensor, clearly indicating the waveform pattern of the artery pulse. Furthermore, both systolic and diastolic phases and peaks were successfully identified (Figure 4.1(f)).

Strong skin adhesion combined with signal amplification of advanced wearable systems may improve noise-free, sensitive, and accurate monitoring of body signals.

More information: https://pi.is.mpg.de/project/bioinspired-microfibers-for-skin-adhesion-and-signal-amplification-of-wearable-sensors

## High-performance multiresponsive flexible paper actuators

*Morteza Amjadi, Metin Sitti*

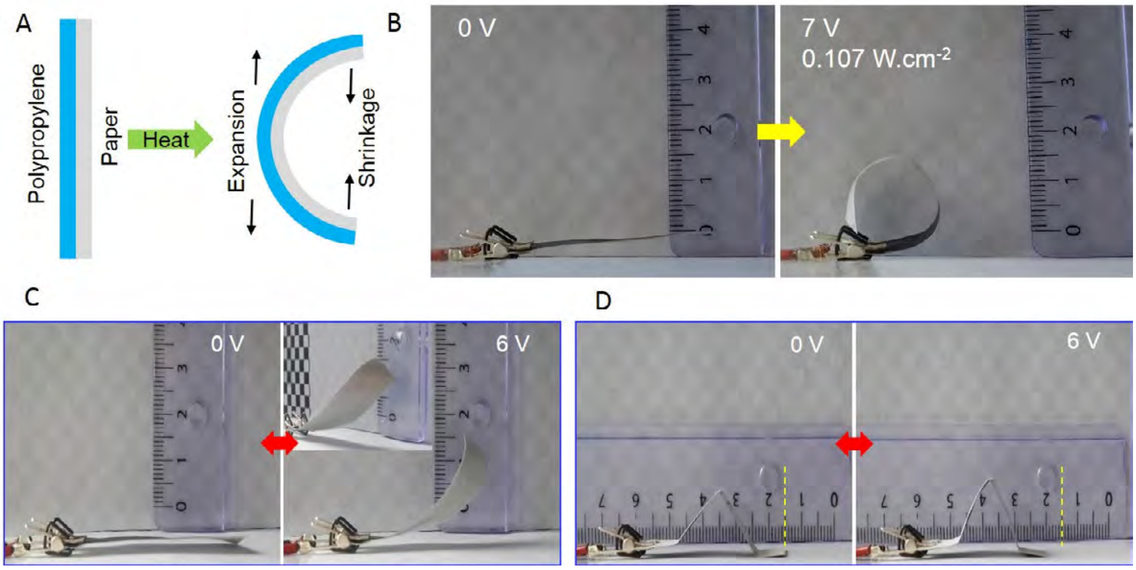


Figure 4.2: (a) Schematic illustration of a bilayer structure composed of two active layers. (b) Shape-changing behavior of an electrically activated paper actuator. (c) Left-handed twisting-bending motion of an actuator due to the anisotropic hygroscopic properties of the normal copy paper. (d) Linear back and forth motion of a folded actuator.

Soft actuator materials are being actively pursued owing to their importance in soft robotics, artificial muscles, biomimetic devices, and beyond. Electrically, chemically, and light-activated actuators are the mostly explored soft actuators. Recently, significant efforts have been made to reduce the driving voltage and temperature of thermoresponsive actuators, develop chemical actuators that can function in air, and enhance the energy efficiency of light-responsive actuators.

We have designed bilayer actuators that operate based on the large hygroscopic contraction of the copy paper and simultaneously large thermal expansion of the polypropylene film upon increasing the temperature. Different from the traditional bilayer designs that consist of an active layer bonded to a passive layer, our bilayer actuators are composed of two active layer materials (Figure 4.2(a)). To locally control the temperature of actuators, we directly deposited electroresistive heaters made of hybrid films of highly conductive silver nanowire network and highly robust PEDOT:PSS conducting polymer on the paper substrate.

As depicted in Figure 4.2(b), electrothermally activated bending actuators can function with low voltages ($\leq$ 8 V), low input electric power per area ($\leq$ 0.14 W cm-2), and low temperature changes ($\leq 35\ ^{\circ}C$). They exhibit reversible shape-changing behavior with curvature radii up to 1.07 cm-1 and bending angles up to 360 °. The paper actuators can produce high output forces and lift objects 53.7 times heavier than their own weight. Considering foldability and high anisotropic hygroscopic behavior of the copy paper, we have demonstrated a series of programmable mechanical motions such as bending, angular, linear, and torsional movements (Figure 4.2(c)(d)). Besides electrical activation, they can be powered by humidity or light irradiation.

More information: https://pi.is.mpg.de/project/paper-actuators

## Inflated soft actuators with reversible stable deformations

*Lindsey Hines, Kirstin Petersen, Metin Sitti*

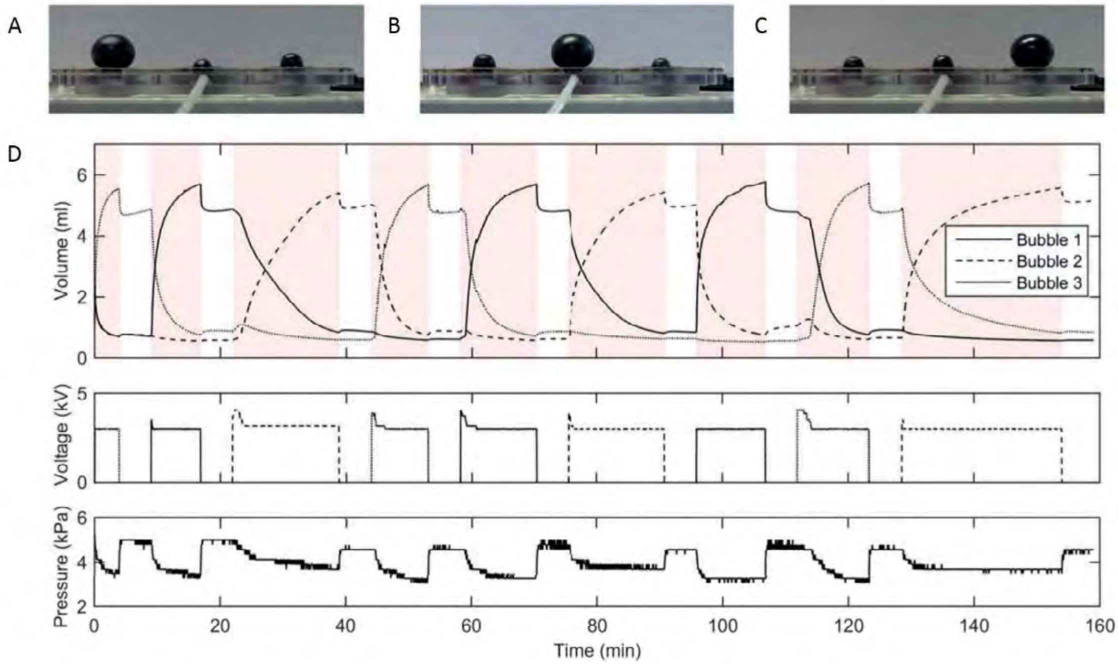


Figure 4.3: (a-c) Enclosed, inflated system with three stable shapes. Three membranes composed of an acrylic elastomer (VHB4910) share a single connected chamber. In (d), the system transitions between its stable configurations using DEAs. The shaded area marks when a DEA is active.

Soft robotic systems and actuators have the potential to make very robust and adaptable platforms and manipulators. While the majority of natural biological systems are soft-bodied, human-made systems are typically designed to be stiff and rigid to allow fast and precise motions. Though well suited for a manufacturing line, they can encounter safety problems when interacting with human, unknown or unpredictable objects and environments. Air or fluid filled systems, in particular, have been of considerable interest to researchers; by pumping or shifting the internal air or fluid, the system can be actuated while maintaining the desired soft flexible body. However, the vast majority remain tethered, relying upon large off board compressors and/or pumps limiting their practical applications.

We have developed soft, inflated actuators, and robots that are completely enclosed, unreliant upon pumps or compressors, easing their use both in the field and in internal medical procedures making them more scalable. Here we rely upon the nonlinear stress/strain characteristics of many hyperelastic materials in order to design inflated systems with several stable states, transitioning between them with the use of soft, dielectric elastomer actuators (DEAs). The hyperelastic films can exhibit 'snap-through' behavior, where after a given strain it becomes easier to induce more. Both before and after this cusp are points at which the balloon can sit stably at different volumes for the same internal pressure. By carefully choosing the film material and designing the inflated structure, one can create different stable shapes that can be transitioned between (Figure 4.3).

More information: https://pi.is.mpg.de/project/soft-inflatable-actuators

## Multifunctional ferrofluid-infused surfaces (FLIPS) with reconfigurable multiscale topography

*Wendong Wang, Dirk Drotlef, Joanna Aizenberg(Harvard University), L. Mahadevan(Harvard University), Metin Sitti*

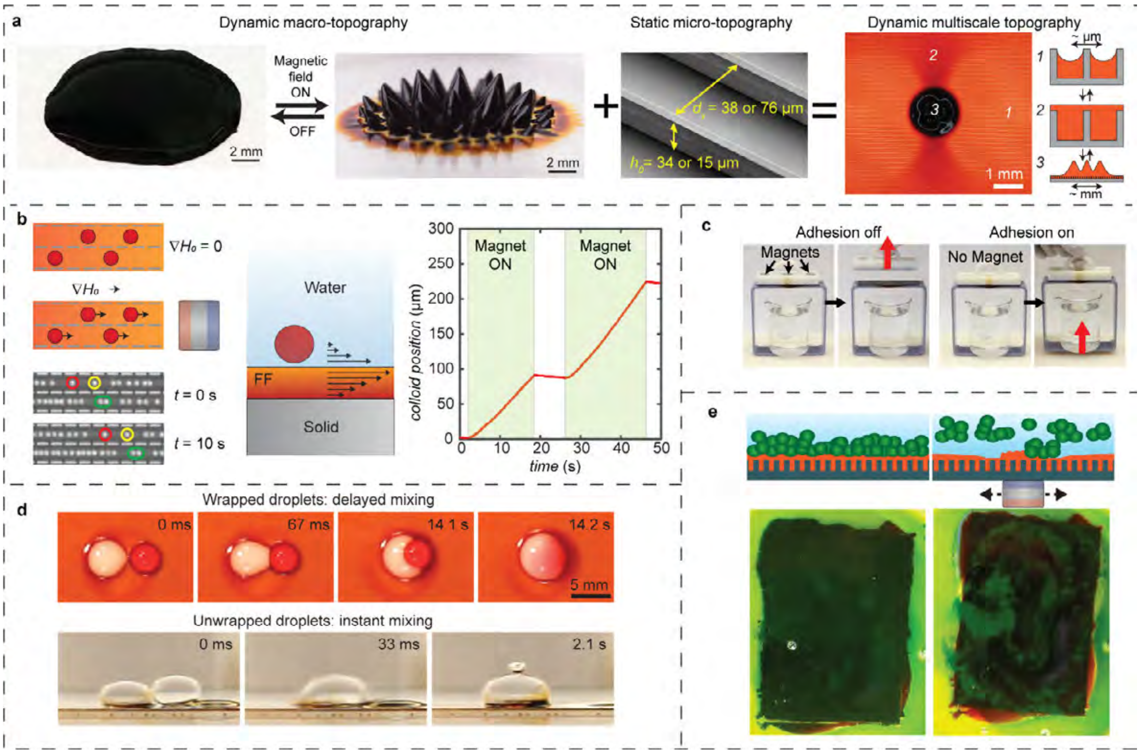


Figure 4.4: Multifunctional ferrofluid-infused surfaces with reconfigurable multiscale topography. (a) A diagram showing the concept of ferrofluid-containing liquid-infused porous surfaces (FLIPS). (b) Manipulation of non-magnetic colloidal particles on FLIPS. (c) Switchable adhesion on FLIPS. (d) Manipulation of droplets on FLIPS. (e) Biofilm removal on FLIPS.

Developing adaptive materials with changing geometries in response to external stimuli provides fundamental insights into the links between the physical forces involved and the resultant morphologies and creates a foundation for technologically relevant dynamic systems. Recently, gels, polymers, composites and slippery surfaces have been explored as reconfigurable surface to control interfacial properties. However, these concepts exhibit a limited range of topographical changes and thus a restricted scope of function.

In collaboration with a team of researchers, led by Prof. Joanna Aizenberg at Harvard University, we have developed a multifunctional ferrofluid-containing liquid infused porous surfaces (FLIPS) with dynamic reconfigurable multiscale topography. FLIPS is a composite surface, made up of two distinct parts: a microstructured solid substrate and a ferrofluid, which consists of magnetic particles suspended in a liquid (Figure 4.4(a)). We demonstrate the applications of FLIPS at multiple length scales, including manipulation of colloidal particles at micron scale (Figure 4.4(b)), manipulation of droplets at millimeter scale (Figure 4.4(d)), switchable adhesion (Figure 4.4(c)) and friction, pumping, and biofilm removal at centimeter scale (Figure 4.4(e)). Each of these applications can be further extended. Our results suggest that FLIPS allows much more diverse combinations of functions and capabilities than surfaces that have only a simple, single-scale topographical response. This could be a platform for a lot of future technologies.

More information: https://pi.is.mpg.de/project/multifunctional-ferrofluid-infused-surfaces-with-reconfigurable-multiscale-topography

## Parallel microcracks-based ultrasensitive and highly stretchable strain sensors

*Morteza Amjadi, Mehmet Turan, Cameron P. Clementson, Metin Sitti*

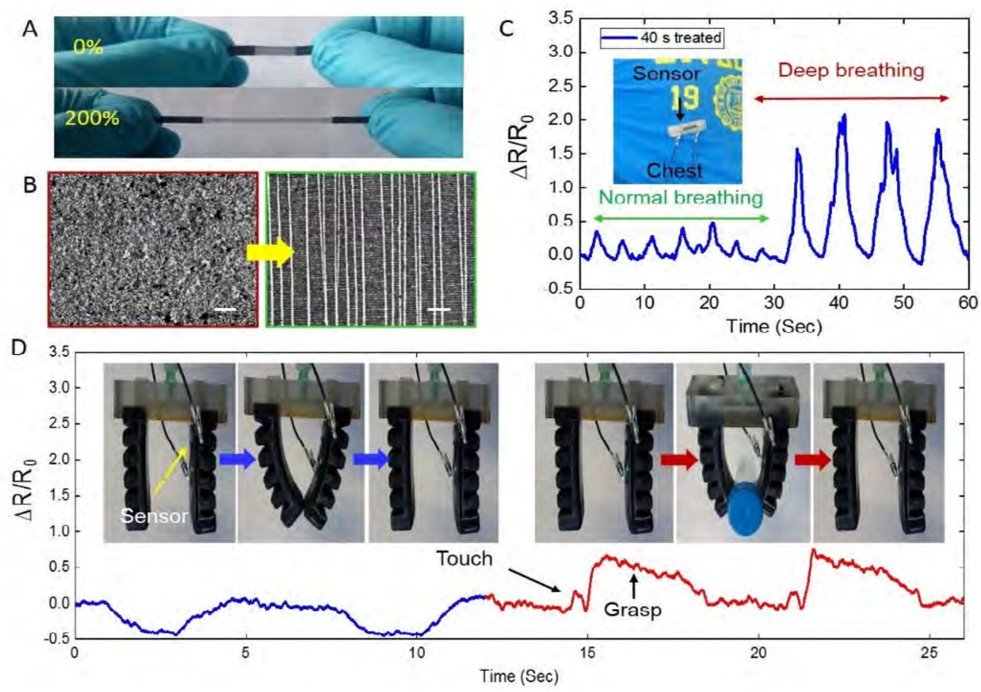


Figure 4.5: (a) Photographs of highly stretchable strain sensors based on graphite conductive films coated on Ecoflex elastomer. (b) Generation of microcracks in graphite films upon stretching. Scale bar: 100 $\mu$m. (c) Respiration monitoring when a strain sensor is attached to the chest. (d) Integration of soft sensors with soft gripper robots.

Flexible electronic devices are a rising interest in both research and industry due to their facile interaction with the human body and surrounding environment. In particular, demand is dramatically increasing for flexible sensors due to their potential applications in personalized health monitoring, human-machine interfaces, and soft robotics. Examples of such electronic devices include sensors for strain, pressure, touch, temperature, and a variety of bio-sensors.

In this project, we have developed highly stretchable strain sensors by combining graphite thin films as strain responsive conductive films and a flexible polymer support (Figure 4.5(a)). To boost the sensitivity, controllable microcracks have been generated in graphite films that can dramatically reduce the conductive passways upon stretching (Figure 4.5(b)). As a result, the strain sensors can simultaneously detect deformation amplitudes ranging from microscale ($\geq$ 10 $\mu$m) to millimeter scale ($\geq$ 10 mm) with ultrasensitivity (GFs $\geq$ 100) and high stretchability ($e \geq 100$).

We have demonstrated the advantages of our strain sensors in various potential applications such as human motion detection, sound visualization, pressure sensing, and soft robotics. Figure 4.5(c) shows the resistance change of a strain sensor mounted onto the chest area for respiration rate and pattern monitoring. As shown, there is an obvious resistance change during deep breathing since larger strain is accommodated by the sensor. Figure 4.5(d) illustrates the integration of our strain sensors with a soft gripper robot as feedback control and sensory skin. The integration enabled touch and grasp sensing through resistance change of the embedded sensors.

More information: https://pi.is.mpg.de/project/parallel-microcracks-based-ultrasensitive-and-highly-stretchable-strain-sensors

## Phase change of gallium enables highly reversible and switchable adhesion

*Zhou Ye, Guo Zhan Lum, Metin Sitti*

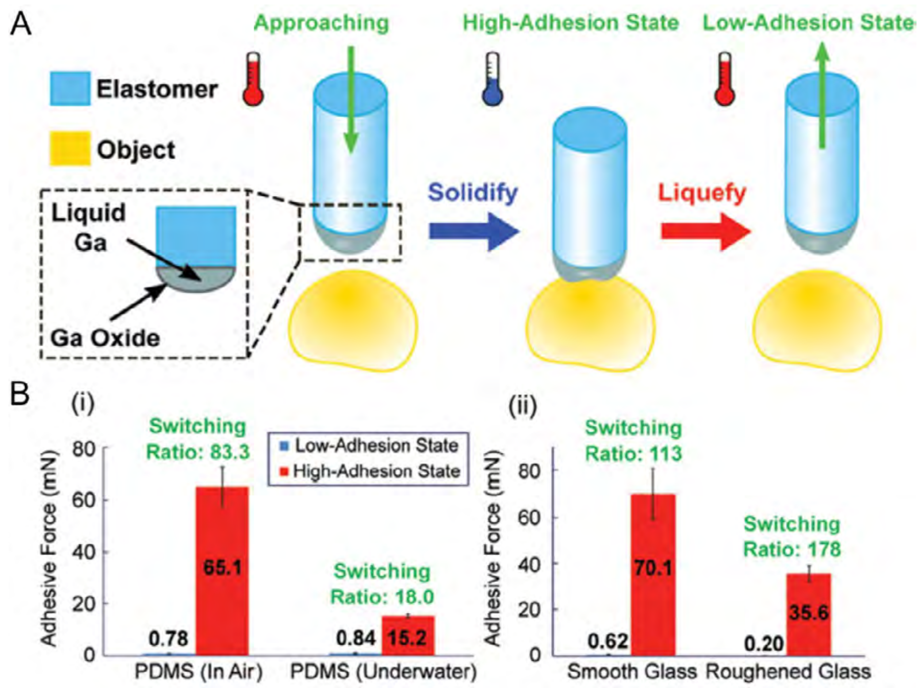


Figure 4.6: (a) Method of reversibly attach a liquid Gallium coated post to a substrate. (b) (i) Adhesion and switching ratios on smooth PDMS under dry and wet conditions and on smooth and roughened glass under dry conditions (ii).

It has remained a great challenge to synthesize a reversible adhesives that can work well on both dry and wet surfaces for a wide range of roughness. Inspired by the reversible adhesion of biological systems, we show how Ga can be used as synthetic reversible adhesive. Ga adhesion is enabled by the phase change ability of Ga near room temperature (30°C) as depicted in Figure 4.6(a). The reversible adhesion is achieved by switching between the liquid and solid phases, and can adhere to even rough and wet surfaces. The advantages of Ga adhesion over other synthetic systems is that it requires low detachment force and therefore exhibits a high maximum-minimum adhesion ratio. While adhesives with high ratios do exist (e.g., 1:35, 1:40, 1:204), they are generally only effective on smooth and dry surfaces. Furthermore, the adhesive properties of Ga were characterized through experiments that measured the adhesion of Ga-coated polydimethylsiloxane (PDMS) posts. Under dry conditions, Ga exhibited a high maximum-minimum adhesion ratio (33-113) across a range of materials with smooth surfaces. The method also worked well under wet conditions (ratio: 18, Figure 4.6b(i)) and on rough surfaces (ratio: 178, Figure 4.6b(ii)). Thus, we demonstrated one of the potential uses of Ga adhesive by performing pick-and-place operations on non-planar objects. We believe Ga adhesion can be implemented in broad adhesive applications including transfer printing, reconfigurable and climbing robots, electronic packaging and biomedicine.

More information: https://pi.is.mpg.de/project/phase-change-of-gallium-enables-highly-reversible-and-switchable-adhesion

## Self-sensing paper actuators based on graphite-carbon nanotube hybrid films

*Morteza Amjadi, Metin Sitti*

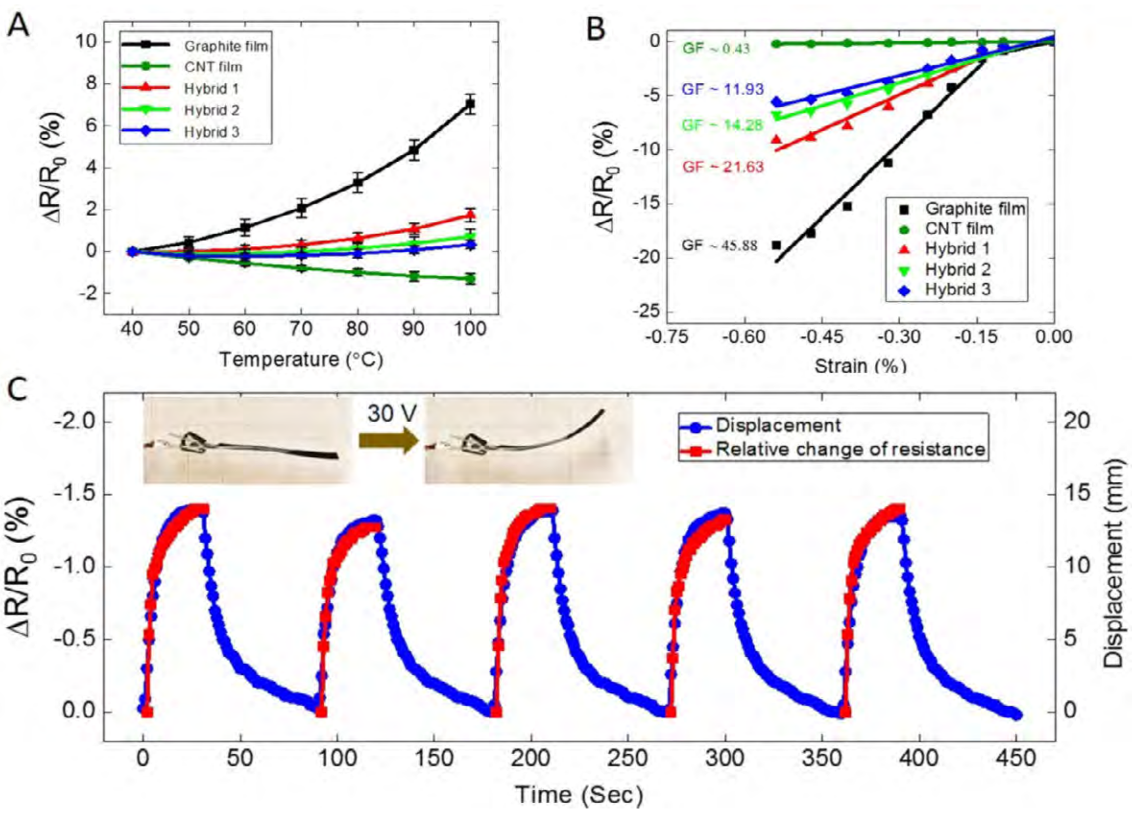


Figure 4.7: (a) Relative resistance change of graphite, carbon nanotubes, and hybrid films upon increasing the temperature. Hybrid films with varying mass ratio of graphite and carbon. (b) Piezoresistivity of graphite, carbon nanotubes, and hybrid films under compressive strain. (c) Self-sensing function of a paper actuator; insets, displacement of the actuator before and after application of 30 V.

In recent years, new emerging functional materials have considerably improved the deformation amplitude, response speed, force generation, and programmable motion output of soft actuators. Aside from efficient shape-changing behavior of soft actuators, real-time motion feedback is essential for their greater functionalities and wider adoption. To date, however, the actuation displacement is typically determined by bulky optical systems and image postprocessing, hindering effective and compact sensing capabilities of soft actuators.

In this project, we have added sensing function to our recently reported multiresponsive actuators composed of copy paper and polypropylene film. We have shown that the combination of functional materials overcomes the self-sensing limitation of current soft actuators. In this study, independent electrothermal stimulation and real-time displacement sensation have been accomplished by the hybridization of graphite microparticles and carbon nanotubes. Given nearly zero thermal coefficient of resistance (Figure 4.7(a)) and high piezoresistivity of hybrid films (Figure 4.7(b)), the signal-to-noise ratio of the proposed self-sensing actuators is significantly boosted to 66. As shown in Figure 4.7(c), relative change of the resistance matches well with the actuator tip displacement. For example, the self-sensing actuator exhibited -1.40% resistance change for around 13.9 mm tip displacement. Thus, unlike previous integrated actuators, the dynamic motion of our actuators can be monitored merely through two input electric terminals, without any need for additional sensing components or input energies.

More information: https://pi.is.mpg.de/project/self-sensing-paper-actuators-based-on-graphite-carbon-nanotube-hybrid-films

## Soft robotic grippers using gecko adhesives

*Sukho Song, Dirk Drotlef, Metin Sitti*

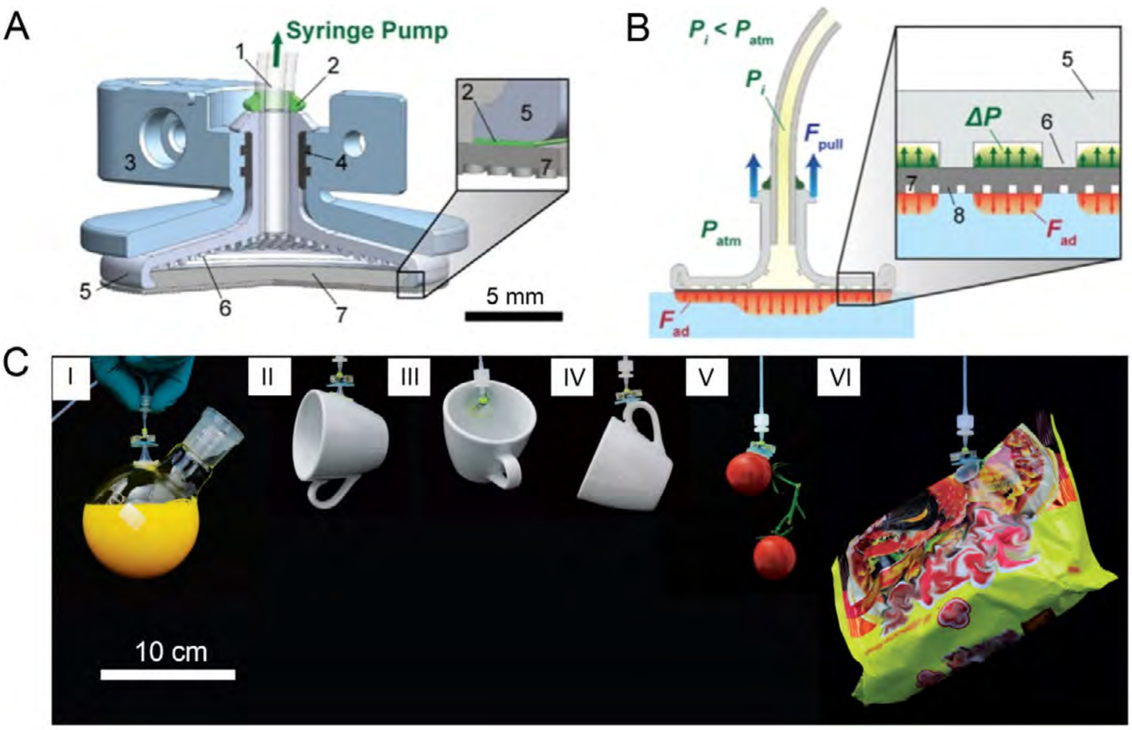


Figure 4.8: (a) Schematic of the adhesion-based soft-gripping system with attached gecko-inspired membrane and Illustration of the equal load sharing mechanism (b). (c) Demonstration of the soft gripper holding various 3D objects (I-VI).

In transfer printing, robotics, and precision manufacturing, adhesion-controlled grasping of complex 3D surfaces is very challenging. Current adhesion systems are limited by a fundamental trade-off between 3D surface conformability and high adhesion strength.

In this work, we overcome this trade-off with an adhesion-based soft-gripping system that exhibits enhanced fracture strength without sacrificing conformability to nonplanar 3D surfaces. Composed of a gecko-inspired elastomeric microfibrillar adhesive membrane supported by a pressure-controlled deformable gripper body (Figure 4.8(a)), the proposed soft-gripping system controls the bonding strength by changing its internal pressure and exploiting the mechanics of interfacial equal load sharing (Figure 4.8(b)).

The soft adhesion system enabling manipulation of various objects with complex 3D geometries (Figure 4.8(c)) such as convex or concave curvatures (I-IV), slightly rough (V) and deformable surfaces (VI). In general, the pull-off force of the soft adhesive system increases when the initial pressure decreases. The soft adhesion system can use up to $\sim 26\%$ of the maximum adhesion of the fibrillar membrane, which is $14\times$ higher than the adhering membrane without load sharing. Our proposed load-sharing method does not only enhance adhesion but also leads to an area scaling law similar to that of the natural geckos′ adhesive system. Such area scalability suggests that improved interfacial load sharing is critical when grasping 3D nonplanar geometries.

The proposed system pushes the upper limit on the maximum adhesion-controlled gripping strength and outperforms previous microfibrillar adhesive systems in handling complex 3D and deformable objects and surfaces.

More information: https://pi.is.mpg.de/project/controllable-load-sharing-for-soft-adhesive-interfaces-on-three-dimensional-surfaces

## Soiled adhesive pads shear clean by slipping: A robust self-cleaning mechanism in climbing beetles

*Guillermo Amador, Thomas Endlein, Metin Sitti*

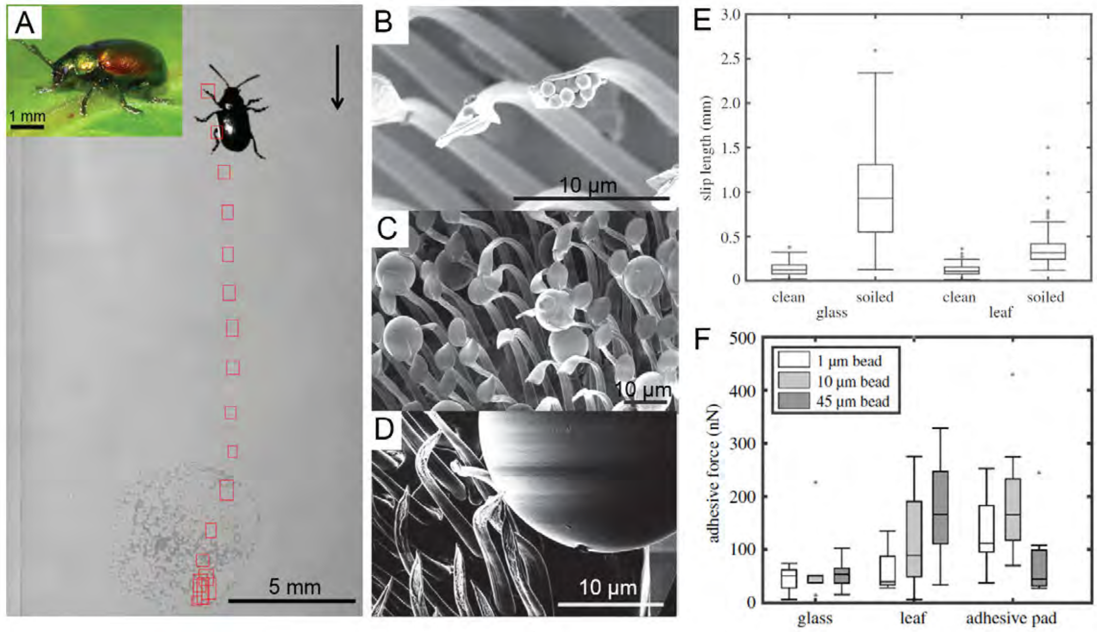


Figure 4.9: (a) Dock beetle climbing up on a clean glass surface after stepping into a patch of microbeads. (b) Soiled adhesive pads with remaining microbeads after 10 steps for (b) 1 $\mu$m, (c) 10 $\mu$m, and (d) 45$\mu$m beads. (e) The slip length for clean and soiled footpads on glass and leaf substrates. (f) Adhesive force on various substrates for AFM cantilevers with attached 1 $\mu$m, 10 $\mu$m and 45 $\mu$m beads.

Many animals, including insects, spiders, geckos and tree frogs, can climb smooth surfaces by using adhesive pads on their feet. Hairy pads exhibit a brush-like array of thin fibres (setae) and are common among spiders, geckos, flies and beetles. As these pads are inherently sticky, they attract dirt particles. Without cleaning, the pads would foul quickly and become unusable. While there may be several mechanisms used for cleaning, like grooming, brushing with legs or flushing with fluids, these mechanisms are time or energy consuming, and so would greatly impede locomotory performance.

Here, a self-cleaning mechanism is proposed whereby soiled feet would slip on the surface due to a lack of adhesion but shed particles in return. Our study offers an in situ quantification of self-cleaning performance in fibrillar adhesives, using the dock beetle as a model organism. After beetles soiled their pads by stepping into patches of spherical beads (Figure 4.9(a)), we found that their gait was significantly affected. Specifically, soiled pads slipped 10 times further than clean pads (Figure 4.9(e)). Like previous studies, we found that particle size affected cleaning performance. Large (45 $\mu$m) beads were removed most effectively, followed by medium (10 $\mu$m) and small (1 $\mu$m) (Figure 4.9(b)-(d)). Consistent with our results from climbing beetles, force measurements on freshly severed legs revealed larger detachment forces of medium particles from adhesive pads compared to a flat surface, possibly due to interlocking between fibers (Figure 4.9(f)). By contrast, dock leaves showed an overall larger affinity to the beads and thus reduced the need for cleaning. Overall, pad slippage and high adhesion of dock leaves were found to promote effective self-cleaning and reduced fouling. Self cleaning through slippage provides a mechanism robust to particle size and may inspire solutions for artificial adhesives, climbing robots or grippers.

More information: https://pi.is.mpg.de/project/soiled-adhesive-pads-shear-clean-by-slipping-a-robust-self-cleaning-mechanism-in-climbing-beetles

## Wrinkling instability and adhesion of a highly bendable gallium oxide nanofilm encapsulating a liquid-gallium droplet

*Muhammad Yunusa, Guillermo Amador, Dirk Drotlef, Metin Sitti*

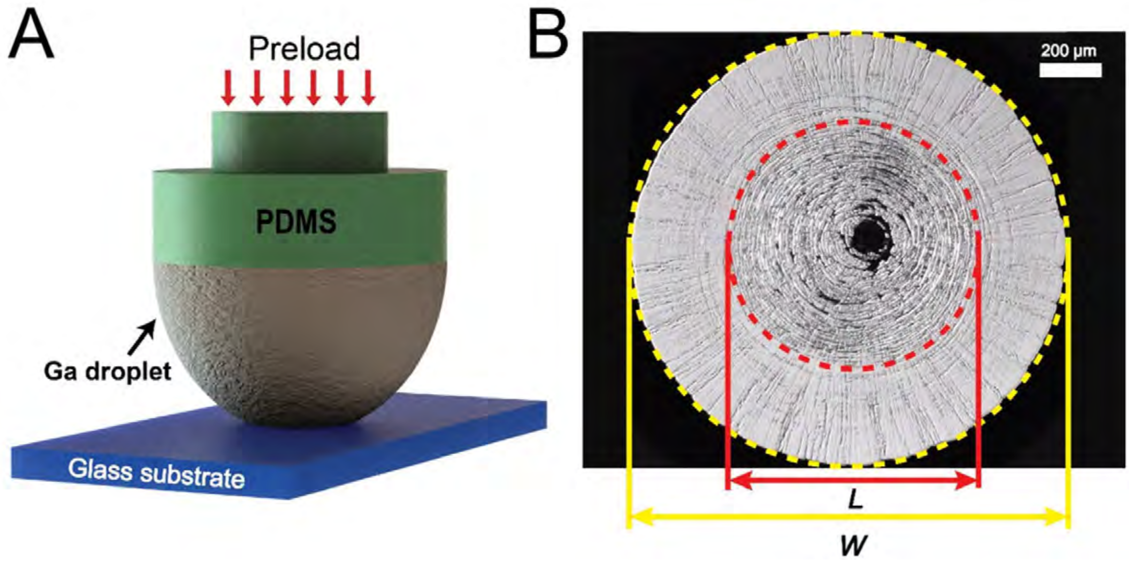


Figure 4.10: (a) A liquid Gallium droplet is attached to an elastomeric PDMS post and brought in contact with a flat smooth glass substrate. (b) The thin Galliumoxide skin is conformed to the smooth glass surface when mechanical load is applied. Wrinkling instability emerges at the contact interface to relieve the applied stress and radial wrinkles evolve under increased load.

Biological system can reversibly adhere to and detach from many uneven terrain. For instance, tree-frog adhesive pad can adhere to smooth, rough, wet, and dry surfaces. Most of the synthetic adhesive systems are only suitable for smooth and dry surfaces. To overcome the challenges facing the synthetic adhesives, Gallium (Ga) liquid metal offers a promising performance due to its phase change property at low temperature (30°C). Ga changes phase from liquid to solid isothermally at room temperature. The mechanism of liquid state adhesion of Ga resembles that of tree-frog pad, which is a very compliant system.

In this study, we present the mechanical behavior of the highly bendable native oxide skin on the surface of Ga droplet when compressed against a rigid flat substrate as it exhibits a fascinating wrinkling phenomenon, and its implications on the adhesion energy necessary to separate the interface. The applied compressive stress at the contact interface is relieved by the $Ga_2O_3$ sheet, transitioning from a circular into a radial wrinkled state (Figure 4.10(b)). The circular wrinkles enhance fracture strength by trapping cracks that propagate along the radial direction. The fundamental response explored in these studies is the out-of-plane wrinkling that occurs to alleviate compressive stresses. Wrinkling has implications in morphogenesis, interfacial-crack trapping similar to frog adhesive pad to enhance adhesion, and functional optical and electronic materials. Ga based liquid metals are attractive for a broad range of applications including flexible or stretchable sensors, robotics, microelectronics, 3D printing, and next generation nanomechanical devices. Additionally, we can predict when the oxide nanofilm ruptures and leaves residue on the surface, which may further facilitate cost effective room-temperature 2D patterning of liquid Ga for future semiconductor applications. Our findings provide the tools for the precise control of liquid metal interfaces.

More information: https://pi.is.mpg.de/project/wrinkling-instability-and-adhesion-of-a-highly-bendable-gallium-oxide-nanofilm-encapsulating-a-liquid-gallium-droplet

## Bacteria-propelled bio-hybrid microswimmers (Bacteriabots)

*Oncay Yasa, Yunus Alapan, Ajay Vikram Singh, Byung-Wook Park, Babak Mostaghaci, Jiang Zhuang, Metin Sitti*

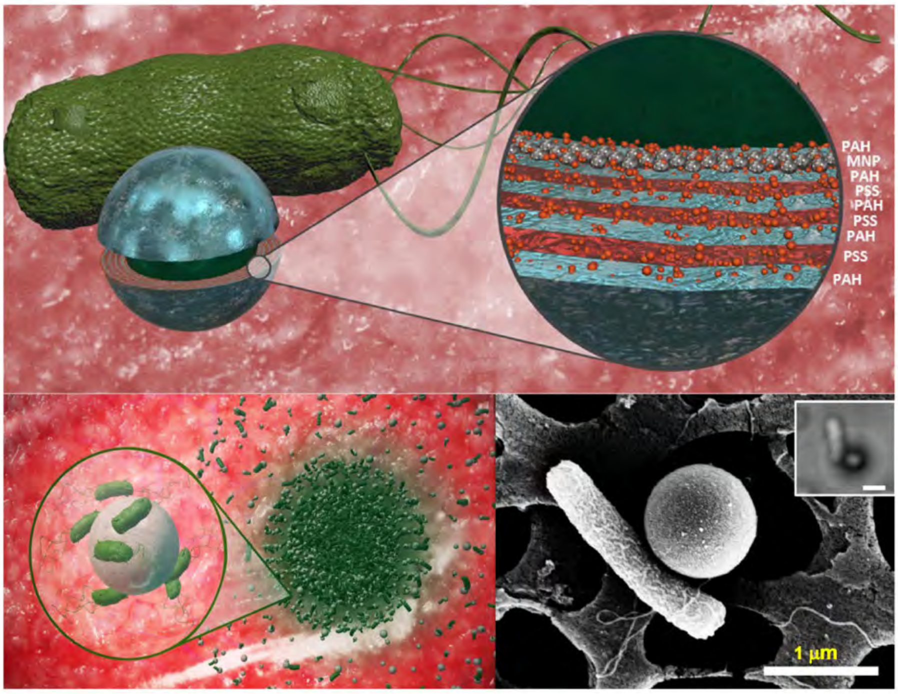


Bacterial microrobots, which are formed by integration of biological entities with synthetic constructs or materials, have been at the forefront of minimally invasive theranostic applications for the past decade. The potential of bacterial microrobots to enable *in vivo* active cargo delivery applications lies in their autonomously functioning biological units, which provide active propulsion and environmental sensing capabilities, and also in their steerability using external magnetic or acoustic fields. Various bacterial species, mainly *Escherichia coli* and *Serratia marcescens*, have been extensively investigated in our department as biological units of bacterial microrobots for the active delivery of cargo moieties. Bacteria, and thus bacterial microrobots, can efficiently swim using nutrients that are already present in the local microenvironment or inside the bacterial cell and are known to have a diverse sensory system, which allows their taxis-based self-guidance *in vitro* through gradients of pH, oxygen, and different attractant chemical molecules. Autonomous propulsion and sensing abilities of bacteria have been exploited for enhancing functions of the bacterial microrobots for their propulsion, guidance, and delivery of various cargos through their synthetic carriers, including bare or coated microbeads, double emulsions, nanofibrous hydrogels, and electropolymerized polypyrrole microtubes. In addition, bacterial microrobots can also be externally guided and selectively brought to target sites via magnetic steering of the microsystem through either magnetic cargo unit or natural magnetotactic bacteria. Overall, microrobot design approaches developed in this department take advantages of both materials and bacterial cells to enable autologous cargo-carrier architectures for the fabrication of biocompatible, biodegradable, and multifunctional microrobots that can be used in future *in vivo* medical applications.

More information: https://pi.is.mpg.de/project/bacteria-propelled-microswimmers

## Biodegradable magnetic microswimmers

*Hakan Ceylan, Immihan Ceren Yasa, Metin Sitti*

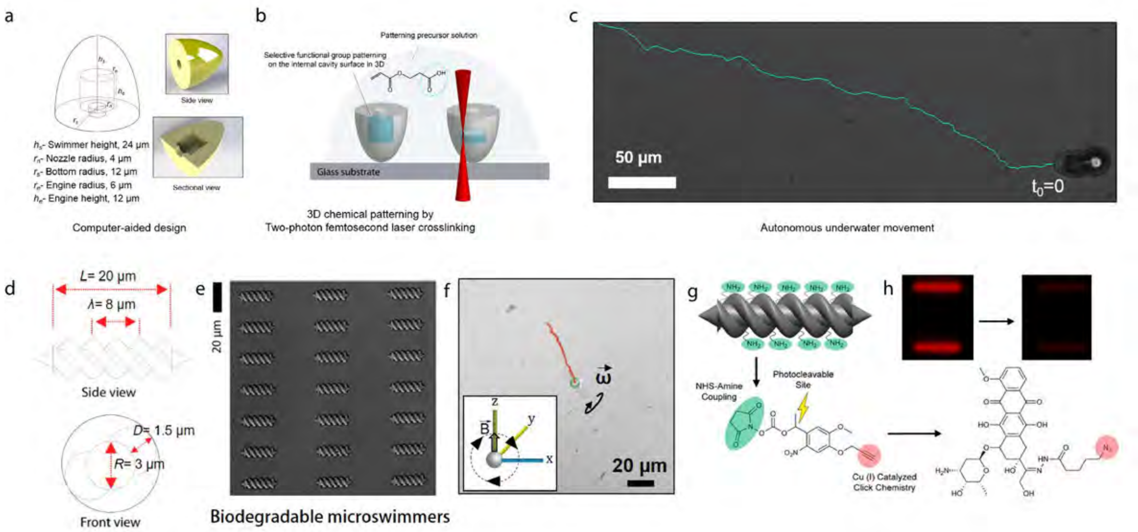


Figure 4.11: 3D chemical patterning of micromaterials for encoding active functionalities. (a) Computer-aided design (CAD) of a low-drag microswimmer with a programmed inner cavity as catalytic bubble-production site, or engine, that opens outside through a nozzle to produce jet bubbles for propulsion. (b) 3D chemical patterning of the inner cavity surface for carboxylic acid group display. (c) Trajectory of an exemplary catalytic microswimmer moving in hydrogen peroxide solution. (d) CAD of a double-helical microswimmer. (e) 3D-printed magnetic hydrogel microswimmers with 90% water content (f) Propulsion of a microswimmer under rotating magnetic fields. (g) Deploying chemotherapeutic drug doxorubicin to the microswimmer through light-cleavage linker. (h) Light-triggered drug release from the microswimmers.

Programmed microscopic carriers that are able to navigate, sense their surroundings, adapt to changing conditions, and perform a set of functions in the physiological environment will revolutionize many clinical practices. The microscopic size makes them unrivalled for accessing small, highly confined and delicate body sites, where conventional medical devices fall short without an invasive intervention. Nevertheless, realization of such aspects in both on-board, *i.e.*, autonomous, and off-board, *i.e.*, externally guided, approaches presents fundamental challenges concerning design, fabrication process, and encoding operational capabilities. Conventional microfabrication techniques mostly provide relatively simple structures with limited design flexibility and function. Realization of complex designs with compartmentalized functionalities in 3D is a daunting task at the micron scale. Our research focuses on the use of additive manufacturing technologies to enable complex microrobots and microactuators.

The use of magnetic fields is a prominent way of remote powering and control of medical microrobots. In contrast to other untethered power transfer alternatives, such as light and chemical fuels, magnetic fields provide a biocompatible source of energy and are able to safely and uniformly penetrate biological tissues. Catalytic microswimmers usually rely on non-biocompatible fuel sources for propulsion, and those that are moving with biocompatible fuels are unable to move inside biologically relevant ionic media. Powering with light is limited with the penetration depth, safely deliverable light intensity and the line-of-sight exposure, so it is not applicable to confined and complex *in vivo* environments. Using acoustic fields is a promising method for off-board propulsion and manipulation of microswimmers. However, the functional design of such a microswimmer and its application in a biological setting is currently insufficient and requires future developments. To rotate the double helix, rotating magnetic fields are needed, which then create a torque on the microswimmer through a magnetic axis defined perpendicular to the helical axis.

Biodegradability, *i.e.*, decomposition over time as a result of the resident biological activity, is a critical aspect of microrobotic design for their safe operation in the living environment. When the prescribed task is accomplished, the safest option for removing the microrobots from the body is to expect their degradation to non-toxic, metabolized products. The use of non-degradable materials can result in serious acute and chronic toxicities, which could require surgical revision, and hence lower the overall desired benefit from the microrobot. As a result, materials that predictably degrade and disappear in a safe manner have become increasingly important for medical applications (Figure 4.11). Microrobotic systems developed so far have not tackled the issue of biodegradability, so it complicates their clinical use due to possible adverse effects in the body.

More information: https://pi.is.mpg.de/project/biodegradable-microswimmers

## Red blood cell based microswimmers (RBCbots)

*Yunus Alapan, Oncay Yasa, Metin Sitti*

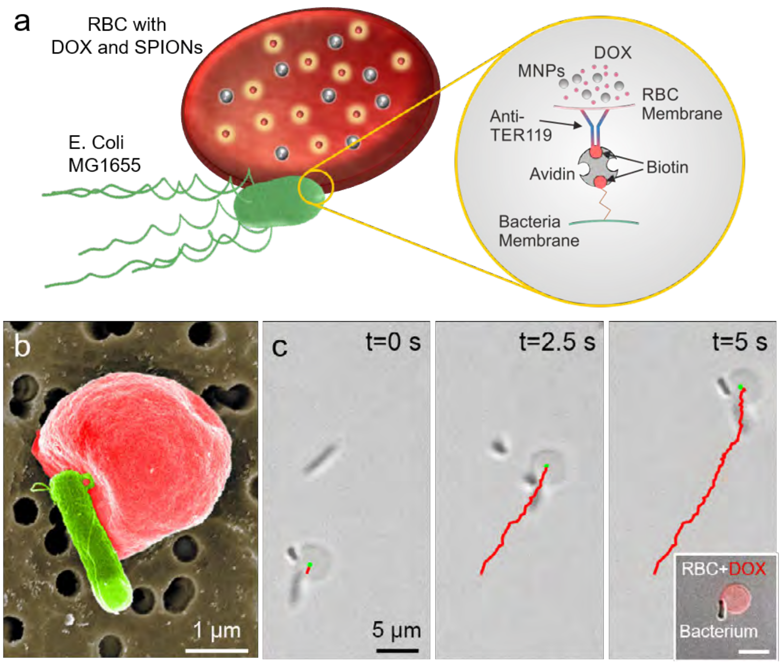


Figure 4.12: A. RBC microswimmers are composed of an RBC, loaded with drug molecules and superparamagnetic iron oxide nanoparticles, bounds to a motile bacterium via biotin-avidin-biotin binding complex. B. Scanning electron microscope image (pseudo-colored red: RBC; pseudo-colored green: bacterium) of an example RBC microswimmer with an attached bacterium. C. Example 2D propulsion trajectories of RBC microswimmers via bacteria over time. The inset displays a bacterium attached to an RBC, loaded with DOX molecules. Scale bar represents 5 $\mu$m.

Bacteria-powered biohybrid microswimmers have recently shown to be able to actively transport and deliver cargos encapsulated into their synthetic constructs to specific regions locally. However, use of synthetic materials as cargo carriers can result in inferior performance in load-carrying efficiency, biocompatibility, and biodegradability, impeding clinical translation of biohybrid microswimmers. In this project, we report construction and external guidance of bacteria-driven microswimmers using red blood cells (RBCs, erythrocytes) as autologous cargo carriers for active and guided drug delivery, shown in Figure 4.12. Multifunctional biohybrid microswimmers were fabricated by attachment of RBCs (loaded with anti-cancer doxorubicin drug molecules and superparamagnetic iron oxide nanoparticles (SPIONs)) to bioengineered motile bacteria, *E. coli* MG1655, via biotin-avidin-biotin binding complex. Autonomous and on-board propulsion of biohybrid microswimmers were provided by bacteria, and their external magnetic guidance was enabled by SPIONs loaded into the RBCs. Furthermore, bacteria-driven RBC microswimmers displayed preserved deformability and attachment stability even after squeezing in microchannels smaller than their sizes, as in the case of bare RBCs. In addition, an on-demand light activated hyperthermia termination switch was engineered for RBC microswimmers to control bacteria population after operations. RBCs, as biological and autologous cargo carriers in the biohybrid microswimmers, offer significant advantages in stability, deformability, biocompatibility, and biodegradability over synthetic cargo-carrier materials. The biohybrid microswimmer design presented in this study transforms RBCs from passive cargo carriers into active and guidable cargo carriers toward targeted drug and other cargo delivery applications in medicine.

More information: https://pi.is.mpg.de/project/bloodcellbasedmicroswimmers

## Dynamic and programmable self-assembly of micro-rafts at the air-water interface

*Wendong Wang, Gaurav Gardi, Vimal Kishore, Metin Sitti*

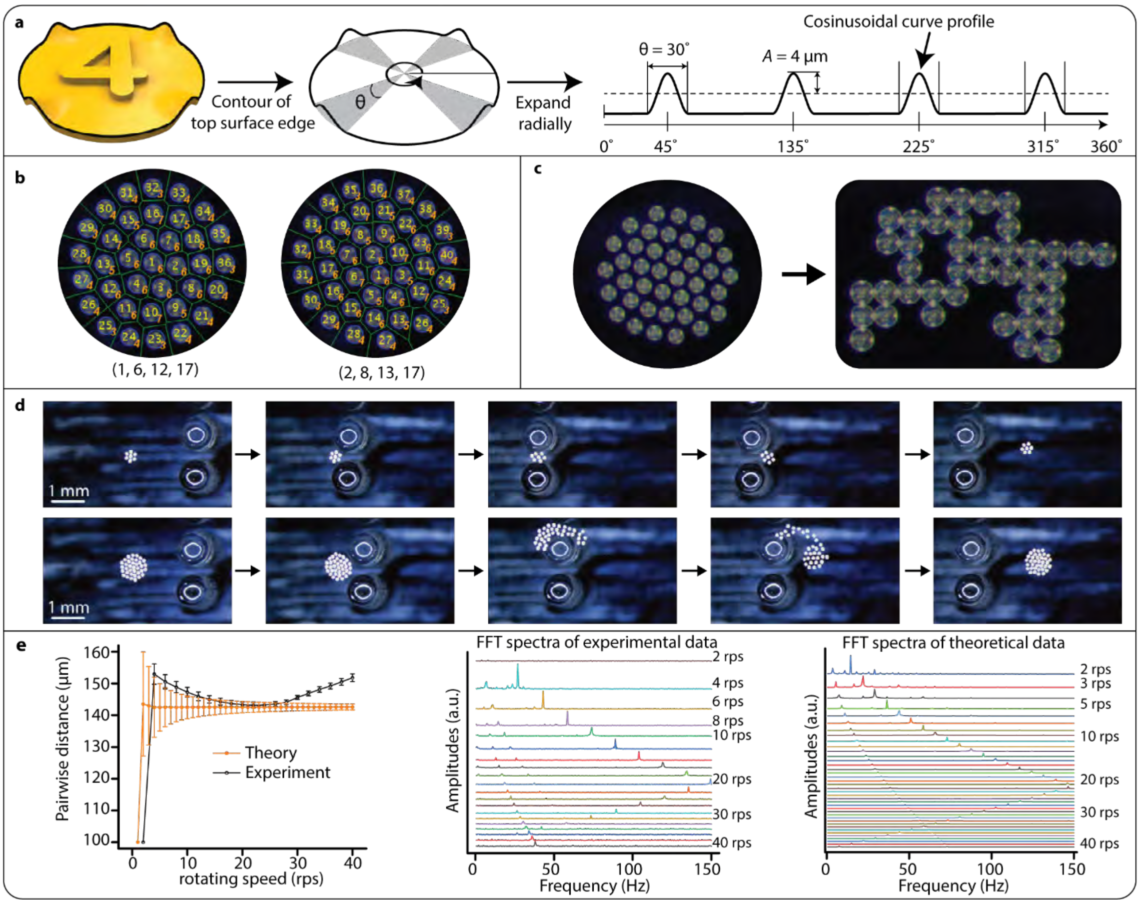


Figure 4.13: Fig. A collective system built from the dynamic and programmable self-assembly of spinning micro-rafts at air-water interfaces. (a) The parametric design of one representative 3D-printed micro-raft. (b) Dynamic patterns of 36 and 40 rafts. (c) Programmable self-assembly of 40 rafts. (d) Channel-crossing of small (7 rafts) and large (36 rafts) assemblies, showing a size-dependent emergent behavior. (e) Mechanistic study of pairwise interactions, showing decent agreement between experimental and theoretical pairwise distance curves.

By linking concepts from chemistry and materials science to techniques in fluid mechanics and robotics, we proposed a platform of collective microrobots based on dynamic and programmable self-assembly of circular magnetic micro-rafts at the air-water interface. The cosinusoidal edge-height profiles of these micro-rafts not only create a net dissipative capillary repulsion that is sustained by continuous torque input, but also enables directional assembly of micro-rafts (Figure 4.13(a)-(c)). We have demonstrated collective behaviors such as channel crossing (Figure 4.13(d)). In addition, in a collaboration with Eric Lauga's group at the University of Cambridge, we have investigated the mechanism of the interactions between rafts, which will help us to simulate and predict collective behaviors in the future. We anticipate that this dynamic and programmable self-assembled materials system will serve as a model system for studying non-equilibrium dynamics and statistical mechanics as well as for use as a novel platform to assemble and control microrobots in the future.

More information: https://pi.is.mpg.de/project/dynamic-and-programmable-self-assembly-of-micro-rafts-at-the-air-water-interface

## Microalga-powered microswimmers toward active cargo delivery (Algabots)

*Oncay Yasa, Pelin Erkoc, Yunus Alapan, Metin Sitti*

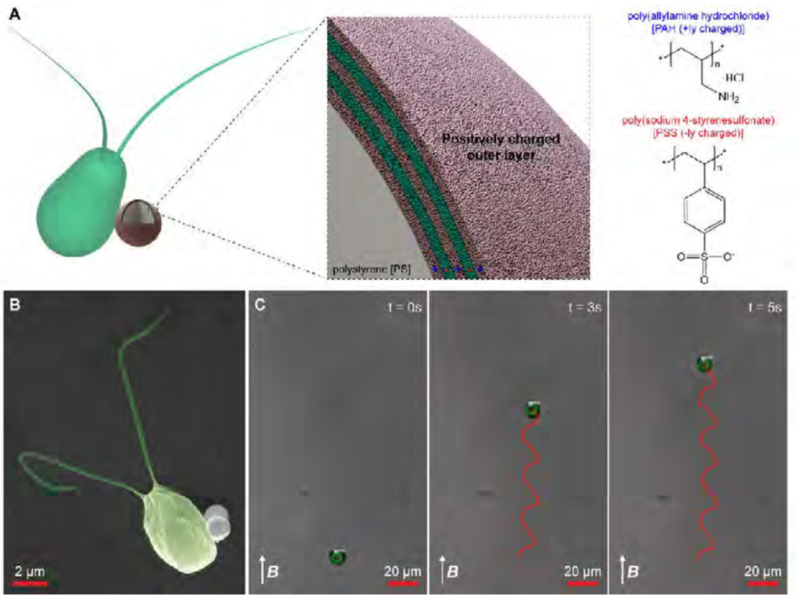


Figure 4.14: Overview of the layer-by-layer polyelectrolyte deposition onto spherical magnetic polystyrene microparticles and fabrication of the microswimmers utilizing electrostatic interactions between the microalga and the functionalized microparticles. Poly (allylamine hydrochloride) [PAH] and poly (sodium 4-styrenesulfonate) [PSS] are utilized as positively and negatively charged polyelectrolytes, respectively. B. SEM image (pseudo colored green, C. reinhardtii) of an example algal microswimmer. C. Example 2D propulsion trajectories of an algal microswimmer under 26 mT uniform magnetic field.

Nature presents intriguing biological swimmers with innate energy harvesting abilities from their local environments. Use of natural swimmers as cargo delivery agents presents an alternative strategy to transport therapeutics inside the body to locations otherwise difficult-to-access by traditional delivery strategies. Even though bacteria are heavily utilized as actuators in biohybrid microswimmer designs for active cargo delivery applications, their possible acute pathogenicity necessitates search for an agile biological swimmer with better biocompatibility, such as microalgae. In this project, we report a biocompatible biohybrid microswimmer powered by a unicellular freshwater green microalga, *Chlamydomonas reinhardtii*. Polyelectrolyte-functionalized magnetic spherical cargoes (1 $\mu$m in diameter) were attached to surface of the microalgae via non-covalent interactions without the requirement for any chemical reaction. Three-dimensional swimming motility of the constructed biohybrid algal microswimmers was characterized in the presence and absence of a uniform magnetic field in the $x$-direction. In addition, motility of both microalgae and biohybrid algal microswimmers was investigated in various physiologically relevant conditions, including cell culture medium, human tubal fluid, plasma and blood. Furthermore, it was demonstrated that the algal microswimmers are cytocompatible when co-cultured with certain healthy and cancerous cells. Finally, fluorescent isothiocyanate-dextran (a water-soluble polysaccharide) molecules were effectively delivered to mammalian cells using the biohybrid algal microswimmers as a proof-of-concept active cargo delivery demonstration. The microswimmer design described here presents a new class of biohybrid microswimmers with greater biocompatibility and motility for targeted delivery applications in medicine (Figure 4.14).

More information: https://pi.is.mpg.de/project/microalga-powered-microswimmers-toward-active-cargo-delivery

## Vampire bat-inspired multimodal locomotion robot (Multimo-Bat)

*Hyun Gyu Kim, Matthew Woodward, Metin Sitti*

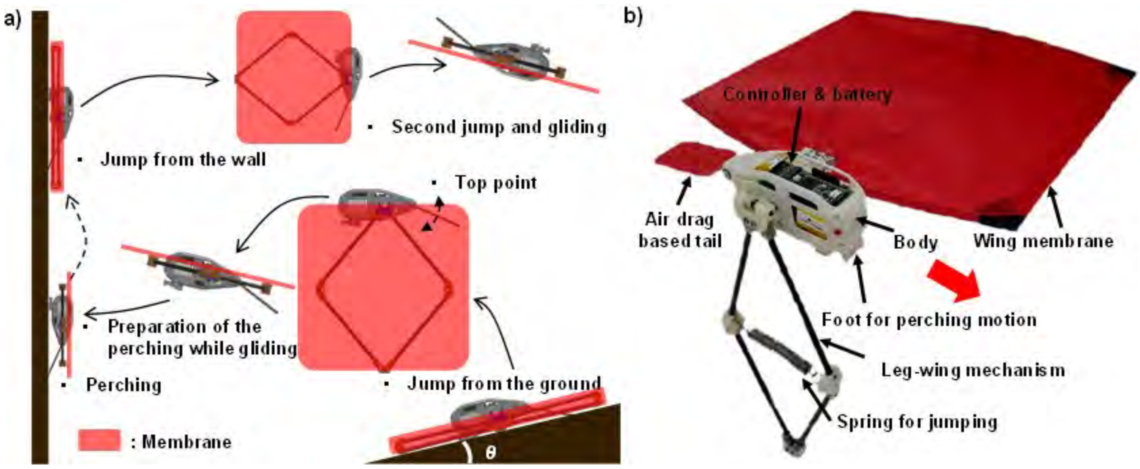


Figure 4.15: The overal design of Multimo-bat. (a) Overall scenario of integrated three motions: jumping, gliding, and perching. The robot jumps from the ground by compressed leg mechanism; the leg mechanism become a wing with a membrane in the air; The robot clings on the wall with additional passive foot mechanism; The robot jumps on the wall for second jump and gliding. (b) Robot prototype. A red arrow denotes moving direction of the robot. The leg-wing mechanism can be rotated to change modes between jumping and gliding motions.

Mobility in unstructured environments is a significant challenge for robotic systems; however, there are systems capable of operation in these environments. For example, flying squirrels, frogs, and snakes are able to jump from a tree and glide to move to other trees. Wood peckers, are able to fly and perch on the bark. These integrated motions provide a maneuverability to avoid danger from predators or obtain food. Even though technology does not allow for replication, with the correct level of abstraction, these systems can inspire designs which can improve the performance of their robotic counterparts.

Inspired by *Desmodus Rotundus* (common vampire bat), who is able to integrate jumping, flying and perching motions by using only one structure, we aim to develop a robot which integrates jumping, gliding and perching motions with minimal necessary components, shown in Figure 4.15. Particularly, we analyze not only this biological system which demonstrates the desired locomotion modes, but also significant levels of integration between the modes. The integration concepts are then abstracted from this organism and applied to the development of a robotic system. The robot has a leg for jumping by storing energy in leg muscles. The leg also functions as a wing with a membrane on its top. The addition of a foot allows for clinging to vertical surfaces. This motion integration creates the potential to improve the mobility of robots in unstructured environments by minimizing the necessary components. These minimal components can result in a smaller, lighter, and higher performance system than that of a system which combines these motions independently.

In summary, the benefit of our approach is that employing multiple locomotion strategies can significantly improve the mobility of systems operating in unstructured terrain. By utilizing an integrated approach for the addition of locomotion modes, the performance of individual modes can be preserved while reducing the additional structure and actuation required, therefore, improving overall system performance.

More information: https://pi.is.mpg.de/project/multimo-bot

## Morphological intelligence counters foot slipping in the desert locust and dynamic robots

*Matthew Woodward, Metin Sitti*

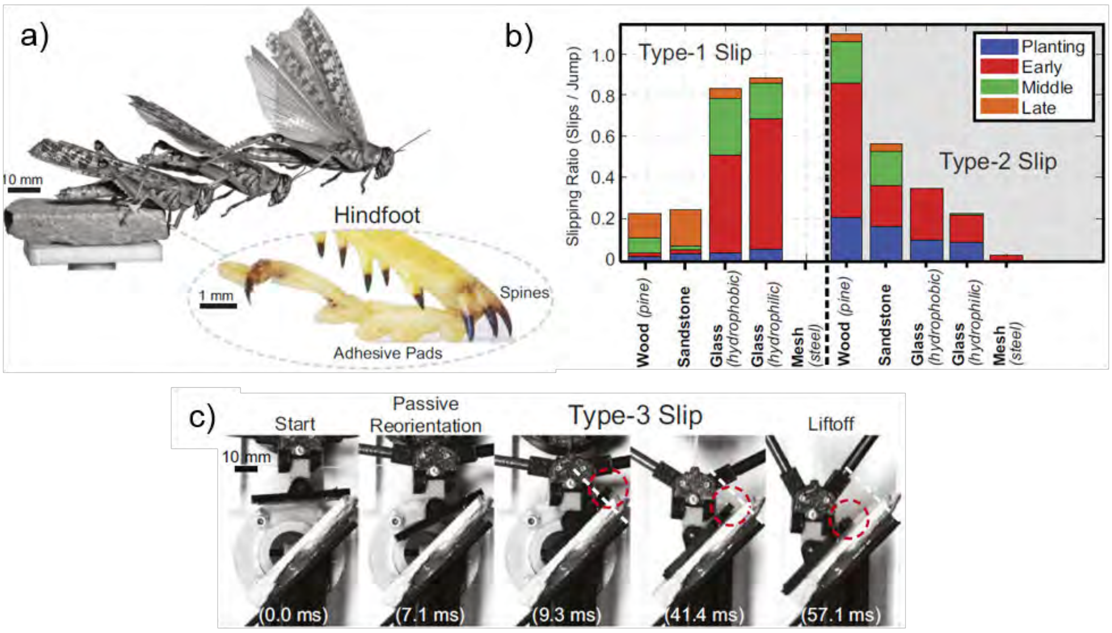


Figure 4.16: The jumping mechanism of desert locust. a) The desert locust's (S. gregaria) jumping behavior and hindfoot. b) Locust experimental results per jump (per-leg rates are half), divided into slip type (type 1 or 2) and energy regions (planting, early, middle, or late). c) Robot noncontact jumps (sandstone), where the foot preload is zero (high-speed video snapshots).

During dynamic terrestrial locomotion, animals use complex multifunctional feet to extract friction from the environment. However, whether roboticists assume sufficient surface friction for locomotion or actively compensate for slipping, they use relatively simple point-contact feet. We seek to understand and extract the morphological adaptations of animal feet that contribute to enhancing friction on diverse surfaces, such as the desert locust (***Schistocerca gregaria***), which has both wet adhesive pads and spines. A buckling region in their knee to accommodate slipping, slow nerve conduction velocity (0.5-3 m/s), and an ecological pressure to enhance jumping performance for survival further suggest that the locust operates near the limits of its surface friction, but without sufficient time to actively control its feet. Therefore, all surface adaptation must be through passive mechanics (morphological intelligence), which are unknown. We report the slipping behavior, dynamic attachment, passive mechanics, and interplay between the spines and adhesive pads, studied through both biological and robotic experiments, which contribute to the locust's ability to jump robustly from diverse surfaces (Figure 4.16). We found slipping to be surface-dependent and common (e.g., wood 1.32+/-1.19 slips per jump), yet the morphological intelligence of the feet produces a significant chance to reengage the surface (e.g., wood 1.10+/-1.13 reengagements per jump). Additionally, a discovered noncontact-type jump, further studied robotically, broadens the applicability of the morphological adaptations to both static and dynamic attachment. Our results demonstrate the potential contribution of morphological intelligence to solving complex dynamic locomotion problems. Furthermore, the concepts discovered can be easily adapted to, for the enhancement of, existing simple miniature and state-of-the-art large-legged terrestrial robots.

More information: https://pi.is.mpg.de/project/morphological-intelligence-counters-foot-slipping-in-the-desert-locust-and-dynamic-robots

## Pill-sized swallowable endoscopic soft capsule robots (MASCE)

*Donghoon Son, Mehmet Turan, Gilbert Hunter, Metin Sitti*

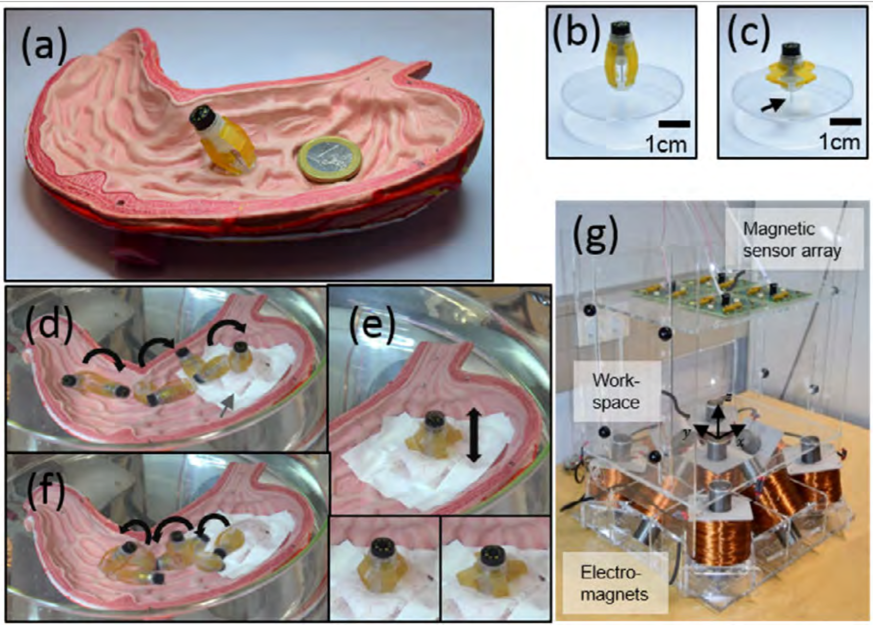


Figure 4.17: The design, setup and performance of the capsule robot. (a) Magnetically-actuated soft capsule endoscopy with fine-needle capillary biopsy (B-MASCE) is developed in the size of a pill (diameter: 12 mm, length: 30 mm). (b-c) Robot can be collapsed based on external magnetic field. The robot exposes the biopsy needle (the black arrow indicates) by the collapse. (d-f) B-MASCE demonstrates fine-needle capillary biopsy on a human anatomy stomach model. The robot takes a biopsy of a porcine fat as a phantom tumor. (g) The actuation and localization system for B-MASCE. The system is composed of a magnetic sensor array and nine electromagnets.

Current medical technologies are converging to minimally invasive diagnosis and therapy. The efforts to reduce pain and discomfort for the diagnosis of gastrointestinal (GI) tract gave birth to a new state-of-art technology, wireless capsule endoscopes (WCEs). However, the lack of the mobility of the WCEs limits its application to only one-dimensional GI tracts relying on peristaltic motions. Due to energy and space limitation of the capsule, magnetically actuated soft capsule endoscopes (MASCEs) have been identified as useful. Developed MASCE can move inside a stomach using the external magnetic field, release drug at a specific position, and implanted.

Currently, there are three tasks in our group for the MASCE (Figure 4.17). First, a magnetic localization method, which is compatible with the magnetic actuation is presented. This method provides reliable feedback for controlling the MASCE in real-time. We are aiming to combine this method with vision algorithms, which enables 3-D mapping of a stomach and localization simultaneously. This will improve the quality of localization and provide a stomach 3-D map for further diagnosis. Second, by exploring the non-uniform magnetic field actuation system and using the localization information, the robot is controlled robustly inside the plastic human anatomical model. The custom designed electromagnetic system is able to generate specific magnetic force and torque for the motion for the navigation and a special diagnostic function, such as biopsy. At last, we demonstrated the fine-needle capillary biopsy *in vitro* using the current robot design platform and the custom designed magnetic system. The robot is able to take a biopsy from a deep tissue, including sub-mucosal tumors. This method enhances the diagnostic accuracy of biopsy capsule robots.

More information: https://pi.is.mpg.de/project/pill-sized-endoscopic-swallowable-robots

## Shape-programmable soft-bodied millirobots with multimodal locomotion

*Wenqi Hu, Guo Zhan Lum, Xiaoguang Dong, Metin Sitti*

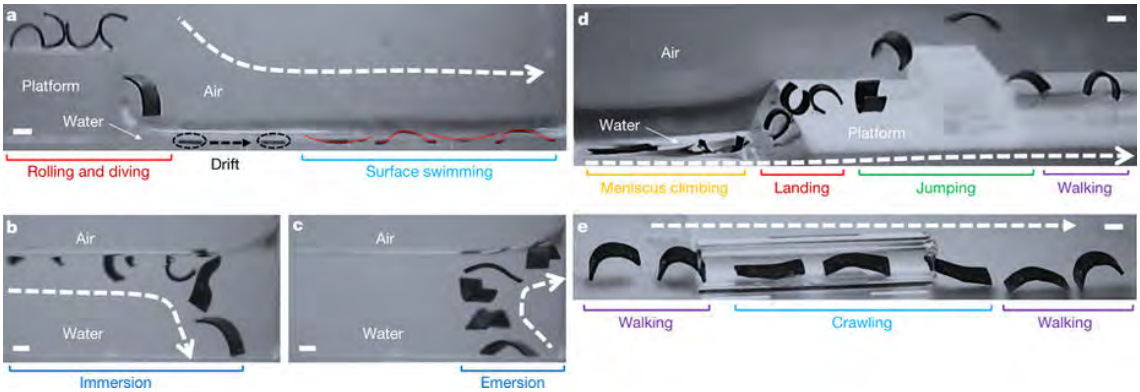


Figure 4.18: Our multimodal locomotion demonstration over a hybrid liquid–solid environment. (a) The soft robot rolls and dives from a solid platform into the adjacent water pool, where it drifts away along the water meniscus. The undulating robot then swims rightwards. (b, c) The robot rotates, disengages from the water surface, sinks, and subsequently swims up from the pool bottom to emerge again at the water–air interface. (d) The robot climbs up a water meniscus, lands on the solid platform, jumps beyond a standing obstacle, and walks away. (e) The robot walks towards a tubular tunnel (diameter: 1.62 mm) that impedes its walking gait. The robot then switches to the crawling mode to cross the tunnel, and finally walks away. The locomotion modes were sequentially captured in four separate videos owing to the restrictions of the workspace. Only one robot is used in this illustration. Scale bars: 1 mm.

At small scales, shape-programmable magnetic materials have significant potential to achieve mechanical functionalities that are unattainable by traditional miniature machines. Unfortunately, these materials have only been programmed for a small number of specific applications, as previous work can only rely on human intuition to approximate the required magnetization profile and actuating magnetic fields for such materials. We proposed a universal programming methodology that can automatically generate the desired magnetization profile and actuating fields for soft materials to achieve desired time-varying shapes. The universality of the proposed method can, therefore, enable other researchers to fully capitalize the potential of shape-programming technologies, allowing them to create a wide range of novel soft active surfaces and devices that are critical in robotics, material science, and medicine.

By using the above methodology, we then addressed a grand challenge facing the existing small-scale robots: *multimodal locomotion in complex terrains*. Previous miniature robots have very limited mobility because they are unable to negotiate obstacles and changes in texture or material in unstructured environments. In our research, we demonstrate magneto-elastic soft millimeter-scale robots (Figure 4.18) that can swim inside and on the surface of liquids, climb liquid menisci, roll and walk on solid surfaces, jump over obstacles, and crawl within narrow tunnels. These robots can transit reversibly between different liquid and solid terrains, as well as switch between locomotive modes. They can additionally execute pick-and-place and cargo-release tasks. We have also developed theoretical models to explain how the robots move. Besides their great biomedical potentials, these soft small-scale robots could be used to study soft-bodied locomotion produced by small organisms.

More information: https://pi.is.mpg.de/project/shape-programmable-soft-bodied-millirobots-with-multimodal-locomotion

## 4.4 Awards & Honors

### 4.4.1 Awards & Honors

#### 2018

**Metin Sitti** received the Rahmi Koç Medal of Science, in Turkey, which is given to one world-wide pioneering Turkish-origin scientist each year.

**Wenqi Hu** received the Günter Petzow Prize.

Visiting Professor **David Gracias** received the Humboldt Research Award.

**Hamed Shahsavan** received the Canadian NSERC postdoctoral fellowship.

**Wenqi Hu**, **Guo Zhan Lum**, and **Metin Sitti** received the Innovator of the Year Award in medical devices topic by the Design & Elektronik Magazine in Germany due to the invention of soft medical robots.

**Ahmet F. Tabak** and **Metin Sitti** were the Best Medical Robotics Award Finalist in the IEEE Robotics and Automation Conference for the research paper entitled, "Mechanical Rubbing of Blood Clots Using Helical Robots Under Ultrasound Guidance".

**Abdon Pena-Francesch**, **Zoey Davidson**, **Ville Liimatainen**, **Xinghao Hu**, **Yubing Guo**, and **Utku Culha** received the Humboldt postdoctoral fellowship.

#### 2017

**Donghoon Son** and **Metin Sitti** were the Best Medical Robotics Award Finalist in the IEEE Robotics and Automation Conference for the research paper entitled, "Magnetically Actuated Soft Capsule Endoscope for Fine-Needle Aspiration Biopsy".

#### 2016

**Yunus Alapan** and **Hunter Gilbert** received the Humboldt postdoctoral fellowship.

### 4.4.2 Faculty Appointments

#### 2018

**ChangKyu Yoon** became an assistant professor at Ewha Womans University, South Korea.

**Pelin Erkoç** became an assistant professor at Bahçeşehir University, Turkey.

**Shuhei Miyashita** became an assistant professor at University of Sheffield, UK.

**Wenqi Hu** became a permanent senior staff scientist at Physical Intelligence Department, Stuttgart, Germany.

**Ajay Vikram Singh** became a senior staff scientist at German Federal Institute for Risk Assessment, Berlin, Germany.

**Seok Kim** became a tenured associate professor at UIUC, USA.

**Burak Aksak** became a tenured associate professor at Texas Tech, USA.

**Morteza Amjadi** will become an assistant professor at Heriot-Watt University, UK.

## 2017

**Hunter Bryant Gilbert** became an assistant professor at Louisiana State University, USA.

**Massimo Mastrangeli** became an assistant professor at Delft University of Technology, Netherlands.

**Ahmet Fatih Tabak** became an assistant professor at Bahçeşehir University, Turkey.

**ByungWook Park** became an assistant professor at Youngstown University, USA.

**Guo Zhan Lum** became an assistant professor at Nanyang Technological University, Singapore.

**Bahareh Behkam** became a tenured associate professor at Virginia Tech, USA.

**Carlo Menon** became a tenured associate professor at Simon Fraser University, Canada.

## 2016

**Kirstin Peterson** became an assistant professor at Cornell University, USA.

**Zeinab Hosseini** became an assistant professor at McMaster University, Canada.

**Alex Spröwitz** became a W2 group leader in our institute's Stuttgart site.

**Sehyuk Yim** became a senior staff scientist at KIST, South Korea.

## 4.5 Equipment

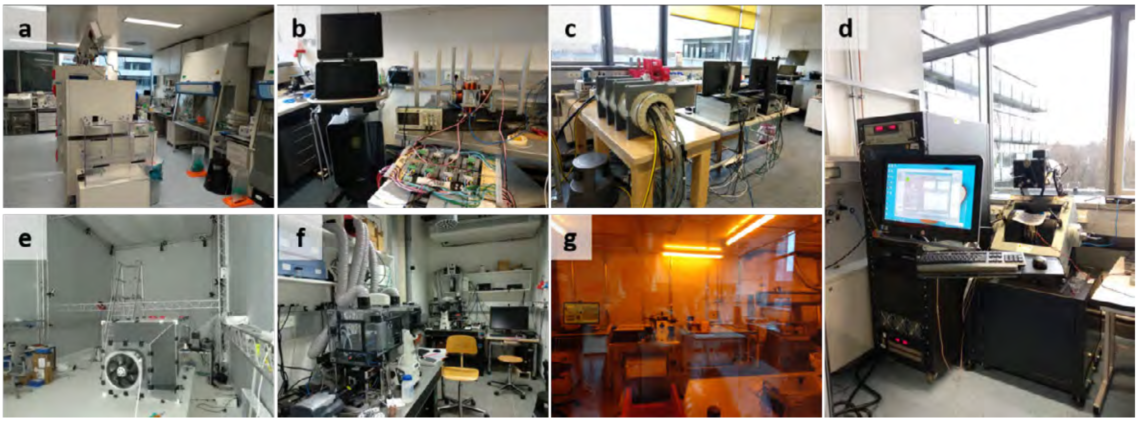

Figure 4.19: A selection of advanced equipment in the PI department. (a) The biological labs with cell culture instruments. (b) Preclinical animal ultrasound and customized coil system. (c) The MRI gradient coil system. (d) Vibrating sample magnetometer (VSM). (e) The cage system equipped with a Vicon system for researching jumping-gliding robot. (f) The two advanced confocal microscopes in the biological characterization lab. (g) The portable cleanroom with Nanoscribe 3D nano-printer and optical lithography system.

To serve our highly multi-disciplinary research, the department is equipped with biological, chemistry, robotic and microfabrication labs. These labs contain general lab instruments such as fume hoods, fridges, safety cabinets for storage, hotplates, stirrers, vortexes, semi analytical balances, precision balances, vacuum ovens and pumps, as well as ultrasonic baths. Besides these general lab supplies, our three major research fields are also equipped with a long range of advanced equipment.

### 4.5.1 Advanced Material

To fabricate different materials, the department owns a small clean room (Class 10,000). The cleanroom has a fume hood with spin coater (Laurell Technologies Corporation) and a mask aligner (MJB-4 Suss MicroTec) to implement standard microfabrication procedure. The lab is also equipped with two Nanoscribe Photonic GT 3D printing systems to manufacture devices of 1-300 $\mu$m size with sub-micron resolution. The same lab hosting the clean room also has a parylene coating system to render devices and robots being fabricated bio-compatible and implantable. For surface treatment of samples such as cleaning of contaminated surfaces, plasma activation of polymer coating or etching for microfabrication the lab utilizes a Diener electronic plasma cleaner. The central cleanroom of the institute is equipped with a roll-to-roll system (Coatema) owned by our department to scale up the production of novel micro/nano-structured materials. To characterize the material being synthesized, the department has a characterization room for scanning electron microscopy (SEM). It includes a SEM (Gemini 500) from Zeiss. It also has a Critical Point Dryer (Leica EM ACE600) and a sputtering machine for pre-treatment of the sample. For magnetic material, a vibrating sample magnetometer (VSM, Microsense EZ7) is available to characterize the hysteresis property. To study soft material, which is a recent key sub-research field in our department, there is a tensile and peel adhesion tester (Instron), as well as a Rheometer (TA Instruments, Discovery HR-3) for studying/measuring the rheological properties. To further characterize the surface property of the material, we have a contact angle goniometer/tensiometer (Kruss DSA100 Drop Shape Analyzer) to measure the contact angle where liquids or vapor meets solid surfaces, and to determine surface-free energy and the wettability of samples. To support our new research momentum in liquid crystal elastomer (LCE), the chemistry lab is equipped with a Zeiss cross-polarized microscope and a pol-scope (Polaviz) for characterizing and mapping LCE molecular orientations. A rheometer from TA instruments is available to measure mechanical properties of LCEs. At last, the department has a dedicated lab to support the research in adhesion and friction properties of artificially produced surfaces. The lab owns a Zeiss inverted compound microscope with a customized force transducer setup for measuring friction and adhesion in the milli-Newton range, and a Keyence Laser Microscope (VK-X200) with the ability to scan surface profiles in 3D with a lateral resolution of 120 $\mu$m and a z-resolution of about 0.5nm, as well as combining optical images with 3D reconstructions. The same lab also has a vibration-free table from Newport, a three-axis motorized stage to move samples with high precision (<1 $\mu$m resolution), I/O-boards for data acquisition, and high-speed Pointgrey cameras to record synchronized with the force measurements.

### 4.5.2 Mobile Millirobots

To manufacture, characterize, and test larger scale rigid and soft robots the department has a mechanical lab equipped with hand tools, a drill press, a Haas vertical CNC mill, an Isel ICV desktop CNC, a manual lathe, and a vertical band saw. The department also has multiple SLA and extrusion-based 3D printers to fast-prototype the robot. For cutting delicate 2D structures, the general robot lab of the department also includes a 60W CO2 laser cutter and engraver from Epilog, and an ultraviolet laser mill from LPKF with 25 $\mu$m resolution for cutting and shaping a variety of materials including plastic, glass, ceramics and composites. To characterize the Multimo-bat robot, the department owns a cage system (5m long x 3.2m width x 4.8m height , shown in the figure) fully equipped with a Vicon system setup to visually track and characterize robot movement patterns. Aero dynamic characteristics of the robot can be measured by a customized wind tunnel (Maximum flow speed and measurable size= 4m/s and 0.3m x 0.3m). To characterize millimeter magnetic robot, the department owns five coil systems. One of them has 5 coils and is integrated with a Zeiss microscope. It is controlled by a real-time linux system; the other four coil systems consist of 5-10 small coils and are controlled by a real-time Linux system from National Instruments. These four coils systems are all compatible with either camera or ultrasound imaging system as the position feedback. In addition to these five coil systems, the department also acquired an MRI gradient coil system for investigating microrobot control inside a medical MRI system. In the coming year, a full MRI system will be established to control and image our robots at the same time. Besides this, more than four high-speed cameras (Phantom) are also available to quantitatively analyze the kinematics and kinetics of the robots being built. In order to monitor the robot in a biomedical environment, the department also owns a preclinical animal ultrasound system with frame rate up to 500 Hz and an X-ray system with down to 100 $\mu$m resolution.

### 4.5.3 Mobile Microrobots

Most of our microscale robots are built and characterized in the biological labs. The department is equipped with four general biological labs for culturing cells and unicellular microorganisms (biosafety level 1) and four biological characterization labs. The general biological labs are all well-equipped with the standard tools (biosafety cabinets, incubators, autoclaves, etc.) to allow cell, various bacteria and microalgae culture. To characterize the mobile microrobots, the biological characterization labs are equipped with a variety of microscopes. This includes a Zeiss inverted microscope (Zeiss Axio Observer. A1), a Nikon Inverted Routine microscope (Nikon ECLIPSE Ts2), a Nikon inverted fluorescence microscope (Nikon ECLIPSE Ti-E), a Nikon Spinning Disk Confocal microscope (Nikon ECLIPSE Ti-E equipped with Yokogawa CSU-X1 Spinning Disk), a Leica Stereomicroscope (Leica M205 FA), and a Leica SP8 Single Point Confocal in combination with Widefield Fluorescence microscope (Leica DMi8 plus SP8). With these microscopes, we are able to perform a large range of experiments, from toxicity tests of robots on cells, to long-term monitoring of robot-cell interaction, to high-speed imaging of bacteria swimming with robots attached, and to high resolution confocal 3D reconstruction of robot-splenocyte interface. We also have a high speed refrigerated bench top centrifuge (SIGMA3-30KS) with gravitational fields up to 60,000 x g to allow for effective separation of micro and nanoparticles used in micro bot fabrication as well as for a number of molecular biology related procedures. For separation of mammalian cells, bacteria and biohybrid microrobots, we use a fluorescence-activated cell sorting (BD Biosciences, LSRFortessa X-20) instrument located in the bio characterization lab. Besides these microscope, researchers also use a Quartz Crystal Microbalance with dissipation (LOT-QuantumDesign) to study sensor devices, including multilayer buildup, drug uptake/release, and bacteria adhesion. To measure the surface property of the material, we own an Atomic Force Microscope (JPK). This is a type of scanning probe microscopy with a resolution on the order of fractions of a nanometer. This tool is also used to perform microfabrication, distinguish samples based on their mechanical properties, and to characterize the properties of living and soft materials such as bacteria, cells, and biopolymers. To monitor and characterize thin film systems and to develop advanced functional materials, the department makes use of an Attenuated Total Reflection Fourier Transform Infrared Spectroscopy (Bruker Tensor II). The same lab also utilizes a potentiostat/Electrochemical Impedance Spectroscopy device (Metrohm Autolab) in virtually all electrochemical applications, including characterization of surface electrical properties and the behavior of biological samples. At last, a Dantec Dynamics stereo micro PIV system was purchased to characterize fluid-robot interaction. Such system is able to cover the fluidic field analysis in from micro- to milli-meter scale and is used to support our research in swimming robot.

## 4.6 Director profile: Metin Sitti

Metin Sitti received a B.Sc. and M.Sc. from Boğaziçi University in 1992 and 1994, respectively, and a Ph.D. degree in electrical engineering from the University of Tokyo, Japan in 1999. During 1999-2002, he was a research scientist at the University of California, Berkeley, USA. During 2002-2016, he was a tenured full professor of Department of Mechanical Engineering and Robotics Institute at Carnegie Mellon University (CMU), Pittsburgh, USA, directing the NanoRobotics Laboratory. During 2011-2012, he was a visiting professor at Harvard University, EPFL and Sorbonne University.

In fall 2014, Metin Sitti founded the Physical Intelligence Department with over 20 PhD students, 21 post-doctoral fellows, 5 technicians, and over 15 visiting scientists/students. He served as the Managing Director of the Max Planck Institute for Intelligent Systems from 2015 to 2017 and is a board member in the Perspectives Committee of the Max Planck Society, the International Max Planck Research School on Intelligent Systems, and the Center for Learning Systems between ETH Zurich and Max Planck Institute for Intelligent Systems. He has been an honorary professor at the University of Stuttgart, Germany since 2017, and a professor of Medical School and Engineering School at Koç University in Istanbul, Turkey since October 2018.

He is a recipient of the Rahmi Koç Science Prize, the National Science Foundation Faculty Early Career Development Award, and the SPIE Nanoengineering Pioneer Award. He is an IEEE Fellow. He was elected as an IEEE Robotics and Automation Society Distinguished Lecturer during 2006-2008. He has been a finalist for and won best paper, video and poster awards at major robotics and materials conferences and journals more than eighteen. His group's recent research breakthroughs have been featured in the popular press, such as New York Times, Wall Street Journal, Le Monde, Economist, Forbes, Science, Science News, Nature News, MIT Technology Review, and IEEE Spectrum Magazine.

Metin Sitti serves as the editor-in-chief of the Progress in Biomedical Engineering since 2018 and Journal of Micro-Bio Robotics since 2008 and as the associate editor of the Extreme Mechanics Letters, Advanced Material Technologies, Biomimetics & Bioinspiration, and Advanced Health Care Technologies. He has published over 430 peer-reviewed papers, over 220 of which have appeared in archival journals. These papers are cited over 19,200 times in Google Scholar (h-index: 75). He has given over 160 invited and keynote/plenary/distinguished talks in universities, conferences, and industry.

He has advised 53 (21 current) PhD students, 55 (3 current) MSc students, and 53 (24 current) postdoctoral fellows since 2002. Over 27 of his lab members became an assistant/associate professor at Cornell University, University of Toronto, UIUC, Oregon State University, Virginia Tech, WPI, Arizona State University, Louisiana State University, Texas Tech, Nanyang Technological University (Singapore), Tampere University of Technology (Finland), Bilkent University (Turkey), University of Sheffield (England), Delft University of Technology (Netherlands), etc. Moreover, some of his lab members are working in industry as senior researchers at Intuitive Surgical, Apple, Intel, Google, Boston-Dynamics, Schlumberger, 3M, Blue Origin, etc. He has recently co-chaired a new international major small-scale robotics conference (MARRS 2018) and a workshop on biomedical applications of micro/nanotechnology, and has ten patents granted and eleven patents pending. He founded a start-up (nanoGriptech Inc.) in 2009 to commercialize his lab's gecko-inspired microfiber adhesive technology as a new disruptive adhesive material (Setex®) for a wide range of industrial applications. Setex has been in the market since 2018 in USA.

Metin Sitti's current research focuses on three topics to understand the underlying principles of creating and controlling small-scale robotic systems for medical applications. His first sub-team works on novel functional micro/nanomaterials for miniature robots, where he has pioneered gecko-inspired microfiber adhesives for robotics and other applications. His second sub-team focuses on mobile millirobotics, where his sub-team has pioneered the soft-bodied magnetic millirobots with multimodal locomotion capability and created new bio-inspired water-walking, water-running, jumping-gliding, flying, and climbing miniature robots and soft capsule endoscopic robots for the gastrointestinal tract medical applications. His third sub-team focuses on mobile microrobotics (scaling down mobile robots down to a few micron scale) for medical applications. This sub-team has pioneered self-directed bacteria-driven and microalga-driven bio-hybrid microswimmers for active drug and other cargo delivery.

## Invited talks (selected from over 160 invited talks)

Zernike Distinguished Lecture, Groningen, Netherlands, 6 December 2018; MRS Fall Conference, Boston, 27 Nov. 2018; MIT, Mechanical Engineering Department, Boston, 11 May 2018; IEEE ROBIO Conference, Malaysia, 14 December 2018 (Plenary Talk); Living Machines, Paris, 18 July 2018 (Plenary Talk); Hamlyn Symposium on Medical Robotics, London, 25 June 2018 (Keynote Talk); Reconfigurable Mechanisms Conference, Delft, Netherlands, 21 June 2018 (Plenary Talk).

## Recent collaborators

Gisela Schütz, Günther Richter, Sebastian Trimpe, Joachim Bill (MPI-IS), Carmel Majidi, Chris Bettinger (CMU, USA), Rahmi Oklu (Mayo Clinic, USA), Joanna Aizenberg (Harvard Univ., USA), Victor Sourjik (MPI for Terrestrial Microbiology, Germany), Arianna Menciassi (SSSA, Italy), Samuel Sanchez (Inst. for Bioengineering of Catalonia, Spain), Marco Dorigo (Université Libre de Bruxelles, Belgium), Seda Kizilel, Mehmet Cengiz Onbasli, Metin Muradoglu, Ihsan Solaroglu (Koç University, Turkey), Islam Khalil (German Univ. of Cairo, Egypt), Eric Lauga (Cambridge University, UK), David Gracias (Johns Hopkins University, USA), Veikko Sariola (Tampere University of Technology, Finland), Roland Siegwart, Mehmet Fatih Yanik (ETH Zurich, Switzerland), and Yasin Temel (Maasricht University, Netherlands).

# 5 HAPTIC INTELLIGENCE

## 5.1 Research Overview

When you touch objects in your surroundings, you can discern each item's physical properties from the rich array of **haptic cues** you experience, including both the tactile sensations arising in your skin and the kinesthetic cues originating in your muscles and joints. For example, reaching out to **grasp a water glass** refines your visual estimates of its location, size, shape, and weight while also making you rapidly aware of the temperature, stiffness, smoothness, and friction of its surfaces. Feeling how these haptic sensations develop in response to your motions enables you to not only perceive the glass's material properties but also manipulate it fluidly, whether your goal is to bring it to your lips to drink, place it upside-down in your dishwasher, or rotate it under a flow of hot water as your other hand scrubs it clean.

Over the course of life, humans leverage their rich sense of touch to master a wide variety of **physical tasks**, from everyday necessities like buttoning a jacket to difficult feats such as sculpting a marble statue or inserting a needle into a patient's vein. Many tasks are challenging when first tried, but practice usually enables one to improve the resulting interaction and optimize one's motions so they become almost automatic. You can gain some appreciation for the complexity of tasks that normally feel effortless, such as slicing bread or brushing your teeth, by trying to complete them with your non-dominant hand. Similarly, even the simplest manual skills become almost impossible if you lose your tactile sensitivity due to local anesthetic or a lack of blood flow. The **crucial role of the sense of touch** is also deeply appreciated by researchers working to create **autonomous robots** that can competently manipulate everyday objects and safely interact with humans in unstructured environments. Such systems rarely take advantage of haptic cues and thus often struggle to match the perception, manipulation, and interaction capabilities of humans.

Although humans experience touch coherently, this sense stems from a wide range of distributed receptors that each responds most strongly to a different type of stimulation, generally broken into the categories of mechanical, thermal, and pain sensations. Much about the sense of touch is understood, and many other aspects still need to be investigated. The most notable differences between touch and the more well-understood senses of vision and hearing are that **exploring the world through touch requires action** and that **what one feels greatly depends on how one moves**. It is also helpful to consider that vision has high spatial acuity and only moderate temporal acuity, while hearing is the opposite; different aspects of haptic perception lie along this spatiotemporal continuum between vision and hearing. Because we don't yet fully understand haptic interaction, few **computer and machine interfaces** provide the human operator with high-fidelity touch feedback or carefully analyze the physical signals generated during an interaction, limiting their usability.

The Haptic Intelligence Department of the Max Planck Institute for Intelligent Systems aims to **elevate and formalize our understanding of haptic interaction** while simultaneously **inventing helpful human-computer, human-machine, and human-robot systems** that take advantage of the unique capabilities of the sense of touch. We pursue this goal by undertaking research projects in the following four main research fields, each of which is more thoroughly described in its own section below:

- **Understanding Tactile Contact:** Haptic perception is tightly coupled to movement (action) via the physics of contact. Despite the complexity of these phenomena, people intuitively learn how best to move to extract desired information while accomplishing the task at hand. We seek to disentangle these elements by instrumenting physical interactions carried out by both natural agents (humans) and artificial agents (robots), analyzing the resulting signals with physics-based models and machine learning.

- **Haptic Interface Technology:** Haptic interfaces are mechatronic systems that modulate the physical interaction between a human and his or her tangible surroundings so that the human can act on and feel a virtual and/or remote environment. How can such systems vividly reproduce the perceptual experience of touching real objects and provide feedback that helps the user improve his or her motor skills? We seek to answer these questions by carefully studying existing technologies and inventing new haptic interfaces.

- **Teleoperation Interfaces:** Commonly used in minimally invasive robotic surgery and hazardous material handling, telerobotic systems empower humans to manipulate items by remotely controlling a robot. How can such systems support the operator to perform tasks with skill and outcomes that are as good as (or even better than) those accomplished via direct manipulation? We work to create new ways to capture operator input, deliver haptic feedback, and otherwise augment the operator's abilities, and we systematically study how these technologies affect the operator.

- **Physical Human-Robot Interaction:** To help humans in unstructured everyday environments like hospitals and homes, robots need better haptic interaction skills as well as increased social intelligence. We are working to discover whether and how physical human-robot interaction can benefit humanity by advancing robotic tactile perception and designing, building, and evaluating new physically interactive robots targeted at particular user populations.

The specific projects we pursue take shape through a **bottom-up process** that draws on the experience and interests of the primary researcher, the expertise and creativity of our director and other members of the department, ideas and expertise from our collaborators both within and outside of MPI-IS, and recent discoveries in haptics and related areas. Most of the resulting projects fit well within one of the above fields; a small number of projects bridge topics or are more remote from these core competencies. Each project typically has a **lead researcher** who is a research scientist, postdoctoral fellow, doctoral student, or masters thesis student. This person will generally be the first author of resulting publications, and Dr. Kuchenbecker will generally be the last author. Depending on the demands of the project, this pair may be supported by one or more other scientists, visiting professors, research engineers, technicians, student research assistants, summer interns, and/or short-term high-school interns.

Members of the HI department enjoy working in our diverse international research environment. Our current team is **gender balanced** and hails from **thirteen countries** around the world. Because we pursue highly **interdisciplinary** projects, we welcome applications from people in a wide range of fields. Many of our current department members have a background in mechanical engineering, biomedical engineering, electrical engineering, computer science, and/or cognitive science. Dr. Kuchenbecker books a **standard one-hour meeting time** for each lead researcher and research engineer every week. This time is spent discussing the recent progress, challenges, and future goals for each of the researcher's projects. When she travels, she assigns pairwise meetings between researchers in lieu of these individual research meetings.

We hold **group meeting** for 90 minutes almost every week in Stuttgart, using video conferencing to enable remote members to participate from Tübingen, Switzerland, the USA, Canada, and many other locations. Our recurring agenda covers personnel changes, publication activities, awards and media attention, third-party funding, internal logistics, recent scientific talks, and upcoming events. Each attendee then spends one to two minutes giving an **individual update** on his or her recent activities, including both successes and challenges. Group members ask each other questions and share suggestions on how to solve the issues others are facing. As a rotating task, a lab member writes down funny things that are said during each group meeting, and our director emails out an edited version of these **quotes** along with her typed **minutes** from the meeting.

Our primary publishing targets are full-length **journal articles** in engineering- and medicine-focused research fields and **top-tier conference papers** in fields related to computer science. We use hands-on **demonstrations** of our technology and peer-reviewed **short papers** (often works in progress, late-breaking reports, or workshop papers) as stepping stones to longer-format papers, beneficially gathering feedback from other researchers early in the process to increase the chances of high-impact contributions. To achieve good visibility for our activities and give our younger scientists experience formally presenting their research, we also frequently publish papers at good conferences in haptics and robotics.

To help people prepare for conference presentations, thesis defenses, and job talks, the HI Department holds **presentation club** most weeks. The presenter practices his or her talk, answers a large volume of questions similar to what they might expect in the target presentation venue, and then receives kind suggestions from the group on how the presentation could be made more effective. In other weeks, presentation club includes reports from people who recently attended a conference, an interactive discussion of a particular paper, a brainstorming session for one of our research projects, tips from lab members on tools helpful for research, or a lesson by Dr. Kuchenbecker on effective writing. Presentation club is coordinated by a lab member who serves a six-month term in this position. Other **departmental leadership roles** that rotate on a six-month basis include Demo Coordinator, Human Subjects Coordinator, Internship Wizard, Poster Coordinator, and Tool Master. Our long-term leadership roles are Director, Department Assistant, IT Wizard, Mechanical Design and Manufacturing Expert, Additive Manufacturing Coordinator, Purchasing Wizard, and Safety Representative. Together, we hope to greatly advance human understanding of touch cues while simultaneously discovering new opportunities for their use in interactions between humans, computers, and machines.

## 5.2 Research Fields



### 5.2.1 Understanding Tactile Contact

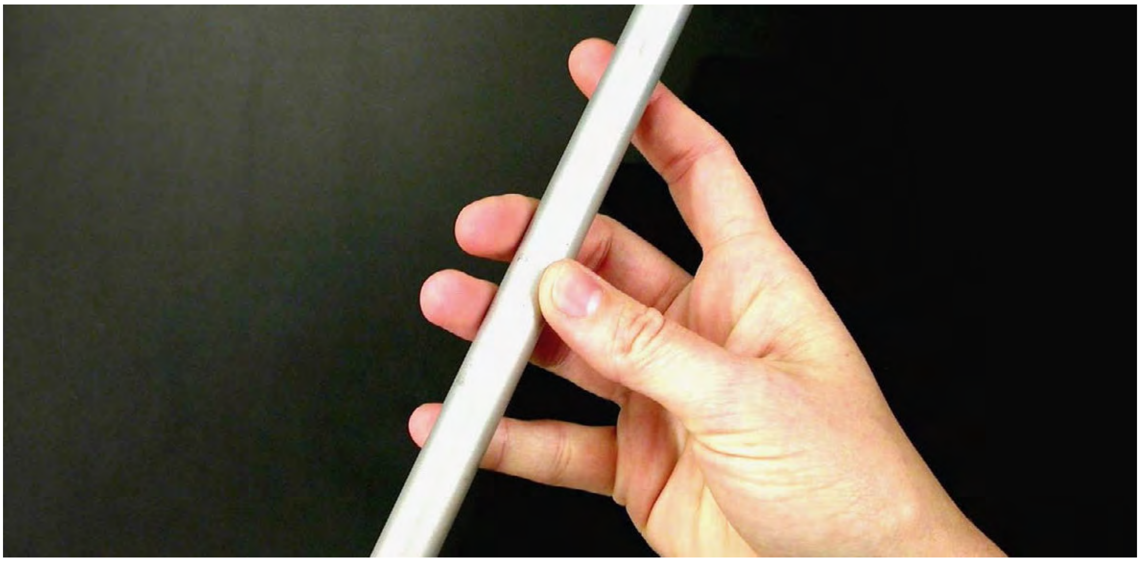

Figure 5.1: A human hand holding an aluminum rod can feel the rod's own hardness, smoothness, friction, and thermal conductivity as well as many detailed haptic properties of objects that the rod touches. How does the exquisite sense of touch work? And how can robots achieve similar levels of haptic intelligence?

Scientists do not yet understand the mechanisms that underpin haptic perception, action, and learning as well as they do the same processes for vision and hearing. Thus we focus a portion of our energy on unraveling the phenomenon of **tactile contact**, wherein an agent interacts with a physical object while feeling the resulting cutaneous and kinesthetic sensations.

In some situations the agent is a healthy **human**, who naturally has an exquisite sense of touch, and in others it is a **robot** endowed with unusually good artificial tactile sensing. Some projects involve both humans and robots touching the same substances at different times, so we can compare their perceptual experiences and try to artificially mimic human capabilities using carefully chosen signal processing and machine learning techniques.

Our projects in this research field often employ high fidelity external motion, force/torque, and acceleration **sensors** that capture the physical details of the interaction as it unfolds. We also occasionally record simultaneous visual data with cameras and audio data with microphones, as both humans and robots are multisensory agents who can benefit from combining cues from different modalities.

Sometimes our primary research goal is to model the **complex physical phenomena** governing contact between human skin and natural objects and/or the perceptual experiences arising from particular haptic stimuli. At other times we want to understand the **cues on which a human relies** to perform a difficult sensory-motor task, so that we can create systems that augment their capabilities. In other scenarios we seek **physical models of relatively new haptic rendering technologies**, so that we can better control their mechanical and perceptual effects on the human user.

More information: https://hi.is.mpg.de/field/understanding-tactile-contact

### 5.2.2 Haptic Interface Technology

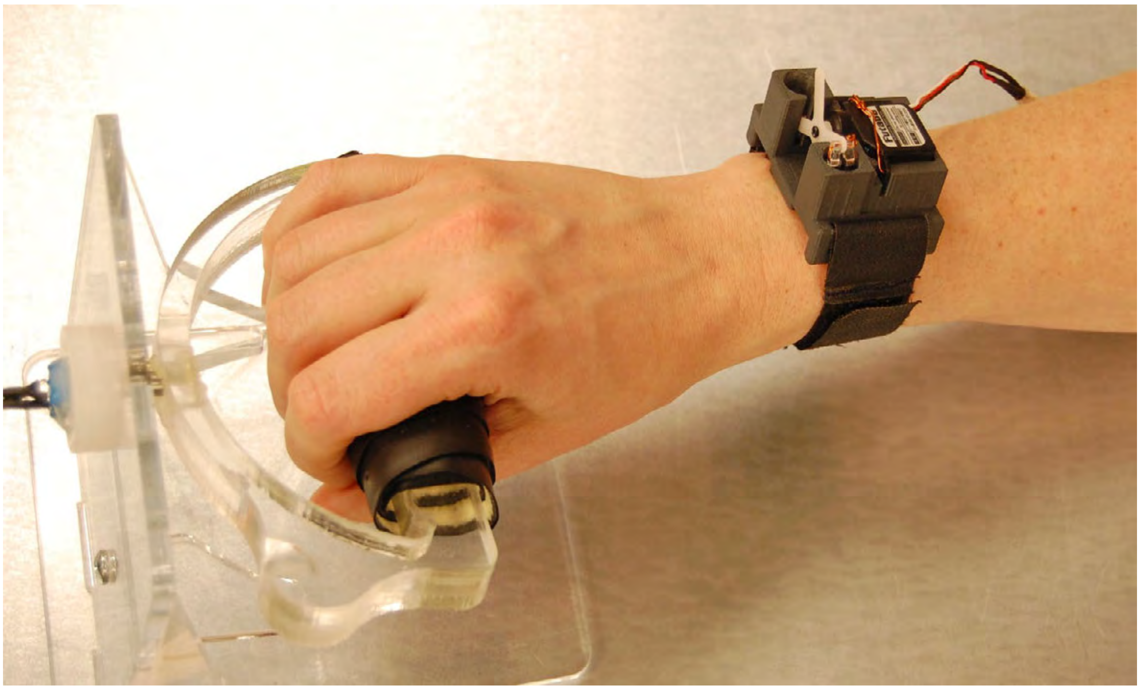

Figure 5.2: A tool for measuring wrist rotation and a wearable haptic feedback device that taps on your skin in response to the wrist movements that you make. Both items were created by undergraduate researcher Andrew Stanley as part of a project that investigated the utility of tactile cues for guiding human movement.

Given the importance of physical contact in human life, scientists and engineers have naturally wondered whether technology could augment or extend such interactions. Indeed, it can! **Haptic interfaces** are mechatronic systems that modulate a human's physical interactions, often to give the user the illusion that they are directly touching things that are at a different location (a remote environment) or that are purely digital (a virtual environment).

Since the start of the field of haptics in the early 1990's, three distinct archetypal **haptic interface categories** have emerged: grounded kinesthetic haptic interfaces, ungrounded haptic interfaces, and surface haptic interfaces. Although they differ in key ways, they all function in the same overall manner: the haptic interface's mechanical, electrical, and computational elements work together to monitor and modify the user's physical interaction with his or her tangible surroundings.

We do research on all three categories of haptic interfaces, as well as on closely related interactive systems that provide the user with visual and auditory feedback based on movement. In the relatively well established area of **grounded force-feedback devices**, we focus on understanding and sharing the diversity of past designs in order to accelerate haptic device innovation and facilitate the use of haptic feedback by interaction designers.

In the newer areas of **ungrounded and surface haptic interfaces**, we aim to expand what is possible by inventing and refining new devices. Here our work often includes hardware design, actuator and sensor selection, calibration, control optimization, application design, and system integration. Most projects also include extensive evaluation by human subjects so that we can understand how our new inventions compare to prior approaches.

More information: https://hi.is.mpg.de/field/haptic-interface-technology

### 5.2.3 Teleoperation Interfaces

Figure 5.3: Paola Forte demonstrates her research to a visitor while three children look on. The Intuitive da Vinci Si surgical system that Paola uses in her research lets a doctor operate on a patient from across the room but normally provides no haptic feedback or other sensory augmentations for the operator.

A human user can carry out tasks in a remote environment by controlling a robot; the interaction site might be just across the room from the user, several kilometers away, deep in the ocean, or orbiting above Earth's surface. Such a scenario is known as **bilateral teleoperation** or **telemanipulation**. The remote robot's job is to represent the user's actions in the remote environment; the user sends these commands and receives multimodal feedback via the **teleoperation interface**, which is where we focus our attention.

Remotely accomplishing complex tasks such as surgical suturing requires a rich bidirectional that is optimized for human capabilities. Because the addition of force feedback tends to drive bilateral teleoperators unstable, most such systems include no haptic cues; the operator thus has to learn to rely on what he or she can see. We work on inventing and refining **clever ways of stably providing haptic feedback** during teleoperation, often by focusing on tactile rather than kinesthetic cues. One main thrust of our work centers on vibrotactile feedback of the robot's contact vibrations, as this approach is both simple and highly effective.

We also study how the addition of haptic cues affects the operator over both **short and long time scales**. Having direct access to physical contact information changes the processing required for one to complete a task; different kinds of feedback have different effects that need to be understood both quantitatively and qualitatively. Our investigations in this domain also show that the haptic signals captured during teleoperation contain significant information about the **manual skill** of the operator currently controlling the robot.

While most of our teleoperation research has considered only manipulation tasks, we also study teleoperation of humanoid robots that have **both task-oriented and social functions**. We are interested in inventing lightweight and effective systems that enable the operator to control the robot's hand gestures and facial expressions with high fidelity.

More information: https://hi.is.mpg.de/field/teleoperation-interfaces

### 5.2.4 Physical Human-Robot Interaction

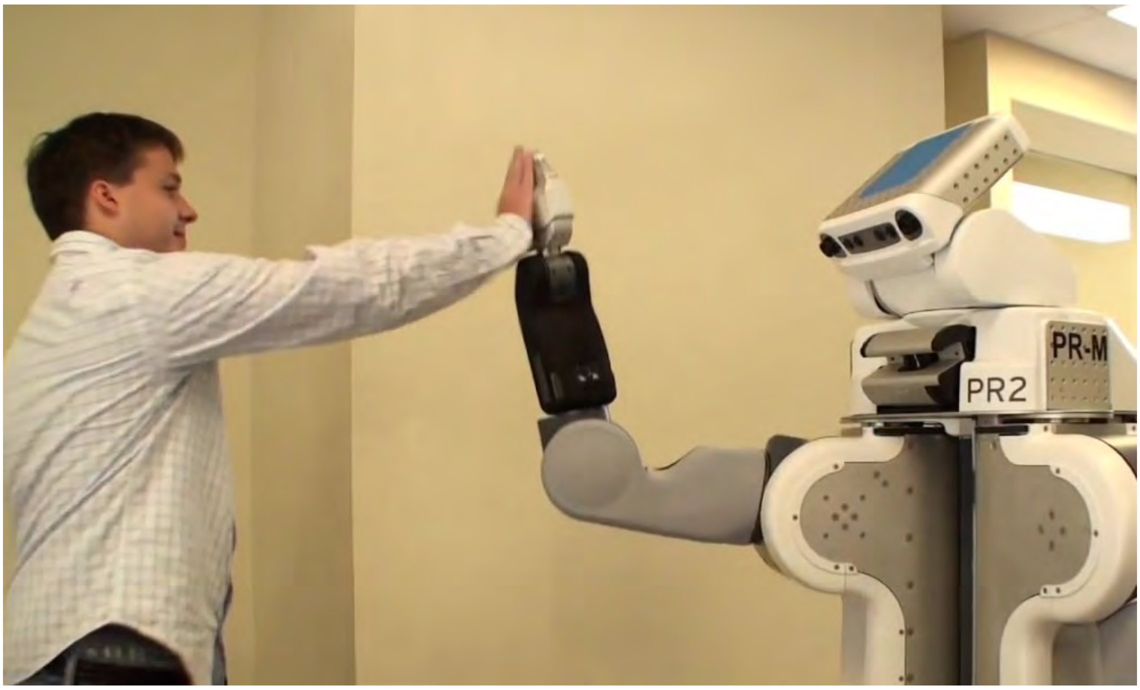


Figure 5.4: Joe Romano, Dr. Kuchenbecker's first Ph.D student at the University of Pennsylvania, high-fives a PR2 robot as part of the pr2-props demo that he created while interning at Willow Garage. Everyone loves to high-five or fist-bump a robot!

While autonomous robots excel at repetitive tasks in controlled environments, the world in which most humans live is messy, constantly changing, and filled with other people. Important opportunities exist for helping humans in these unstructured environments, particularly as our population ages, but new approaches are needed for robots to be as successful inside homes, clinics, and hospitals as they already are in factories.

To enable the completion of useful tasks around humans, we are interested in robots that **physically interact** with both objects and people. Most robots have no haptic sensing, as commercial tactile sensors tend to be expensive and rather limited in size, robustness, sensitivity, and/or reliability. We are thus working to create **tactile sensors** that could be easily manufactured to cover all exposed surfaces of a robot and provide useful contact information. We also frequently investigate new ways of using existing sensors to increase a robot's haptic intelligence.

Interestingly, good physical skills alone are not enough for a robot to succeed at helping humans in everyday environments. Because humans are social creatures, robots also need good **social skills**. Many physical interactions that transpire between humans, such as object handovers and hugs, have strong social dynamics that have been carefully studied. Robots that skillfully take part in such interactions may be able to work more effectively with and around humans. We are interested in how such robots should behave, and how people react to them in different scenarios.

**Social-physical human-robot interaction** (spHRI) may have a special role to play in clinical and therapeutic settings, where patients need to perform repetitive physical activities to improve their skills. We envision robots that function as an exercise partner or coach, interacting with the user both physically and socially in a natural and engaging manner. We are also interested in algorithms that enable a human to teach a robot new manual tasks, which is again a task that has both physical and social aspects.

More information: https://hi.is.mpg.de/field/physical-human-robot-interaction

## 5.3 Selected Research Projects

## Feeling With Your Eyes: Visual-Haptic Surface Interaction

*Alex Burka, Katherine J. Kuchenbecker, Siyao Hu, Trevor Darrell (University of California, Berkeley)*

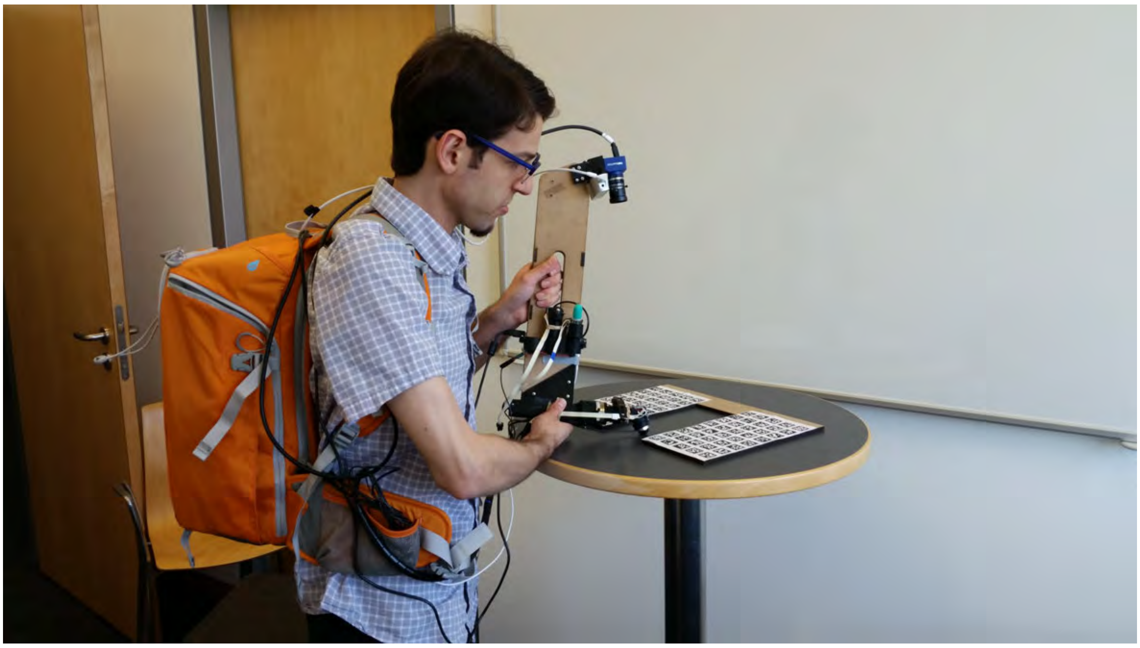

Figure 5.5: An experimenter operates the Proton Pack, a portable visuo-haptic surface interaction recording device.

Humans draw on their vast life experience to make inferences about the objects in their environment; this enables one to make haptic judgments before actually touching things. For example, when selecting the correct grip force for picking up a delicate object, or to choose a gait that will enable safe walking across a patch of ice. Robots struggle to **"feel with their eyes"** in the same way, and we believe that this problem can be solved through machine learning on a large dataset [579].

Modern robots are often equipped with advanced vision hardware that typically provides 3D data (via a stereo or depth camera). However, multimodal perception is still in development, especially in the area of haptics. Not many robots have haptic sensors, and even when they do, they aren't standardized, and we don't always know how to interpret the data.

In this project, we designed and built the Proton Pack, a portable, self-contained, human-operated visuo-haptic sensing device [579]. It integrates visual sensors with an interchangeable haptic end-effector. The major parts of the project are as follows:

- **Calibration**: to relate all the sensor readings to each other, several calibration processes determined the relative camera positions and the mechanical properties of each end-effector [576].
- **Surface classification**: to ensure that collected data is relevant to haptic perception, and to validate hardware improvements along the way, we did some proof-of-concept classification tests among small sets of surfaces [577].
- **Data collection**: we used the Proton Pack to collect a large dataset of end-effector/surface interactions [563]. Both the Proton Pack and several sets of material samples traveled across the Atlantic for this project [566].
- **Machine learning**: we attempted to train a computer to "feel with its eyes" (that is, predict haptic properties from visual images) using the dataset we collected [550].

This work is supported by NSF grant 1426787 and NSF Complex Scene Perception IGERT program, award number 0966142.

More information: https://hi.is.mpg.de/project/feeling-with-your-eyes

## Learning Haptic Adjectives from Tactile Data

*Ben Richardson, David Schultheiss, Katherine J. Kuchenbecker*

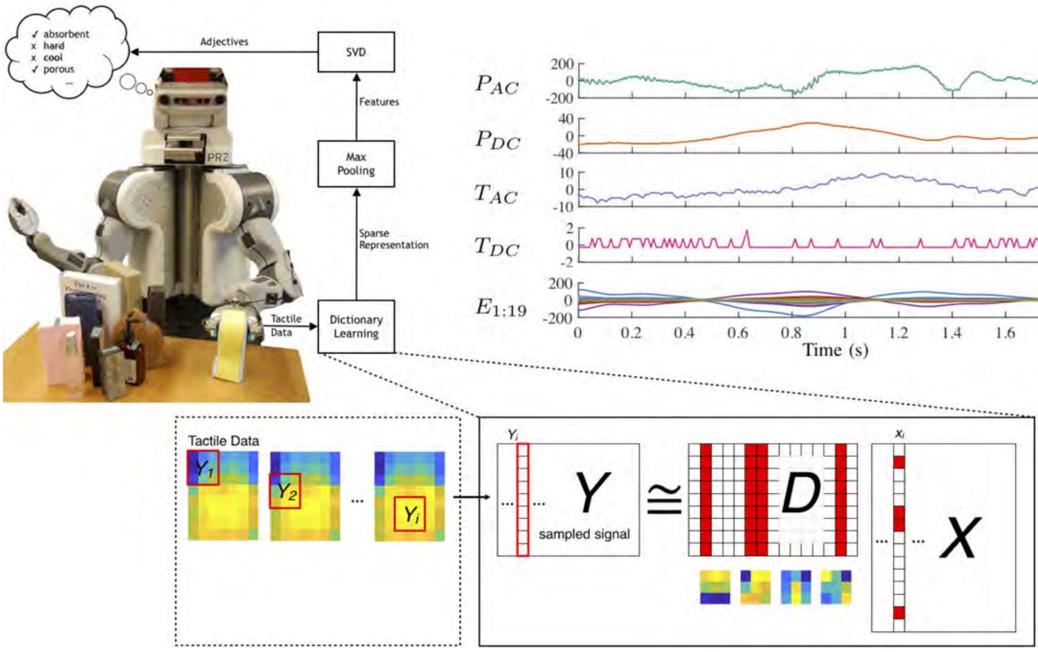


Figure 5.6: As the robot explores objects, raw tactile data is gathered and partitioned into an observation matrix *Y*. Unsupervised dictionary learning extracts a dictionary of basis vectors *D* from the data. *Y* is reconstructed by sparsely combining elements of *D* using the sparse representation matrix *X*. The elements of *X* are max pooled, and the resulting feature vectors are used in adjective classification.

Humans can form an impression of how a new object feels simply by touching its surfaces with the densely innervated skin of the fingertips. Recent research has focused on endowing robots with similar levels of **haptic intelligence**, but these efforts are usually limited to specific applications and make strong assumptions about the underlying structure of the haptic signals. By contrast, unsupervised machine learning can discover important underlying structure without biases for specific applications; such representations can then be effectively applied to a variety of learning tasks.

Our work applies **unsupervised feature-learning methods** to **multimodal haptic data** (fingertip deformation, pressure, vibration, and temperature) that was previously collected using a Willow Garage PR2 robot equipped with a pair of SynTouch BioTac sensors. The robot used four pre-programmed exploratory procedures (EPs) to touch sixty unique objects, each of which was separately labeled by blindfolded humans using haptic adjectives from a finite set.

We are particularly interested in how representations derived from this data set depend on EPs and sensory modalities, and how accurately they classify across the set of adjectives. Do representations learned from certain EPs and data streams classify texture-related adjectives better than others? Can variability in performance be predicted by properties of the learned representations? Is it necessary to learn a different representation for each EP, or does enough shared information exist to use a combined representation?

Our work shows that learned features are far superior to hand-crafted features. Additionally, the features learned from certain EPs and sensory modalities provide a better basis for adjective classification than those from others. We are also investigating dependencies between EPs and sensory modalities, relationships between learned representations, and ordinal (rather than binary) adjectives.

More information: https://hi.is.mpg.de/project/learning-haptic-adjectives-from-tactile-data

## Tangential Force vs. Friction Coefficient

*David Gueorguiev, Katherine J. Kuchenbecker, Julien Lambert (Université catholique de Louvain), Jean-Louis Thonnard (Université catholique de Louvain)*

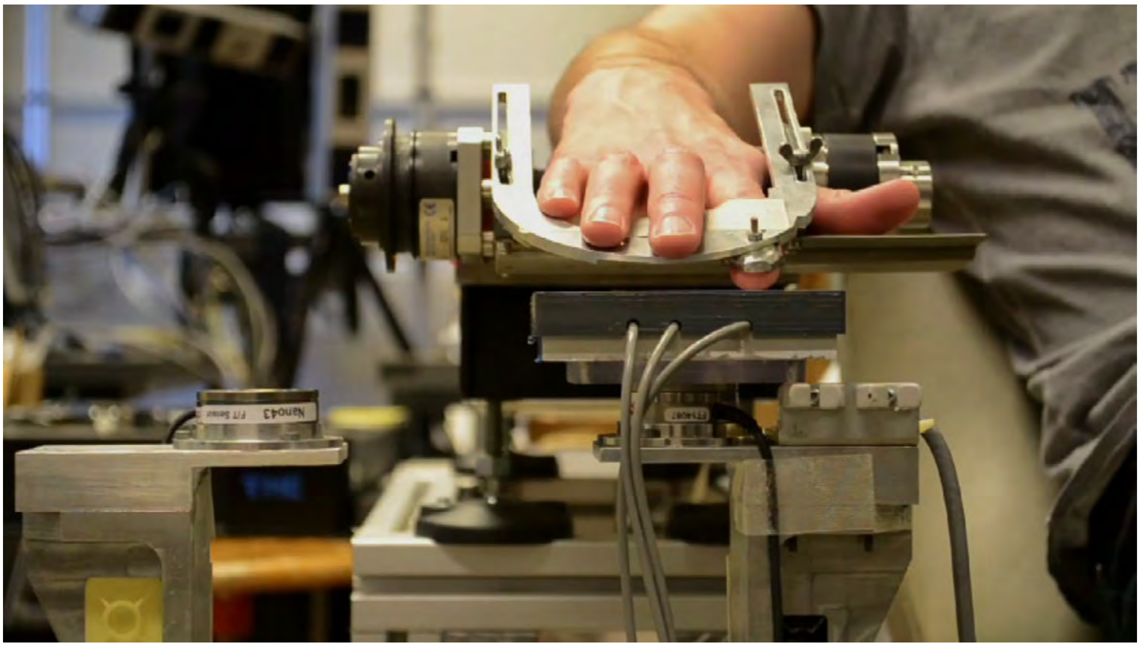

Figure 5.7: The experimental setup combines a force-controlled robotic platform for passive stimulation of the finger and an ultrasonic friction modulating device (STIMTAC) with the aim to better understand the human perception of frictional cues.

The frictional forces we experience when our body interacts with objects provide essential sensory cues that help adapt our behavior. We rely on these sensory cues daily, for example when we feel the smoothness of fabric before buying clothing, or when we slide our finger against the screen of a smartphone to move between photographs. These frictional signals are induced by the complex contact mechanics that occur at the skin-fingertip interface during tactile interaction.

Humans have been shown to be very sensitive to frictional changes while sliding the fingertip across a flat surface. Despite these observations, it is unclear how the mechanical deformations related to frictional strains are processed by different skin receptors and the brain. This project aims to *quantify the respective contributions of the dynamic friction coefficient and the lateral force to the human perception of frictional cues*. Although it is known that humans can easily detect changes in lateral force, little is known about human sensitivity to the coefficient of dynamic friction.

In this collaborative project, we aim to induce **variations of the normal force** with a force-controlled robotic platform during passive stimulation of the finger while simultaneously **modulating the lateral force with an ultrasonic haptic display** that can reduce the finger-surface friction. By coupling these two technologies, we can keep the coefficient of dynamic friction constant while changing the lateral force and vice versa.

The results of these experiments will show *which tactile cues are used by humans to perceive changes in friction*. This knowledge is important because of growing interest in friction-based haptic technologies for manipulation (e.g., in virtual reality) and for texture rendering on flat screens. These growing applications make it essential to understand which parameters are the most efficient at conveying frictional cues to the human brain.

More information: https://hi.is.mpg.de/project/tangential-force-vs-friction-coefficient

## Perception of Ultrasonic Friction Pulses

*David Gueorguiev, Katherine J. Kuchenbecker*

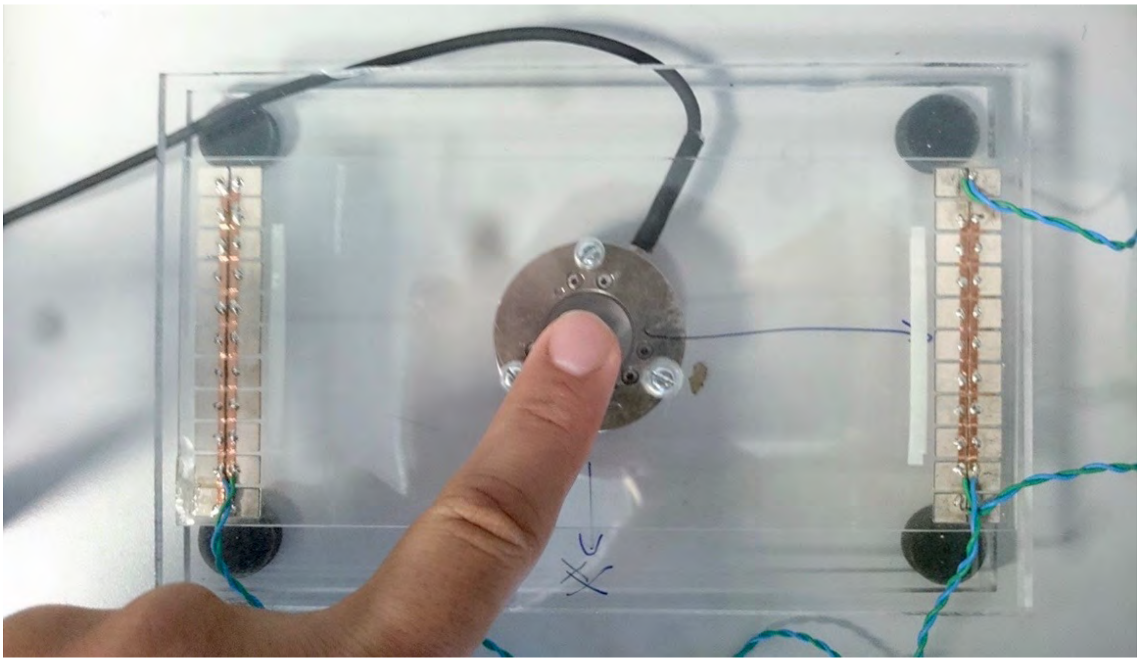

Figure 5.8: Vibrating a glass screen with high-frequency ultrasonic waves reduces the finger-surface friction during contact interactions. We can use this phenomenon to create virtual tactile elements on the flat screen, such as edges, holes and bumps. As shown in the image, piezoceramic actuators glued on the left and right sides of the screen are used to generate the ultrasonic waves.

The demand for **natural haptic feedback on touch screens** such as smartphones and car control panels has been rapidly growing in recent years. One such feedback technology, vibrotactile stimulation, is already incorporated into many consumer devices but provides only a general vibration sensation to the user's fingers and does not work well on mechanically grounded screens.

Friction-based tactile feedback solutions have recently been demonstrated as a good alternative actuation approach. The development of these novel interfaces has raised interest in touch-based human-machine interactions while also highlighting the need for appropriate **high-fidelity strategies for tactile rendering**. Such new approaches face the limitation that little is known about the sensory mechanisms that mediate human perception of frictional cues.

As shown in the figure, virtual geometric features and complex textures (e.g., buttons, edges, patterns, fabrics) can be rendered on flat screens via **ultrasonic vibrations that modulate the friction force** experienced by the user. However, optimizing these tactile sensations requires one to understand which components of the frictional signal are critical and how the intensity of each component should be scaled according to the dynamics of the interaction.

In this project, we investigate how ultrasonic vibration waveforms with different durations and slopes are perceived by humans, as well as what kind of high-level percepts (such as edges, bumps and holes) they generate. To that end, we physically measure the frictional patterns that the finger experiences during ultrasonic stimulation, and we then correlate these results with the verbal reports of the users.

The long-term goals of this project are to *understand how to generate tactile features that are perceived unambiguously by users* and to leverage this knowledge to *guide the design of friction-based tactile stimuli* on future tactile displays that provide haptic feedback.

More information: https://hi.is.mpg.de/project/perception-of-ultrasonic-friction-pulses

## Understanding the Physics Behind Electroadhesion

*Yasemin Vardar, Katherine J. Kuchenbecker, Eberhard Goering*

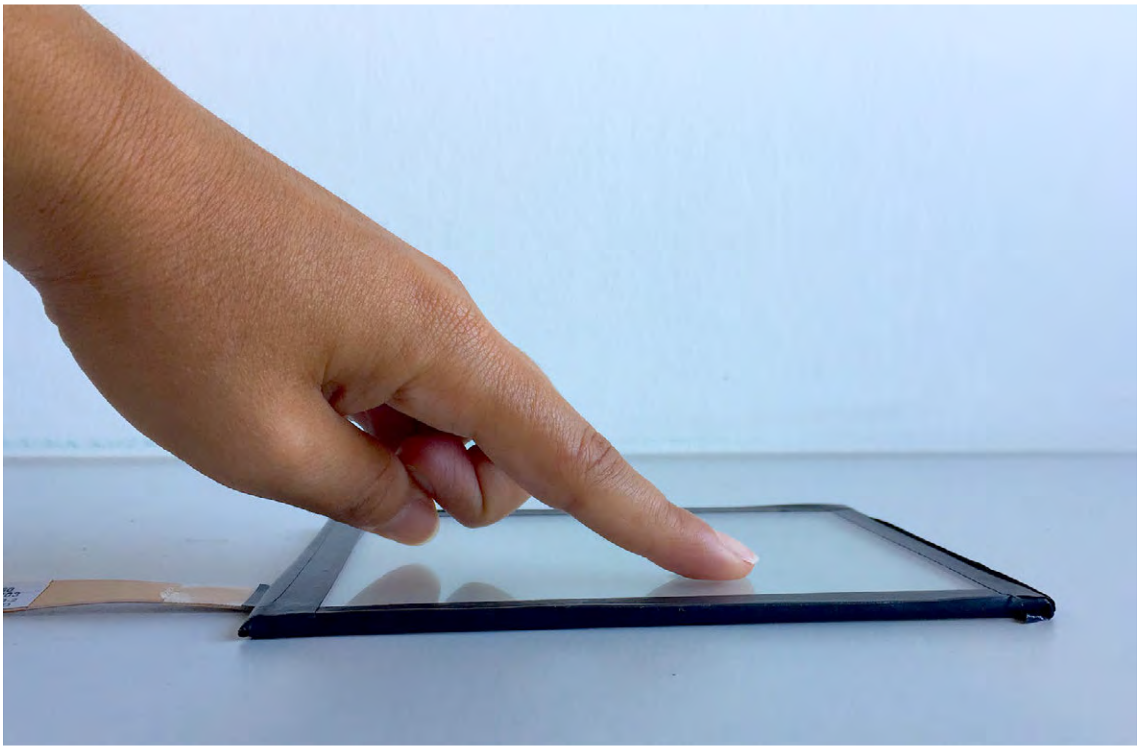

Figure 5.9: A user interacts with a surface haptic display. The friction force between the fingertip and the surface is modulated via electrostatic forces.

Researchers worldwide want to discover how to generate **compelling tactile sensations** on touchscreens to increase the usability of mobile devices and other interactive computer systems. One approach for generating such sensations is to control the friction force between the screen and the finger-pad of the user via electrostatic actuation.

When an alternating voltage is applied to the conductive layer of a touch screen, an attractive force is generated between its surface and the user's finger. This computer-controllable force modulates the friction between the surface and the skin of the finger moving on it, so that one can create various haptic effects. This haptic rendering approach is commonly called **electroadhesion** or **electrovibration**.

Even though it is straightforward to generate various haptic stimuli via this method, there is limited understanding of the physics behind electroadhesion, as well as its perceptual effects. This research project aims to fill this gap by conducting psychophysical experiments in tandem with physical measurements while human subjects interact with an electroadhesively actuated surface haptic display.

More information: https://hi.is.mpg.de/project/understanding-the-physics-behind-electroadhesion

## Visual and Haptic Perception of Real Surfaces

*Yasemin Vardar, Elisa Loeffler, Katherine J. Kuchenbecker, Christian Wallraven (Korea University)*

Figure 5.10: Four surfaces from the Penn Haptic Texture Toolkit; each material has different properties that can be both seen and felt.

Both vision and touch play important roles in human perception of real surfaces. Judging material properties based on only one modality may not give reliable results. For example, many of us have had the experience that something we ordered online did not match with our expectations when it arrived in the mail. On the other side, we are often confused about the objects that we touch when trying to navigate an unfamiliar place in the dark. These undesirable situations occur because we are being forced to make perceptual decisions without information from all of our relevant senses.

Although there have been many studies characterizing the information gathered by the individual senses, it is still unclear how each modality's information is processed and integrated. In this research, we aim to answer these questions for vision and touch. Will humans perceive surfaces similarly or differently through these two senses? And what visual and haptic properties correlate most strongly with human perception?

This project is currently focused on the 100 isotropic and homogeneous surfaces that make up the Penn Haptic Texture Toolkit. We are conducting psychophysical and physical experiments with a subset of these materials to understand the roles that visual and haptic cues each play in the perception of real surfaces.

More information: https://hi.is.mpg.de/project/visual-and-haptic-perception-of-real-surfaces

## Understanding Fingerpad Moisture and Friction

*Saekwang Nam, Katherine J. Kuchenbecker*

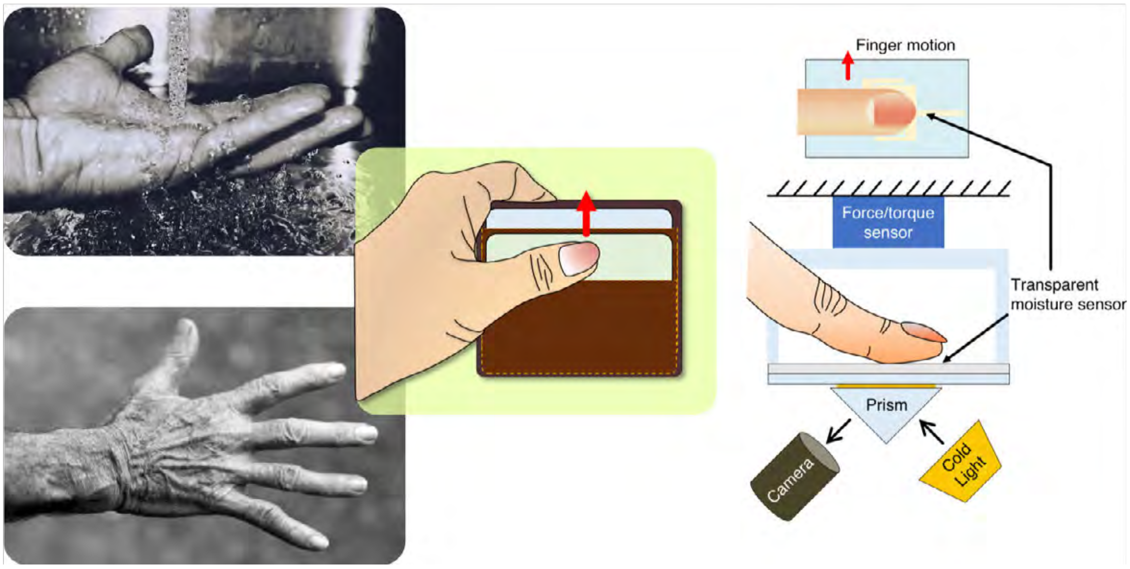


Figure 5.11: This project aims to understand the relative contributions of contact area, force, and moisture of the finger in the interaction between a finger-pad and a physical surface. The proposed apparatus shown to the right measures all three phenomena simultaneously using a camera system, a force/torque sensor, and a transparent moisture sensor.

The interaction between a human finger-pad and a physical surface generates not only the tangential friction needed for gripping objects but also a wide variety of perceptual experiences. Finger-surface contact behavior is known to depend on the actual **contact area**, the **normal force** with which the finger and the surface are pressed together, and the presence of **liquids** (such as sweat) at the interface; however, the relative contributions of these phenomena are not fully understood, nor are their interactions.

Let's think about a situation in which you are trying to pull a plastic credit card from your wallet. When your finger is dry, you may want to press on the card more to increase frictional force and prevent slip. On the other hand, less pressing force is required if your fingerpad is moist, as this moisture creates additional friction.

To investigate this problem, we are preparing a system that can **simultaneously measure** all three quantitites (contact area, force and moisture) over time with high resolution and frame rate. We plan to conduct experiments wherein a human subject presses his or her finger on the transparent surface at the middle of the apparatus. A light illuminates the contact from below the surface, so that a camera located there can record the contact area. A force/torque sensor mounted above the surface measures the normal and tangential forces exerted by the subject. The quantity of liquid on the fingerpad is measured by a custom transparent moisture sensor that is embedded in the transparent surface at the point of fingertip contact.

By analyzing the relative contributions of the three factors, we aim to deepen our understanding of the physical phenomena that affect human manipulation and perception through the hands.

More information: https://hi.is.mpg.de/project/understanding-fingerpad-moisture-and-friction

## Instrumented Needle Insertion Robot

*Rachael Lorsa, Katherine J. Kuchenbecker, Chris Macnab (University of Calgary), Kourosh Zareinia (Ryerson University)*

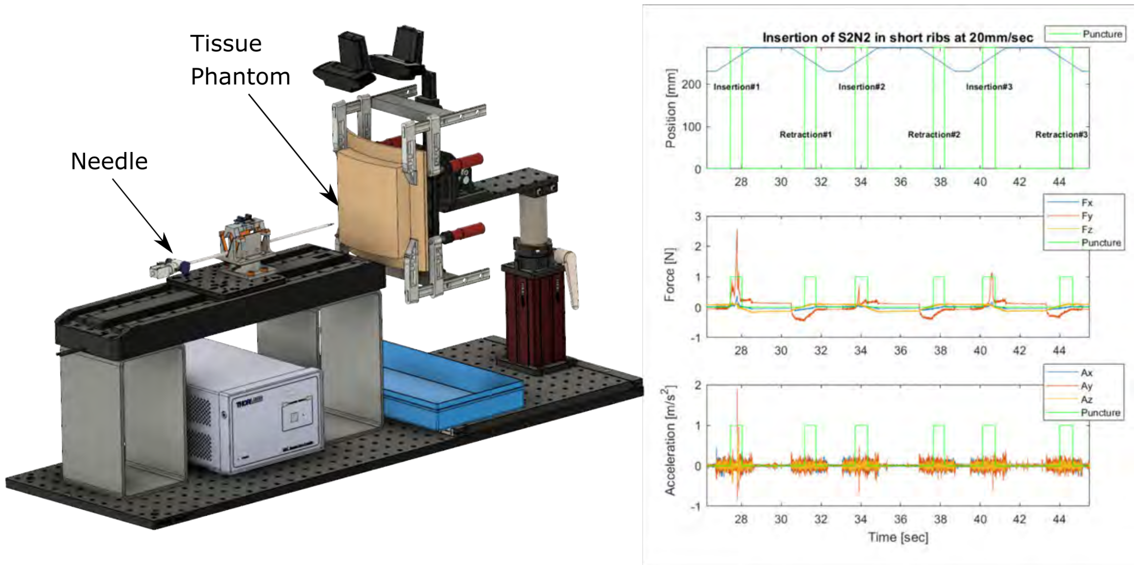


Figure 5.12: How much force is needed to puncture tissue? Sensors on the needle-insertion device (NID0) capture the needle position, force, and acceleration during insertions into and retractions from soft tissue phantoms and food-grade pork ribs.

Human soft-tissue properties vary widely: factors such as the patient's age and health can make enormous differences. As a result, junior medical professionals often have difficulty learning how much force is needed to cut through or puncture a particular type of tissue.

For example, during chest tube insertion, the doctor must use a needle to access the pleural space of the patient's body, puncturing through the pleural membrane but going no further. Surgeons typically perform these procedures blind, relying on a combination of anatomical knowledge and their sense of touch to guide the needle. Surgeons report feeling pleural punctures strongly in some patients, but not in others. Thus, even well trained surgeons may have difficulty, being limited by the resolution of their own human sensory organs.

Given these widely varying tissue properties, complications can occur for percutaneous subtasks within lifesaving procedures like chest-tube insertion. These complications occur more often with inexperienced medical professionals. An area of particular concern is that manually-inserted needles unintentionally perforate vital organs such as the lungs, spleen, and even the heart.

The development of realistic robotic training devices would allow surgeons to learn this difficult technique before their first real procedure. Augmenting their sense of touch via sensorized needles, or using such a needle to warn them of impending punctures, also has the potential to improve surgical outcomes.

In this research we measure tissue properties during needle insertions and will use this data to develop a predictive software model of needle-tissue interactions. We started by developing a custom needle-insertion device (NID0); preliminary experiments with it indicate that our force sensor is of sufficient quality, and we are currently in the process of developing predictive software models based on the captured data.

Ultimately our techniques should allow the development of robotic surgical-training devices that will react realistically (unlike current simulators and tissue phantoms), as well as sensorized needles that could augment a surgeon's sense of touch during a real insertion.

More information: https://hi.is.mpg.de/project/instrumented-needle-insertion-robot

## Modeling Hand Deformations During Contact

*David Gueorguiev, Dimitris Tzionas, Michael Black, Katherine J. Kuchenbecker, Claudio Pacchierotti (CNRS, Irisa and Inria Rennes Bretagne Atlantique)*

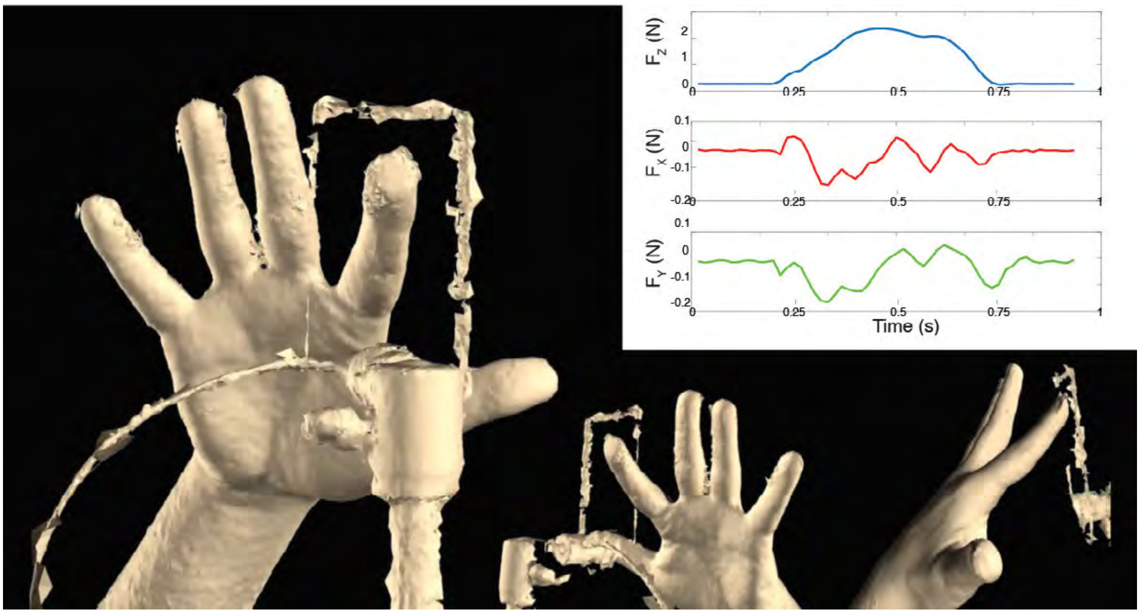


Figure 5.13: 4D capture (in space and time) of a finger pressing on a 2-mm-thick glass surface mounted on a 3-axis force sensor.

Little is known about the **shape and properties of the human finger during haptic interaction**, as such situations are difficult to instrument. Interestingly, these parameters are essential for designing and controlling wearable haptic finger devices and for delivering realistic tactile feedback with such devices. For example, there is currently no standard model of the fingertip's shape and its potential variations across the population. Interaction-dependent deformations have also not yet been modeled.

This project explores a framework for four-dimensional scanning (3D shape over time) and modelling of **finger-surface interactions**, aiming to capture the motion and deformations of the entire finger with high resolution while simultaneously recording the interfacial forces at the contact. We are currently capturing the deformations of the fingertip during active pressing on a rigid surface, which is a first step toward an accurate characterization of the shape and deformations of the physically interacting human hand and fingers [557],[556].

In the future, we would like to optimize the scanner configuration and capture a large number of surface-finger interactions from a range of participants that is representative of the general population. We believe that these data will enable the creation of a **statistical model** that simulates the natural behavior of the interacting finger.

An accurate model of the variations in finger properties across the human population could enable one to infer the user's fingertip properties from scarce data obtained by lower resolution scanning. It may also be relevant for inferring the physical properties of the underlying tissue from observing the surface mesh deformations, as has been shown for body tissues. Further applications of this research include simulation of human grasping in virtual worlds that is consistent with our manipulation of real objects and estimation of the contact forces generated by the tactile interaction from the 4D data alone.

More information: https://hi.is.mpg.de/project/modeling-hand-deformations-during-contact

## Haptipedia

*Hasti Seifi, Farimah Fazlollahi, Gunhyuk Park, Katherine J. Kuchenbecker, Michael Oppermann (University of British Columbia), John Sastrillo (University of British Columbia), Jessica Ip (University of British Columbia), Ashutosh Agrawal (IIT Guwahati), Karon E. MacLean (University of British Columbia)*

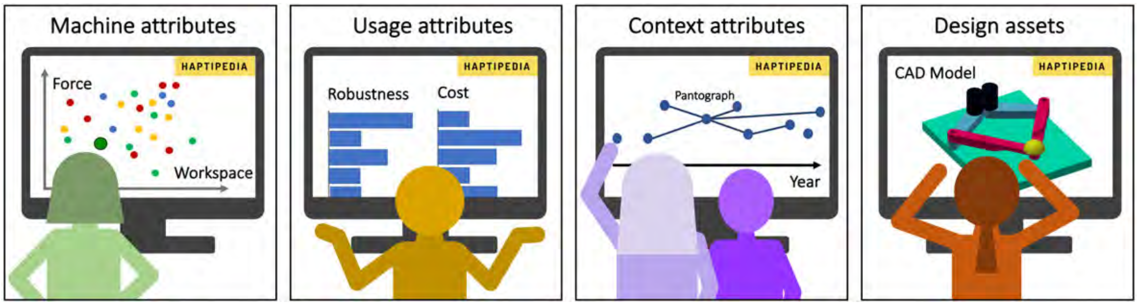


Figure 5.14: Haptipedia is a taxonomy, database, and interactive visualization of 105 haptic devices. The Haptipedia visualization enables designers with varying backgrounds and purposes to browse this device database according to the attributes that matter to them.

Creating haptic experiences often entails inventing, modifying, or selecting specialized hardware. However, experience designers are rarely engineers, and 30 years of haptic inventions are buried in fragmented literature that describes devices mechanically rather than by purpose.

We conceived of **Haptipedia** to unlock this trove of examples: Haptipedia provides a **practical taxonomy**, **database**, and **visualization** to efficiently navigate this fragmented corpus based on metadata that matters to both device and experience designers [547]. Through it, designers can search and browse a database of 105 grounded force-feedback devices with an online visualization, examine their design trade-offs, and repurpose them into novel devices and interactions.

To design Haptipedia, we asked: what taxonomy of attributes can best delineate haptic devices for both device and experience designers? What subset of attributes is most informative? Which other attributes, missing from the literature, do users care about? Which missing attributes can be estimated by an expert designer?

Haptipedia's design was driven by both systematic review of the haptic device literature and rich input from diverse haptic designers. We focused on grounded force-feedback (GFF) devices, as the earliest subset of haptic technology with rich device variation and considerable maturity in both research and commercial settings. From simple haptic knobs to robotic arms with a dozen degrees of freedom, GFF devices typically measure the user's motion and output force and/or torque in response. We iteratively developed Haptipedia by screening 2812 haptics publications, selecting 105 papers that described a haptic device, extracting device attributes from documentations, building a GFF taxonomy, database, and visualization, and evaluating them with users. Device and experience designers provided input on three major iterations of Haptipedia's components (taxonomy, database, and visualization) during haptics conference demonstrations [554],[553], focus group sessions, and in-depth individual interviews.

We show that 1) the Haptipedia taxonomy provides a framework and lexicon for describing various aspects of GFF devices, 2) our taxonomy and visualization let users assess device trade-offs and overall hands-on feel, and select relevant devices from a large corpus, and 3) our process can inform design of future platforms for other technology subsets (e.g., 3D fabrication technology or head-mounted displays).

In our ongoing work on this project, we plan to evaluate salient features and patterns in the current dataset, investigate ways of crowdsourcing the entry and validation of data for new haptic devices, create Haptipedias for other types of haptic interfaces, and assess Haptipedia's utility and impact in haptics and the larger designer community over time.

You can interact with the Haptipedia visualization at http://haptipedia.org/

More information: https://hi.is.mpg.de/project/haptipedia

## 6DOF Tactile Fingertip Display

*Eric Young, Katherine J. Kuchenbecker*

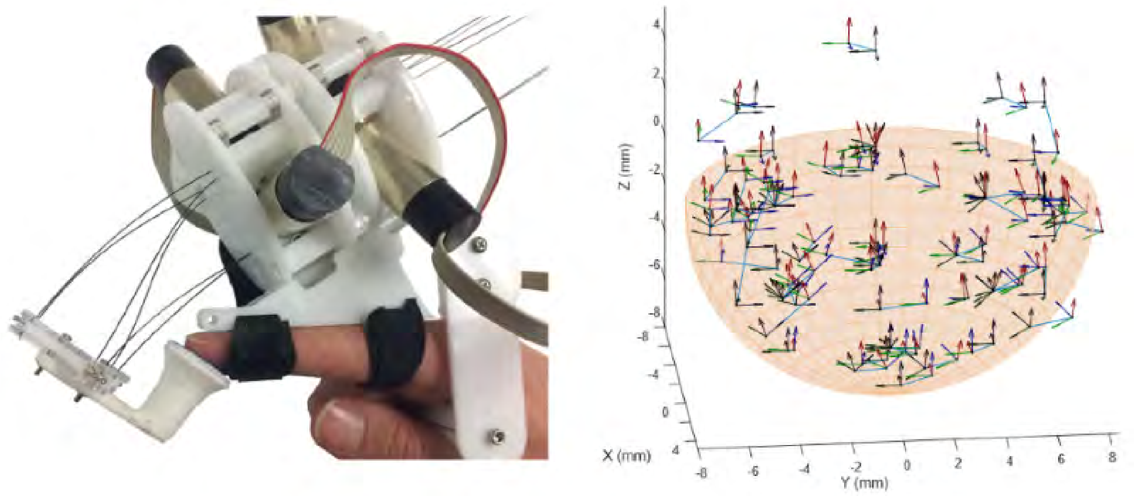


Figure 5.15: The left image shows the complete 6-DoF fingertip device. The position and orientation of the contact surface can be controlled by changing the lengths of the six Nitinol legs independently. The right image shows the difference between commanded and measured poses around a simulated fingertip.

Although a human can move his or her fingertip with six degrees of freedom (three for position and three for orientation), displaying **6DOF haptic cues** continues to escape the capabilities of body-grounded fingertip haptic displays. All six degrees of freedom have been displayed in smaller subsets, so the limiting factor seems to be the additional volume, weight, and complexity of combining all six degrees of freedom in one device.

Meanwhile, robotics research in non-haptics areas has shown the potential of parallel continuum manipulators (PCMs) to be strong and compact. PCMs have six flexible wires that run from a stationary base platform to a free-moving distal platform, allowing the position and orientation of the distal platform to be controlled by adjusting the six wire lengths independently. This project aims to *design and evaluate a **wearable haptic device** that is capable of delivering 6DOF tactile cues to a user's fingertip by leveraging the compact and lightweight nature of PCMs*.

After refining the design of a parallel continuum manipulator specifically for fingertip haptic applications [573], we designed and constructed the **motorized prototype** seen in the figure above [551]. A preliminary user study using a feed-forward position control scheme has shown that users can discern **contact location** quite well throughout the workspace of the device and that they discern **shear cues** moderately well depending on the contact location. This approach thus seems to have potential for delivering rich haptic cues to the fingertip of a user for interactions in augmented reality (AR) and virtual reality (VR).

Looking forward, we hope to use the created device to deliver better haptic cues by compensating for the precise shape and position of a user's fingertip, to explore the addition of force or position feedback via the contact surface, and to evaluate further the perception of high-dimensional haptic cues at the fingertip.

More information: https://hi.is.mpg.de/project/6dof-tactile-fingertip-display

## Gait Propulsion Trainer

*Katherine J. Kuchenbecker, Siyao Hu, Erin V. Vasudevan (SUNY Stony Brook University), Krista Fjeld (SUNY Stony Brook University )*

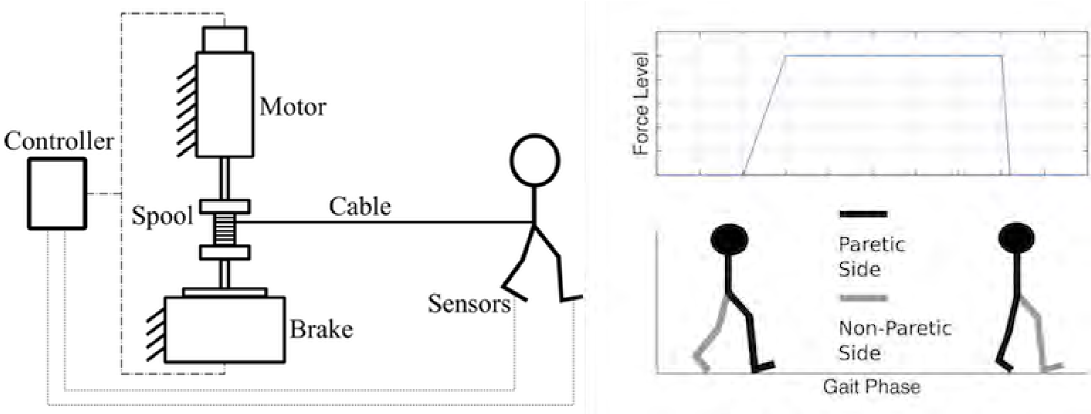


Figure 5.16: The left image shows a schematic of the proposed device: a person is tethered to a spool, which is connected to a motor and a brake, and data from the sensors underneath the shoes is transmitted wirelessly to the controller. The right image shows the asymmetric resistive force with the corresponding gait phases.

Walking speed and symmetry are high priorities for people with hemiparesis from stroke. We developed the **Gait Propulsion Trainer (GPT)** to help such individuals improve their walking abilities by increasing the propulsive force generated by the paretic leg. The GPT centers on a cable spool attached to a stand at waist level. The end of the cable attaches to the center back of a waist harness worn by the user. As the person walks away from the spool, a rotary brake on the spool's shaft resists its rotation between paretic mid-stance (non-paretic toe off) and paretic swing onset (paretic toe off). This periodic resistance was designed to make the participant push off harder than usual on the paretic side when advancing forward. A motor helps cancel the brake's friction when the resistance is not active and also reels in the cable between trials [574].

After building and characterizing the device, we conducted a preliminary study at SUNY Stony Brook University, USA, with a 24-year-old female with left-side hemiparesis and gait asymmetry following pediatric traumatic brain injury. We found that ***GPT resistance increased paretic leg propulsive forces generated in late stance by 25% over baseline values***. Importantly, increased paretic propulsion **persisted** when GPT resistance was removed in post-braking trials, even though the participant walked with the GPT for a short training period (only ten 10 m trials).

Given these encouraging results, we are currently recruiting more participants for an evaluation study at SUNY Stony Brook University. This project is funded by the US National Institutes of Health (NIH) via grant #R03HD092822; Erin Vasudevan and Katherine J. Kuchenbecker are this grant's co-principal investigators.

More information: https://hi.is.mpg.de/project/gait-propulsion-trainer

## Immersive VR for Phantom Limb Pain

*Katherine J. Kuchenbecker, Elisabetta Ambron (University of Pennsylvania), Alexander Miller (University of Pennsylvania), Laurel J. Buxbaum (Moss Rehabilitation Research Institute), H. Branch Coslett (University of Pennsylvania)*

Figure 5.17: The right inset shows a lower-limb amputee using an early version of the created system. He wears an immersive VR headset and motion-tracking sensors on the upper and lower segments of both legs. The larger image shows his view of the virtual world, including the game scene and two intact legs that track the motions of his real legs.

Up to 90% of individuals who undergo amputation experience a persistent sensation of the missing limb, which is called a phantom limb A substantial subset of these people feel intense pain in the missing extremity; this **phantom limb pain (PLP)** often responds poorly to medications or other interventions and significantly interferes with quality of life. There is an urgent need for better treatments for PLP, particularly for people with lower-limb amputation.

Other researchers have published a small number of studies showing that virtual reality (VR) can alleviate phantom limb pain for some people, but the strength of the effect varies greatly, and some people show no response. The research objective of this project is to test the hypothesis that ***limitations in the verisimilitude of the sensory feedback provided by current PLP therapies limit their efficacy***.

Specifically we aim to determine whether PLP after leg amputation is reduced by high-quality, multi-modal feedback provided through immersive VR technology that tracks the participant's real leg motions and shows him or her a pair of intact legs moving in the same way. Immersive games that require a large range of leg motions add to the engagement and long-term interest of such a treatment option.

Pilot data collected from two individuals with PLP using an early version of this system strongly support the hypothesis in question [536]. After improving the system's hardware and software, we are now conducting a proof-of-concept study to serve as the next step in the translational pipeline.

This study is being conducted in Philadelphia, USA, and this project is funded by an NIH R21 grant to H. Branch Coslett.

More information: https://hi.is.mpg.de/project/immersive-vr-for-phantom-limb-pain

## TouchTable: A Musical Interface with Haptic Feedback for DJs

*Adam Spiers, Katherine J. Kuchenbecker*

Figure 5.18: The TouchTable prototype with headphones

DJing is a musical activity that involves the blending of songs (mixing) and the rythmic manipulation of sounds (scratching). Traditionally, DJs used vinyl records as their music sources. The vibration of a stylus (needle) in the groove of these records produced not just sound, but also subtle haptic sensations that could be felt by the DJ as they manipulated the record by hand. Modern DJ technology is digital in nature and no longer provides this haptic cue.

The TouchTable is a prototype device that augments digital music files with rich haptic feedback delivered by a vibrotactile actuator. This system can not only replicate the literal ’feel’ of vinyl but also provide additional sensations related to the music being played. We are investigating whether this tactile augmentation of music can be beneficial to the performance of novice and skilled DJs.

More information: https://hi.is.mpg.de/project/touchtable

## Effects of Vibrotactile Feedback on Residents Learning Robot-Assisted Surgery

*Haliza MatHusin, Katherine J. Kuchenbecker, Ernest (Ted) Gomez (Hospital of the University of Pennsylvania)*

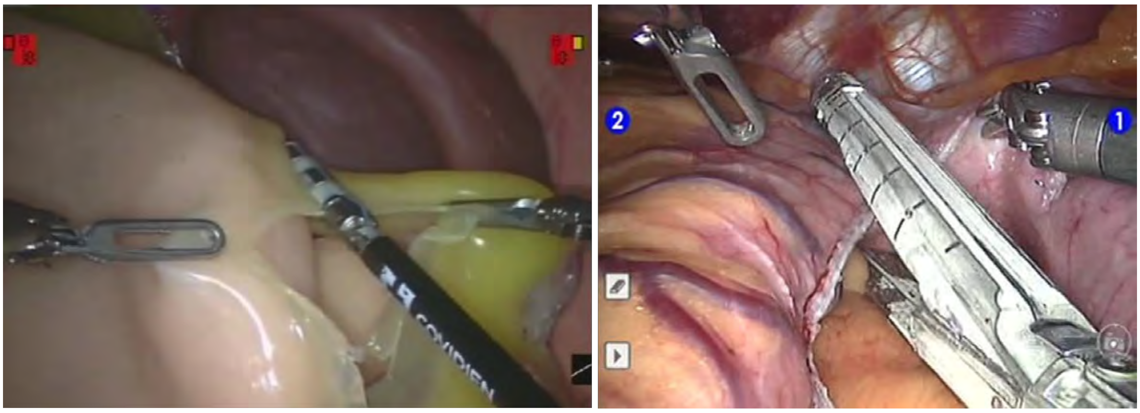

Figure 5.19: Screen captures of robot-assisted sleeve gastrectomy: (Left) Simulated surgery on our custom organ model. (Right) Clinical surgery during the stapling of the stomach.

The lack of haptic feedback is a potential limitation of existing robotic surgical systems. Members of Dr. Kuchenbecker's group at Penn previously invented a haptic feedback system named **VerroTouch** that is able to deliver the vibrations of surgical instruments to the surgeon's hands as he or she remotely operates a laparoscopic surgical robot. To record the instrument vibrations, the system uses MEMS-based accelerometers mounted on the arms of the surgical robot.

Our recent work has shown that instrument contact vibration is a potential quantitative measure of **technical surgical skill**. Compared to novices, experienced surgeons complete robotic surgery tasks with significantly lower instrument vibration magnitudes, shorter completion times, and fewer instrument contacts. These vibration metrics might be a better approach to current surgical skills assessment through manual video evaluation, which is time consuming and subject to rater bias.

Our aim is to continue to study the effect of haptic feedback on surgeons during robotic surgery. In this project, we are investigating the effect of haptic feedback on intraoperative performance of residents in robot-assisted surgery. Twelve resident surgeons were included in this study; eleven of them are right-hand dominant. Residents were assigned to a series of simulation training and were randomized to receive either haptic feedback through VerroTouch or no haptic feedback (the standard operating method). The performance of these residents was then tracked as they performed robotic surgery both in simulation and on live human patients. We hypothesize that *haptic feedback of instrument vibrations during robotic surgery practice will significantly improve resident skill in the operating room.*

This study was conducted at the Hospital of the University of Pennsylvania, USA.

More information: https://hi.is.mpg.de/project/effects-of-vibrotactile-feedback-on-residents-learning-robot-assisted-surgery

## Effects of Vibrotactile Feedback on Learning for Telerobotic Manipulation

*Gunhyuk Park, Katherine J. Kuchenbecker, Ernest Gomez (Department of Otorhinolaryngology – Head and Neck Surgery, Hospital of the University of Pennsylvania), Thamolwan Surakiatchanukul (Haptics Group, GRASP Laboratory, University of Pennsylvania), Michael Gosselin (Haptics Group, GRASP Laboratory, University of Pennsylvania)*

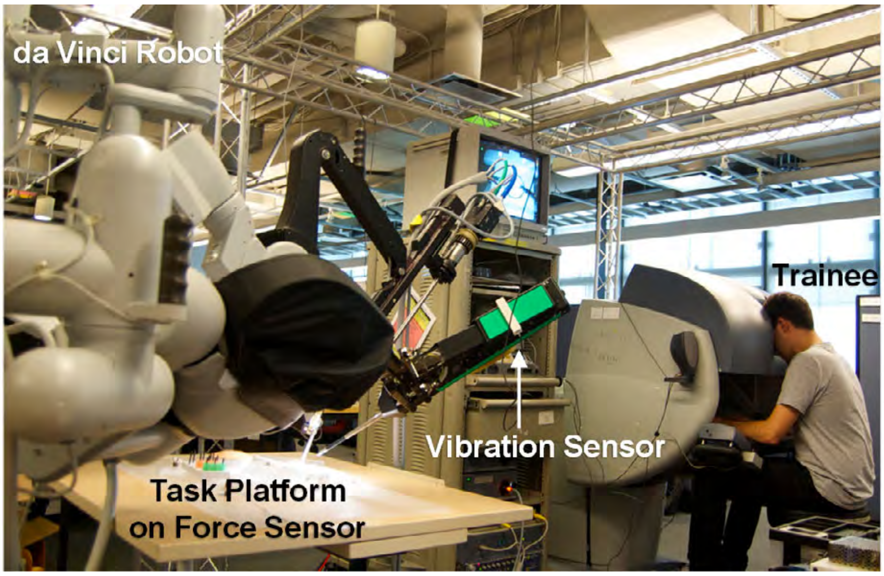


Figure 5.20: The overall configuration of the learning curve study, wherein surgical trainees learned to perform the peg transfer training task using a da Vinci surgical robot either with or without vibrotactile haptic feedback. A single 3-axis vibration sensor is attached to each arm of the da Vinci robot. A force sensor located below the platform measures the force that the trainee exerts with the robot.

Robotic minimally invasive surgery systems such as the Intuitive Surgical da Vinci system physically separate the surgeon from the surgical tools. As touch cues are known to play a critical role in manipulation tasks, the resulting loss of the sense of touch may affect the speed and skill with which one can learn to perform surgery through such a system.

We hypothesize that **tactile feedback of instrument vibrations** could accelerate trainee learning of key robotic surgery skills. Specifically we quantify task performance in terms of completion time, the force applied to the task materials, and the vibrations induced in the surgical instruments. We chose to test this hypothesis by application to a peg transfer task, as it is a well-known training task in minimally invasive surgery that is easy for novices to perform but difficult for them to master.

We designed a human-subject experiment that enrolled 16 fourth-year medical students intending to pursue a career in surgery. We then tested our hypothesis by evaluating the process by which each participant learned to perform peg transfer; half learned with vibrotactile feedback, and the other half learned without. Our overall findings show that this form of haptic feedback does indeed accelerate the learning process for peg transfer.

More information: https://hi.is.mpg.de/project/effects-of-vibrotactile-feedback-on-the-learning-curve-for-telerobotic-manipulation

## Digital VerroTouch

*Katherine J. Kuchenbecker, Haliza MatHusin, Bernard Javot, Ecda Erol*

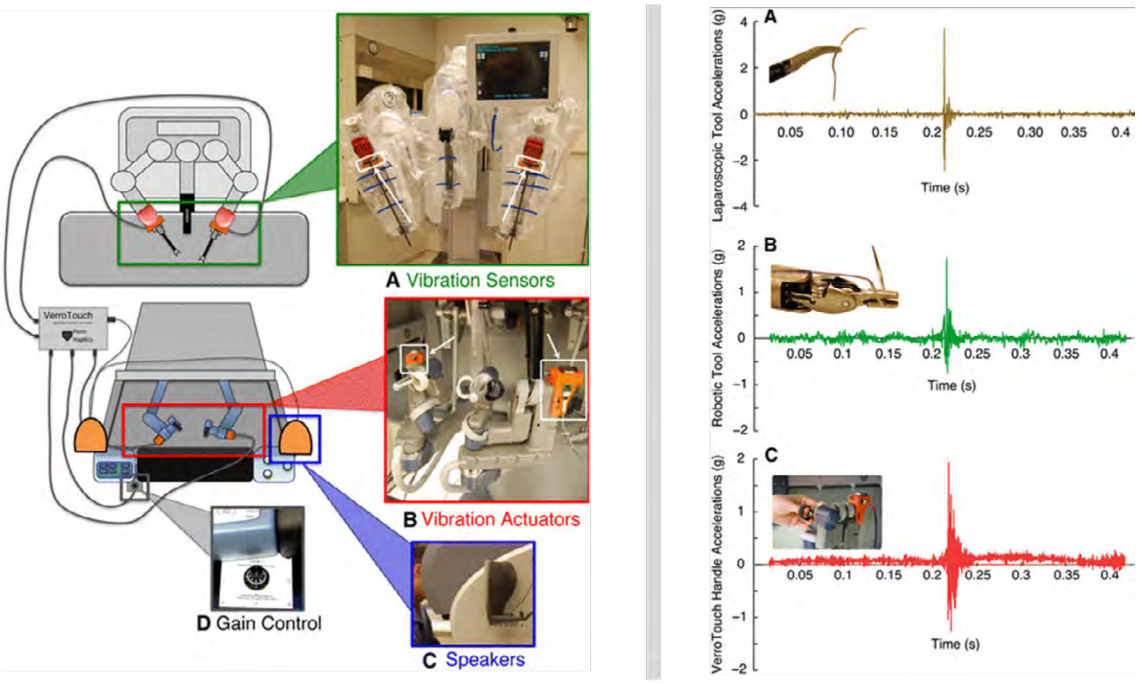


Figure 5.21: (Left) Overview of the current version of the VerroTouch system for providing haptic feedback in robotic surgery. (Right) High-frequency accelerations measured when grasping a needle using (a) a laparoscopic tool and (b) a robotic surgery instrument. (c) Accelerations produced by VerroTouch at the surgeon's control handle when the corresponding instrument experiences the accelerations shown in (b).

When performing minimally invasive robotic surgery, surgeons must currently rely only on their visual sense, as commercially available robotic surgery systems provide no touch feedback. Dr. Kuchenbecker and other members of the Penn Haptics Group previously invented a vibrotactile haptic feedback system that accurately replicates vibrotactile events for robotic surgical platforms.

Called **VerroTouch**, this system senses the contact vibrations of the robotic instruments using 3-axis accelerometers mounted on the robot's arms. It then reproduces these vibrations for surgeons to feel through a voice coil actuator that is attached to each handle of the surgeon's console. VerroTouch allows surgeons to perform robotic surgery in a teleoperated environment with sensations that are closer to those felt during normal surgery.

A robust, low-cost system with high accuracy is critically important for clinical settings. Therefore, this research project aims to create a new integrated version of the VerroTouch system that is able to sense, transmit and process vibrations both digitally and wirelessly. Moving from analog circuitry to a digital platform will not only improve system performance but will also enable us to study a wide range of new scientific questions in this domain, such as the optimal way of transforming 3D vibrations to 1D, or whether overlaying virtual cueswith real vibrotactile signals enables new capabilities.

More information: https://hi.is.mpg.de/project/digital-verrotouch

## Reducing 3D Vibrations to 1D

*Gunhyuk Park, Katherine J. Kuchenbecker*

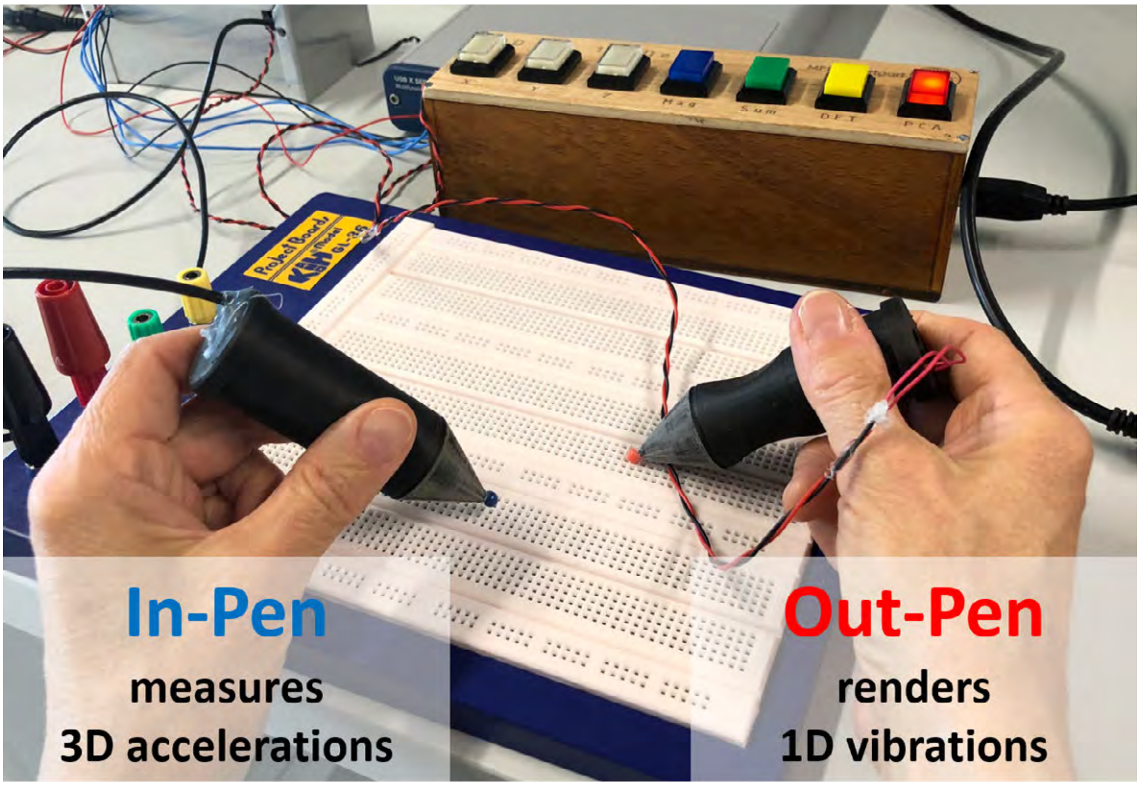


Figure 5.22: The overall system structure of the demonstration. 3D accelerations measured from the In-Pen are captured every 20 ms and converted to corresponding 1D vibration signals. When a user interacts with surrounding objects using the In-Pen, the user simultaneously feels the converted vibration from the Out-Pen.

Many scenarios arise wherein the high-frequency accelerations of a tool need to be captured and either portrayed for a human to feel or analyzed by a computer system. For example, this approach provides a simple way for a surgeon to feel tactile information from a remotely controlled surgical robot. Similarly, vibrations are the main signal of interest for the modeling of haptic textures.

In most scenarios, 3D accelerations are captured from real contact interactions and used as a realistic vibration source. Because humans cannot perceive the direction of high-frequency vibrations, haptics researchers usually **reduce these 3D accelerations into 1D signals** and render them using a single-axis vibration actuator. This dimensional reduction can be performed in many ways, and the chosen approach has a substantial impact on the resulting experience.

This research project implements a real-time conversion system that simultaneously measures 3D accelerations using an In-Pen and renders the corresponding 1D vibrations using an Out-Pen. The user can freely interact with various objects using the In-Pen, which contains a 3-axis high-bandwidth accelerometer. The captured accelerations are converted to a single-axis signal, and the Out-Pen renders the reduced signal for the user to feel using a Tactile Labs Haptuator.

Our system can quickly switch between seven different conversion methods ranging from the simple use of a single-axis signal to applying principal component analysis (PCA). After gathering informal feedback via conference demonstrations [561],[548], we are now investigating both the quantitative signal similarity and the qualitative perceptual similarity between the 3D accelerations and the reduced vibrations to determine which method is optimal for each scenario where such an algorithm is needed.

More information: https://hi.is.mpg.de/project/reducing-3d-vibrations-to-1d

## Robust Visual Augmented Reality in Robot-Assisted Surgery

*Maria Paola Forte, Katherine J. Kuchenbecker*

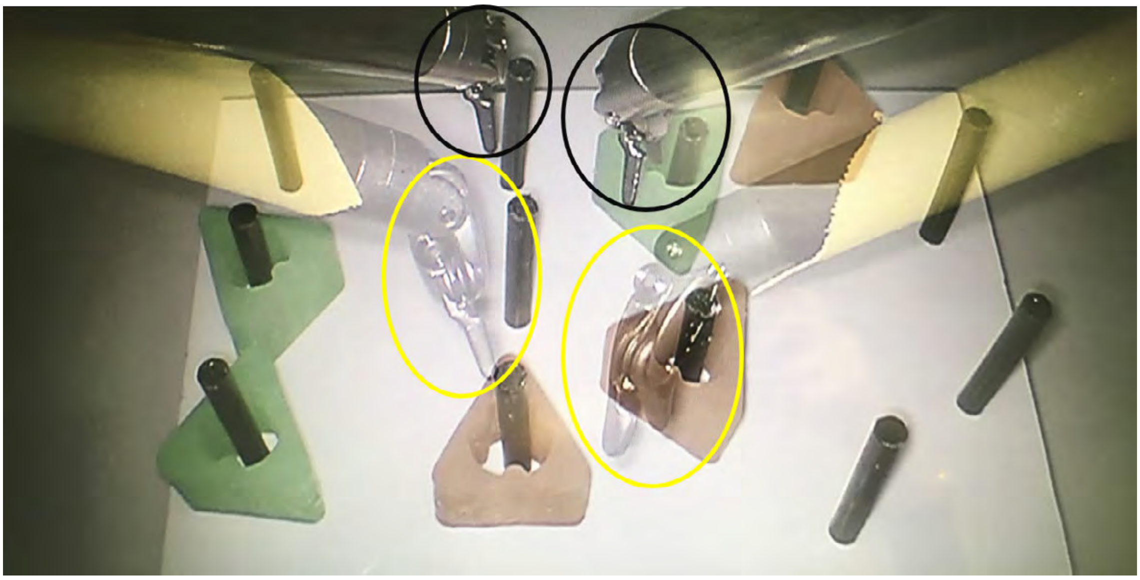

Figure 5.23: Augmented reality (AR) could be used to help surgeons learn and also perform minimally invasive surgery. In the sample scenario shown above, a trainee manipulates the real tools of a da Vinci robot (circled in yellow for visibility). This real tool view is overlaid on a pre-recorded stereoscopic view of an expert's instruments (circled in black for visibility), from which the trainee can learn.

Out of the many technologies being applied to healthcare delivery today, **augmented reality (AR)** is one of the most exciting. We believe AR can be applied to surgical robots to provide surgeons with preoperative data, intraoperative data, and useful computational tools directly in their field of vision. To have a universal AR system that does not require special access to the robot's data streams, we have chosen to implement our AR system by applying computer vision techniques to the robot's stereo camera images.

The research objective of this project is to test the hypothesis that *real-time stereo video analysis and AR can increase safety and task efficiency in robot-assisted surgery*. Specifically, we aim to develop a **robust AR system** that delivers the envisioned feedback to a surgeon while he or she controls a surgical robot identical to those used on human patients [580].

In addition to writing all of the system's **software** for **real-time execution**, we will design clinically relevant simulated surgical tasks. We will then conduct a **human-subject study** to gather both quantitative and qualitative feedback from potential users so that we can definitively evaluate the future potential of this new approach to AR in surgery.

More information: https://hi.is.mpg.de/project/robust-visual-augmented-reality-in-robot-assisted-surgery

## Teleoperating Max’s Head and Arms

*Mayumi Mohan, Katherine J. Kuchenbecker*

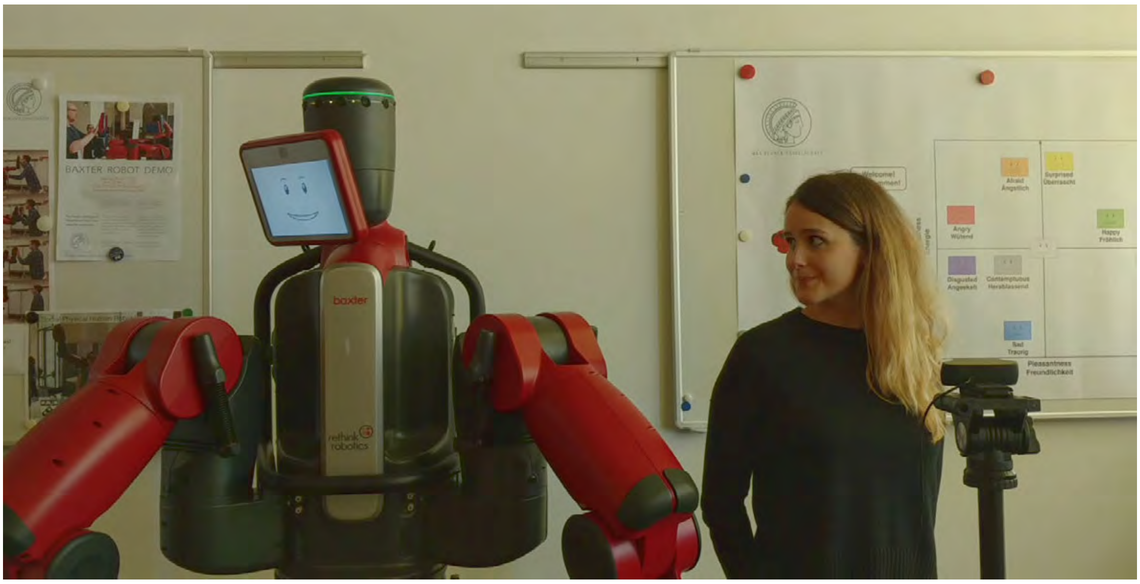


Figure 5.24: A human teleoperates the head of Max, our Baxter robot, via camera-based pose estimation and facial emotion recognition.

Many situations arise where it is beneficial for a human to control the movements of a robot at a distance, such as handling hazardous materials, doing surgery deep inside the human body, or taking part in remote meetings with other people. In these scenarios, the **robot’s control interface** has a significant influence on the effectiveness of the interaction.

This project aims to invent an intuitive and user-friendly interface for remotely controlling Max, our Baxter Research Robot (developed by Rethink Robotics). Given our ongoing research in **social-physical human-robot interaction**, we are particularly interested in an interface that will allow the operator to control the robot in an expressive and physically interactive manner. We envision scenarios where the remote-controlled robot acts as an exercise partner or coach, performing collaborative tasks or playing interactive games with a human. Such a scenario may take place in an elderly care home or a physical rehabilitation center.

We are in the process of assembling a control interface suitable for the envisioned use case [546]. In our present design, Max’s **arm and head movement** are controlled via an inertial-sensor-based motion-capture suit (XSens MVN) worn by the operator. The same person simultaneously controls the **robot’s face** via camera-based facial emotion recognition, as shown in the figure above. We will refine and evaluate this interface and then use it to prototype a wide range of social-physical interactions between a human and a robot.

More information: https://hi.is.mpg.de/project/teleoperating-max-s-head-and-arms

## Hierarchical Structure for Learning from Demonstration

*Siyao Hu, Katherine J. Kuchenbecker*

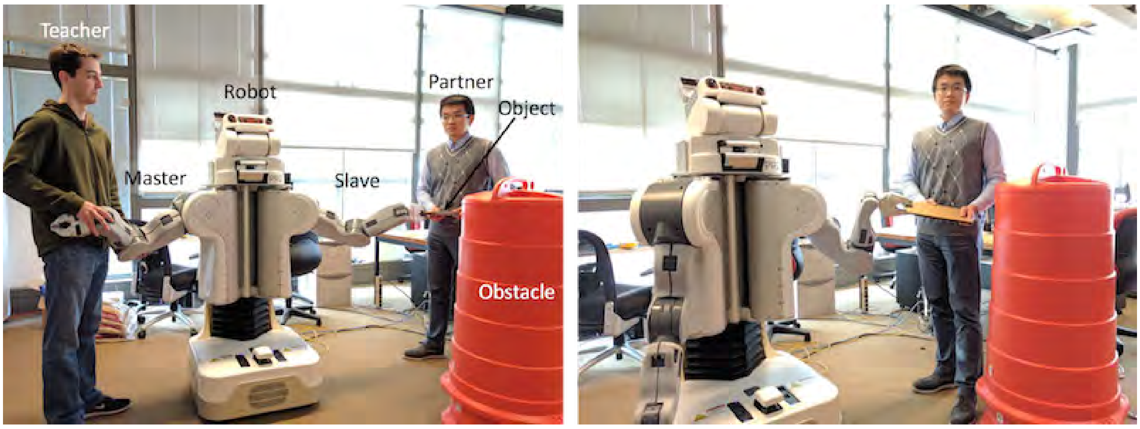


Figure 5.25: The left image shows the training scenario for collaborative manipulation with a Willow Garage PR2, where the teacher guides the robot to manipulate an object with the partner. The right image shows the testing scenario, where the robot automatically collaborates with the partner in the same task.

Many modern humanoid robots are designed to operate in human environments, like homes and hospitals. If designed well, such robots could help humans accomplish tasks and lower their physical and/or mental workload. As opposed to having an operator devise control policies and reprogram the robot for every new situation it encounters, **learning from demonstration (LfD)** provides a direct method for users to program robots.

We present a hierarchical structure of task-parameterized models for learning from demonstration: ***demonstrations from each situation are encoded in the same model, and models are combined to create a skill library***. This structure has four advantages over the standard approach of encoding all demonstrations together: users can easily add or delete existing robot skills, robots can choose the most relevant skill to use for new tasks, robots can estimate the difficulty of these new test situations, and robot designers can apply domain knowledge to design better generalization behaviors.

We elucidate the proposed structure via application to a collaborative task where a robot arm moves an object with a human partner, as shown in the figure above. The demonstrations were collected for three situations using a novel kinesthetic teleoperation method [569] and were encoded using **task-parameterized Gaussian mixture models** (TP-GMM). Naive participants then collaborated with the robot while it was controlled by a passive model with only gravity compensation, a single TP-GMM encoding all demonstrations, and the proposed structure. Our method achieved significantly **better task performance and subjective ratings** in most tested scenarios.

More information: https://hi.is.mpg.de/project/hierarchical-structure-for-learning-from-demonstration

## Exercise Games with Baxter

*Katherine J. Kuchenbecker, Mayumi Mohan, Naomi T. Fitter (University of Southern California), Michelle J. Johnson (University of Pennsylvania)*


Figure 5.26: A user playing the exercise games with Max, our Baxter robot. The six games on the left involve physical interaction between the user and Max, whereas Flamenco and Roboga require motion but no contact.

Improvements in healthcare have led to an increase in human life expectancy. Members of this aging population want to stay healthy and active, but many forms of exercise and physical therapy are expensive, boring, or inefficient. Past research has shown that **well-designed robots** can play a vital role in motivating users to perform regular exercise and physical therapy.

To discover how people respond to physical exercise interactions with a robot, we have developed **eight human-robot exercise games for Max**, our Baxter Research Robot (developed by Rethink Robotics): six of these games involve some form of physical contact with the robot, and two involve performing movements as directed by the robot, which has been the standard approach in prior work. These games were developed with the input and guidance of experts in game design, therapy and rehabilitation, as well as through extensive pilot testing [565],[568]. The viability of the games was then formally evaluated in a user study conducted at the Rehabilitation Robotics Laboratory at the University of Pennsylvania.

Our subject group included 20 younger and 20 older adult users. Participants of both age groups were willing to enter Baxter's workspace and physically interact with the robot through all of these games [555]. Additionally, participating in the experiment caused a significant increase in **user trust** and **confidence** in Baxter [555]. Careful analysis of the human-robot interactions that occurred throughout the study provided us with detailed feedback on the **usability** of all of the games. These results support the potential use of bimanual humanoid robots for social-physical interaction in exercise and will help guide our ongoing efforts in this research domain.

More information: https://hi.is.mpg.de/project/Exercise-Games-Baxter

## How Should Robots Hug?

*Alexis Block, Katherine J. Kuchenbecker*


Figure 5.27: A participant hugging the PR2 with its custom Soft-Warm outfit during the experiment.

Hugs are one of the first forms of contact and affection that humans experience. Not only are hugs a common way to provide comfort, support, or affection, but they have also been shown to have significant health benefits. Hugs can lower blood pressure, increase oxytocin levels (the hormone that makes you happy), lower cortisol levels (the hormone that makes you stressed), and improve your immune system.

As roboticists who study human interaction, we are interested in creating robots that can hug humans as seamlessly as people hug each other. The purpose of this first **HuggieBot** project was to evaluate human responses to different physical characteristics and hugging behaviors exhibited by a hugging robot. Specifically, we aimed to test the hypothesis that *a soft, warm, touch-sensitive PR2 humanoid robot can provide humans with satisfying hugs by matching both their hugging pressure and their hugging duration* [570],[562],[559].

After a brief introduction and opening survey, thirty relatively young participants with mostly technical backgrounds experienced and evaluated twelve hugs with the robot. Each person started the study with three randomly ordered trials that focused on physical robot characteristics, wherein the robot was Hard-Cold, Soft-Cold, or Soft-Warm. The study then proceeded to nine randomly ordered behavioral trials with all combinations of low, medium, and high hug pressure and low, medium, and high hug duration.

Analysis of the results showed that participants significantly preferred soft, warm hugs over hard, cold hugs [534]. Furthermore, users preferred hugs that physically squeeze them and then release immediately, when they are ready for the hug to end. Taking part in the experiment also significantly increased positive user opinions of robots and robot use.

More information: https://hi.is.mpg.de/project/huggiebot-how-should-robots-hug

## HuggieBot 2.0: A More Huggable Robot

*Alexis Block, Bernard Javot, Katherine J. Kuchenbecker, Roger Gassert (ETH Zurich), Otmar Hilliges (ETH Zurich)*

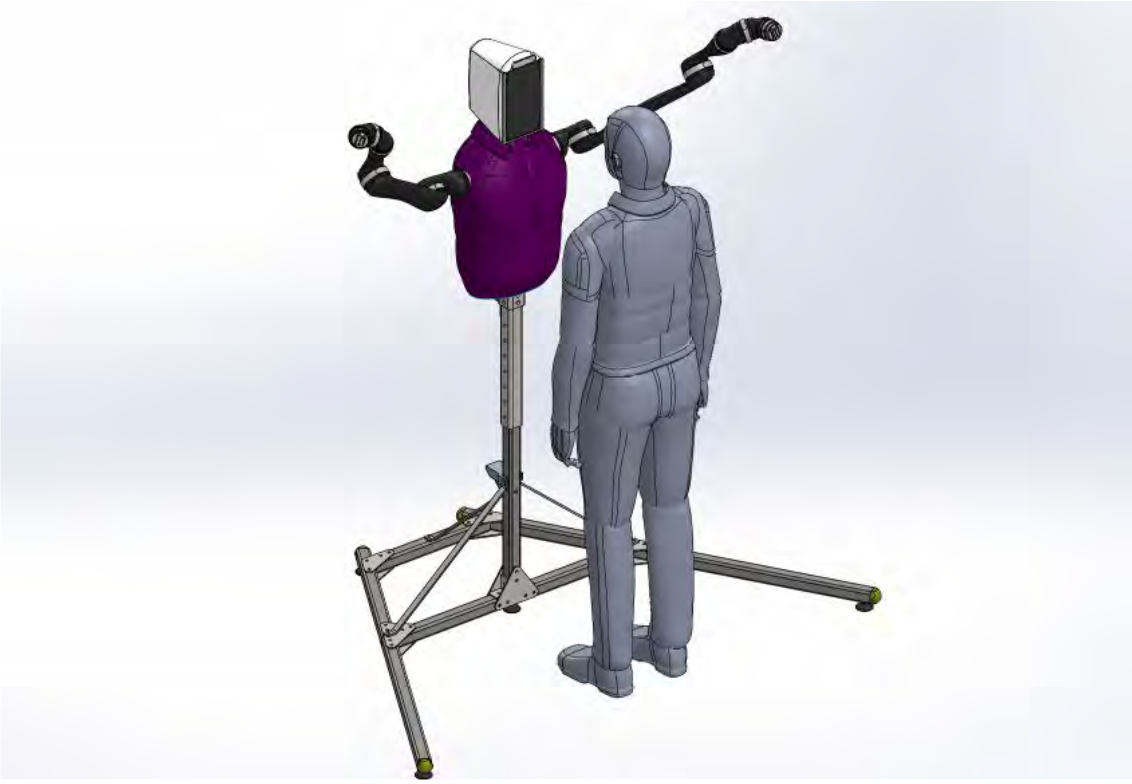

Figure 5.28: A CAD model of HuggieBot 2.0. This robot features two Kinova JACO 6-DOF arms, a customized v-shaped base for close interactions, an adjustable stand to adapt to user height, an inflatable pressurized chest to measure hug tightness, and a screen for showing either an animated face or a pre-recorded video message.

Our first HuggieBot research project showed that people are generally accepting of soft, warm robot hugs that squeeze the user and release promptly. In the future we would like to see if robot hugs can provide the same known health benefits as human hugs. Testing this hypothesis requires a new version of HuggieBot that allows more control over hugs than our original robot test platform; thus, we are now designing, building, and programming HuggieBot 2.0.

The HuggieBot 2.0 custom robot design builds off critiques of the Willow Garage PR2 robot, which was used in our previous experiment [534]. A v-shaped stand will allow participants to get closer to the robot without having to lean over a large base, which was a pain point with the PR2.

This new platform uses two Kinova JACO 6-DOF arms, which are slimmer, quieter, and smoother than the arms of the PR2. To accommodate various hugging arm placements and to streamline the hugging process, we are currently creating an inflatable chest that includes haptic sensors; the chest will additionally be soft and heated because users responded favorably to these two features during the previous study. A final element of HuggieBot 2.0 is a face screen that can show both animated faces or pre-recorded video messages to make the hug experience more personal, more enjoyable, and less mechanical for the user.

To test if HuggieBot 2.0 can provide similar health benefits as human hugs, we plan to carefully induce mental stress on study participants and then offer either a human hug, an active robot hug, a passive robot hug, or self-soothing. We will use objective measures such as heart rate, facial expression, and cortisol and oxytocin levels to evaluate users' physical responses to the different hugs.

More information: https://hi.is.mpg.de/project/huggiebot-2-0-a-more-huggable-robot

## General Tactile Sensor Model

*Hyosang Lee, Katherine J. Kuchenbecker*

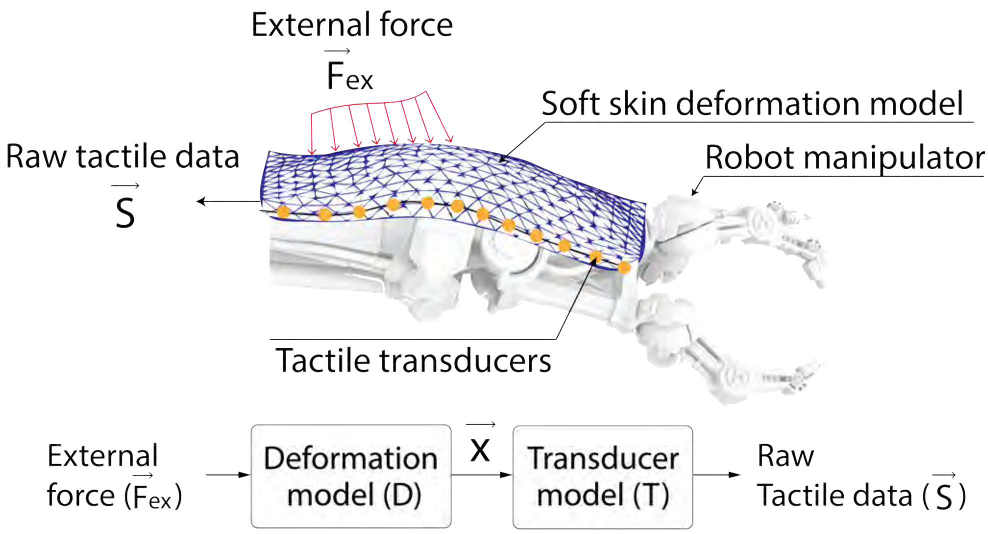


Figure 5.29: This schematic shows the generalized process of obtaining raw tactile data when external forces are applied to the sensor. The bottom block diagram represents a calibration model separated by a deformation model (D) and a transducer model (T).

Robots working in unstructured environments and alongside people need to be able to sense contact information from both intentional and unintentional interactions. Soft and skin-like tactile sensors can provide a robot with physical contact information from the surrounding environment and beneficially also supply a compliant interface for conformal contact and impact protection.

Raw data from soft tactile sensors is normally nonlinear and has high dimensionality. These characteristics make calibration of soft tactile sensors very complicated, such that one often cannot convert the readings into physical values. The research objective of this project is to *implement a calibration method that is applicable to a wide range of soft tactile sensors*.

We aim to develop a **parameter identification method** to obtain the calibration model by combining model-based and data-driven approaches. Specifically, mechanical deformation of the soft body of the sensor can be estimated through solid mechanics. The transducer model, which produces raw data from mechanical deformation, is obtained by a data-driven method.

The future outcome of this ongoing project will be to provide a generalized method that could be used to calibrate a wide range of soft tactile sensors so that their outputs can be confidently interpreted in physical units such as pressure.

More information: https://hi.is.mpg.de/project/general-tactile-sensor-model

## Large-Scale Fabric-Based Tactile Sensor

*Hyosang Lee, Katherine J. Kuchenbecker, Kyungseo Park (Korea Advanced Institute of Science and Technology), Jung Kim (Korea Advanced Institute of Science and Technology)*

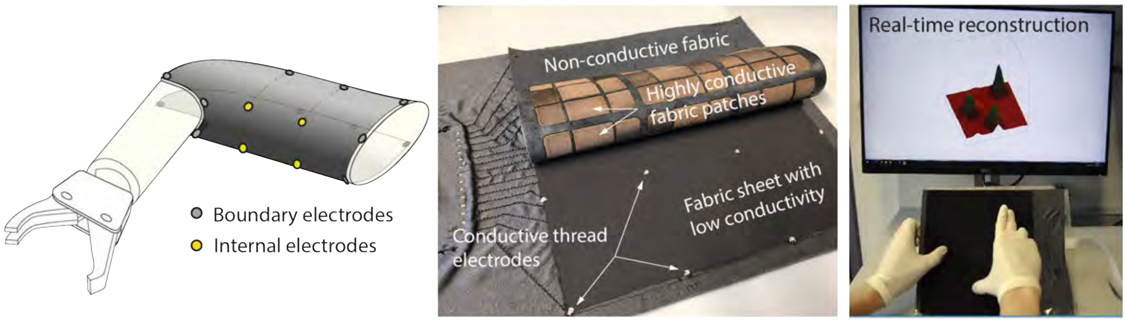


Figure 5.30: The left image shows the large-scale tactile sensor wrapped over the body of a robot manipulator. The sensor is able to measure contact forces using distributed electrodes. In the center, a prototype of the fabric-based tactile sensor is presented. The right image shows a demonstration of simultaneous sensing of multiple contacts.

Being able to cover the entire body of a robot with soft tactile sensors has become an attractive concept in intelligent robotics. Soft, stretchable materials can conform around surfaces and also absorb impacts, which is beneficial for safety around humans and in unpredictable environments [558].

However, such a sensor has not yet been developed because such systems are challenging to design and manufacture. In the traditional approach to tactile sensors, a high number of transducers must be placed on soft and stretchable substrates and then wired together. This strategy results in large wiring bundles, rigid elements, and many potential points of failure.

A geophysical imaging method called Electrical Resistance Tomography (ERT) has recently shown promising results in tactile sensing, simultaneously achieving large-area coverage, ease of manufacturability, and robustness. ERT is a reconstruction method that estimates the internal conductivity distribution of a conductive medium using electrodes that are typically located only around the boundary. These electrodes are used to inject electrical currents into the medium and also measure the resulting voltage potentials. The internal conductivity distribution is reconstructed from these measurements and calibrated to external stimuli.

This project aims to develop a fabric-based ERT tactile sensor that can cover a large area [552]. To achieve this goal, we focus on two aspects. First, we develop a low-cost and reliable sensor manufacturing method that involves fabric layers and sewing. Second, we focus on overcoming the limitations of conventional ERT-based sensors to improve sensing resolution and framerate.

In the future, we plan to optimize the sensor design and conductivity reconstruction techniques to improve tactile sensor's capabilities. Overall this project aims to provide a systematic method to develop whole-body tactile sensors for any robot.

More information: https://hi.is.mpg.de/project/large-scale-fabric-based-tactile-sensor

## Haptic Empathetic Robot Animal (HERA)

*Rachael Burns, Katherine J. Kuchenbecker*

Figure 5.31: A conceptual design and initial prototype of the robot animal - a NAO robot enclosed inside a soft koala form.

**Autism spectrum disorder (ASD)** is a complex condition that impacts many systems in the body, from neurological aspects to physical co-morbidities. Children with autism often endure sensory overload, may be nonverbal, and may have difficulty understanding and relaying emotions. These combined experiences can cause the child heightened stress when interacting with others. As the rate of autism diagnosis continues to rise, there is an urgent need to develop mechanisms to help autistic children cope with exciting or unfamiliar situations.

**Deep touch pressure therapy**, the standard method of intervention in the American special needs education system, uses tools such as hand brushes or weighted vests to reduce stress. However, these tools provide the children neither with control over the contact nor with opportunities to socialize. **Animal-assisted intervention** can improve the stress levels and behavior of autistic children. However, live animal companions are not always feasible given a number of factors, such as the resources of caregivers. **Robot therapy** also shows promise - children with ASD appear to respond more favorably to robots than human strangers, possibly due to the robot's safely predictable behavior.

Our research objective is to test the **hypothesis** that *an appropriately designed robot animal can reduce stress in children with autism and empower them to engage in social interaction* [560].

We enclose NAO, a commercially available humanoid research robot, inside a soft koala suit to serve as the initial robot animal prototype. The robot animal (which has been named 'HERA') will use a system of custom tactile sensors to identify and react to the users' touches. As autism is a spectrum with low to high functioning patients, HERA's personality and behavior will be customizable to generate reactions to tactile interactions that are appropriate to the user.

More information: https://hi.is.mpg.de/project/haptic-empathetic-robot-animal-hera

## 5.4 Awards & Honors

### 2018

**Mayumi Mohan** is selected to participate in the HRI Pioneers 2019 Workshop

**Alexis Block** wins a 2,500 USD research grant from the IEEE Technical Committee on Haptics

**Yasemin Vardar** is selected to participate in the Sign up! career-building program for excellent female postdocs in the Max Planck Society

**Hasti Seifi** wins the 2017 EuroHaptics Society prize for the best Ph.D. thesis in the field of haptics

Visiting Professor **Brent Gillespie** receives a Humbold Research Award

**Hyosang Lee** is selected to participate in the RSS Pioneers 2018 Workshop

**Hyosang Lee** receives a postdoctoral fellowship from the National Research Foundation of South Korea

**Hasti Seifi** receives a two-year postdoctoral fellowship from the National Sciences and Engineering Research Council (NSERC) of Canada

**Rachael L'Orsa** receives a four-month scholarship from the Deutscher Akademischer Austauschdienst (DAAD)

Ph.D. student **Alexis Block** is elected as co-general chair of the HRI Pioneers 2019 Workshop

### 2017

**Alexis Block** is selected to participate in the HRI Pioneers 2018 Workshop

A paper by **Eric Young** and **Katherine J. Kuchenbecker** is a finalist for the Best Poster Paper Award at the 2017 IEEE World Haptics Conference

**Rachael Bevill Burns** is selected to receive a grant from the Whitaker International Program

### Faculty Appointments

**Gunhyuk Park**, a postdoc in the Haptic Intelligence department, is appointed as a Assistant Professor at the Gwangju Institute of Science and Technology in South Korea. He begins his new position on February 1, 2019.

## 5.5 Director profile: Katherine J. Kuchenbecker

Katherine J. Kuchenbecker received her B.S. (2000), M.S. (2002), and Ph.D. (2006) all in Mechanical Engineering from Stanford University. After ten months as a postdoctoral researcher at the Johns Hopkins University, she joined the faculty of the University of Pennsylvania in the Department of Mechanical Engineering and Applied Mechanics (MEAM) as the Skirkanich Assistant Professor of Innovation. She earned tenure and was promoted to Associate Professor in 2013, at which point she also obtained a secondary appointment in the Department of Computer and Information Science. She served as MEAM Undergraduate Curriculum Chair from July 2013 to June 2016 and held the Class of 1940 Bicentennial Endowed Term Chair from July 2015 to December 2016.

In January 2017 she joined the Max Planck Society as a Scientific Member and Director at the Stuttgart site of the Max Planck Institute for Intelligent Systems. She is the IMPRS-IS Spokesperson, and she has served as Managing Director for the institute's Stuttgart site since January 2018. Dr. Kuchenbecker's research centers on haptics, which involves human and robot interaction with physical objects through the sense of touch. She is most well known for her work on data-driven haptic textures, fingertip tactile displays, haptic feedback in robotic surgery, and tactile information processing for robots.

Dr. Kuchenbecker received a US National Science Foundation CAREER Award in 2009 and gave the RSS Early Career Spotlight Talk the same year. She received the 2012 IEEE Robotics and Automation Society Academic Early Career Award. In 2012 she also delivered a widely viewed TEDYouth talk on haptics, and she now gives several keynote talks about her research every year. Her team has been honored with numerous best paper and best demonstration awards, and she has won awards as an associate editor, scientific reviewer, advisor, and teacher. She co-chaired the IEEE Haptics Symposium in 2016 and 2018 and the IEEE Technical Committee on Haptics from 2015 to 2017. Five of her previous trainees are now assistant professors at top research universities.

She is a co-founder and Chief Scientific Advisor of two start-up companies that have licensed technology that she and her team invented at the University of Pennsylvania. Tactai focuses on delivering realistic haptic cues in consumer products, and VerroTouch Medical aims to bring haptic feedback to robotic surgery.

# Dr. Katherine J. Kuchenbecker

## Appointments

| | |
|---|---|
| 01/2017 – present | Director at the Max Planck Institute for Intelligent Systems |
| 01/2017 – present | Spokesperson of the IMPRS-IS, a new joint Ph.D. program between MPI-IS, the University of Stuttgart, and the University of Tübingen |
| 01/2018 – present | Managing Director of the MPI for Intelligent Systems, Stuttgart |
| 01/2019 – present | Adjunct Professor, MEAM Department, University of Pennsylvania |

## Awards & Honors (Selected)

| | |
|---|---|
| 2017 | Finalist for Best Poster (with co-author), IEEE World Haptics Conference |
| 2015 | TCPW Award for Excellence in Undergraduate Advising, Univ. of Pennsylvania |
| 2014 | Lindback Award for Distinguished Teaching, University of Pennsylvania |
| 2013 | Best Demo (by committee vote, with co-authors), SIGGRAPH Asia |
| 2013 | Best Cognitive Robotics Paper (with co-authors), IEEE ICRA |
| 2013 | Best Demo (by audience vote, with co-authors), IEEE World Haptics Conf. |
| 2012 | IEEE Robotics and Automation Society Academic Early Career Award |
| 2012 | Best Demo (three-way tie, with co-authors), IEEE Haptics Symposium |
| 2009 | US National Science Foundation CAREER Award |
| 2007 | Best Haptic Technology Paper (with co-author), IEEE World Haptics Conference |

## Selected Organization and Community Service (2014–2019)

| | |
|---|---|
| 2019 | Program Committee, IEEE World Haptics Conference |
| 2018 | Co-organizer of the "Haptipedia" workshop at AsiaHaptics |
| 2018 | General Co-Chair, IEEE Haptics Symposium |
| 2017 | Program Committee, IEEE World Haptics Conference |
| 2016 | General Co-Chair, IEEE Haptics Symposium |
| 2015 | Co-organizer of the "Cutaneous Feedback for Teleoperation in Medical Robotics" workshop at the IEEE World Haptics Conference |
| 2014 | Publications Chair, IEEE Haptics Symposium |

## Memberships (2014–2019)

Integrative Computational Design and Construction for Architecture (IntCDC, new DFG Cluster of Excellence centered at the University of Stuttgart), one of 22 Principal Investigators since 2018
Institute for Electrical and Electronic Engineers (IEEE), Senior Member since 2018
Max Planck ETH Center for Learning Systems (CLS), Member since 2017
Intel Network on Intelligent Systems (NIS), Member since 2017

## Startup Activity (2014–2019)

Tactai, Inc., Boston, USA, Co-founder and Chief Scientific Advisor, 2014 – present
VerroTouch Medical, Inc., Philadelphia, USA, Co-founder and Chief Scientific Advisor, 2016 – present

## Selected Keynote Talks and Seminars (2014–2019)

Plenary, *IEEE/RSJ International Conf. on Intelligent Robots and Systems (IROS)*, Macau, China, 2019
Keynote, *RehabWeek (six conferences combined)*, Toronto, Canada, 2019
Keynote, *AsiaHaptics*, Songdo, South Korea, 2018
Keynote, *Hand, Brain and Technology*, Monte Veritá, Switzerland, 2018
Keynote, *Robotics: Science and Systems (RSS)*, Pittsburgh, USA, 2018
Seminar, Institute of Robotics and Intelligent Systems (IRIS), ETH Zurich, Switzerland
Keynote, *IEEE Virtual Reality (VR)*, Reutlingen, Germany, 2018
Seminar, Department of Neuroscience, University of Tübingen, Germany, 2018
Seminar, Institute for Systems Theory and Automatic Control, University of Stuttgart, Germany, 2017
Keynote, *Annual Meeting of the Vision Sciences Society (VSS)*, St. Pete Beach, USA, 2017

# 6 AUTONOMOUS LEARNING

## 6.1 Research Overview

Our mission is to make robots learn in a developmental fashion - similar to children. Why should that be useful? Already now robots are present in many areas of our life, taking over tasks that are too dangerous, too repetitive, or require too high precision for humans. However, their area of application is bound because preprogrammed robots are not able to successfully interact with our complex and constantly changing world.

In living beings, this problem is solved by an interplay of learning, self-organization and innate information. Self-organization—ubiquitous in nature—offers promising perspectives for practical applications because it is based on local interactions and typically scales well to large fault-tolerant systems. In our research, we build the theoretical basis for self-organized robot control. We study how dynamical systems and information theory can be used to generate sensory-motor coordination in robots, see project *Self-exploration of Behavioral Primitives*.

In order to make robots successful in learning new skills, they have to extract as much information as possible from experience. Learning a model of its body and the environment can capture the information needed to master unknown future tasks. Classical regression models are not good at predictions in new situations. We investigate how an algorithm can identify the underlying relationships in data, thereby obtaining a model that can extrapolate well, see project *Equation Learner for Extrapolation and Control*. In a first study we have shown that such an approach can be very efficient in learning a simple control task. We are currently investigating more challenging robotic systems.

For actually learning, improving and controlling goal-oriented behavior, we use reinforcement learning and optimal control methods. Our aim is to make reinforcement learning so data-efficient that it can be applied to real systems easily. We currently take two complementary routes. One route is to combine optimal control methods and planning with reinforcement learning, see project *Data-driven robust and optimal control in complex systems*. The other route is to learn many tasks at the same time using hierarchical learners that are intrinsically motivated to explore appropriate goal spaces and understand relational structures in the environment. We obtained first results on this challenging task [582] and are intensifying our efforts in this direction.

Machine Learning methods are the working horses behind our recent developments in model learning, reinforcement learning, and representation learning. While investigating the state of the art methods we realized a potential for speeding up stochastic gradient descent in practical applications, see project *Stepsize adaptation for stochastic optimization*.

Autonomous Learning systems need to build the right representations for good generalization capabilities and efficiency in learning new skills. We investigate unsupervised learning of independent representations using information theory and variational inference.

Our work on principles of developmental learning is accompanied by our efforts in creating a suitable hardware for validating our algorithms. We have assembled and modified a humanoid robot named “Poppy”, which is an open-source hardware project of a robot at the size of a 4-year old child. In order to give the robot a sense of touch, which is important for learning to interact with the environment, we developed a novel large surface haptic system, see *Robust and Affordable Haptic Sensation*. Our system is exploiting machine learning to infer touch information from a few deformation sensors inside the structure making it highly robust.

In the following we provide more detail on above mentioned projects.

## 6.2 Selected Research Projects



### Self-exploration of Behavioral Primitives

*Georg Martius, Cristina Pinneri*

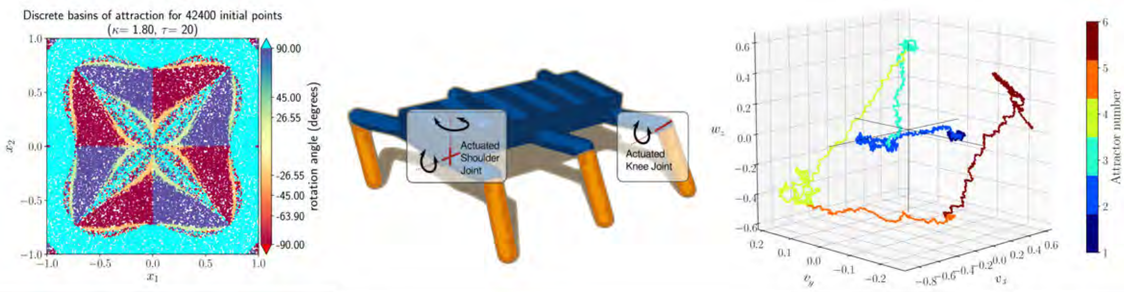


(a) (b) (c)

Figure 6.1: (a) Discrete basins of attraction for a 2 DoF system. Different colors indicate different kinds of attractors, showing a rich variety of behaviors even in a low-dimensional system. (b) Hexapod simulated robot with 18 DoF. The robot is initialized in a random environment and thanks to the repelling potential is able to execute several behaviors in succession. The sweeping through the behavioral landscape is shown in the velocity space (c).

One of the primary questions of this project is understanding how to generate structured behaviors for a robotic system, without specifying any reward nor objective function. The systems we are interested in should be able to exploit the embodiment among brain, body and environment to self-explore a wide range of behaviors and automatically extract a suitable controller for each of them. In order to bootstrap this goal-free explorative process, we use a biologically plausible synaptic mechanism for self-organizing controllers, in particular, differential extrinsic plasticity (DEP) (Der & Martius, PNAS 2015), which has proven to enable embodied agents to self-organize their individual sensorimotor development and generate highly coordinated behaviors during their interaction with the environment.

We use a dynamical systems framework to describe a behavior as an attractor in the brain-body-environment system using DEP. The behaviors self-organize within a few seconds of live interaction and are specific to the embodiment of the robot. Each behavior corresponds to a potentially useful motion primitive.

The behavioral landscape generated by DEP is then explored thanks to a "repelling potential" which allows the system to actively explore all its attractor behaviors in a systematic way. With a view to a self-determined exploration of goal-free behaviors, our framework enables switching between different motion patterns in an autonomous and sequential fashion. Our algorithm is able to recover all the attractor behaviors in a toy system and it is also effective in two simulated environments. A spherical robot discovers all its major rolling modes and a hexapod robot learns to locomote in 50 different ways in 30min [588].

More information: https://al.is.mpg.de/project/sweeping-behaviorspace

## Equation Learner for Extrapolation and Control

*Georg Martius, Michal Rolinek, Jia-Jie Zhu, Subham S. Sahoo, Christoph H. Lampert*

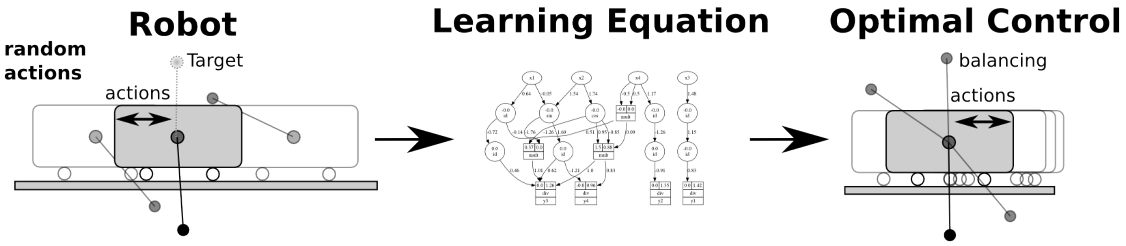


Figure 6.2: Example of the equation learning framework for control. The robot is first controlled with some random action sequence to excite the system. Using this data an equation describing the system dynamics is learned. Utilizing this model, a planning based control strategy (MPC) can make the robot perform a task, here swinging up and balancing the pole.

Intelligent systems today, including physical robots, are capable of amazing feats compared to their early predecessors. However, much of the planning and control relies on the accurate identification of the robotic system, often in the form of known physics equations derived by experienced engineers. In the case when the system model is unknown, various function approximations have been used in the context of model-based reinforcement learning. There is a typical trade-off between the explicability and model capacity when applying function approximation, e.g., deep neural networks may model complex system dynamics but are hard to interpret and optimize. Typically, control and planning tasks require much higher accuracy compared to regular machine learning tasks. In addition, it naturally occurs that unseen parts of the state space are visited. So we wish to address the problem of learning models that are highly accurate and also able to extrapolate. Instead of resorting to black-box function approximation we designed a system to *identify the system equation from a limited data range* using a gradient-based neural network training [587]. In this way, we solve a symbolic regression problem using a continuous optimization method. The crucial aspect of the method is to find the simplest solution that still explains the data well—implementing Occam's razor. The identified equations are suitable for extrapolation to unseen domains which we demonstrate using several synthetic and real-world systems. Most importantly we show how to use these learned equations to obtain a highly efficient learning algorithm for control and planning. In our first experiment, a cart-pendulum system learns to swing up from scratch after two random rollouts. Ongoing work is to improve the efficiency and applicability of the method and applying it to challenging robotic systems. The code is available at github/martius-lab repository.

More information: https://al.is.mpg.de/project/equation-learner

## Data-driven Robust and Optimal Control in Complex Systems

*Georg Martius, Jia-Jie Zhu*

We are interested in data-driven approaches to optimal and robust control, with applications robotics. To tackle various challenges arisen in complex systems, we combine optimal control – a principled way of decision-making and control, with reinforcement learning for control designs in robotic systems. This on-going project currently has the following focus areas.

Reinforcement learning (approximate dynamic programming) for robust resource-aware control. Real-world feedback control is typically resource-heavy. We are interested in controlling systems with limited resources while achieving successes in control tasks. In the heart of the approach, we use a deterministic policy gradient to design robust controllers without knowledge a priori. The resulting control design achieves model-free control with significant resource saving. The first publication appeared in the proceedings of 2018 IEEE Conference on Decision and Control (CDC), see [586].

The second focus area is control design in hybrid systems. Hybrid systems consist of continuous dynamical systems and discrete events. They can be used to model real-world systems well, such as in legged locomotion of robots. We apply optimal control (e.g. trajectory optimization, model predictive control) and probabilistic inference for planning and control under uncertainty. The goal of the project is robust control design that can plan and control under uncertain hybrid systems, such as robotic contact models.

## Stepsize Adaptation for Stochastic Optimization

*Georg Martius, Michal Rolinek*

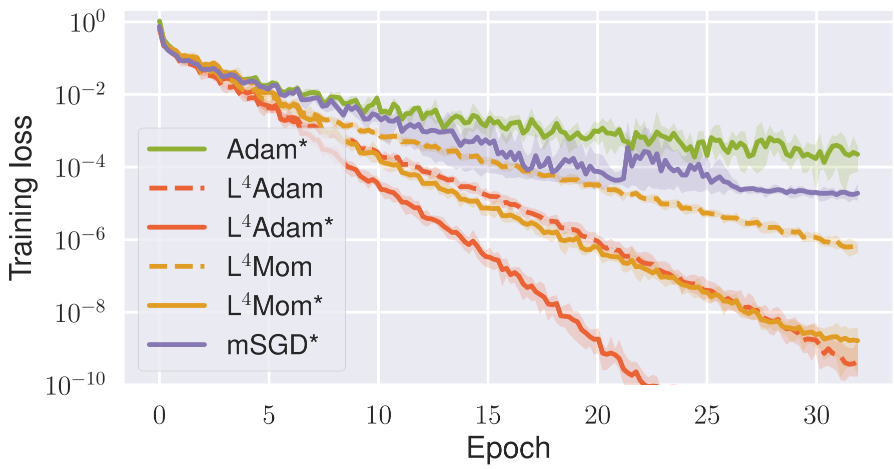


Figure 6.3: Training progress of multilayer neural networks on MNIST. Average (in log-space) training loss with respect to five restarts with a shaded area between minimum and maximum loss (after log-space smoothing).

Stochastic gradient methods are the driving force behind the recent boom of deep learning. As a result, the demand for practical efficiency as well as for theoretical understanding has never been stronger. Naturally, this has inspired a lot of research and has given rise to new and currently very popular optimization methods such as Adam, AdaGrad, or RMSProp, which serve as competitive alternatives to classical stochastic gradient descent (SGD).

However, the current situation still causes huge overhead in implementations. In order to extract the best performance, one is expected to choose the right optimizer, finely tune its hyperparameters (sometimes multiple), often also to handcraft a specific stepsize adaptation scheme, and finally combine this with a suitable regularization strategy. All of this, mostly based on intuition and experience.

We propose [585] a stepsize adaptation scheme for stochastic gradient descent. It operates directly with the loss function and rescales the gradient in order to make fixed predicted progress on the loss. We demonstrate its capabilities by conclusively improving the performance of Adam and momentum optimizers. The enhanced optimizers with default hyperparameters consistently outperform their constant stepsize counterparts, even the best ones, without a measurable increase in computational cost. The performance is validated on multiple architectures including dense nets, CNNs, ResNets, and the recurrent Differential Neural Computer on classical datasets MNIST, fashion MNIST, CIFAR10 and others. The results are presented at the Neural Information Processing Systems (NeurIPS) 2018. The code is available at github/martius-lab repository.

More information: https://al.is.mpg.de/project/stepsize-adaptation-for-stochastic-optimization

## Robust and Affordable Haptic Sensation

*Georg Martius, Huanbo Sun*

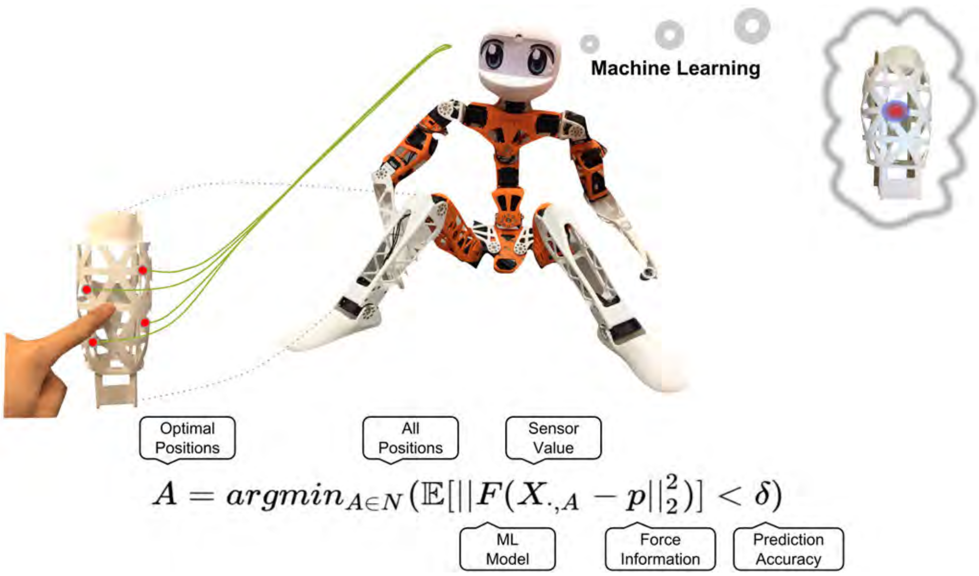


Figure 6.4: Design of a 3D haptic sensation system. Instead of applying a dense sensor array, we propose to infer interaction forces with high accuracy using machine learning from few, optimally placed, deformation sensors inside the structure.

During the rapid development of robot technologies, actuators and sensors have become increasingly compact and powerful. Nevertheless, robots are still far from matching human capabilities especially when it comes to touch sensation. For this, haptic sensors have to be robust to sustain long-lasting experiments. Besides robustness, another important aspect of robotic hardware is its price, availability, and performance. A low cost makes robotic technologies widely accessible and thus facilitates research.

In this project, we aim at providing a low-cost, robust and sufficiently precise method for inferring haptic forces on the surface of 3D robot limbs. Instead of applying dense array-shaped sensors, we opt for a small number of sensors measuring internal deformation. This offers a couple of conceptual advantages: (1) The system is robust to environmental impacts because sensors are placed internally; (2) The surface shape can be freely designed; (3) Only a few channels have to be read out which reduces both the energy consumption and the data rate. On the downside, a measurement of the sensors does not directly correspond to the impacting force. Instead, an inference mechanism is required to estimate the force. We propose a data-driven approach using machine learning algorithms to perform this inference efficiently. To require as few sensors as possible, we employ several optimization schemes to determine an optimal sensor placement.

The contributions of this project [584] are as follows: On the theory side, we propose a new way of implementing a whole surface haptic sensor and provide a method for determining the optimal number and position of sensors using finite element method. On the application side, we provide a method to assemble the strain gauges, designed a hardware system to systematically collect data and demonstrated the sensing system on a robotic limb.

More information: https://al.is.mpg.de/project/haptics

## 6.3 Awards & Honors

### 2018

Jia-Jie Zhu is awarded the **Marie Skłodowska-Curie Individual Fellowship.**

Michal Rolinek is awarded with a **CVPR Honorable Mention**

## 6.4 Research group leader: Georg Martius

### Dr. Georg Martius

| | |
|---|---|
| 2017 – present | Max Planck Research Group Leader, MPI for Intelligent Systems, Tübingen |
| 2015 – 2017 | IST-Fellow at Institute of Science and Technology, Austria |
| 2010 – 2015 | Post-doctoral position at Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany. Group for Cognition and Neurosciences, Prof. Jost |

### Education

| | |
|---|---|
| 2005 – 2009 | Graduate student at Bernstein Center for Computational Neuroscience Göttingen, Germany in the PhD program for Theoretical and Computational Neuroscience. Graduated 2009 with Dissertation: *Goal-oriented control of self-organizing behavior in autonomous robots*. |
| 2003 – 2005 | University of Leipzig, Faculty of Mathematics and Computer Science. Diploma in Computer Science, graduated with mark very good (1.3). Diploma Thesis: *Active Learning in the Sensorimotor Loop*. |
| 2002 – 2003 | University of Edinburgh, Division of Informatics, UK. Visiting student. |
| 1999 – 2002 | University of Leipzig, Faculty of Mathematics and Computer Science, Pre-diploma. |

### Awards

| | |
|---|---|
| 2016 | Max Plank Research Group Leader grant. |
| 2016 | Distinghuised oral presentation award on IEEE ICDL 2016, France |
| 2014 | IST-Fellowship from Institute of Science and Technology Austria. |
| 2014 | Best paper award on SAB 2014, Castellon, Spain. |
| 2010 | Best paper award on SAB 2010, Paris, France. |
| 2009 | Dissertation *Summa cum laude* (highest praise), Göttingen, Germany |
| 2005 | Stipend of Graduate school on *Analysis, Geometry and their Connection to Natural Sciences*, Leipzig, Germany |

# 7 AUTONOMOUS VISION

## 7.1 Research Overview

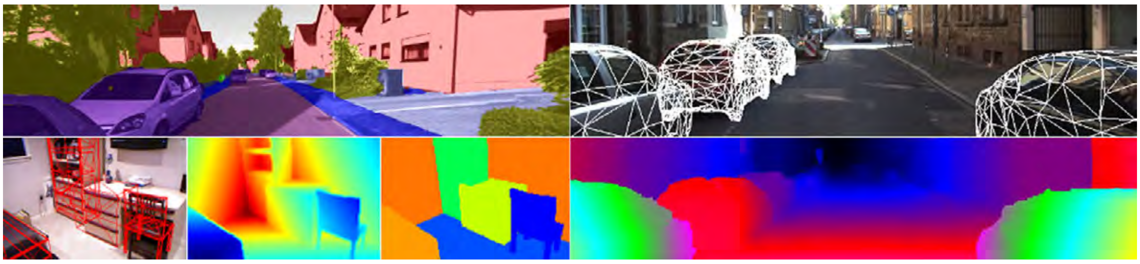

The **Autonomous Vision Group** is focused on 3D scene understanding, reconstruction, motion estimation, generative modeling and sensori-motor control in the context of autonomous systems. Our goal is to make artificial intelligent systems such as self-driving cars or household robots more autonomous, efficient, robust and safe. By making progress towards these goals, we also strive for uncovering the fundamental concepts underlying visual perception and autonomous navigation. Our research is guided by the following principles.

**(1) Representations:** Our world is inherently three-dimensional as all physical processes (including image formation) occur in 3D and not in the 2D image plane. Thus, we strive for inferring compact 3D representations of our world from 2D or 3D measurements. To this end, we develop novel scalable spatio-temporal representations, learning frameworks, reconstruction as well as motion estimation algorithms.

**(2) Prior Knowledge:** Visual perception is a highly ill-posed task with many explanations for a single observation. We therefore investigate how prior knowledge (e.g., about the shape of objects, image formation or driving laws) can be incorporated into visual perception and autonomous navigation to make both tasks more robust. We also develop probabilistic representations which capture uncertainty in the output.

**(3) Learning from Little Data:** High-capacity models such as deep neural networks require large amounts of annotated training data which limits scalability. To address this problem, we develop novel techniques for learning from little annotated data, including self-supervised models for geometry and motion estimation, methods for transferring labels across domains, techniques which incorporate a-priori knowledge into the structure of neural networks and approaches that explicitly model invariances in the data.

**(4) Generative Models and Simulation:** Generative models are at the core of understanding the fundamental processes underlying vision, building robust models and generating large amounts of training data for discriminative models. We investigate generative models (such as GANs and VAEs) from a theoretical perspective and apply them to tasks in the context of data generation for autonomous driving and beyond.

**(5) Empirical Risk Minimization:** Many state-of-the-art computer vision models (e.g., 3D reconstruction) or sensori-motor control systems (e.g., self-driving cars) are trained using auxiliary loss functions instead of the actual task loss due to difficulties in representation or computational limitations. We work on end-to-end trainable models for these tasks which exploit task information in a data-efficient manner.

**(6) Datasets and Evaluation:** We strongly believe that research is a collective effort that is only possible by sharing research results, code and data. We are therefore committed to publishing our results and code. Moreover, we construct novel datasets like KITTI or ETH3D to foster progress in the field and across research areas.

## 7.2 Selected Research Projects



### 3D Reconstruction

*Simon Donne, Yiyi Liao, Despoina Paschalidou, Gernot Riegler, David Stutz, Osman Ulusoy, Andreas Geiger*

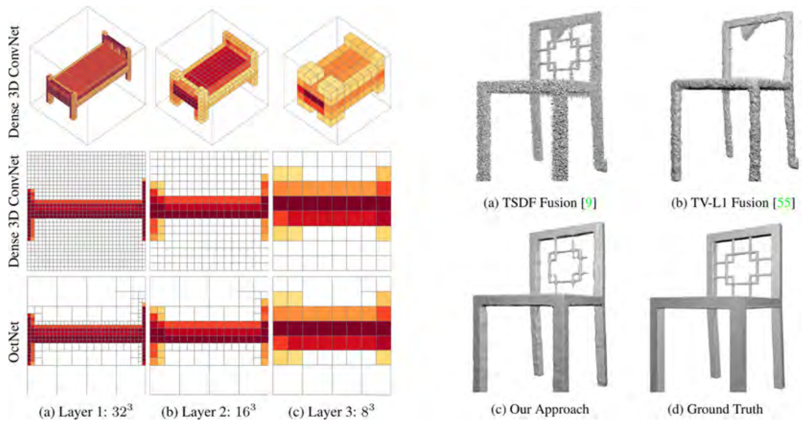


Figure 7.1: 3D deep learning suffers from a cubic increase in memory, hence limiting current approaches. To address this problem, we have developed a data-adaptive framework for discriminative and generative convolutional neural networks which allows 3D recognition, segmentation and reconstruction at resolutions up to $256^3$ voxels.

We live in a three-dimensional world, thus understanding our world in 3D is important. A first step towards this 3D understanding is to accurately reconstruct our environment based on 2D image measurements or sparse 3D point clouds. We have developed novel representations, algorithms and benchmarks for this task.

As 3D reconstruction is an ill-posed inverse problem, solving it requires strong a-priori knowledge or models that incorporate 3D shape information from training data. Towards this goal, we have developed OctNet [389], one of the first frameworks for scaling notoriously memory-hungry 3D deep learning techniques to high-resolution input and output spaces, allowing for detailed volumetric 3D reconstructions at $256^3$ voxels and beyond [604] as illustrated in Fig. 7.1.

While generic convolutional neural networks are agnostic to the image formation process, classical approaches to 3D reconstruction leverage explicit knowledge about 3D geometry and light propagation [382, 395]. With RayNet [599], we have presented the first approach that integrates this knowledge into a deep 3D reconstruction model by unrolling a high-order CRF as layers in a convolutional neural network. In similar spirit, we have presented a technique to reconstruct geometry and semantics jointly using a deep variational reconstruction approach [592].

To go beyond voxel representations, we have presented Deep Marching Cubes [598], a deep neural network that outputs mesh representations of arbitrary topology. In contrast to prior work, our model can be trained end-to-end to predict meshes without resorting to auxiliary representations (e.g., TSDFs) and loss functions.

Point clouds are another popular output representation. However, they lack connectivity and topology. To tackle this problem, we have presented novel approaches to depth map completion [605] and 3D shape completion [597].

Access to 3D datasets and benchmarks is crucial for driving progress in the field. Beyond our popular KITTI dataset, we have therefore proposed novel large-scale datasets for single image depth prediction and depth map completion [605] as well as two-view and multi-view 3D reconstruction in indoor and outdoor environments [610]. Evaluation servers with held-out ground truth provide the basis for a fair comparison.

More information: https://avg.is.mpg.de/field/3d

## Motion Estimation and Scene Understanding

*Aseem Behl, Benjamin Coors, Fatma Güney, Joel Janai, Moritz Menze, Axel Sauer, Andreas Geiger*


Figure 7.2: Understanding scenes in motions is crucial for intelligent systems. We have developed novel self-supervised approaches for learning optical flow from multiple images without ground truth (left). We have further developed a novel model for learning driving affordances from video sequences using limited supervision (right).

Intelligent systems as well as many computer vision tasks benefit from an accurate understanding of object motion as well as the interplay between scene elements. However, motion is even less constrained than geometric structure and obtaining ground truth is difficult. Furthermore, besides pure 2D image motion, systems operating in the 3D world require access to 3D motion information. Our research addresses all of these aspects and combines them for sensori-motor control tasks such as autonomous driving.

Optical flow is the problem of estimating 2D motion in the image plane [391]. While deep learning has led to significant progress, large amounts of training data are required. We have addressed this problem by developing space-time tracking techniques for generating training data using high-speed cameras [380]. Moreover, we have proposed a new model for self-supervised flow estimation from multiple frames, explicitly accounting for occlusions [359] (Fig. 7.2, left).

To recover 3D motion [601], we have developed a state-of-the-art 3D scene flow estimation technique which exploits recognition (bounding boxes, instance segmentation and object coordinates) to support the challenging matching task [606]. Furthermore, we have also investigated how optical flow can be extracted from spherical imagery [593] and empirically evaluated the effectiveness of various motion estimation algorithms for down-stream tasks such as action recognition from video sequences [354].

Intelligent systems not only require the relative motion of objects around them, but typically also a precise global location [609] with respect to a map, i.e., for planning or navigation tasks. With LOST [351], we have demonstrated that localization solely based on map information is feasible. In [596], we further showed that semantic and geometric information can significantly improve visual localization, allowing for localizing wrt. the opposite driving direction or in the presence of strong environmental changes.

Moreover, we have developed a novel model for sensori-motor control [590] which learns driving affordances from video sequences using only very limited supervision (Fig. 7.2, right). To foster new research on 3D scenes in motion, we have created a new dataset for 3D urban scene understanding, annotated at the object level [394].

More information: https://avg.is.mpg.de/field/stereo-flow

## Generative Models and Image Synthesis

*Hassan AlHaija, Lars Mescheder, Sebastian Nowozin, Michael Oechsle, Carsten Rother, Andreas Geiger*

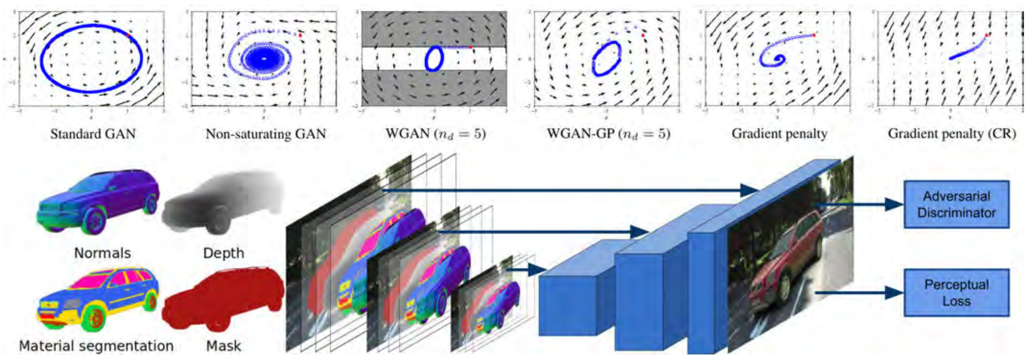


Figure 7.3: Generative models are important for understanding the fundamental processes underlying vision, building robust models and training high-capacity deep neural networks. We have worked on fundamental theoretical questions including optimization (top) and applied generative models to complex image synthesis tasks (bottom).

While deep learning undeniably achieves impressive results for numerous different computer vision tasks, the theoretical foundations behind the success of these methods are often not well understood. One particular focus of our group is on deep generative models which are important for understanding the principles of vision, for building robust models and for training high-capacity deep neural networks.

Variational Autoencoders (VAEs) represent a particular class of latent variable models that can be used to learn complex probability distributions from training data. However, the quality of the results crucially relies on the expressiveness of the inference model. In [611], we introduced Adversarial Variational Bayes (AVB), a technique for training Variational Autoencoders with arbitrarily expressive inference models by introducing an auxiliary discriminative network that allows to rephrase the maximum-likelihood-problem as a two-player game.

We have also analyzed the optimization problems underlying recent Generative Adversarial Networks (GANs). In [602], we analyzed the gradient vector field associated with the GAN training objective using the formalism of smooth two-player games. Using our findings, we developed a new algorithm with better convergence properties. In [594], we showed that the requirement of absolute continuity is necessary for convergence of unregularized GAN training using a simple counterexample. We further analyzed which GAN training methods converge (Fig. 7.3 top) and proved local convergence for GAN training with simplified gradient penalties.

We leveraged our insights on training generative models for tackling image synthesis tasks, with the ultimate goal of generating large amounts of training data at limited cost. We first demonstrated that data augmentation is a viable alternative to annotating real images or synthesizing entire scenes from scratch [589, 607]. Towards this goal, we augmented real images from the KITTI dataset with photo-realistically rendered 3D car models, yielding significant performance improvements when used for training deep neural networks on recognition tasks.

In [608], we demonstrated that the rendering process itself can be learned. In particular, we used the output of OpenGL (depth, normals, materials) as input to a neural network which generates a rendering of the respective real-world object (Fig. 7.3 bottom). Importantly, our generative model handles ambiguity in the output (e.g., cars might have different color) and learns to add realistic transparency, reflection and shadowing effects to the augmented objects.

More information: https://avg.is.mpg.de/field/theoretics-of-deep-learning

## 7.3 Research group leader: Andreas Geiger

### Prof. Dr. Andreas Geiger

Andreas Geiger is a full professor at the University of Tübingen and a group leader at the Max Planck Institute for Intelligent Systems. Prior to this, he was a visiting professor at ETH Zürich and a research scientist in the Perceiving Systems department of Dr. Michael Black at the MPI-IS. He studied at KIT, EPFL and MIT and received his PhD degree in 2013 from the Karlsruhe Institute of Technology. His research interests are at the intersection of 3D reconstruction, 3D motion estimation and visual scene understanding with a particular focus on integrating rich prior knowledge and deep learning for improving perception in intelligent systems. In 2012, he has published the KITTI vision benchmark suite which has become one of the most influential testbeds for evaluating stereo, optical flow, scene flow, detection, tracking, motion estimation and segmentation algorithms.

His work on stereo reconstruction and optical flow estimation has been ranked amongst the top-performing methods in several international competitions. His work has been recognized with several prizes, including the IEEE PAMI Young Investigator Award, the Heinz Maier Leibnitz Prize of the German Science Foundation DFG, the German Pattern Recognition Award, the Ernst-Schoemperlen Award and the KIT Doctoral Award. In 2013, he received the CVPR best paper runner up award for his work on probabilistic visual self-localization. He also received the best paper award at GCPR 2015 and 3DV 2015 as well as the best student paper award at 3DV 2017. He is an associate member of the Max Planck ETH Center for Learning Systems and the International Max Planck Research School for Intelligent Systems, and serves as an area chair and associate editor for several computer vision conferences and journals (CVPR, ICCV, ECCV, PAMI, IJCV).

### Appointments

| | |
|---|---|
| 03/2018 – present | Professor (W3), University of Tübingen |
| 06/2016 – present | Independent Max-Planck Research Group Leader, Max Planck Institute for Intelligent Systems |
| 06/2016 – 02/2018 | Visiting Professor, ETH Zürich |
| 06/2013 – 05/2016 | Research Scientist, Max Planck Institute for Intelligent Systems |
| 06/2013 – 05/2016 | Research and Teaching Assistant, Karlsruhe Institute of Technology |

### Awards & Honors (Selected)

| | |
|---|---|
| 2018 | IEEE PAMI Young Researcher Award |
| 2017 | Best Student Paper Award, International Conf. on 3D Vision (3DV) |
| 2017 | German Pattern Recognition Prize, German Conf. on Pattern Recognition (GCPR) |
| 2017 | Heinz Maier-Leibnitz Prize, Deutsche Forschungsgemeinschaft (DFG) |
| 2015 | Best Paper Award, International Conf. on 3D Vision (3DV) |
| 2015 | Best Paper Award, German Conf. on Pattern Recognition (GCPR) |
| 2013 | Best Paper Runner Up Award, Conf. on Comp. Vision and Pattern Recog. |

### Selected Keynote, Conference, Workshop, and Public Talks (2016-2018)

| | |
|---|---|
| 2018, Munich | 3 Invited Talks, *European Conference on Computer Vision (ECCV)* |
| 2018, Renningen | Invited Talk, *Bosch Center for Artificial Intelligence (BCAI)* |
| 2018, Salt Lake City | 4 Invited Talks, *Comp. Vision and Pattern Recog. (CVPR)* |
| 2018, Beijing | Invited Talk, *Baidu ApolloScape Workshop* |
| 2017, Venice | 2 Invited Talks, *International Conference on Computer Vision (ICCV)* |
| 2017, Basel | Award Lecture, *German Conf. on Pattern Recognition (GCPR)* |
| 2017, London | Invited Tutorial Lecture, *British Machine Vision Conference (BMVC)* |
| 2017, Honolulu | Invited Talk, *Comp. Vision and Pattern Recog. (CVPR)* |
| 2016, Singapore | Invited Talk, *National University of Singapore (NUS)* |
| 2016, Amsterdam | Invited Talk, *European Conference on Computer Vision (ECCV)* |

# 8 DYNAMIC LOCOMOTION

## 8.1 Research Overview

We aim to understand animal legged locomotion by researching principles that enable robust, yet agile and energy-efficient locomotion in non-structured, natural environments. We are extending legged locomotion research in biomechanics and neurocontrol with novel physical and experimental systems: legged robots and their simulated models.

At the Dynamic Locomotion Group we are focusing on the following scientific topics: design, experimentation and characterization of bipedal and quadrupedal robotic systems, bipedal locomotion and trunk stabilization during walking and running, developing novel frameworks to evaluate integrated mechanical and controller designs for legged locomotion, and bioinspired sensing and actuation. To evaluate our concepts, we cooperate with scientists at the Royal Veterinary College at the University of London, and at Universities in Stuttgart, Tübingen, and Munich. The Dynamic Locomotion Group was founded as an independent Max Planck Research Group in October 2016 by Alexander Spröwitz, Dr. Sc.

## 8.2 Selected Research Projects



### Quadruped Research Platform To Evaluate Bioinspired Locomotion Mechanisms

*Alexander Spröwitz, Felix Ruppert, An Mo, Alborz Aghamaleki Sarvestani, Steve Heim, Benedikt Gyoerfi*

Animals outperform robotic walking machines albeit their limitations in muscle power density, neuromuscular conductivity velocity [More et al. 2010, 2018], and tissue properties, by running faster, more robustly, and with better energy efficiency. On the engineering side, technological progress recently led to power dense actuators, gigahertz fast computation, accurate and fast sensing, and stiff but light-weight materials. Eventually, engineered components individually have the potential to outperform their biological counterparts. But we are starting to learn that legged animals are not just the sum of their components. As an example, no current legged robotic system reaches agility and low cost of transport of an animal of equivalent size.

We expect that both a smart and adapted biomechanical integration of design and control causes the superior biological performance and robustness, other than individual components with better properties.

It is feasible that mechanical mechanisms evolved and adapted specifically for legged locomotion, leveraging and utilizing resources like springy and damping muscles, and effective moment arms. Similar adaptations are expected on for neuromuscular control. These adaptations have the potential to expand working range, power amplification, and efficiency of a legged system.

In our bottom-up approach **we often utilize biomechanical and neuromuscular mechanisms hypothesized in Biology**. Specific examples are spring loaded, pantographic leg designs[1] and neuromuscular controllers regulating gait pattern generation and adaptation[2]. We transfer these mechanisms into robotic hardware and control. We then can test and characterize individual mechanisms, before merging and integrating them into a complex legged system.

Specifically, we investigate mechanisms in leg design and neurocontrol that potentially **reduce the control effort in animals, bioinspired robots[3], exoskeletons, and prostheses**. One focus is on leg mechanics. Leg spring stiffness, segmentation ratios, and tendon parameters influence the performance the biological and the bioinspired robotic leg. We eventually collect and analyze biomechanical and robotic data, with biomechanical tools like force plates and motion capture systems.

By porting mechanisms from animals into robots[4,5] we are able to systematically investigate design parameters [612, 613, 619, 621]. Potentially, our approach allows us finding real functional couplings, instead of correlations. One goal is then to create a design language, aiding integration of sensing and actuation of legged structures. We want to understand why animals developed certain anatomical structures, and how these influence dynamic locomotion.

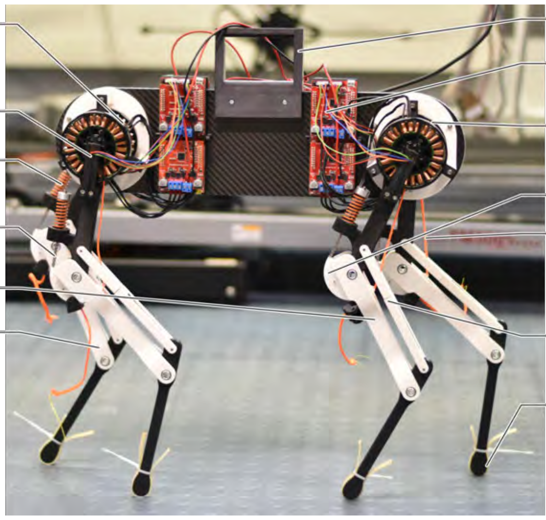


Figure 8.1: Our custom developed quadruped robot features pantograph, compliant legs. A tendon mechanism actively shortens leg length through powerful, brushless motors. The tendons also serve as an actuator disengagement, such as at step-down perturbations in unstructured environments.

More information: https://dlg.is.mpg.de/project/quadruped-research-platform-for-evaluation-of-bioinspired-locomotion-mechanisms

[1] A. Spröwitz, A. Tuleu, M. Vespignani, M. Ajallooeian, E. Badri, et al. Towards Dynamic Trot Gait Locomotion: Design, Control, and Experiments with Cheetah-cub, a Compliant Quadruped Robot. *The International Journal of Robotics Research* **32** (8): 932–950, 2013.

[2] A. Spröwitz, M. Ajallooeian, A. Tuleu, A. J. Ijspeert. Kinematic primitives for walking and trotting gaits of a quadruped robot with compliant legs. *Frontiers in Computational Neuroscience* **8** (27): 1–13, 2014.

[3] F. Moro, A. Spröwitz, A. Tuleu, M. Vespignani, N. G. Tsagakiris, et al. Horse-Like Walking, Trotting, and Galloping derived from Kinematic Motion Primitives (kMPs) and their Application to Walk/Trot Transitions in a Compliant Quadruped Robot. *Biological Cybernetics* **107** (3): 309–320, 2013.

[4] P. Eckert, A. Spröwitz, H. Witte, Ijspeert, Auke Jan. Comparing the effect of different spine and leg designs for a small bounding quadruped robot. In *Proceedings of ICRA*, pages 3128–3133, 2015.

[5] M. Khoramshahi, A. Spröwitz, A. Tuleu, M. N. Ahmadabadi, A. J. Ijspeert. Benefits of an active spine supported bounding locomotion with a small compliant quadruped robot. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 3329–3334, 2013.

## Shaping the Reward Landscape without Shaping the Reward

*Alexander Spröwitz, Steve Heim, Felix Ruppert, Alborz Aghamaleki Sarvestani*

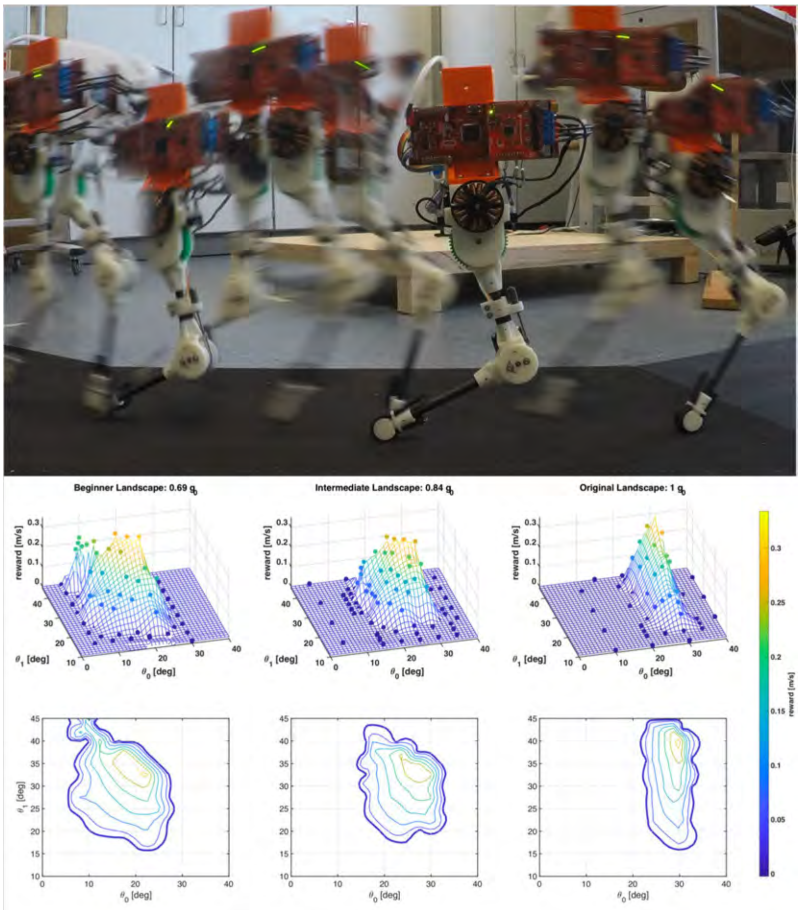


Figure 8.2: We can empirically map the true reward landscape (here: robot velocity, color bar) of our monoped hopping robot for two control parameters. We then observe the change of the landscape due to the modified effective weight of the robotic system.

In reinforcement learning, tasks that are difficult to learn are often made more amenable by shaping the reward or cost landscape. This is typically done by adjusting the reward signal $R$ in the Markov Decision Process, composed of $(S, A, P, R, \gamma)$, where $S$ is the state-space, $A$ is the action-space, $P$ is the probability transition matrix, i.e. the system dynamics, $R$ is the reward signal and $\gamma$ is the discount factor.

We formalize and show the effectiveness of changing other parts of the MDP, in particular the dynamics $P$ and the initial state conditions $S_0$. **The concept of training wheels**, first formalized by [Randløv, 2001], uses instead a **temporary adjustment of the system dynamics**. As a practical example, we show that mechanically adjusting the weight of a hopping robot influences its reward landscape. This results in learning that can occur more reliably and with less initial tuning [616].

In a simple simulation, we show that state initialization can play an important role. Indeed, we show that **state initializations that are normally doomed to fail can be exploited to learn effective policies**. Utilizing information from these normally discarded initializations **can result in more reliable and quicker learning** [615].

More information: https://dlg.is.mpg.de/project/shaping-the-reward-landscape-without-shaping-the-reward

## Trunk Influence in Bipedal Locomotion

*Alexander Spröwitz, Oezge Drama, Alborz Aghamaleki Sarvestani*

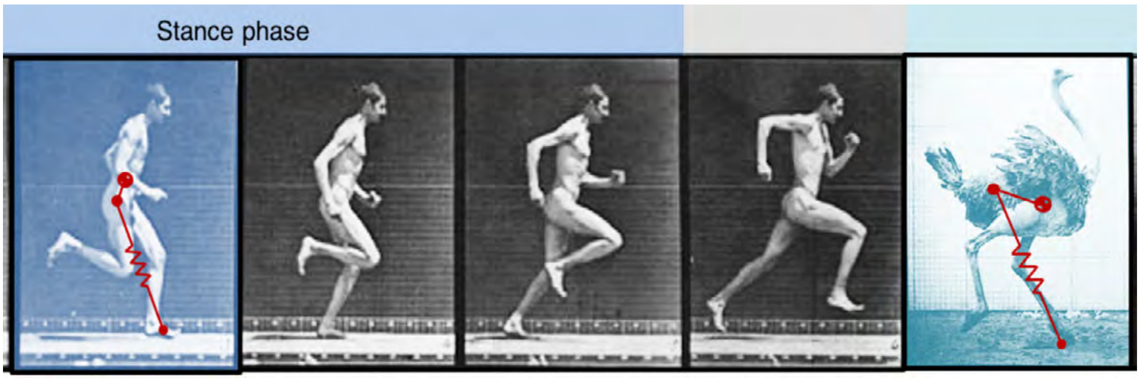


Figure 8.3: Bipedalism evolved multiple times. It is currently widely expressed in running, flightless birds, and also in humans. Besides the difference in leg and foot design of both species, also orientation, position and moment of inertia of the trunk vary: a human trunk is balanced almost vertically, while a bird carries the trunk in front of its hip joint.

We investigate bipedal locomotion, and its consequences for function and control in different species: the upright, humanoid configuration, and the flight-less, leg-locomoting bird-like configuration. This research is tightly coupled to postural stability in locomotion tasks like walking and running.

Bipedal locomotion is a multi-objective control task. For postural stability the trunk must be carried and balanced by the actuated hip joint, maintaining trunk posture, orientation, and range of motion. Equally, hip torques are provided by the same hip joint, but in opposite direction. Due to leg segmentation, hip torques and femur rotation further influence leg forces and leg length kinematics[6]. These tasks are directly and physically coupled by and in the hip joint. The bipedal system becomes more complex through the interplay of two legs acting in a phase-shifted manner, while the trunk is balanced, within the alternating sequence of stance and flight phase.

Albeit this complexity, bipedalism evolved in lineages ranging from reptiles, avians, theropods to primates. Diverse morphologies exhibit a large and viable parameter range for limb segment lengths, trunk masses, and trunk orientations. Each yields different locomotion characteristics.

Amongst bipeds, birds demonstrate exceptional agility [Daley et al. 2010], locomotion efficiency and terrain traversability despite having limitations on actuation and sensory delays [Alexander, 1992, More et al. 2010, 2018]. We aim to understand the necessary components to generate versatile locomotion and we analyze the requirements they impose on control policies.

The animal trunk effects both locomotor kinetics and kinematics. It accounts for almost 50% of total body mass in humans, and 70-80% in birds. Even higher ratios are possible in bipedal robots[7]. Stabilizing a heavy trunk with small support base is a major challenge in bipeds. Hence, trunk orientation is integral for motion generation. We consider two distinct animal trunk postures. Orthograde (upright) posture such as humans, and pronograde (horizontal) posture such as in birds, or extinct theropods.

In order to isolate the functional effect of the trunk, we utilize a spring loaded inverted pendulum model, extended with a rigid trunk similar to T-SLIP model [Maus et al, 2008]. We focus on the effect of trunk orientations on the torque and ground reaction profiles, for different control strategies.

More information: https://dlg.is.mpg.de/project/trunk-influence-in-bipedal-locomotion

[6] D. Renjewski, A. Spröwitz, A. Peekema, M. Jones, J. Hurst. Exciting Engineered Passive Dynamics in a Bipedal Robot. *IEEE Transactions on Robotics and Automation* **31** (5): 1244–1251, 2015.

[7] C. Hubicki, J. Grimes, M. Jones, D. Renjewski, A. Spröwitz, et al. ATRIAS: Design and validation of a tether-free 3D-capable spring-mass bipedal robot. *The International Journal of Robotics Research* **35** (12): 1497–1521, 2016.

## Viability, Learning and Robust Natural Dynamics

*Alexander Spröwitz, Steve Heim*

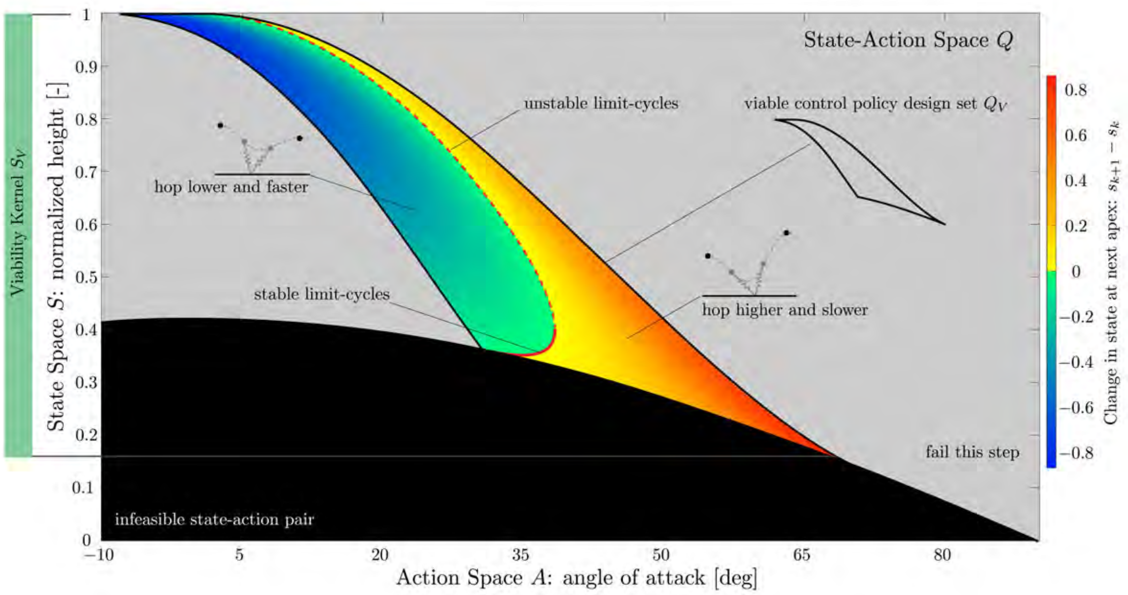


Figure 8.4: We can find the structure of the natural dynamics of a system by moving into the state-action space. Shown is the state on the y-axis and action on the x-axis. The colored region shows the space available for designing a viable control policy.

**To quantify the effect of natural dynamics** on the **policy learning problem**, we formalize our work using the **concept of viability**. Unlike traditional control, viability analysis starts by defining a set of failures instead of a target. The viable set, also known as the viability kernel, is then the set of all states which can remain inside the set, and therefore avoid failure for all time.

While finding the viable set does not shed any information on convergence or optimality, **it requires no definition of an objective, the reward function** or the **policy parameterization**. This **allows us to begin quantifying robustness to failure for a system design prior to designing or learning the actual control policy**.

We currently observe that systems that are more robust to noise in action space are also more amenable to learning control policies, and allow more flexibility. This allows us to compare different designs of the mechanical system as well as low-level controllers such as reflexes[8].

More information: https://dlg.is.mpg.de/project/viability-learning-and-robust-natural-dynamics

[8] S. Heim, A. Spröwitz. Beyond Basins of Attraction: Evaluating Robustness of Natural Dynamics. *arXiv:1806.08081 [cs] (submitted December 2018).*

## 8.3 Research group leader: Alexander Spröwitz

### Alexander Spröwitz, Dr. sc.

Alexander Spröwitz leads the Dynamic Locomotion group at the Max Planck Institute for Intelligent Systems in Stuttgart. After his PhD, he was a postdoctoral researcher at the Biorobotics Laboratory at EPFL, Lausanne between December 2010 and December 2012, and worked with large bipedal robots at the Dynamics Robotics Laboratory at Oregon State University, Corvallis until June 2013. He researched biomechanics and functional anatomy of ground running birds at the Structure and Motion Lab at the Royal Veterinary College, University of London until August 2014, and robotic legged mechanisms at the Physical Intelligent Department at MPI IS between October 2014 and June 2016. He studied Mechatronics at the Technical University Ilmenau (Diploma, 2006), and received his Doctorate in Science from the Ecole Polytechnique Fédérale de Lausanne (Switzerland) in 2010. Dr. Spröwitz holds a W2 independent Max Planck Research Group Leader position since July 2016. His current research focus is on functional morphology and locomotion control of animals and robots, with focus on bipedal and quadrupedal legged systems.

### Employment

| | |
|---|---|
| 2016 – present | Independent Max Planck Research Group Leader, MPI-IS, Germany |
| 2014 – 2016 | Postdoctoral Researcher, Physical Intelligence Department (Metin Sitti), MPI-IS, Germany |
| 2013 – 2014 | Postdoctoral Researcher, Dynamics Robotics Laboratory (Jonathan Hurst), Oregon State University, USA and at Structure & Motion Lab (Monica Daley), Royal Veterinary College, UK |
| 2010 – 2012 | Postdoctoral Researcher, Biorobotics Laboratory (Auke Ijspeert), Ecole Polytechnique Fédérale de Lausanne, Switzerland |

### Education

| | |
|---|---|
| 2010 | Doctorate in Science, Ecole Polytechnique Fédérale de Lausanne, Switzerland. Supervisor: Prof. A. J. Ijspeert |
| 2006 | Diploma in Mechatronics, Technische Universität Ilmenau, Germany |

### Publications and Talks

| | |
|---|---|
| Publications | 11 refereed journal, 14 refereed DOI conference<br>h-index: 20, i-10 index: 32, 1319 citations (src: Google Scholar, 12/18) |
| Talks | 3 invited talks, 20 regular talks in international conferences<br>8 invited talks in universities and research institutes |

### Supervision Activities

| | |
|---|---|
| 2016 – present | Supervision of 6 PhD candidate researchers (1 since 2016, 3 since 2017, 2 since 2018) |
| 2006 – present | (Co)Supervision of 9 MSc and 11 BSc Theses |

### Professional Activities (selected)

| | |
|---|---|
| 2018 – present | Executive board, International Max Planck Research School for Intelligent Systems |
| 2017 – present | Faculty, International Max Planck Research School for Intelligent Systems |
| 2016 – present | Associated Member, Max Planck ETH Center for Learning Systems |
| 2015 – 2016 | Associated Fellow, Max Planck ETH Center for Learning Systems |
| 2013 – present | (Co)Organization and scientific committee of 4 international scientific meetings/workshops |

### Public Outreach (selected)

| | |
|---|---|
| 2018 – present | Open-sourcing blueprints of Oncilla robot: c4science online repository |
| 2013 | Cheetah-cub robot presented at London Science Museum exhibition: Robot Safari EU |
| 2010 – present | Various national and international media appearances of projects Roombots, Cheetah-cub and related: BBC, DPA, Stuttgarter Zeitung, 1.3 Million views on YouTube channel ‘epflnews’ etc. |

### Awards

| | |
|---|---|
| 2016 | Max Plank Research Group Leader grant |

# 9 EMBODIED VISION

## 9.1 Research Overview

Intelligent agents such as robots require the ability to learn and adapt within their environment. Our group investigates novel methods for learning to understand dynamic 3D scenes and their functioning, and use this knowledge to perform complex tasks such as autonomous navigation and object manipulation. Traditional approaches often integrate perception and control components engineered for specific tasks and scenarios. In contrast, we aim at systems that learn to act and perceive from raw sensor measurements such as images or tactile information and action experience acquired in their environment. We investigate computer vision methods and end-to-end trainable architectures for learning task-relevant representations that allow agents to plan their actions.

## 9.2 Selected Research Projects

## Model Learning for Scene Understanding, Control and Planning in Dynamic Environments

*Jörg Stückler, Jan Achterhold, Michael Strecke, Deyao Zhu*

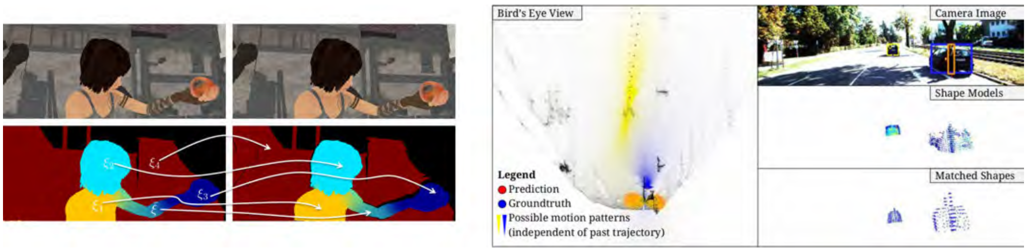


Figure 9.1: Left: RGB-D motion segmentation and scene flow estimation[1]. Right: Prediction model learned for objects in dynamic street scenes[2].

Complex everyday environments pose a challenge to the implementation of autonomous robots that purposefully act in these environments. A classical approach for developing autonomous robots is to hand-craft object and environment models that are used for control and planning. These models, however, are typically limited by the implicit assumptions the engineers make about the properties and functioning of the environment, such that the designed system might not generalize to unseen tasks and situations.

In this project, we investigate the learning of causal models of dynamic environments for scene understanding, control and planning. We aim at learning approaches that acquire such models from raw sensory observations such as images or tactile information and action experience in the environment. For object manipulation, we start our investigation with physical scene understanding in simple manipulation scenes that involve multiple rigid and non-rigid objects. Robots not only need the ability to discover and perceive the objects in a scene that can be manipulated (e.g. through object detection[3] or motion segmentation[4]), but also need models of the effect of actions on these objects. By interacting with the scene, the robot shall learn "intuitive physics" on a Newtonian scale which enables the robot to predict the future evolution of the scene conditioned on actions. For robot navigation, we research approaches for learning appropriate priors for simultaneous localization and mapping such as predictive models of dynamic objects[5].

We will also develop suitable control and planning methods that use the learned models. Differentiable control and planning methods allow for learning and adapting scene models end-to-end during the execution and exploration of tasks. Tasks have to be described using a cost function or a goal configuration for which we also research the learning of reward models and approaches for describing goal states based on our scene models.

More information: https://ev.is.mpg.de/project/model-learning-for-dynamic-scene-understanding

[1]M. Jaimez, M. Souiai, J. Stueckler, J. Gonzalez-Jimenez, D. Cremers. Motion Cooperation: Smooth Piece-Wise Rigid Scene Flow from RGB-D Images. In *Proc. of the Int. Conference on 3D Vision (3DV)*, 2015

[2]D. Klostermann, A. Osep, J. Stueckler, B. Leibe. Unsupervised Learning of Shape-Motion Patterns for Objects in Urban Street Scenes. In *British Machine Vision Conference (BMVC)*, 2016

[3]M. McElhone, J. Stueckler, S. Behnke. Joint detection and pose tracking of multi-resolution surfel models in RGB-D. In *Proc. of the European Conference on Mobile Robots (ECMR)*, pages 131–137, 2013.

[4]J. Stueckler, S. Behnke. Efficient Dense Rigid-Body Motion Segmentation and Estimation in RGB-D Video. *International Journal of Computer Vision (IJCV)* **113** (3): 233–245, 2015.

[5]D. Klostermann, A. Osep, J. Stueckler, B. Leibe. Unsupervised Learning of Shape-Motion Patterns for Objects in Urban Street Scenes. In *British Machine Vision Conference (BMVC)*, 2016.

## Vision-based Navigation for Autonomous Robots

*Jörg Stückler, Michael Strecke, Paul Sanzenbacher*

Figure 9.2: Left: Simultaneous localization and mapping with stereo vision and inertial sensing in an urban environment[6]. Right: Vision-based autonomous navigation of a multicopter using semi-dense monocular SLAM[7].

We develop navigation approaches for autonomous robots that are primarily based on computer vision. Vision-based approaches to localization and mapping are a popular choice for systems that are constrained in sensor and computing resources such as multicopters or handheld devices. A popular approach is to complement vision with inertial sensing to achieve more robust and accurate camera motion tracking.

We investigate visual-inertial simultaneous localization and mapping approaches that enable robots to navigate in challenging dynamic environments. In previous work[8] we propose a direct method for visual-inertial SLAM that achieves state-of-the-art results on the EuRoC MAV benchmark. This dataset has been recorded with a microcopter that flies through indoor environments including various speeds and lighting conditions. Recently, we also proposed a new benchmark for visual-inertial odometry with monocular and stereo cameras[9], that includes an accurate calibration of sensor intrinsics and extrinsics, ground-truth trajectories recorded with a motion capture system, and challenging new indoor and outdoor sequences.

In previous work[10] we also demonstrate an approach which applies semi-dense direct SLAM methods for exploration and navigation by multicopters. A main challenge here is that the monocular SLAM method cannot observe 3D information in textureless areas. Instead the robot needs to take specific maneuvers to uncover the unoccupied free-space in the environment by the semi-dense measurements of well textured surfaces and contours.

In ongoing work, we investigate vision-based navigation approaches that not only localize, map and plan with respect to the static part of the environment, but also incorporate dynamic objects. We aim at novel algorithms that combine deep learning-based priors and object segmentation with classical approaches to visual SLAM to perceive dynamic objects which can subsequently be used for motion control and planning.

More information: https://ev.is.mpg.de/project/vision-based-navigation

[6] V. Usenko, J. Engel, J. Stueckler, D. Cremers. Direct Visual-Inertial Odometry with Stereo Cameras. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016

[7] L. von Stumberg, V. Usenko, J. Engel, J. Stueckler, D. Cremers. From Monocular SLAM to Autonomous Drone Exploration. In *European Conference on Mobile Robots (ECMR)*, 2017

[8] V. Usenko, J. Engel, J. Stueckler, D. Cremers. Direct Visual-Inertial Odometry with Stereo Cameras. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.

[9] D. Schubert, T. Goll, N. Demmel, V. Usenko, J. Stueckler, et al. The TUM VI Benchmark for Evaluating Visual-Inertial Odometry. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2018.

[10] L. von Stumberg, V. Usenko, J. Engel, J. Stueckler, D. Cremers. From Monocular SLAM to Autonomous Drone Exploration. In *European Conference on Mobile Robots (ECMR)*, 2017.

## Visual Simultaneous Localization and Mapping

*Jörg Stückler, Michael Strecke, Pia Ana Cuk, Paul Sanzenbacher*

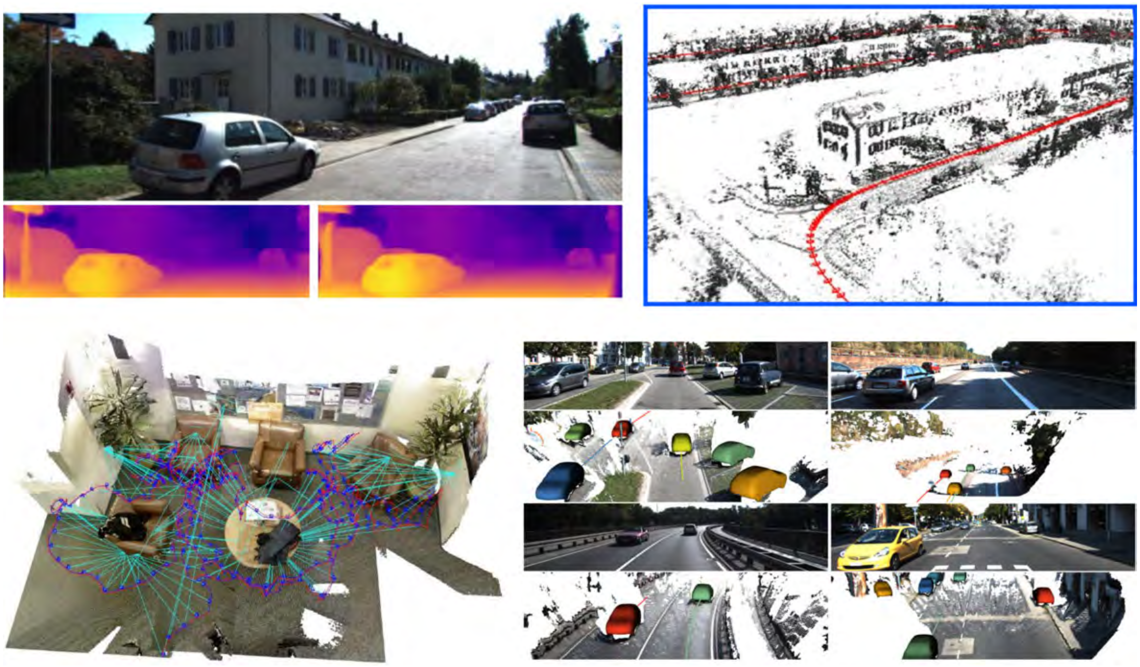

Figure 9.3: Top: Monocular visual odometry using deep-learning based monocular depth estimation[11]. Bottom left: RGB-D SLAM with maps that combine keyframes with planar objects[12]. Bottom right: Shape and motion reconstruction of vehicles in a dynamic street scene[13].

Simultaneous localization and mapping or structure-from-motion is a long-standing problem in computer vision and robotics. In this project, we investigate direct visual SLAM approaches that enable robots to acquire 3D maps of the environment and localize in the maps in real-time. We are interested in combining learning-based priors with classical optimization-based methods to SLAM in order to tackle challenging problems such as dense and detailed reconstruction with monocular cameras or SLAM in dynamic environments and to increase robustness and accuracy.

Direct methods avoid hand-crafted interest point detectors and descriptors. Instead they estimate image correspondences and camera motion directly on the pixel level. The estimation of camera trajectory and 3D reconstruction from a video is formulated as an optimization problem. It uses the estimated depth of pixels and camera motion to project pixels between images and measure their photoconsistency. Based on this principle, real-time SLAM approaches, for instance, for stereo cameras can be implemented[14]. The approach can also be complemented with high-frame rate measurements of linear accelerations and rotational velocities by inertial measurement units[15]. This significantly reduces the accumulation of visual errors and overcomes limitations of pure visual odometry in textureless areas or bad lighting conditions.

Recently, we proposed Deep Virtual Stereo

[11] N. Yang, R. Wang, J. Stueckler, D. Cremers. Deep Virtual Stereo Odometry: Leveraging Deep Depth Prediction for Monocular Direct Sparse Odometry. In *European Conference on Computer Vision (ECCV)*, 2018

[12] L. Ma, C. Kerl, J. Stueckler, D. Cremers. CPA-SLAM: Consistent Plane-Model Alignment for Direct RGB-D SLAM. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016

[13] F. Engelmann, J. Stueckler, B. Leibe. SAMP: Shape and Motion Priors for 4D Vehicle Reconstruction. In *IEEE Winter Conference on Applications of Computer Vision, WACV*, 2017

[14] J. Engel, J. Stueckler, D. Cremers. Large-Scale Direct SLAM with Stereo Cameras. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2015.

[15] V. Usenko, J. Engel, J. Stueckler, D. Cremers. Direct Visual-Inertial Odometry with Stereo Cameras. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.

Odometry[11], a hybrid method that uses a learned prior model for monocular depth prediction for direct monocular SLAM in a classical optimization pipeline. The model predicts dense depth from single images which is used to initialize depth and regularize the depth estimation during optimization.

Direct methods are also well suited for SLAM with RGB-D cameras[16,17]. In previous work[18], we combine keyframe-based direct SLAM with tracking and mapping of planar objects in a single optimization framework. The addition of planar objects provides semantic information in the map, reduces drift during camera tracking and improves the estimated map and trajectory estimate. In[19], we have investigated 3D reconstruction and motion estimation of dynamic objects. The approach uses a motion model and a 3D shape prior to reconstruct moving vehicles in stereo image sequences of street scenes. In our ongoing work we aim at integrating deep learning priors for dense real-time SLAM and SLAM in dynamic environments.

More information: https://ev.is.mpg.de/project/visual-simultaneous-localization-and-mapping

[16] C. Kerl, J. Stueckler, D. Cremers. Dense Continuous-Time Tracking and Mapping with Rolling Shutter RGB-D Cameras. In *IEEE International Conference on Computer Vision (ICCV)*, 2015

[17] R. Maier, J. Stueckler, D. Cremers. Super-Resolution Keyframe Fusion for 3D Modeling with High-Quality Textures. In *International Conference on 3D Vision (3DV)*, 2015

[18] L. Ma, C. Kerl, J. Stueckler, D. Cremers. CPA-SLAM: Consistent Plane-Model Alignment for Direct RGB-D SLAM. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.

[19] F. Engelmann, J. Stueckler, B. Leibe. SAMP: Shape and Motion Priors for 4D Vehicle Reconstruction. In *IEEE Winter Conference on Applications of Computer Vision, WACV*, 2017.

## 9.3 Research group leader: Joerg Stueckler

### Dr. Joerg Stueckler

Joerg Stueckler received his Diploma (equiv. M.S.) in Computer Science from the University of Freiburg (2007), and his Ph.D. in Computer Science from Bonn University (2014). As a postdoc he performed research in the Computer Vision groups of the Technical University of Munich and RWTH Aachen University. In 2017/2018, he spent one semester as a visiting professor in the Department of Computer Science at the Technical University of Munich for the Chair of Computer Vision and Artificial Intelligence. Since April 2018, he is leader of the independent Max Planck Research Group on Embodied Vision at the Max Planck Institute for Intelligent Systems in Tuebingen. Dr. Stueckler investigates research topics in computer vision, machine learning and robotics. In computer vision and robotics, he is most known for his research on localization, mapping and object perception. Dr. Stueckler is the recipient of the 2015 Georges Giralt Award for the best PhD thesis in European robotics.

### Academic Positions

| | |
|---|---|
| 04/2018 – present | Max Planck Research Group Leader at the MPI for Intelligent Systems |
| 10/2017 – 03/2018 | Visiting Professor, Department for Computer Science, Technical University of Munich |
| 10/2015 – 10/2017 | Postdoctoral Researcher, Computer Vision Group, RWTH Aachen University |
| 09/2014 – 09/2015 | Postdoctoral Researcher at Chair for Computer Vision and Pattern Recognition, Technical University of Munich |
| 04/2008 – 08/2014 | Research Associate at Chair for Autonomous Intelligent Systems, University of Bonn |
| 10/2007 – 04/2008 | Research Associate in Research Group on Learning Humanoid Robots, University of Freiburg |

### Awards & Honors (Selected)

| | |
|---|---|
| 2015 | Georges Giralt PhD Award by euRobotics aisbl |
| 2011 – 2013 | 1. Place RoboCup@Home, Team NimbRo@Home |
| 2011 | Finalist KUKA Service Robotics Best Paper Award at IEEE ICRA 2011 |

### Selected Organization and Community Service

| | |
|---|---|
| 2018 | Demo Chair European Conference on Computer Vision |
| 2015 – 2016 | Associate Editor IEEE Int. Conf. on Intelligent Robots and Systems |
| 2015 | Advisory Committee of the 1st Amazon Picking Challenge 2015 |

### Memberships

MPI-ETH Center for Learning Systems, Associated Member since 2018
International Max Planck Research School on Intelligent Systems, Faculty Member since 2018
Advisory Committee of the 1st Amazon Picking Challenge 2015

# 10 INTELLIGENT CONTROL SYSTEMS

## 10.1 Research Overview

*Research in the Intelligent Control Systems (ICS) group focuses on decision making, control, and learning for autonomous intelligent systems. We develop fundamental methods and algorithms that enable robots and other physical intelligent systems to interact with their environment through feedback, autonomously learn from data, and interconnect with each other to form collaborative networks. Turning mathematical and theoretical insight into enhanced autonomy and performance of real-world physical systems is an essential and driving facet of our work.*
*The ICS group was founded as an independent Max Planck and Cyber Valley Research group in February 2018 and is headed by Dr. Sebastian Trimpe.*

The Intelligent Control Systems (ICS) group aims to develop machine learning and decision algorithms for machines in the physical world. Our research often starts with fundamental theoretical questions on learning and control, leading us to develop new methods and algorithms, which we finally implement and demonstrate on physical machines such as robots, vehicles, and other autonomous systems.

When learning on physical machines, some special challenges arise, which are different from other machine learning domains typically involving pure software or computer systems. For example, learning in the real world often has to cope with imperfect and relatively small data sets, because physical systems cannot be sampled arbitrarily and exhibit high-dimensional and continuous state-action spaces. A constant stream of data (e. g. from sensors) requires online and lifelong learning, but often on embedded hardware with limited computational resources. Finally, theoretical guarantees on safety, robustness, and reliability are essential for physical learning systems, but often not available in standard machine learning. These are some of the fundamental challenges that arise when artificial intelligence meets the physical world – and that drive our research.

In addition to learning, control, and decision making for a single physical system, we are also interested in *distributed* and *networked* problems, for example, where multiple intelligent agents cooperate to achieve a common goal. How can a team of robots efficiently coordinate their actions? What information should they exchange, and when? And how to design for limited embedded resources such as bandwidth, computation, or energy? These are some of the questions that we address in this research direction.

As we seek to bridge computational and physical intelligence, research at ICS is highly interdisciplinary. In particular, we combine and intersect the disciplines of machine learning, systems & control theory, applied mathematics, and robotics. The main directions of current research at ICS can be summarized as:

- **Learning-based control:** machine learning and control theory for learning on physical systems with guarantees;
- **Distributed intelligence:** learning, control, and cooperation across multi-agent and cyber-physical networks;
- **Resource efficiency:** achieving high performance control and learning with limited resources (embedded computation, small data, communication bandwidth, energy).

### Research projects

Research projects at ICS are interdisciplinary and span some or all of the above-mentioned directions. Current projects and related publications are listed below. A selection of these projects is presented in more detail on the following pages.

- Controller Learning using Bayesian Optimization [24, 31, 58, 623]
- Event-based Wireless Control of Cyber-physical Systems [43, 586, 628]
- Learning Probabilistic Dynamics Models [20, 23]
- Adaptive Locomotion of Soft Microrobots [627]
- Event-triggered Learning [202, 629]
- Gaussian Filtering as Variational Inference [10, 62]
- Learning-based Model Predictive Control [624, 626]
- Model-based Reinforcement Learning for PID Control [32]
- Networked Control and Communication [3, 5]

### Group development

Within this reporting period, the Intelligent Control Systems (ICS) group has transitioned into an independent research group. The ICS group was established as an independent Max Planck Research Group (MPRG) at MPI-IS Stuttgart in February 2018 within the Cyber Valley Initiative. Previously, Sebastian Trimpe was leading a sub-group within the Autonomous Motion Department in Tübingen, which was the starting point for the new group.

With becoming an independent MPRG, the ICS group has grown significantly. At the end of 2018, the group included 7 PhD students, 5 Master students, 2 student assistants, and 1 administrative assistant. Primary funding of the group is through the Cyber Valley Initiative. We have also attracted additional third-party funding and positions, especially DFG funding (1 PhD position) and two industry projects with Cyber Valley industry partners *Bosch* and *IAV Automotive Engineering* (2 PhD positions).

The ICS group has also established multiple research collaborations with national and international academic partners including University of Stuttgart, University of Tübingen, TU Dresden, ETH Zurich, KTH Stockholm, University of Toronto, and Stanford University. Moreover, the PI and the group are actively participating and contributing to the institute's wider research collaborations such as the *International Max Planck Research School on Intelligent Systems*, the *Max Planck ETH Center for Learning Systems*, the *Cyber Valley*, and the *SimTech Cluster of Excellence* at the University of Stuttgart. Finally, ICS has several active research collaborations within MPI-IS, both in Tübingen and Stuttgart (e.g., Empirical Inference Department, Physical Intelligence Department, Probabilistic Numerics Group, Autonomous Learning Group, and Micro, Nano, and Molecular Systems Group).

## 10.2 Selected Research Projects

## Controller Learning using Bayesian Optimization

*Alonso Marco Valle, Philipp Hennig, Alexander von Rohr, Jeannette Bohg, Stefan Schaal, Sebastian Trimpe*

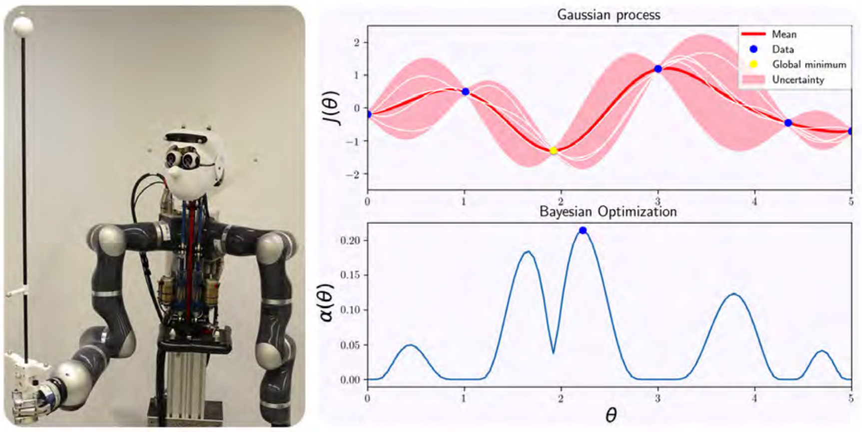


Figure 10.1: Left: Humanoid robot Apollo learning to balance an inverted pole using Bayesian optimization. Right: One-dimensional synthetic example of an unknown cost J($\theta$) modeled as a Gaussian process for controller parameter $\theta$, conditioned on observed data points. The next controller to evaluate is suggested by the Bayesian optimizer where the acquisition function $\alpha(\theta)$ finds its maximum.

Autonomous systems such as humanoid robots are characterized by a multitude of feedback control loops operating at different hierarchical levels and time-scales. Designing and tuning these controllers typically requires significant manual modeling and design effort and exhaustive experimental testing. For managing the ever greater complexity and striving for greater autonomy, it is desirable to tailor intelligent algorithms that allow autonomous systems to learn from experimental data. In our research, we leverage automatic control theory, machine learning, and optimization to develop automatic control design and tuning algorithms.

In [58] we propose a framework where an initial controller is automatically improved based on observed performance from a limited number of experiments. Entropy Search (ES)[1] serves as the underlying Bayesian optimizer for the auto-tuning method. It represents the latent control objective as a Gaussian process (GP) (see above figure) and sequentially suggests those controllers that are most informative about the location of the optimum. We validate the developed approaches on the experimental platforms at our institute (see figure).

We have extended this framework into different directions to further improve data efficiency. When auto-tuning real complex systems (like humanoid robots), simulations of the system dynamics are typically available. They provide less accurate information than real experiments, but at a cheaper cost. Under limited experimental cost budget (i.e., experiment total time), our work [31] extends ES to include the simulator as an additional information source and automatically trade off information vs. cost.

The aforementioned auto-tuning methods model the performance objective using standard GP models, typically agnostic to the control problem. In [24], the covariance function of the GP model is tailored to the control problem at hand by incorporating its mathematical structure into the kernel design. In this way, unforeseen observations of the objective are predicted more accurately. This ultimately speeds up the convergence of the Bayesian optimizer.

Bayesian optimization provides a powerful framework for controller learning, which we have successfully applied on very different settings: humanoid robots [58], micro robots [627] and automotive industry [623].

More information: https://ics.is.mpg.de/project/cont-learn-bayes-opt

[1] P. Hennig, C. Schuler. Entropy Search for Information-Efficient Global Optimization. *Journal of Machine Learning Research* **13**: 1809–1837, June 2012.

## Event-based Wireless Control of Cyber-physical Systems

*Dominik Baumann, Fabian Mager, Harsoveet Singh, Marco Zimmerling, Sebastian Trimpe*

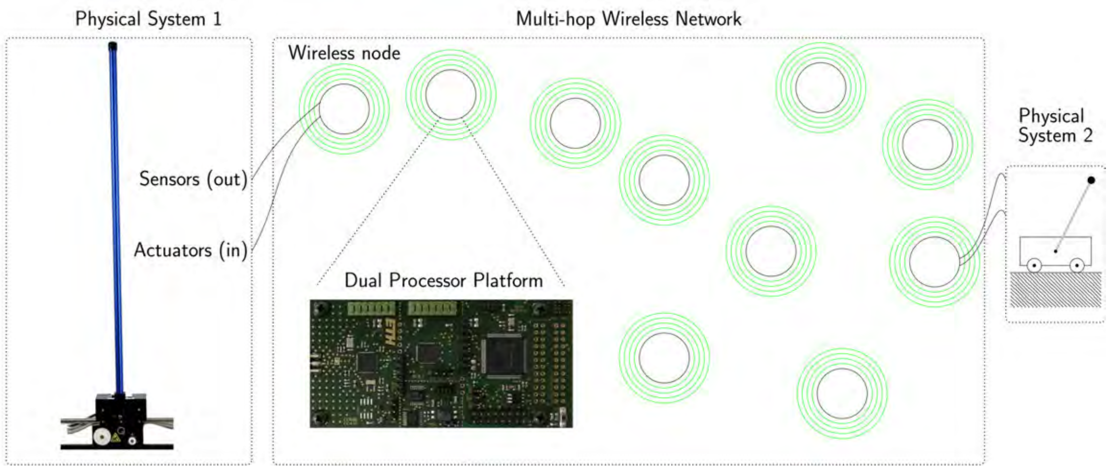


Figure 10.2: Schematic of the cyber-physical testbed developed within this project to investigate novel approaches for feedback control over wireless networks. A wireless multi-hop network connects multiple physical cart-pole systems. Each of the wireless nodes is implemented by a Dual Processor Platform, which combines an application and a communication processor for end-to-end-real-time operation. The testbed can be flexibly used to study, for example, stabilization and synchronization over wireless, which are relevant for control in CPSs in general.

Cyber-physical systems (CPS) tightly integrate physical processes with computing and communication, thus, enabling emerging applications such as coordinated flight of autonomous vehicles or controlling factory automation machinery over wireless networks. The adoption of wireless technology offers unprecedented flexibility in sharing data between these systems, for example, to increase collective information or take collaborative action, but also comes with severe challenges. Wireless networks are orders of magnitude less reliable than wired communication, while feedback control typically poses strict requirements on reliability and timeliness of data.

Within this project, novel feedback control strategies are developed in order to realize the full potential of future CPSs. This is done by a tight co-design of network protocol and control system. The protocol tames natural imperfections of wireless communication to the extent possible. The control design then takes the remaining imperfections (delays and packet losses) into account to realize stable closed-loop control over low-power multi-hop networks. Stability is proven through a formal analysis and demonstrated on the cyber-physical testbed shown above that includes multiple, spatially distributed physical processes controlled over a multi-hop low-power wireless network [628].

A further challenge in wireless CPS is that they can easily be overloaded when multiple systems communicate over the same network. In order to realize energy savings and optimal system-level performance, communication should occur only if necessary to achieve desired performance. Current research activities within this project focus on the development of novel event-triggering mechanisms [43],[3],[5] for the integration of control with the communication system and on learning-based concepts for further resource savings [586],[629].

More information: https://ics.is.mpg.de/project/distributed-and-event-based-wireless-control-of-cyber-physical-systems

## Learning Probabilistic Dynamics Models

*Andreas Doerr, Christian Daniel, Duy Nguyen-Tuong, Stefan Schaal, Marc Toussaint, Sebastian Trimpe*

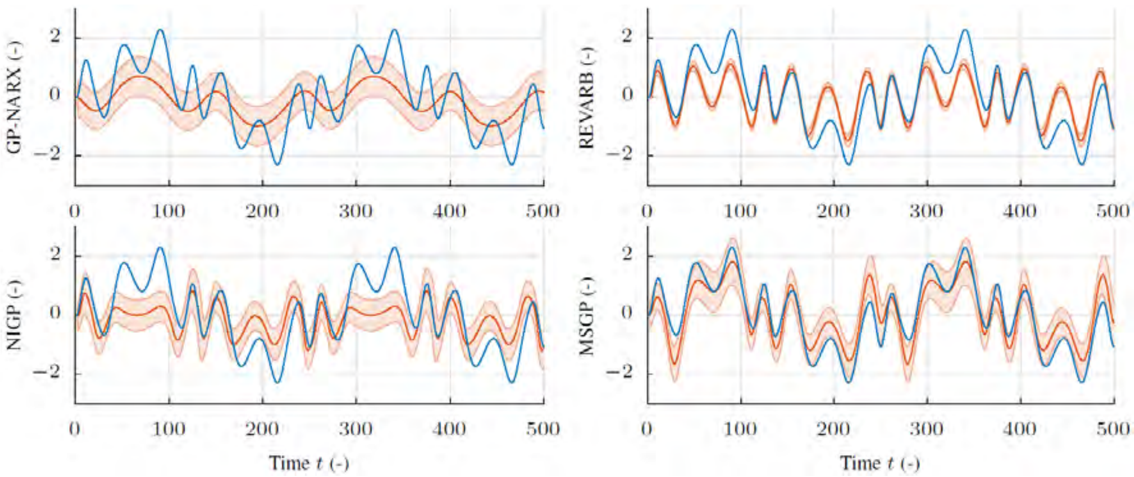


Figure 10.3: Visualization of long-term model predictions from four state-of-the-art probabilistic model learning methods. The observed system output time series (blue) is displayed together with the model's predictive distribution (mean and +/- 2 std in red). Results from the proposed Multi-Step Gaussian Process (MSGP) [23] optimization scheme are in the bottom right.

In Reinforcement Learning (RL), an agent strives to learn a task solely by interacting with an unknown environment. Given the agent's inputs to the environment and the observed outputs, model-based RL algorithms make efficient use of all available data by constructing a model of the underlying dynamics. Data-efficiency has been shown to greatly improve over model-free (e.g. policy gradient) or value function based methods. At the same time, incorporation of uncertainty is essential to mitigate effects of sparse and non-iid data and to prevent model bias.

Learning probabilistic predictive models from time-series data on real systems is however a challenging task, for example, because of imperfect data (e.g. noise and delays), unobserved system states, and complex, non-linear dynamics (e.g. joint friction and stiction). This research aims for high quality, probabilistic, and long-term predictive models, in particular for the use in RL.

In [23], we exploit three main ideas to improve model learning by leveraging structure from the subsequent RL problem:

1. Optimize for long-term predictions.
2. Restrict model learning to the input manifold reachable by the specific policy.
3. Incorporate the approximations made for computing the expected discounted cost into the model learning.

The proposed model learning framework *Multi Step Gaussian Processes (MSGP)* [23] was shown to enable robust, iterative RL without prior knowledge on a real-world robotic manipulator. At the same time, state-of-the-art predictive performance is demonstrated in a benchmark of synthetic and real-world datasets [23].

Oftentimes in practice, the underlying system state cannot be directly measured, but must be recovered from observed input/output data. In our work on *Probabilistic Recurrent State-Space Models (PR-SSM)* [20], we lift ideas from deterministic Recurrent Neural Networks (RNN) into the realm of probabilistic Gaussian Process State-Space Models (GP-SSMs). The resulting inference scheme is derived as approximate Bayesian inference using variational techniques to robustly and scalably identify GP-SSMs from real-world data.

More information: https://ics.is.mpg.de/project/learning-probabilistic-dynamics-models

## 10.3 Research group leader: Sebastian Trimpe

### Dr. Sebastian Trimpe

Since February 2018, Sebastian Trimpe is a Max Planck Research Group Leader (W2) at MPI-IS Stuttgart. He leads the independent Max Planck Research Group on Intelligent Control Systems, which has been established within the Cyber Valley Initiative and focuses on fundamental research at the intersection of control, machine learning, distributed systems, and robotics. Trimpe obtained his PhD (Dr. sc.) degree in Dynamic Systems and Control from ETH Zurich in 2013 with Raffaello D'Andrea as his advisor. Previously, he received a B.Sc. degree in General Engineering Science in 2005, a M.Sc. degree (Dipl.-Ing.) in Electrical Engineering in 2007, and an MBA degree in Technology Management in 2007, all from Hamburg University of Technology. In 2007, he was a research scholar at the University of California at Berkeley. Trimpe is recipient of the General Engineering Award for the best undergraduate degree (2005), a scholarship from the German Academic National Foundation (2002-2007), the triennial IFAC World Congress Interactive Paper Prize (2011), and the Klaus Tschira Award for achievements in public understanding of science (2014).

### Appointments

| | |
|---|---|
| 02/2018 – present | W2 Independent Max Planck Research Group Leader, MPI-IS |
| 10/2016 – 01/2018 | Senior Research Scientist & Group Leader, AMD, MPI-IS |
| 09/2013 – 09/2016 | Research Scientist, Autonomous Motion Deptartment (AMD), MPI-IS |
| 03/2013 – 08/2013 | Postdoctoral Researcher & Lecturer, ETH Zürich |

### Education

| | |
|---|---|
| 2013 | PhD (Dr. sc.), ETH Zürich, Switzerland, advisor: Raffaello D'Andrea |
| 2007 | MSc (Dipl.-Ing.), Electrical Engineering, TU Hamburg, Germany |
| 2007 | MBA, Technology Management, TU/NIT Hamburg, Germany |
| 04/2007 – 11/2007 | Visiting Student Researcher, University of California Berkeley, USA |
| 2005 | BSc General Engineering Science, TU Hamburg, Germany |

### Academic Memberships and Positions (selected)

| | |
|---|---|
| 2018 – present | Co-speaker Project Network 4, Cluster of Excellence Data-integrated Simulation Science (SimTech), University of Stuttgart, Germany |
| 2018 – present | Faculty Member, International Max Planck Research School for Intelligent Systems (IMPRS-IS), Stuttgart/Tübingen, Germany |
| 2017 – present | External Institute Member and Lecturer, Institute for Systems Theory and Automatic Control, University of Stuttgart, Germany |
| 2016 – present | Associated Member, Max Planck ETH Center for Learning Systems, Switzerland/Germany |
| 2016 – present | IEEE Control Systems Society, TC on Intelligent Control |
| 2015 – present | IEEE Control Systems Society, TC on Networks and Communications |

## Awards & Honors (selected)

| | |
|---|---|
| 2016 | Top Four Finalist for Best Student Paper Award (as advisor and co-author), International Workshop on Discrete Event Systems |
| 2014 | KlarText! Klaus Tschira Award for achievements in public understanding of science |
| 2011 | IFAC Congress Interactive Paper Prize (best out of 450 interactive papers) |

## Speaking Engagements (selected)

| | |
|---|---|
| 2018 | Keynote speaker, 50th International Symposium on Robotics (ISR), Munich |
| 2017 | Invited plenary speaker, VDI Conference on Humanoid Robots, Munich |
| since 2009 | More than 60 invited seminars and talks (to academic, industry, and general audiences) |

## (Co-)Organization of Scientific Workshops and Events (selected)

| | |
|---|---|
| 2019 | 1st IEEE Workshop on Cyber-Physical Networking, Las Vegas, USA |
| 2018 | DFG Priority Program (SPP 1914), Plenary meeting, Tübingen |
| 2018 | 1st Workshop on Benchmarking Cyber-Physical Networks and Systems (CPSbench), Porto, Portugal |
| 2016 – 2018 | Invited Session on Learning-based Control at annual IEEE Conference on Decision and Control (CDC) (3 times) |
| 2015, 2018 | Max Planck ETH Workshop on Learning Control, Tübingen/Zürich |

## Teaching and Advising (selected)

| | |
|---|---|
| 2016 – present | (Co-)Lecturer, *Statistical Learning and Stochastic Control*, U. of Stuttgart |
| 2013 | Lecturer, *Recursive Estimation*, ETH Zürich |
| since 2015 | Supervision of 5 PhD students as main advisor, 3 PhD as co-advisor (1 graduated) |
| since 2008 | Supervision of 20 Master students and 18 Bachelor/other students |

## Third-party Funding (selected)

| | |
|---|---|
| 2018 – 2022 | Cyber Valley Initiative, funding for Max Planck Research Group |
| 2018 – 2021 | IAV GmbH, funding of 1 joint PhD-project *Interpretable Learning Control* |
| 2017 – 2019 | DFG Priority Program 1914, tandem project *Event-based Wireless Control for Cyber-physical Systems* (equal share with Co-PI at TU Dresden) |
| 2015 – 2019 | Robert Bosch GmbH, funding of 1 joint PhD-project *Model-based reinforcement learning* |

# 11 LOCOMOTION IN BIOROBOTIC AND SOMATIC SYSTEMS

## 11.1 Research Overview

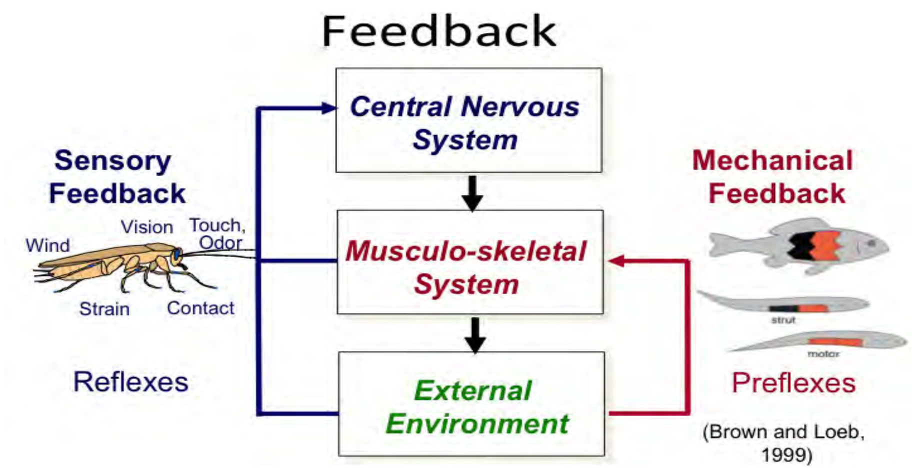


Figure 11.1: Overview of feedback loops present in all animals. Sensory feedback and responses from the mechanical system, using a fish as an example, can modify a gait. Dickinson et al. 2000

Despite remarkable advances in robotics, a substantial capability gap persists vis-a-vis natural systems. The biggest difference between human-made devices and biological organisms is robustness. A key capability of organisms is perturbation-response: their ability to adjust their locomotion to external disturbances and challenges. Organisms respond to perturbations by combining behavior and materials-based responses in ingenious ways often not yet dreamt of by robotics engineers. The role of feedback in adaptive locomotion (Fig. 1) includes both the linkage between sensory and motor control systems and also the mechanical coupling between the body and an animal's external environment. The research goal of the group is to understand animal locomotion as a system, bringing together multiple levels of analysis from neuromechanics to physical models based experimental validation. Leveraging advances in novel soft sensors and actuators, my team complements theoretical models with physical modelling, facilitating control experiments for comparison with live animal and tissue measurements. The group will take the next step in understanding how neural sensory feedback is integrated with the mechanical responses of appendages to simplify control. In the study of locomotion we explore how motion systems get from 'A' to 'B'. We would like to know how animals move and what we might learn from nature when it comes to effective movement in human-made devices such as legged robots. The group's strengths reside in biomechanics, dynamics, engineering robotics, control of locomotion, and bio-inspired design. Based on discoveries pertaining to the locomotor function of lizard tails as control and inertial appendages, allowing them to overcome obstacles and slippery surfaces during climbing, as well as assisting them in aerial righting, we test hypotheses on the tail's role in stability of locomotion and maneuverability using physical models. The findings from locomotor tail function in lizards have inspired the design of biologically informed robots (Fig. 2) with capabilities that could ultimately help test hypotheses regarding robust locomotion in complex, irregular terrain.

## 11.2 Selected Research Projects



### Integrative Systems Biomechanics

Deciphering robustness in perturbation response requires insight in the multiple neuromechanical layers required for locomotion, at levels ranging from somatic neural system circuits to whole organisms. My group's overall goal is to understand animal locomotion as a system, bringing together multiple levels of analysis from neurophysiology to mechanical computation. Our research aims to provide a quantitative understanding of the selective pressures and adaptive tradeoffs that have influenced vertebrate body shapes and locomotor behaviors. We will take the next step in understanding how neural sensory feedback is integrated with the mechanical responses of a ppendages to simplify control. The initial model system we selected for study is the gecko, an arboreal lizard. They possess some of the most spectacular arboreal specializations among nature's elite climbers. To achieve their remarkable climbing performance lizards with adhesive pads strongly rely on the unparalleled ability of their sticky feet, which cling to substrates by means of intermolecular forces. Each one of its toes can easily support the animal's entire body weight. Secure footholds are important because geckos running up vertical surfaces that provide a good grip counterbalance the tendency to pitch back by pulling their head toward the wall with their fore leg on each step, relying on van der Waals forces of their hairy feet.

### Transitions with Biologically Informed Control

By challenging single footholds in wall-running geckos, we discovered a control structure the significance of which had not been previously recognized. Although the remarkable climbing performance of geckos has traditionally been attributed to specialized feet, we showed that a gecko's tail functions as an emergency fifth leg to prevent falling during rapid climbing. Published in *PNAS*, the experimental technique (Figure 11.2 A) of inducing perturbations allowed the discovery of a tail reflex initiated by slipping, which causes the tail to flex ventrally and the tip to push against the vertical surface, thereby preventing pitch-back of the upper body (1) and ultimately catastrophic falls. These experiments suggested that the secret to the gecko's arboreal acrobatics robustness includes an active tail. When geckos encountered an unsurmountable gap, we found that when the lizard finds itself falling with its back to the ground (like the falling cat), a swing of its tail induces the most rapid air-righting reflex yet measured. These findings suggested that large, active tails can function as effective control appendages (2). After developing a three-dimensional analytical model on the aerial righting reflex in lizards, we performed a comparative analysis (3) to contrast the falling cat phenomenon, providing the first measurements on two major phyla of the animal kingdom, with reptiles and invertebrates previously unstudied. Taking this a step further, we found a rapid pitch inversion behavior during incline running (4).

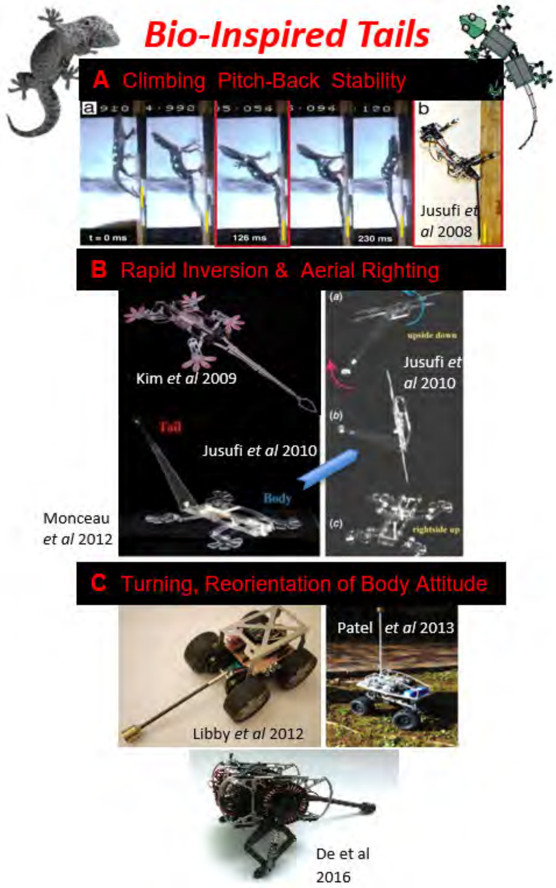


Figure 11.2: Addition of a Tail to a robot can help simplify control of locomotion. (A) Tail use in climbing lizard and legged climbing robot in response to a large pitch-back. From Jusufi et al 2008. Climbing robots use tails as a point of support as in the example of climbing robot RiSE. (B) Aerial righting on a prototype of StickyBot. (C) Tail-assisted body attitude control and rapid turning.

## Systems Biophysics and Biorobotics

In a study published in Nature, we found that a sub-Saharan lizard can use its tail during running jumps to control body attitude by temporarily redirecting excess energy from body to tail (5), which shed new light on the role of caudal reflexes in perturbation response capacity. Rather than inducing caudal autotomy in the lizards, the control experiments were facilitated by a robotic physical model, which allowed for the testing of the tail ventral flexion hypothesis of the running jump (Movie). The cooperative research that we have carried out with colleagues has inspired the design and locomotion capabilities of several robots (Figure 11.2). As shown in (1) the gecko, nature's elite climber, uses its tail to avert catastrophic falling during rapid vertical climbing (Figure 11.2 A), as well as mid-air righting during free fall. It was used passively in Stickybot for climbing (Kim et al 2009) and in RightingBot (Figure 11.2 B). Moreover, tails were also found to have critical functions for body attitude control during transient reorientations in lizards and robots (Figure 11.2 C) by Libby and Jusufi and colleagues (4). Meanwhile, it has also been found to help with turning (Patel et al. 2013). Since, tails have been used so stabilize hopping robots Penn Jerboa (De et al. 2016) and even designed for water-running robots. Drawing from experience with locomotion experiments of nearly 15 years, the team has developed considerable know-how at the cross-disciplinary intersection of experimental robotics and biology. Physical models are critical because they meet the real environment and can inform our mathematical models by revealing new parameters and serve to test our hypotheses about the natural systems.

## Soft robotics

*Ardian Jusufi*

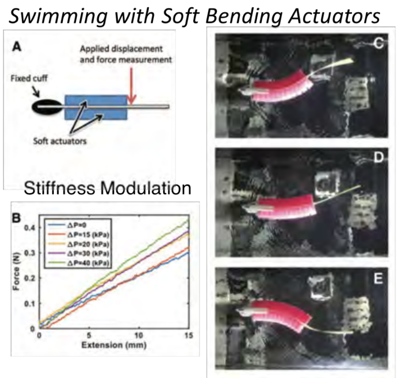


Figure 11.3: Soft robotic fish for study of Co-Contraction facilitated by a physical model (1) with soft artificial 'muscles' swimming in flow channel (c-e) under modification of body stiffness (b).

Soft Robotics is a new frontier that combines flexible, active materials and soft sensors to unlock novel capabilities in robots. Soft active materials allow for a more life-like movement patters that are not afforded by conventional rigid robots. To implement this vision my team will leverage our experience with soft composite manufacturing. Having discovered that geckos maintain lateral undulation of the torso in all gaits from climbing to running, we find they also do this even when racing across the water surface (A). The caudal musculature responsible for the aforementioned tail behaviors, from the muscle activation patterns to the underlying neural control remains under-explored, particularly as it pertains to co-contraction of the axial musculature can modulate the stiffness of body and tail. In the following we describe how our team will utilize soft actuators to discover how axial co-contraction can be used to maximize locomotion performance. Undulatory motion of the body is the dominant mode of locomotion in many reptiles, amphibians, and fishes. Numerous studies of body kinematics and muscle activity patterns have provided insights into the mechanics of swimming. However, it has not been possible to investigate how key parameters such as the extent of bilateral muscle activation affect propulsive performance due to the inability to manipulate muscle activation in live, freely-swimming fishes. To study how co-contraction affects body stiffness control, we have built a soft robotic fish (6) which addresses longstanding hypotheses of undulatory swimming mechanics. This experimental platform (Figure 11.3) demonstrates that a soft physical biorobotic model can answer questions on how body stiffness modulation via co-contraction enhances swimming performance. We discovered that thrust is maximized with a small amount of co-contraction.

Leveraging knowledge in manufacturing of soft sensors and actuators, my team complements theoretical models with physical modelling, facilitating control experiments for comparison with live animal and tissue measurements. To study how co-contraction can be used to modify axial body stiffness during swimming (Movie), we have developed a soft robotic physical model (6). To gain more insight into undulatory locomotion and mechanisms for body stiffness control my team will perform a parameter sweep of contraction phasing and frequency to find the emergent optimum conditions of hydrodynamic stiffening for the highest performance. Two to four pneumatic actuators are attached on each side of a flexible panel with stiffness comparable to that of a fish body. This active pneumatic model is capable of producing substantial trailing edge amplitudes with maximum excursion equivalent to 1.4 foil lengths, and of generating considerable thrust. Altering the extent of bilateral co-contraction ranging from 17% to -22% of the cycle period reveals that maximum thrust is generated with some amount of simultaneous bilateral co-contraction of approximately 3% to 5% of the cycle period; thrust was substantially reduced for conditions of greatest antagonistic overlap in left-right actuation, and also for the largest latencies introduced when the system is exposed to water flow. Our group seeks to collaborate with scientists overseas to determine if they can corroborate the findings from our physical model with electromyographic measurements on Bluegill Sunfish showing bilateral co-activation when the animals accelerate during swimming. We will follow this up with cooperative experiments and visits, also comparing his measurements of the lateral line in lampreys with our soft sensor based lateral line to test hypotheses related to swimming performance. Our group is actively pursuing a soft biorobotic physical model in cooperation with Prof. Koh Hosoda of Osaka University in Japan.

## 11.3 Research group leader: Ardian Jusufi

### Dr. Ardian Jusufi

Dr. Ardian Jusufi received his Ph.D. from the University of California at Berkeley, USA. After graduating, he went to Cambridge University, UK, for a postdoctoral position at Queens' College. Ardian then moved to Harvard University, USA, for a second postdoctoral position in Professor Robert Wood's Harvard Microrobotics Lab at the Wyss Institute for Biologically-Inspired Engineering.
Ardian positions his research in the fields of both engineering and biology, which has been published in leading journals including Nature, PNAS, Current Biology, Soft Robotics, and PloSOne. Dr. Jusufi is a Cyber Valley Max Planck Group Leader and Faculty of the International Max Planck School of Intelligent Systems as well as Associate Faculty of the Max Planck ETH Center for Intelligent Systems. Dr. Jusufi serves as an Associate Editor for the Proceedings of IEEE Robosoft.

### Education

| | |
|---|---|
| 2007 - 2013 | PhD, University of California at Berkeley, USA. |
| 2013 - 2017 | Postdoc, University of Cambridge, UK, and Harvard University, USA. |

### Awards and Honors

| | |
|---|---|
| 2017 | Best Poster Award. 8th International Symposium on Adaptive Motion of Animals & Machines. Hokkaido, Japan. |
| 2015 - 2016 | Fellowship Early Postdoc.Mobility - Swiss National Science Foundation. |
| 2013 - 2014 | Post-Doctoral Research Associate Award. Queens' College. University of Cambridge, UK. |
| 2012 | Outstanding Teaching Award for Graduate Student Instructors – U.C. Berkeley. |
| 2009 - 2011 | Fellowship for Prospective Researchers - Swiss National Science Foundation. |
| 2009 | Best Student Paper Competition, Runner-Up, Society of Integrative & Comparative Biology. |
| 2008 | William V. Power Award – University of California, Berkeley. |

### Leadership and Service (selected)

| | |
|---|---|
| 2018 - present | Associate Editor. Proceedings of IEEE RoboSoft. |
| 2018 - present | Program Committee Member, IEEE 2nd International Conference on Soft Robotics, Korea. |
| 2013 - 2018 | Representative of Postdoctoral Researchers, University of Cambridge, UK. |

### Invited presentations (selected)

| | |
|---|---|
| 2018 | Robotics: Science Systems RSS 2018. Chaired Session in workshop on multi-modal or multi- functional uses for limbs, tails, and other body parts. |
| 2011 | Invited Symposium Presentation. Society of Integrative and Comparative Biology. Symposium on the Biomechanics of Bio-aerial Locomotion, USA. |
| 2010 | Invited Brown Bag Lecture. University of Zurich, Prof. Rolf Pfeifer. Department of Computer Science, Artificial Intelligence Institute, Switzerland. |

### Publications (selected)

A. Jusufi, D. Vogt, R. Wood, G. Lauder. Undulatory Swimming Performance and Body Stiffness Modulation in a Soft Robotic Fish-like Physical Model. *Soft Robotics*, 2017

J. Mongeau, B. McRae, A. Jusufi, P. Birkmeyer, A. Hoover, et al. Rapid Inversion: Running Animals and Robots Swing like a Pendulum under Ledges. *PLoS One*, 2012

T. Libby, T. Moore, E. Chang, D. Li, D. Cohen, et al. Tail-assisted pitch control in lizards, robots and dinosaurs. *Nature*, 2012

A. Jusufi, Y. Zeng, R. Full, R. Dudley. Aerial righting reflexes in flightless animals. *Integ. Comp. Biol.* 2011

A. Jusufi, D. Kawano, T. Libby, R. Full. Righting and turning in mid-air using appendage inertia: reptile tails, analytical models and bio-inspired robots. *IOP Bioinsp. Biomim.* 2010

A. Jusufi, D. Goldman, S. Revzen, R. Full. Active tails enhance arboreal acrobatics in geckos. *PNAS* **105** (11): 4215–4219, 2008

# 12 MOVEMENT GENERATION AND CONTROL

## 12.1 Research Overview

*The ERC Starting Grant Research Group on robotic motion planning and control headed by Ludovic Righetti has been established in June 2015. Since September 2017, L. Righetti is an associate professor at New York University and retains a 50% affiliation at MPI-IS.*

**Research Statement** What are the algorithmic principles that would allow a robot to run through a rocky terrain, lift a couch while reaching for an object that rolled under it or manipulate a screwdriver while balancing on top of a ladder? By answering these questions, we try to understand **the fundamental principles for robot locomotion and manipulation** that will endow robots with the robustness and adaptability necessary to efficiently and autonomously act in an unknown and changing environment.

Our research assumption is that understanding how robots should move is a necessary step towards autonomy and **a comprehensive theory of robot movement is needed**. This theory should have at least three important properties: 1) it can be used to control any robot with legs and arms for both manipulation and locomotion tasks, 2) it allows robots to constantly improve their performances as they experience the world and 3) it is fully automated, (i.e. no need for time-intensive engineering each time a new robot is used or a new task needs to be performed).

With this goal in mind, our research agenda follows several complementary directions that define a consistent research program for the generation of movements in autonomous robots. In particular, we explore problems related to high performance torque control, contact interactions, reactive motion planning and movement learning and we apply our research to both locomotion and manipulation problems.

## 12.2 Selected Research Projects

## Learning contact dynamics

*Ludovic Righetti, Julian Viereck, Avadesh Merduri, Felix Grimminger, Majid Khadiv*

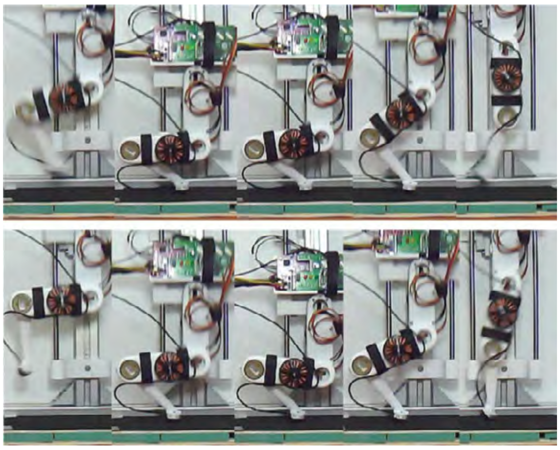

Figure 12.1: Demonstration of the motions learned with very dynamic impacts on a real robot. The robot was designed at the MPI-IS as part of a Grassroots project.

In this project, we are interested in using machine learning techniques to handle robot behaviors with hard impact dynamics and contacts, such as walking, jumping or catching a ball. We investigate reinforcement learning techniques to learn dynamic motions under contact. We also use statistical learning tools to find representations for contacts that can be efficiently incorporated in estimation and control algorithms.

More information: https://mg.is.mpg.de/project/learning-contact-dynamics

## Multi-contact trajectory optimization

*Ludovic Righetti, Alexander Herzog, Andrea Del Prete, Maximilien Naveau, Brahayam Ponton, Majid Khadiv, Sean Mason, Nick Rotella*

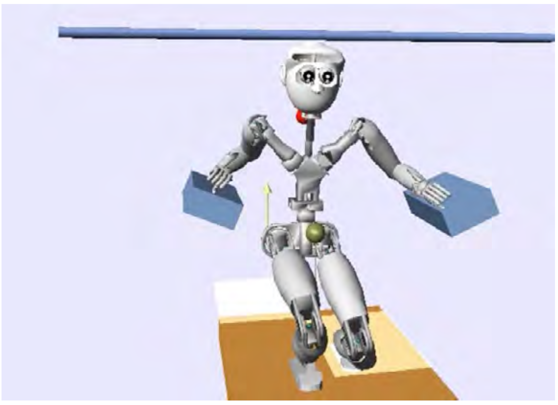

Figure 12.2: Example of motions optimized with our algorithms. Full kino-dynamic plans involving complex multi-contact situations can be computed close to real time.

In this project, we investigate the problem of computing movements for robots with arms and legs in intermittent contact with their environment that are physically correct. Our approach uses trajectory optimization and optimal control techniques. In particular, we study how the physics of mechanical systems impose a structure to the underlying optimization problem. This can be leveraged to find computationally efficient algorithms that optimize for the centroidal-momentum dynamics of the robot, contact timing, locations and forces as well as whole-body motions. Our ultimate goal is to demonstrate such techniques in real-time on complex multi-contact tasks.

More information: https://mg.is.mpg.de/project/multi-contact-trajectory-optimization

## Sensor fusion for legged robots

*Ludovic Righetti, Nick Rotella, Sean Mason, Alexander Herzog, Stefan Schaal*

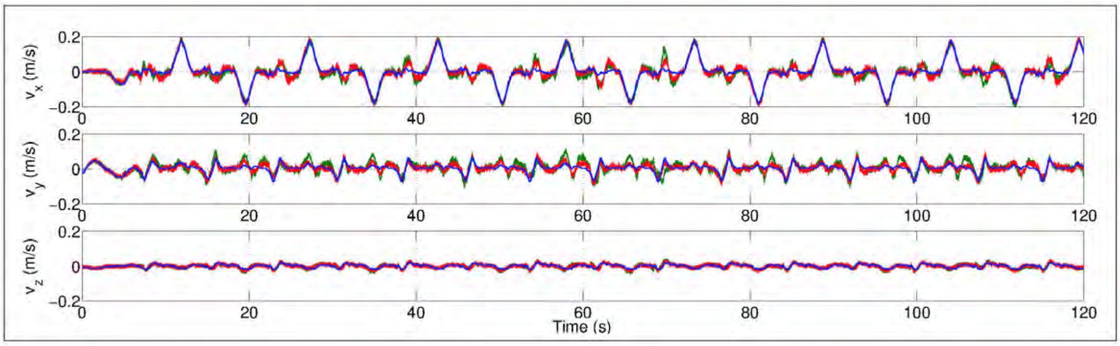


Figure 12.3: The plots show estimated and ground truth velocities for the base of a humanoid robot when using our novel contact-aware base state estimator. It was fundamental to realize the experiments presented in the next section.

In this project, we explore the problem of fusing sensor information from inertial, position and force measurements to recover quantities fundamental for the feedback control of legged robots. Our final goal is to find a systematic way of fusing multiple sensor modalities to improve the control of legged robots.

- Our theoretical work studies the observability properties of nonlinear estimation models to better understand the fundamental limitations of sensor fusion approaches for legged robots
- We design novel estimators combining multiple sensor modalities (inertial, position and force) to provide accurate estimates of important quantities, such as the position and orientation of the robot in space, its center of mass and angular momentum or any external forces applied on the robot. We demonstrate through numerical simulation and real robot experiments that our estimators can significantly improve the control performance.
- We also investigate machine learning techniques to automatically extract sensory information, for example, to learn how to classify the contact states of a robot using unsupervised learning

More information: https://mg.is.mpg.de/project/sensor-fusion-for-legged-robots

## Whole-body control of legged robots

*Ludovic Righetti, Alexander Herzog, Nick Rotella, Sean Mason, Felix Grimminger, Stefan Schaal*

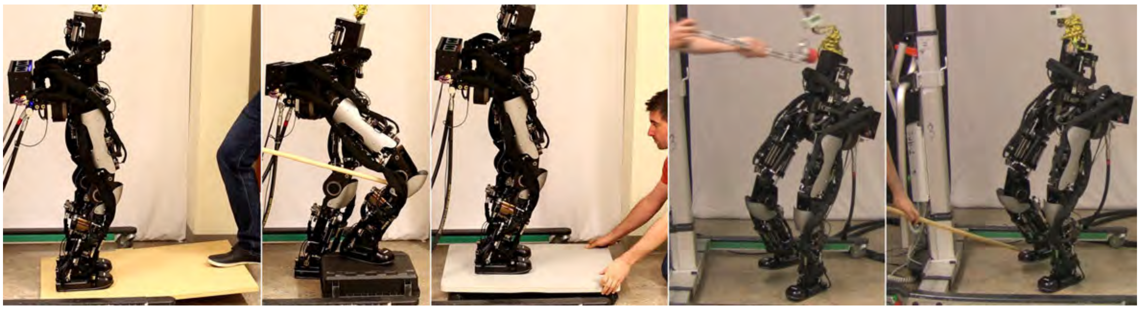

Figure 12.4: The robot is balancing while different kinds of disturbances are applied to its mechanical structure. At each millisecond an optimization problem is solved that decides which joint torques are to be applied in order to realize feedback controllers that a) stabilize the robot after a push while b) generating admissible forces on the ground. With our whole-body controller the robot remains in balance despite comparably strong pushes.

Legged robots are expected to locomote autonomously in an uncertain and potentially dynamically changing environment. Active interaction with contacts becomes inevitable to move and apply forces in a goal-directed way and withstand unpredicted changes in the environment. Therefore, we need to design algorithms that exploit interaction forces and generate desired motions of the robot leading to robust and compliant interaction with the environment. In this context, the choice of a control strategy for legged robots is of primary importance as it can drastically improve performance in face of unexpected disturbances and therefore open the way for agile robots, whether they are locomoting or performing manipulation tasks.

In the theoretical part of our work, we develop algorithms based on dynamic models to control any robots with arms and legs. Our controllers allow for simultaneous control of robot motion and contact interactions in a unified framework. This allows, for example, optimization of contact forces while walking on uneven terrain. In the experimental part of our work, we implement these algorithms on state-of-the-art machines and show the capabilities and limitations of the algorithms when dealing with imperfect models, sensor noise and actuation limitations.

Our experiments show that our humanoid robot can balance with very good performance. It withstands strong pushes and rapid displacements of the ground support. Our experiments demonstrate that our algorithms are robust to dynamic model inaccuracies, sensor noise and actuator bandwidth limitations and it opens the way to more general controller formulations through optimization.

More information: https://mg.is.mpg.de/project/model-based-control-of-floating-based-robots

## 12.3 Awards & Honors

### 2018

- Ludovic Righetti was designated by New York University Provost's Office, together with the Tandon School of Engineering as the Vivian G. Prins Global Scholar at New York University for 2018-2019.

### 2017

- Cédric de Crousaz receives an ETH Medal 2017 for his excellent master thesis (supervision by L. Righetti and S. Trimpe).
- Julian Viereck receives the ETH Medal for his outstanding master thesis on "Learning to Hop Using Guided Policy Search" (Handed out to < 2.5% of all master theses per year at ETH Zurich). He did his master thesis at the Movement Generation and Control Group under the supervision of A. Herzog and L. Righetti

### 2016

- The paper "A Convex Model of Humanoid Momentum Dynamics for Multi-Contact Motion Generation" by Brahayam Ponton, Alexander Herzog, Stefan Schaal and Ludovic Righetti was Finalist for the Best Interactive Session Award at the 2016 IEEE-RAS International Conference on Humanoid Robotics. 10 finalists out of 283 submitted / 186 accepted papers.
- The paper "Stepping Stabilization Using a Combination of DCM Tracking and Step Adjustment" by Majid Khadiv, Sébastien Kleff, Alexander Herzog, Ali Moosavian, Stefan Schaal and Ludovic Righetti was Finalist for the Best Paper Award at the 4th RSI International Conference on Robotics and Mechatronics (ICROM).
- Ludovic Righetti receives the Heinz Maier-Leibnitz Prize 2016. It is awarded by the German Research Foundation (DFG) and the German Federal Ministry of Education and Research. It is considered the most prestigious award in Germany for early career researchers across all disciplines who have established an independent scientific career since having gained their doctorates
- Ludovic Righetti receives the IEEE-RAS Early Career Award (Academic) for "his contributions to the theory of, and experiments in, robot locomotion and manipulation". It is awarded by the IEEE Robotics and Automation Society to "individuals in the early stage of their career who have made an identifiable contribution or contributions which have had a major impact on the robotics and/or automation fields"

## 12.4 Research group leader: Ludovic Righetti

### Prof. Dr. Ludovic Righetti

Ludovic Righetti is an Associate Professor in the Electrical and Computer Engineering Department and in the Mechanical and Aerospace Engineering Department at the Tandon School of Engineering of New York University and a Research Group Leader at the Max-Planck Institute for Intelligent Systems (MPI-IS) in Tübingen, Germany. He studied at the Ecole Polytechnique Fédérale de Lausanne (Switzerland) where he received an engineering diploma in Computer Science (eq. MSc) in 2004 and a Doctorate in Science in 2008 under the supervision of Professor Auke Ijspeert. Between March 2009 and August 2012, he was a postdoctoral fellow at the Computational Learning and Motor Control Lab with Professor Stefan Schaal (University of Southern California).

In September 2012 he started the Movement Generation and Control Group at the Max-Planck Institute for Intelligent Systems in Tübingen, Germany where he became a W2 Independent Research Group Leader (eq. associate professor) in September 2015. His research focuses on the planning and control of movements for autonomous robots, with a special emphasis on legged locomotion and manipulation but he is more broadly interested in problems at the intersection between decision making, automatic control, optimization, applied dynamical systems and machine learning.

#### Employment

| | |
|---|---|
| 09/2017 – present | Associate Professor, New York University, USA |
| 09/2017 – present | Independent Research Group Leader, MPI-IS, Germany |
| 2015 – 2017 | W2 Independent Research Group Leader, MPI-IS, Germany |
| 2012 – 2015 | Group Leader, Autonomous Motion Department, MPI-IS, Germany |
| 2012 – 2017 | Visiting Researcher, University of Southern California, USA |
| 2011 – 2012 | Postdoctoral Researcher, Autonomous Motion Department, MPI-IS, Germany |
| 2009 – 2012 | Visiting Researcher, Computational Neurosciences Laboratories, ATR, Japan |
| 2009 – 2012 | Postdoctoral Researcher, University of Southern California, USA |

#### Education

| | |
|---|---|
| 2008 | Doctorate in Science, Ecole Polytechnique Fédérale de Lausanne, Switzerland<br>Supervisor: Prof. A. J. Ijspeert - Received the Georges Giralt PhD Award |
| 2004 | Diploma in Computer Science, Ecole Polytechnique Fédérale de Lausanne |

#### Fellowships and Awards (selected)

| | |
|---|---|
| 2018 | Vivian G. Prins Global Scholar at New York University (awarded by NYU's Provost Office) |
| 2016 | Heinz Maier-Leibnitz Prize (given by the German Research Foundation (DFG) & the German Ministry of Education and Research (BMBF) |
| 2016 | IEEE Robotics and Automation Society Early Career Award |
| 2014 | ERC Starting Grant, 5 year project CONT-ACT (1.49M€ ) |

| | |
|---|---|
| 2011 | Best Paper Award, IROS 2011 (790 accepted papers/2459 submissions) |
| 2010 | Georges Giralt PhD Award for the best PhD thesis in Europe in the field of robotics (given by the European Robotics Research Network). |
| 2009 | Young Researcher Scholarship from the Swiss National Science Foundation |

## External Research Funding at MPI-IS

| | |
|---|---|
| 2018 – 2021 | *Memmo: Memory of Motion*, European Union Horizon 2020 Research and Innovation Program, ICT-2017-1 Research Grant, Amount: €4,000,000 (€595,250 for co-PI Righetti). |
| 2015 – 2020 | *CONT-ACT: Control of contact interactions for robots acting in the world*, European Research Council Starting Grant, Amount: €1,490,000. |
| 2016 – 2017 | *WALK3D: From walking in 3D environments to multi-contact locomotion*, German Academic Exchange Service, Program PROCOPE France, Amount: € 8,000. |

## Publications and Talks

| | |
|---|---|
| Publications | 14 Refereed Journals, 50 Refereed Conferences, 2 Book Chapters<br>h-index: 31, i-10 index: 50, 3517 citations (Google Scholar 11/28/18) |
| Talks | 34 Invited Seminars at Universities/ 22 Invited Conference Talks |

## Teaching/Supervision Activities

| | |
|---|---|
| | Supervision of 10 PhD students (3 graduated in 2017/2018) |
| | Supervision of 34 Master student projects/theses |
| 2017 – present | Graduate classes Optimal and Learning Control for Robotics, Sensor-Based Robotics and Networked Robotic Systems (NYU) |
| 2014 | Graduate Course *Optimization in Robotics* (with J. Bohg), Univ. of Tübingen |
| 2004 – 2007 | Invited Lecturer (Nonlinear Sys.) and Teaching Assist. (Logical Sys.), EPFL |
| 2014 – present | Thesis committee member for 13 PhD Theses |

## Professional Activities (selected)

| | |
|---|---|
| 2017 – present | Member of the IEEE Robotics and Automation Research and Practice Ethics in Robotics and Automation Committee (RARPEC) |
| 2017 – present | Faculty of the Int. Max-Planck Research School on Intelligent Systems |
| 2015 – present | Associate Editor IEEE RA-Letters |
| 2015 – present | Associated Member of the Max Planck ETH Center for Learning Systems |
| 2014 – present | Advisory Board Member, IEEE/RAS TC on Whole-Body Control |
| 2013 – present | Organization of 7 international scientific meetings/workshops |

## Public Service and Outreach (selected)

| | |
|---|---|
| 2014 – present | Expert for several Non-Governmental Organizations discussing ethical implications of Robotics and AI and autonomous weapon systems including the International Committee of the Red Cross and the Stockholm International Peace Research Institute |
| 2016 | Profile (2 page feature) in the Deutsche Universitats Zeitung (Germany). |
| 2015 | Interview for Deutsche Presse-Agentur on the ERC Starting Grant project CONT-ACT, Article featured in German national news medias including Die Welt, N-TV, Berliner Morgenpost, Reutlinger General-Anzeiger, Hannoverlche Allgemeine, NOZ |
| 2015 | Invited Speaker, Berlin Tech Open Air |

# 13 MICRO, NANO, AND MOLECULAR SYSTEMS

## 13.1 Research Overview

*The independent Max Planck Research Lab headed by Peer Fischer, who is also a Full Professor of Physical Chemistry at the Univ. of Stuttgart, combines research on the physics and chemistry of active matter, with the development of new nano and micro-scale fabrication methods, alongside original engineering approaches to realize novel "Micro Nano and Molecular Systems". The group currently consists of 17 physicists, chemists, material scientists, and biomedical, electrical and mechanical engineers (PhDs and Postdocs). The research group has discovered new effects and developed patented instruments and unique fabrication capabilities. The lab has won major competitive research grants (EU ERC, Volkswagen Foundation, MPG-Fraunhofer, Baden Württemberg Foundation, BMBF, DFG, German Israeli Foundation) with >€3.5Mio in the past three years.*

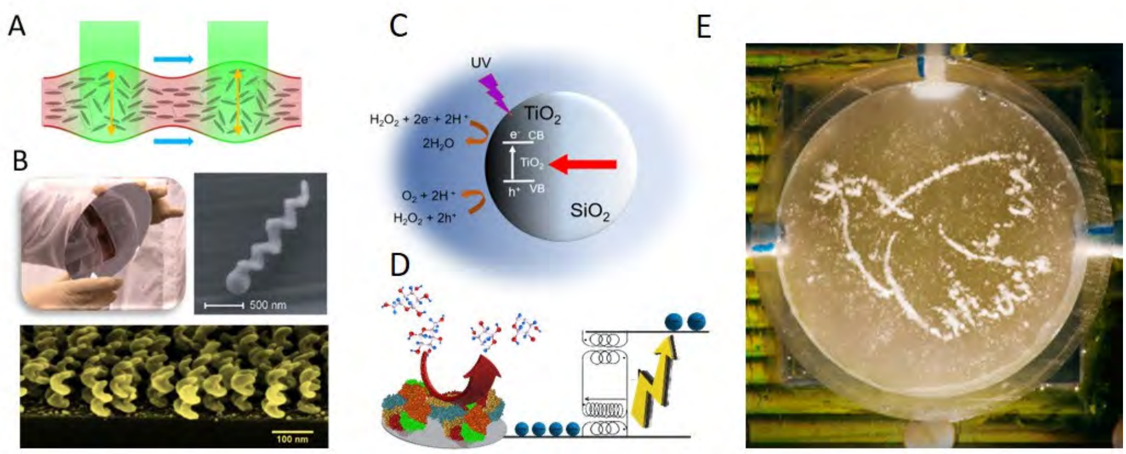


Figure 13.1: Overview image: A) Liquid crystal elastomer microswimmer, B) Wafer scale fabrication of complex colloids and nanostructures, C) Self-propelled chemical nanomotor, D) Boosting the energy from enzymatic reactions for biofuel cells, and E) the invention of the acoustic hologram permits particle assembly 'by design'.

Nature has evolved micro-organisms that operate without neurons, yet exhibit remarkably complex behaviors. These are wonderful examples of sophisticated "micro nano and molecular systems" whose function and interaction with the environment is entirely governed by the chemistry and physics at small scales. There is thus no fundamental reason why we should not also be able to build similar 'bio-inspired' *synthetic* motors, machines, and robotic systems. This will require the ability to synthesize and engineer complex 3D parts and ways to organize or assemble these components. Another challenge is to provide energy for activation, which calls for strategies that are entirely different to the many engineering solutions that have been devised at large scales. Progress in these endeavors promises the development of new materials, fabrication methods, potential biomedical technologies, and novel devices. Crucially, natural organisms encompass dynamical systems that are driven away from equilibrium by constant energy input. The living world is thus special in that it is made from "active" building blocks that transduce energy from the environment. "Active" non-equilibrium components are, therefore, an essential ingredient for future synthetic autonomous systems.

Since many chemical processes take place in solution, one of the research topics of the group concerns diffusion and swimming at the microscale, where the property of a fluid is governed by low Reynolds number hydrodynamics. Here, the lab has made several fundamental contributions including the demonstration that in it is possible to overcome the scallop theorem in biological fluids. This led to the first reciprocal micro-scallop swimmer[1]. In 2016 we were able to demonstrate the smallest microrobots that swim through fluid with internally-generated body-shape changes (see 13.1A)[669]. The lab has also developed unique fabrication methods at the MPI that permit the growth of hundreds of billions of designer nanostructures and colloids at the wafer scale[2]. These are for instance of interest as photonic materials [668]. The growth method is also used to fabricate the world's smallest nanodrills (see 13.1B) that are so small that they can efficiently penetrate the macromolecular networks found in biological fluids and tissues[34]. A recent highlight is that it was possible to demonstrate the first microrobotic system that could be steered efficiently through real tissue over cm-distances[639] (see Projects). Surprisingly, the design of micro-propellers and drills could be further simplified, by exploiting the underlying symmetries of rotation-translation coupling (see Projects). Another ongoing research direction is to realize synthetic chemical motors – structures with localized catalytic reactions, which are "active". The lab's fabrication processes can be used to grow chemical nanomotors [5], which self-propel (13.1C) and give rise to photogravitaxis and dynamic self-organization [648, 655]. An important fundamental question [6] is whether similar mechanisms can be observed at the scale of enzymes. Several projects in the lab address this important and still unanswered question, with high precision diffusion measurements of catalytically active enzymes using optical and NMR methods [642] (see Projects). Only in a system-based approach is it possible to understand how chemical reactions can lead to complex behaviors and potentially give rise to synchronization, muscle-like action, or even self-propelling drug-carriers that can autonomously *chemotax*[644].

In an international collaboration the lab could demonstrate how one can use abundant, simple enzyme reactions to power energy-hungry reactions and biofuel cells. The major implication is that it is now possible to obtain sufficiently high voltages in small biofuel-cell implants. The free enthalpy (Gibbs energy) sets fundamental limits on the energy that can be obtained from a spontaneous reaction. The electrochemical potential is thus restricted. A hybrid chemical-electronic completely passive autonomous device that does not need any outside power, can be coupled to the reaction to overcome the limits set by thermodynamics. It first stores the energy released in many spontaneous (low energy) enzyme reactions and then bundles the energy to drive a second reaction that needs more energy and would therefore not proceed spontaneously (13.1D)[643].

While, dynamic non-equilibrium chemically-powered systems and assembly is a focus, the lab has recently also patented a major engineering advance that permits ultrasound fields to be used for directed assembly. The invention permits ultrasound fields to be formed that are orders of magnitude more sophisticated and complex than what has been possible to date (see Projects) [663]. The method promises a "one shot" 3D fabrication method (13.1E)[650]. The lab invented and patented a second new ultrasound technology to develop a new miniaturized multi-degree of freedom robot arm that can be used for minimally invasive endoscopy[651] (see Projects).

The lab has licensed technology to a start-up company, and has collaborations with a number of international research groups and with research groups and departments at the MPI-IS (e.g. Empirical Inference on machine learning for acoustics (MPG-FhG funded), Intelligent Control Systems for soft-micro-robot control, Haptic Intelligence on organ phantoms for robotic surgery, Modern Magnetic Systems, and Theory of Inhomogeneous Condensed Matter).

In the following pages some of the lab's recent projects will be described in greater detail.

[1] T. Qiu et al. Nat. Commun. 5:5119 (2014).
[2] A.G. Mark et al.,Nature Materials 12, 802, (2013)
[3] D. Schamel et al.,ACS Nano 8, 8794–8801, (2014).
[4] D. Walker et al.,Science Advances 1 (11), e1500501 (2015)
[5] T.-C. Lee et al.,Nano Letters 14, 2407–2412, (2014).
[6] Special Issue edited by A. Sen, A. Balazs, and P. Fischer, Accounts of Chemical Research 51 (9), (2018).

## 13.2 Selected Research Projects



### Most advanced ultrasound fields and "one shot" additive manufacturing

*Kai Melde, Tian Qiu, Eunjin Choi, Peer Fischer*

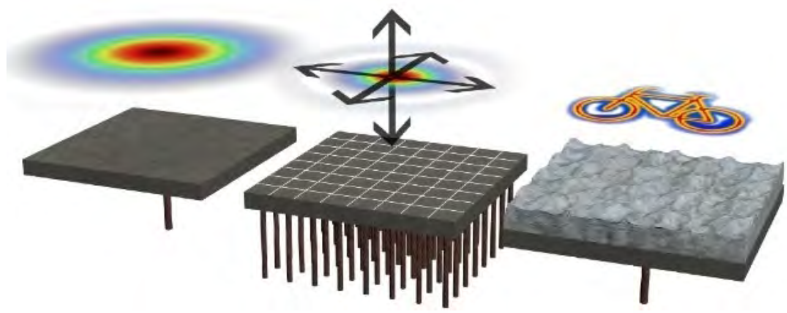

Figure 13.2: Comparison of different sound shaping strategies and achievable degrees of freedom (DOF). Higher DOF leads to more complex sound patterns. From left to right: Single transducer with natural focus (DOF=1), Phased array transducer (as is so far used in medical imaging) with steerable focal point but limited DOF<1000, our simple single transducer with hologram achieves ultrasound images with DOF>10,000.

**Introduction and problem** Additive manufacturing, such as 3D printing, opens new avenues for mechanical design, as it allows fabrication of geometries not possible with classical subtractive machining. However, established methods are serial and therefore very slow. Why is it not possible to assemble matter in the desired shape "in one shot"? Ultrasound waves are potentially very interesting for this application, as they are benign and versatile and as they allow contact-free particle manipulation over many length scales. Acoustic waves exert forces when they interact with matter, which has been used to trap and levitate individual microparticles and cells. The acoustic contrast between the material and the medium, and the spatial variation of the ultrasound field determine the interaction. Resonators and arrays of a few hundred transducers have thus far been used to generate the sound fields, but the former only yields highly symmetrical pressure patterns, and the latter cannot be scaled to achieve sufficiently complex fields.

**Our approach** A radically new approach to generate ultrasound fields with orders of magnitude higher complexity than what has been possible to date, has been invented by our group [663]. We could demonstrate the first acoustic hologram, which is the acoustic analog to the optical kinoform [663]. This is a diffractive element that defines the phase of a passing wave at each point across its wavefront. Theoretically it is now possible to achieve up to 100% diffraction efficiency. The device encodes the desired wave front in the thickness distribution of a thin slab. Due to the difference in the speed of sound between the hologram and the surrounding medium (e.g. fluid) the passing wave will locally lead or lag at each point depending on the local thickness. The phase distribution that is needed to reconstruct an arbitrary target field can be computed by algorithms, that can be based on established optical holography methods. We adapted the theory for physical acoustics.

**Results** We designed and fabricated holograms with an aperture of 50 mm, which effectively contains 15,000 pixels, and is thus one to two orders of magnitude larger than what has been achieved to date using expensive and cumbersome phased array technology. We showed that commercial printers can be used for the rapid fabrication of the diffractive elements (holograms), which means that holograms can be

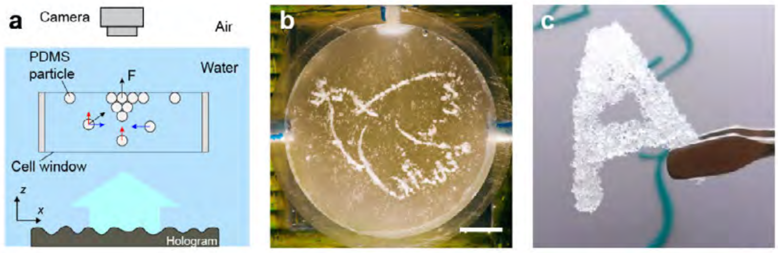


Figure 13.3: Acoustic particle assembly and fabrication. (a) Schematic of the experimental setup (side view). (b) Assembly of microparticles to form a dove (top view). (c) Chemical fusion of the particle assembly results in mechanically stable, self-supporting objects (and potentially much faster than 3D printing).

easily and cheaply made. We then demonstrated parallel assembly of silicone microparticles into complex shapes defined by the projected acoustic image from the hologram (13.3). We expanded this further by a chemical fixation method to allow instant fabrication of objects that can be removed from the bath, while the surrounding bulk liquid remains unaffected [650].

The large DOF of our holograms enables not only control over the scalar amplitude distribution in the acoustic image. It even allows control over the phase, which is the in-plane component of the wave vector. This can be used for advanced and seemingly dynamic object propulsion techniques, such as the surfing particles shown in Fig. 13.4. There the acoustic image is projected onto the water-air interface where the radiation pressure results in crests on the surface. Rubber discs of several millimeters are trapped on the crests due to surface tension. The hologram can now project elongated tracks with a defined phase distribution such that the trapped particles are continuously propelled around the track as long as the sound stays on.

The complex acoustic fields also need to be experimentally validated. The conventional method works by scanning a hydrophone point-by-point through the region of interest and therefore scales badly with size. We therefore developed a very fast method for mapping 3D sound fields using a thin membrane that is continuously imaged by a thermal camera [641]. The evolution of the film's surface temperature can provide the pressure map of large volumes – orders of magnitude faster than what is possible with a hydrophone.

**Discussion and Outlook** The advanced control over ultrasound fields – enabled by acoustic holograms – is a major advance for ultrasound research, useful for many applications in science, medicine and industry. Holography is a general concept and the algorithms and processes we develop are not limited to static diffractive elements. They can be applied to phased arrays if devices with larger DOF become available.

We filed two patent applications on acoustic fabrication and medical therapy, respectively. Parts of this work have been reported as Letters in *Nature*, *Advanced Materials*, and *Applied Physics Letters* and received considerable attention in the media and the scientific community. In 2018 we also won two major grants: An ERC Advanced Grant (*HoloMan*, 2019-2023), and a Max Planck - Fraunhofer cooperation grant (*Akustogramme*, 2018-2022). We collaborate with the Empirical Inference Department at the MPI-IS and the Fraunhofer Institute for Biomedical Engineering.

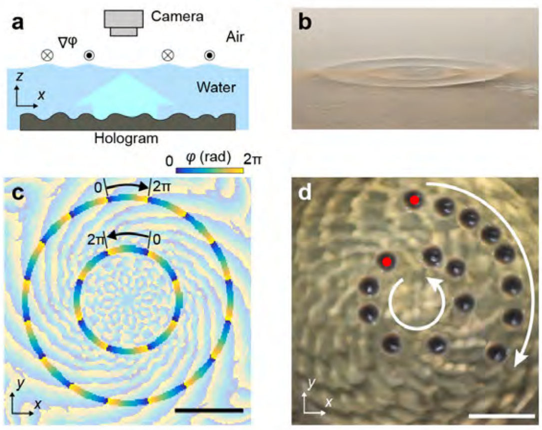


Figure 13.4: Acoustic propulsion of objects along closed tracks with opposing directions. (a) Schematic of the setup (side view). (b) Photo of the water crests. (c) Simulated phase distribution. (d) Time-lapse image showing motion of discs

More information: https://pf.is.mpg.de/project/holoman

## New wireless ultrasonic actuator and its application in miniaturized endoscopy

*Tian Qiu, Kai Melde, Fabian Adams, Andrew Mark, Stefano Palagi, Peer Fischer*

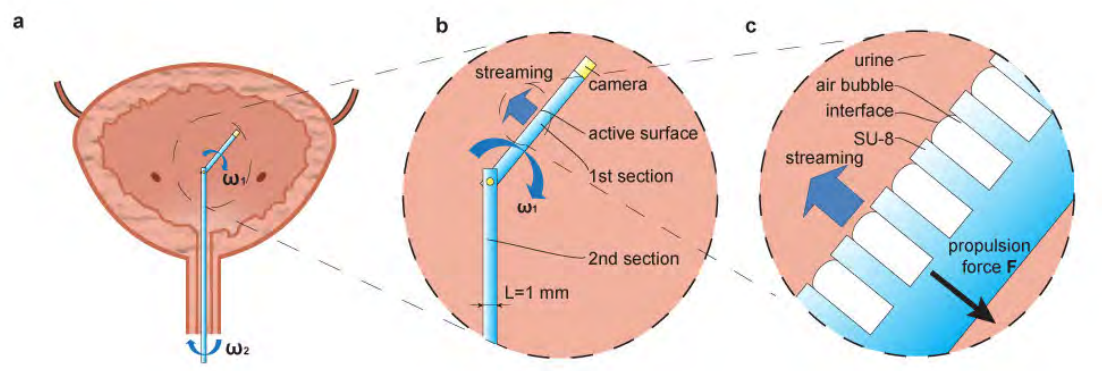


Figure 13.5: Schematic of the miniaturized endoscope driven wirelessly by ultrasound. A surface is microfabricated (c) such that it supports an array of micro-bubbles. The oscillation of the air-liquid interface at the bubble surface causes liquid streaming and thus generates a propulsion force (due to the recoil) (b). The robot arm can now be powered and moved wirelessly to position a camera (a). The robot arm is only 1mm in diameter and thus much smaller than current endoscopes.

**Introduction and problem** Endoscopy enables minimally invasive procedures in many medical fields, such as urology. However, current endoscopes are normally cable-driven, which limits their dexterity and makes them hard to miniaturize. Indeed, current urological endoscopes have an outer diameter of about 3 mm and still only possess one bending degree of freedom. In order to avoid anesthesia and to enable routine inspections, it would be advantageous to develop an endoscopic camera arm that needs no cables or wires. Is it possible to actuate an endoscope that is much smaller than existing instruments? This calls for a new kind of actuator – one that has a simple structure and is powered wirelessly such that it can be miniaturized and integrated more efficiently. Ultrasound has been widely used in medicine and it has also been proposed as a means to transfer power wirelessly. However, common approaches first convert mechanical vibrations to electrical power using the piezoelectric effect, and then in a second step utilize the electricity to power a device. This conversion process adds additional complexity and suffers from low efficiency. In this project, we have developed a wireless miniaturized actuator that can directly convert ultrasound into mechanical work. Our new patented actuation scheme allowed us to build the first ultrasound powered endoscope.

**Our approach** Our idea is to design functional surfaces that are excited by ultrasound and directly generate a force. For this we exploit acoustic streaming, where a resonant bubble can cause fast fluid flow. We fabricate surfaces that contain thousands of resonant microcavities, which we show can actuate a millimeter-scale robot arm in a fluid (Fig. 13.5). We demonstrate that under resonant ultrasound excitation, the functional surface, which we call a Bubble Array Streaming Surface (BASS), generates a strong streaming flow, and thus provides a propulsive force that can be used for actuation. We achieve independent control of different BASS surfaces by frequency-division multiplexing, which allows us to develop a robot arm with many degrees of freedom. We also develop a method to increase the force of the actuator, such that the BASS actuator can even propel mini-robots at the centimeter-scale.

**Results** The functional surfaces, consisting of two-dimensional arrays of cylindrical cavities, were fabricated by a photolithography process. When remotely powered by an acoustic field, the surfaces provide highly directional propulsive forces in fluids *via* acoustic streaming. The size of the microcavities and hence bubbles determines the resonance frequency. This can be exploited to independently address different surfaces, simply by changing the frequency of the

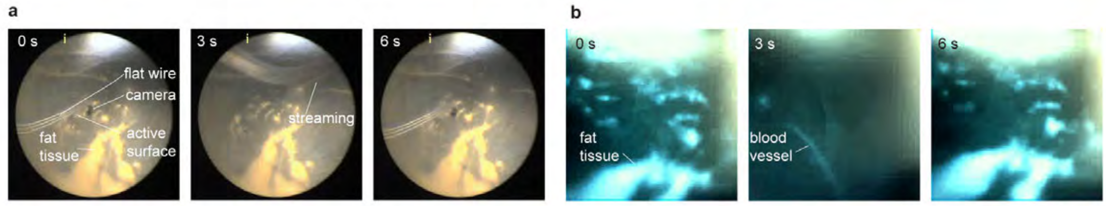


Figure 13.6: Bladder imaging by the miniaturized flexible endoscope. (a) Snapshots of a video showing the movement of the endoscope inside a rabbit bladder. (b) Respective snapshots by the mini-camera installed on the tip of the flexible endoscope.

ultrasound. We easily obtain more than 10 degrees of freedom, thus permitting the wireless actuation of a complex motion of a flexible endoscope, which has a diameter of only 1 mm (Fig. 13.7a). The flexible endoscope successfully performed an endoscopic examination in both a high-fidelity kidney phantom that we developed in our lab [659], as well as in a rabbit bladder *ex vivo* (Fig. 13.6) [651]. The flexible endoscope entered the tool channel of a rigid cystoscope. The structure of the fat tissue can be seen as a marker (image at 0s in Fig. 13.6). When the ultrasound is applied, streaming from the active surface bends the flexible endoscope arm to the right (3 s in Fig. 13.6). A blood vessel on the bladder sidewall is clearly visualized by a 1x1mm camera afixed at the end of the endoscope arm.

**Discussion and Outlook** We developed a novel functional surface that is wirelessly powered by ultrasound in water and permits the actuation of miniaturized, multiple degrees-of-freedom endoscopes and robot-arms. The functional surfaces consist of arrays of micro-bubbles. Streaming can be induced with standard ultrasound sources from the outside, which ensures wireless operation. The surfaces can readily be integrated in devices. In contrast to bulk actuators, no specific consideration is needed for the mechanical coupling and integration of the surface actuators, which also minimizes the volume and weight required for actuation, thus making them particularly well-suited for miniaturization. Moreover, if the surface actuators are made of a thin polymeric material, they can also be soft and compliant with a complex three-dimensional object. The adoption of micro-structured surfaces as wireless actuators opens new possibilities in the development of miniaturized devices and tools for fluidic environments that are accessible by low intensity ultrasound fields. The miniaturized endoscope arm we developed promises a significant reduction of surgical trauma and thus reduces unwanted complications for the patient.

A patent has been filed and the work has been published in Applied Physics Letters [662], Advanced Materials Interfaces (cover article) [654], as well as ACS Applied Materials & Interfaces [651]. The work was highlighted by Applied Physics Letters and by the media.

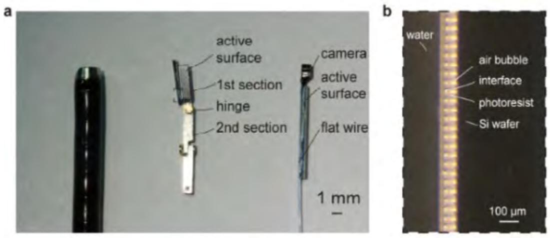


Figure 13.7: (a) Comparison of a commercial endoscope (left) with our ultrasonic mini-arm (middle) and the flexible endoscope (right). (b) Microscope image of the active surface under water, which consists of an array of micro-bubbles.

More information: https://pf.is.mpg.de/project/ultrasonic-actuator

## Nanorobotic system for propulsion through biological tissue

*Tian Qiu, Zhiguang Wu, Jonas Troll, Hyeon-Ho Jeong, Peer Fischer*

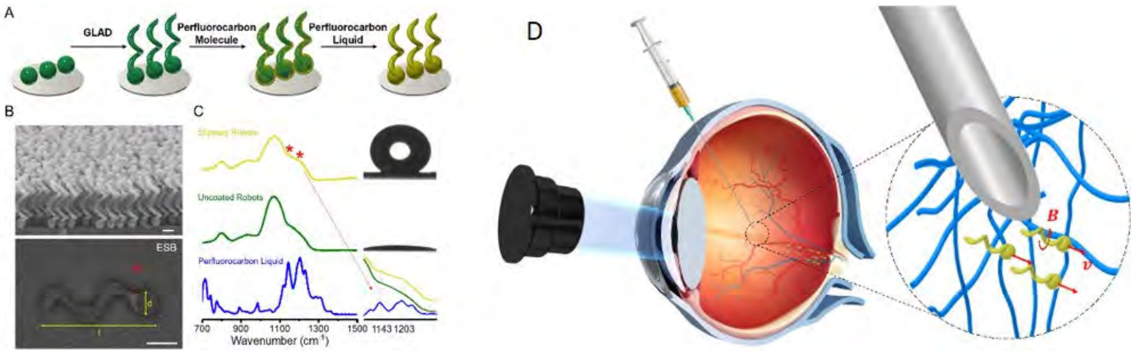


Figure 13.8: Schematic of the targeted delivery procedure used for the slippery nanopropellers: (Left) Fabrication of the perfluorocarbon coated "slippery" nanopropellers. (A) Schematic of the fabrication process. (B) SEM images of the nanopropellers. Scale bars, 500 nm. (C) FTIR of the nanopropellers without coating and with coating. The contact angles of the wafer with coated helices (top) and uncoated helices (bottom). (Right) Injection of the nanopropellers into the vitreous humour of the eye (D); the magnetically-driven long-range propulsion in the vitreous towards the retina; The observation of the nanopropellers at the target region on the retina by OCT.

**Introduction and problem** Targeted drug delivery through dense biological tissue is very challenging. Traditional delivery methods rely on the random, passive diffusion of molecules, which is often a slow and ineffective process. The systemic application of pharmaceuticals requires larger drug volumes and can negatively impact healthy tissue. It has therefore always been a dream to develop nano-vehicles that can actively move through tissue and precisely target a region. Actively propelled micro- and nanoparticles provide a novel potential pathway for the targeted delivery of drugs. A few synthetic microparticle systems have been developed that can be actively propelled in biological media. These have predominately focused on propulsion in liquids or fluid-filled cavities in the body, and include structures that can move in the fluids of the stomach, or that have been propelled through blood *in vitro*. However, the propulsion of micro- and nanoparticles through dense biological tissues has not been demonstrated before. Here, we use the vitreous humour of the eye as a model tissue, and demonstrate that sufficiently small structures that are wirelessly actuated are able to efficiently move through the vitreous humour over cm-distances, when they are equipped with a suitable surface chemistry. Without a coating most particles 'become stuck', as biological tissue contains hydrophobic and hydrophilic as well as positively and negatively charged regions. Adhesion has therefore been a major challenge.

**Our approach** Our research group has developed novel 3D nanofabrication tools that allow us to grow billions of cork-screw like nano and microparticles from a wide range of materials[7]. We have in our previous work shown that the nanopropellers, which mimic the locomotion of flagella, can be magnetically actuated and precisely steered[8]. Recently, we have shown that the surface of the propellers can be functionalized with the same enzymes the bacterium *H. pylori* uses for propulsion through the otherwise impenetrable mucus of the stomach [9]. We have also made and actuated nanopropellers that have an outer diameter so small (< 100nm) that they can easily move through the gaps in a macromolecular polymeric network[10]. However, very small structures can potentially carry only a small load. We therefore fabricate helical propellers with a diameter that is as large as the mesh size of the biopolymeric network of the vitreous, and use a novel perfluorocarbon surface coating to min-

[7] A.G. Mark, J.G. Gibbs, T.-C. Lee, P. Fischer, Nature Materials 12, 802, (2013). Int. media coverage, ISI highly cited.
[8] A. Ghosh, and P. Fischer, Nano Letters 9 (6), 2243, (2009). Highlighted in Nature, ISI highly cited.
[9] D. Walker, H.-H- Joeng, B. Käsdorf, O. Lieleg, P. Fischer, Science Advances 1 (11), e1500501 (2015)
[10] D. Schamel, A.G. Mark, J.G. Gibb, C. Miksch, K.I. Morozov, A.M. Leshansky, P. Fischer, ACS Nano 8, 8794–8801, (2014). Featured cover article & int. media

imize the interaction of the propellers with the vitreous (Fig. 13.8).

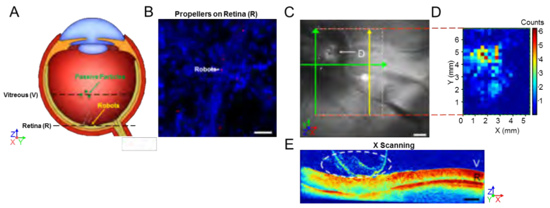

Figure 13.9: Movement of the slippery nanopropellers in the complete porcine eye. (A) Schematic illustrating the movement of the nanopropellers in the vitreous towards the retina. (B) Fluorescent microscopy image of the incised retina. Nanopropellers (red fluorescent) are observed on the retina. Scale bar, 20 µm. (C) Image of the retina near the optic disc (labeled as D). Scale bar, 1 mm. (D) 3D reconstruction of the OCT scans showing the nanopropellers concentrate near the optic disc. (E) OCT image of X scanning. Scale bars are 500 µm in (C) and (E).

The coating is inspired by a liquid layer found on the carnivorous Nepenthes pitcher plant, which presents a slippery surface on the peristome to catch insects. The non-toxic silicone oil and fluorocarbon coatings are medically approved. We expect that the intravitreal injection, long-range propulsion, and non-invasive monitoring with a clinically approved instrument will enable future ophthalmological therapies (Fig. 13.8).

**Results** The fabrication of the slippery nanopropellers consists of two main steps: the preparation of helical microstructures and their coating. The nanopropellers were fabricated using physical vapor shadow growth via glancing angle deposition (GLAD)[7, 8]. The surface functionalization consisted of a gas phase deposition of a perfluorosilane, and the subsequent fusion with a slippery perfluorocarbon liquid layer. Under wireless actuation with an external magnetic field, the coated propellers not only show controllable propulsion, but can also be driven as a large swarm over centimeter-distances through the eye ball and reach the retina within 30 min. To investigate the propulsion of nanopropellers in an intact eye (Fig. 13.9A), the propellers were loaded with fluorescent nanodiamonds. Red fluorescence was observed on the expected target region on the retina (Fig. 13.9B). We showed that standard clinical optical coherence tomography (OCT) can bes used for the observation of the nanopropellers in a non-invasive and label-free manner. The OCT scans of the vitreous-retina boundary reveal the swarm of nanopropellers 13.9E-D). The results indicate that the propulsion of the micropropellers shows an excellent directionality under the control of the external magnetic field.

**Discussion and Outlook** We demonstrate that two major criteria should be fulfilled for the propulsion in biological media: 1) matched size of the particles to the macromolecular network – which calls for sub µm-sized structures, and 2) a means to reduce the interaction between the propellers and the biopolymer network. We have succeeded in demonstrating both. Most biological tissues are porous, for example, the extracellular matrix has a pore size of approximately 100-300 nm, the mucus has a pore size of 100 nm to several micrometers. The propulsion mechanism reported here is general and can therefore be applied to various porous biological media.

This work has just been published in Science Advances [639], and the work has been highlighted by Science News (including Science Video) and other international media. The work also led to several patents, which have been licensed a start-up company in the US.

More information: https://pf.is.mpg.de/news/nanorobots-propel-through-the-eye

## Are all propellers chiral?

*Johannes Sachs, Peer Fischer*

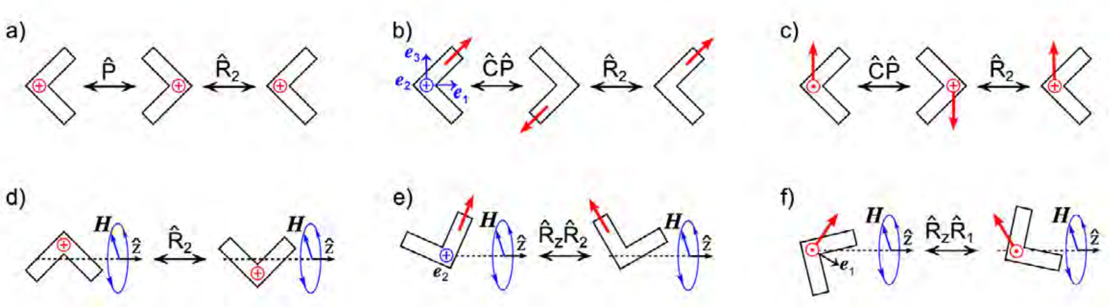


Figure 13.10: Schematic diagrams illustrating symmetries of a planar V-shaped object carrying a magnetic dipole moment (red) of different orientations with respect to parity, $\widehat{P}$, and charge conjugation, $\widehat{C}$. (a) $\widehat{P}$- (and $\widehat{C}\widehat{P}$-) even object; (b) $\widehat{C}\widehat{P}$-even (and $\widehat{P}$-odd) object; (c) $\widehat{C}\widehat{P}$- (and $\widehat{P}$-) odd object. Principal axes of rotation $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ are shown in (b). (d-f) diagrams illustrating the pairs of solutions corresponding to (a-c).

**Introduction and problem** Swimming at the micro scale is subject to low Reynolds (Re) number hydrodynamics. The scallop theorem dictates that any symmetric (time-reciprocal) motion cannot result in any net translation. Strictly, this symmetry requirement is only true in Newtonian fluids. We were the first to show that the scallop theorem can be overcome in non-Newtonian fluids [11]. Nevertheless, all swimming microorganisms, such as bacteria or ciliates, have devised efficient schemes to break reciprocity in their motion and hence satisfy the scallop theorem. In artificial microsystems, however, one can apply external forces and torques, and then the scallop theorem no longer holds. For instance, a particularly powerful propulsion scheme at low Re is to rotate a propeller or drill using external fields and thus cause translation. This raises the question if a propeller has to be chiral to exhibit efficient rotation-translation coupling? Is it possible for highly symmetrical and achiral objects to propel, or does one need a corkscrew-propeller, which is commonly regarded to be the most efficient propeller? Surprisingly this fundamental question has thus far not been answered. The motivation is to find shapes that are simple to fabricate, yet propel well.

**Our approach** Together with our collaborators at the Technion in Israel, we have just published a comprehensive analysis of the symmetries that underlie rotation-translation coupling [637]. We fabricated V-shaped structures that possess a magnetic moment. The V-shape has two mutually orthogonal symmetry planes and is therefore highly symmetrical. We developed a cm-scale model as well as micron-sized structures (see Fig. 13.11a, b). The macroscopic model allowed us to precisely fix the orientation of the dipole moment with respect to the body. This is essential as the dipole orientation has important consequences on the symmetry and the associated propulsion gait of the object, a fact that has been overlooked in previous experimental studies. The micro-scale structure was fabricated using our unique glancing angle deposition (GLAD) technique. We also fabricated V-shaped structures that are electrically polarizable (see Fig. 13.11c) and could therefore be driven with an electric instead of a magnetic field. This allowed us to identify the differences between the underlying symmetries and propulsion characteristics.

**Results** The first surprising result is that a highly symmetric V-shape can propel. Our rigorous symmetry analysis then shows – and this is the second surprising result – that a V-shape can be chiral! The latter is found when one correctly accounts for the symmetry of the magnetic moment in the structure. This has so far been overlooked in the literature. A truly remarkable result

[11] T. Qiu, T.-C. Lee, A. G. Mark, K. I. Morozov, R. Münster, O. Mierka, S. Turek, A. M. Leshansky, P. Fischer, Nat. Comm. 5, 5119 (2014)

is that when an electric moment is used to rotate the structure, then the V-shape can propel even though it is achiral! For rotation-translation coupling to occur, the V-shape must precess around the external fields' rotation axis. This is only possible, when the dipole moment is angled such that it does not rotate around one of the principal axes of rotation. The V-shapes propulsion gait is critically dependent on the orientation of the dipole moment. We could show that the V-shape can propel like a corkscrew.

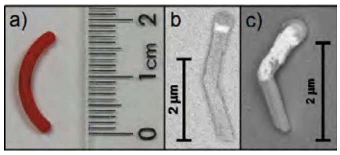


Figure 13.11: A cm-sized V-shape was fabricated via 3d printing (a), whereas micron-sized V-shapes where grown with the GLAD technique (c, b). These structures are made of $SiO_2$ and had either a ferromagnetic Nickel segment (white part in b) or one arm coated with gold (white part in c). The latter one could therefore be driven with an electric instead of a magnetic field.

**Discussion and Outlook** We describe the symmetry that is needed to propel a structure (e.g. in a simple microrobotic system) by rotation-translation coupling. We show that chirality is not a prerequisite and highly symmetrical shapes can be propulsive too. We realized the first truly achiral propeller. This is useful as it allows one to identify simple shapes that can be used as propellers, including shapes that do not have to be chiral. We believe that this is important, because microfabrication of structures with high symmetry (including 2.5 D shapes) is generally easier than having to fabricate 3D chiral shapes. The work has been published in the journal Physical Review E [637].

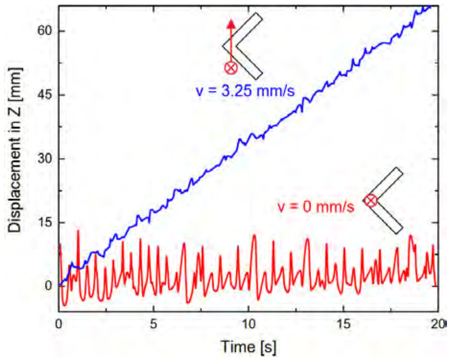


Figure 13.12: Displacement upon rotation in an external field plotted over time for two cm-sized V-shapes with different orientations of their magnetic dipole moment. While the upper object resembles the same propulsion gait as a corkscrew propeller (blue) the lower structure (red) rotates without notable net translational motion.

More information: https://pf.is.mpg.de/project/are-all-propellers-chiral

## Can enzymes swim and show chemotaxis?

*Jan-Philipp Guenther, Peer Fischer*

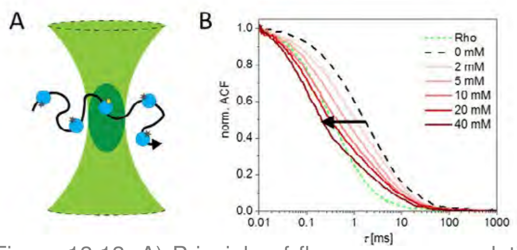


Figure 13.13: A) Principle of fluorescence correlation spectroscopy; the transitions of a single enzyme (blue) through a confocal volume (green) is observed due to the fluorescence of the dye (asterisk); B) for the enzyme alkaline phosphatase an apparent decrease in diffusion can be observed with increasing concentrations of its substrate[12]; However, we showed that an apparent decrease in transition time can also be caused by fluorescence quenching [642].

**Introduction and problem** Chemical reactions on the surface of microparticles can cause local concentration gradients of substrate or product molecules around the colloid which in turn give rise to fluid flows around the particle. Momentum conservation then requires that the particle moves in the opposite direction to these flows. Energy is consumed, which powers the self-propelled particles. The particles are therefore a model system for “active matter” and are also known as chemical motors, which can autonomously swim and self-organize [648, 655]. Whether these or related mechanisms can also translate to the molecular scale is, however, not yet fully understood. The question whether the reaction of enzymes can cause their self-propulsion has potentially important implications for biochemistry. Similarly, the question whether these systems can show chemotaxis has far reaching implications as this would suggest a radically new mechanism how one may realize drug delivery, e.g. by means of active self-propelled carriers that can steer themselves to the location they are needed [644].

**Our approach** Previous reports on enzyme motion used conventional fluorescence correlations spectroscopy (FCS) to determine the diffusion of fluorescently labeled biomolecules. Together with our collaborator M. Boersch we have performed precise measurements to examine if enzymes can self-propel. We used a more advanced optical setup to elucidate how the fluorescence dye interacts with its environment and the enzyme it is coupled to. We also introduced in situ fluorescence lifetime analysis via time-correlated photon counting (TCPC), and we measured the anisotropy of the fluorescence signal. The former provides information regarding the quantum states of the dye and the latter can provide information about its binding to the enzyme. Additionally we considered serval biochemical factors that affect FCS signals.

**Results** We first focused on the enzyme alkaline phosphatase for which the highest diffusion enhancement has been reported[12]. Enzyme assays were used to monitor the enzymes’ activity throughout the studies. Importantly, we could show that a significant decrease in the transition time through the confocal volume of the FCS setup occurs (Fig. 13.13B), but the in situ lifetime measurements suggest that quenching of the fluorophore by the substrate of alkaline phosphate is the real reason for the change in the diffusion constant. Quenching adds additional fluctuations to the fluorescence signal, which have been misinterpreted by others as a sign for enhanced diffusion. Through careful control experiments we were able to completely rule out any diffusion enhancement by alkaline phosphatase related to its activity [642].

**Discussion and Outlook** The insight that enzymes do not exhibit activity-enhanced diffusion provides size limits for active matter and self-propelled swimmers. Interestingly, many papers claim that enzymes self-propel and are a form of active matter[12,13]. Our work therefore suggests that these experiments must be revisited. Our ongoing studies using Nuclear Magnetic Resonance lend further support to our findings that previous experiments are in error. We are currently investigating how small nanostructures can be and still efficiently self-propel. Additionally, we are designing autonomous nanomotors that can also chemotax.

More information: https://dx.doi.org/10.1021/acs.accounts.8b00276

[12] C. Bustamante et al., Nature 517, 227, (2015).
[13] A. Sen et al., Nano Lett. 17, 4807 (2017); A. Sen et al., Nat. Chem. 10, 311 (2018); S. Granick et al., PNAS 115, 14 (2018).

## 13.3 Equipment

The group operates several custom thin-film deposition setups with unique computerised substrate manipulation capabilities to grow 3D designer nano-structures at the wafer-scale. An optical and spectroscopy lab is used to observe single molecules, nano-structures, and enzymes. Electron microscopy and cell experiments are performed in house in collaboration. The group has custom magnetic manipulation instruments, as well as a fully equipped chemistry lab for synthesis and functionalisation.

**Physics and nano-fabrication:** High vacuum electron beam setups for glancing angle deposition (three setups, one dual e-beam, one e-beam combined with Knudsen cells, one home-built), plasma chamber.

**Chemistry and material characterisation:** Organic chemistry lab, $CO_2$ laser cutter, 3D printer, Langmuir Blodgett trough and deposition, flow field fractionation and chromatography, contact angle measurement, circular dichroism spectrometer, nanoparticle analysis light scattering and Zeta-potential, rheometer.

**Optics and spectroscopy:** Two inverted microscopes (TIRF, fluorescence, and dual imaging setups), spectroscopy, and optics setups, including lasers, digital micro-mirror array, polarisation modulators Pockels cells and PEMs, cameras, precision stages, laser tweezer setup for femtonewton force measurements.

**Electromagnetics:** 1.8 Tesla electromagnet, custom amplifiers, ultrasound generators, network analyser, precision electronics, and a 12 core workstation for simulation work.

## 13.4 Awards & Honors

### 2018

**Kai Melde** won the "Excellent Communication Award" at the International Conference on Ultrasonics 2018 in Caparica Portugal.

**Peer Fischer** wins an ERC Advanced Grant.

### 2017

**Eunjin Choi** and **Hyeon-Ho Jeong**, respectively, received the 2017 scholarship awards from the Korean Scientists and Engineers Association in the FRG (VeKNI e.V)

**Peer Fischer** is named Steinhofer lecturer at the University of Freiburg and he receives the associated prize of the Steinhofer-Foundation.

**Tian Qiu, Stefano Palagi, and Fabian Adams** who won the Best Oral Paper Award, at the 10th Hamlyn Symposium on Medical Robotics, London 2017.

**Jan-Philipp Günther** and **Vincent Mauricio Kadiri** won the first poster prize on the NanoBioMater 2017 Summer School International Conference for our poster with the title "Nanorheology and Nanopropellers in Biological Fluids".

**Hyeon-Ho Jeong** was awarded this year's Graduate Student Award at the 2017 E-MRS Spring Meeting in Strasbourg.

### 2016

**Peer Fischer** has received the World Technology Award 2016. Together with his group, he has developed new 3D nanofabrication methods and nanorobots, made the first reciprocal microswimmer, and realized the first swimming soft microrobot that moves using only body shape changes.

**Hyeon-Ho Jeong** was awarded the MRS Gold Graduate Student Award for a presentation on Chiral Plasmonic Nanosensors at the 2016 MRS Fall Meeting in Boston.

**Peer Fischer** has been elected a Fellow of the Royal Society of Chemistry (FRSC).

**Tian Qiu** has been granted the "National Award for Outstanding Self-financed Chinese Students Abroad 2015" by the China Scholarship Council (April, 2016).

## 13.5 Research group leader: Peer Fischer

### Prof. Dr. Peer Fischer

| | |
|---|---|
| 06/2013 – present | Professor of Physical Chemistry, Univ. of Stuttgart |
| 09/2011 – present | Max Planck Institute for Intelligent Systems, Head, Micro Nano and Molecular Systems Lab |
| 2009 – 2011 | Fraunhofer Society, Independent Attract Group Leader |
| 2004 – 2009 | Harvard University, The Rowland Institute Principal Investigator and Rowland Fellow |
| 2000 – 2004 | Cornell University, NATO Postdoctoral Fellow and Postdoctoral Associate, Department of Chemistry and School of Applied & Engineering Physics |

### Qualifications

| | |
|---|---|
| 1995 – 1999 | Ph.D., Dep. of Chemistry, University of Cambridge. Supervisor: Prof. A. D. Buckingham, FRS CBE |
| 1992 – 1995 | B.Sc. Physics, (First Class Honours), Imperial College London |

### Academic Offers

| | |
|---|---|
| 2013 | Professor of Physical Chemistry (W3 Professorship), Univ. of Stuttgart |
| 2011 | Max Planck Institute for Intelligent Systems, independent W2 with tenure |
| 2011 | Assistant Professor Physical Chemistry, Boston University, (declined) |
| 2009 | Assistant Professor Physical Chemistry, Yale University, (declined) |
| 2009 | Assistant Professor Physical Chemistry, University of Oregon, (declined) |

### Awards and Fellowships

| | |
|---|---|
| 2018 | ERC Advanced Grant (start date 2/2019; € 2,500,000) |
| 2017 | Steinhofer Lecturer and Prize, Steinhhofer Foundation, Dep. of Chemistry, Univ. of Freiburg, Awarded for "fundamental contributions to the controlled 3D fabrication of artificial nanostructures and their biomedical application". |
| 2016 | Winner, World Technology Award, (IT Hardware), Los Angeles |
| 2016 | Fellow of the Royal Society of Chemistry (FRSC). |
| 2012 – 2018 | ERC Starting Grant (as consolidator) “Chiral Nanostructured Surfaces and Colloidal Microbots” (€ 1,480,000) |
| 2009 | Attract Award, Fraunhofer Society, The *Fraunhofer Attract* grant is the excellence stipend programme of Fraunhofer. *Fraunhofer Attract* invites "outstanding researchers to develop their ideas towards innovation". (5 yrs., approx. € 2,500,000, stayed until 9/2011). |
| 2004 | Rowland Research Fellow, Harvard University. "Fellows are selected for their scientific achievement, their creativity, their resourcefulness as experimentalist". (5 yrs., approx € 1.200,000) |
| 2000 – 2002 | DAAD (NATO) Postdoctoral Fellowship |

## Professional Activities (selected, recent)

| | |
|---|---|
| 2018 | Guest Editor, Accounts of Chemical Research, Special Issue on *Fundamental Aspects of Self-Powered Nano- and Micromotors* (with A. Sen and A. Balazs). |
| 2018 – present | Editorial Advisory Board: MDPI Robotics. |
| 2018 | Guest Editor, Advanced Functional Materials, Special Issue on *Micro- and Nanomachines on the Move* (with M. Pumera and J. Wang) |
| 2017 | Guest Editor, MDPI Micromachines, special issue on *Microswimmers* |
| 2016 – present | Editorial board: (Science/AAAS) *Science Robotics* |
| 2015 – present | Steering committee Int. Conf. on Manipulation, Automation and Robotics at Small Scales |
| 2013 – present | University of Stuttgart research strategy board member |
| 2010 – present | Editorial Board: Wiley journal Chirality |
| 2011 – present | Scientific advisory committee: International conference on chiroptical spectroscopy |

## Lecture Courses (selected, recent)

| | |
|---|---|
| 2018 – 2019 | Univ. of Stuttgart, Master degree course, special lecture *Nanosystems*, (3 lectures/week). |
| 2017 – 2019 | Univ. of Stuttgart: Physical Chemistry, Introduction to Chemistry, Bachelor degree course (in German). WS17 and WS18 (4 lectures/week & exercises, PC/AC/OC block lecture). |
| 2015 – 2018 | Univ. of Stuttgart: Master degree course; Surfaces and Colloids (multiple lecturers, one unit with lab experiments, 2 lectures/week). WS15 and WS17. |
| 2015 – 2017 | Univ. of Stuttgart: Physical Chemistry, PCII; Atoms, Molecules and Spectroscopy, Bachelor/Master degree course (in German). WS15 and WS16 (4 lectures/week & exercises). |
| 2018 | Graduate lectures: Summer school, *Current topics in active fluids*, Imperial College London, 16-20 July. |

## Invited Talks (selected, recent)

| | |
|---|---|
| 2018 | Plenary lecture, *Bioinspired Magnetic Systems*, BIMS2018, Exeter, UK |
| 2018 | Keynote lecture, *Nanocolloids on the move*, 92nd ACS Colloids and Surface Science Symposium in State College, Pennsylvania, USA |
| 2018 | Plenary lecture, *Physical Vapor Deposition at the Nanoscale*, Glancing Angle Deposition Films and their Application, IIT Delhi, India |
| 2017 | Keynote lecture, *Synthetic microswimmers*, summer school, Dresden |
| 2017 | Keynote lecture, *Chiral plasmonic nanostructures*, META 2017, 8th Int. Conf. on Metamaterials, Incheon-Seoul, South Korea |
| 2017 | Keynote lecture, *Multifunctional chiral nanostructures*, Int. Conf. on Chiroptical Spectroscopy CD2017, Rennes, France |
| 2017 | Keynote lecture, *Multifunctional nanohelices*, Chiroptics 2017, Munich |
| 2016 | Keynote lecture, *Untethered Micro- and Nano-robots for Potential Biomedical Applications*, The Hamlyn Symposium on Medical Robotics, London, England |
| 2016 | Keynote lecture, *Micro- and Nanorobots: Fabrication, actuation and control*, 47th Int. Symposium on Robotics, Munich |
| 2016 | Keynote lecture, Plenary talk, *Active micro- and nanosystems*, Institute of Microstructure Technology (IMT), KIT Karlsruhe |

# 14 PHYSICAL REASONING AND MANIPULATION

## 14.1 Research Overview

Fig. 14.1 illustrates hallmarks of what is considered intelligent behavior in animals: *Betty the Crow* (Fig. 14.1 top-left) demonstrated the ability to use a sequence of hooks to retrieve a piece of food[1]. Koehler's apes stacked crates to reach a hanging bunch of bananas (bottom-left)[2]. Humans perform such sequential manipulation planning tasks naturally and flexibly (bottom-right). As these behaviors have not been trained or learned from massive data—as would be needed with the current methodologies of Reinforcement Learning—the question is what prior models and inference mechanisms allow for such systems to conceive and execute such solutions. To date, there are no satisfying theories or methods to model or reproduce such capabilities in artificial systems. In the project we follow up on the idea of an **intuitive physics engine**[3] that has been proposed in the cognitive sciences as a foundation to enable general purpose physical reasoning. Such a mental physics engine was proposed to internally simulate the outcome of an action or tool manipulation[4]. Dedicated neural architectures near the motor cortex were proposed as implementation of this capability[5].

However, a physics engine is a pure forward model. Brute-force forward sampling (Monte-Carlo) in principle allows for inference and goal-directed reasoning, and was the basis for the mentioned work on intuitive physics[6]. However, from our understanding of the complexities of (and relations between) control theory, probabilistic inference, motion planning, and AI task planning[7] we know that pure Monte-Carlo with a physics engine cannot scale.

In this project we address how general purpose physical reasoning and manipulation can be enabled, assuming we had a (potentially stochastic) physical simulation of our environment available. We approach this question from the perspective of recent advances in AI, ML, and robotics, in particular the embedding of differentiable reasoning mechanisms in computation graphs (neural networks), as well as recent advances in optimization (ADMM, Alternating Direction Multiplier Method) and AI planning (learning to plan). We structure our research program around the following topics, as detailed in the work program:

1. Differentiable Physics $\longleftrightarrow$ Invertible Physics?
2. Symbolic Structure in the Inverse Problem of Physics
3. Learning to Think
4. Probabilistic Modeling and Inference with Physics
5. Reactive Execution under Uncertainty

[1] J. Wimpenny, A. Weir, L. Clayton, C. Rutz, A. Kacelnik. Cognitive processes associated with sequential tool use in new caledonian crows. *PLoS One, 4(8)*, 2009.

[2] M. Koehler. Cntelligenzprüfungen am Menschenaffen. *Springer, Berlin (3rd edition, 1973)*, 1917.

[3] P. Battaglia, J. Hamrick, J. Tenenbaum. Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences, 110(45):18327–18332*, 2013.

[4] F. Osiurak, A. Badets. Tool use and affordance: Manipulation-based versus reasoning-based approaches. *Psychological review, 123(5):534*, 2016.

[5] J. Fischer, J. Mikhael, J. Tenenbaum, N. Kanwisher. Functional neuroanatomy of intuitive physical inference. *Proceedings of the national academy of sciences, 113(34):E5072–E5081*, 2016.

[6] P. Battaglia, J. Hamrick, J. Tenenbaum. Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences, 110(45):18327–18332*, 2013.

[7] S. LaValle. Planning Algorithms. *Cambridge university press*, 2006.

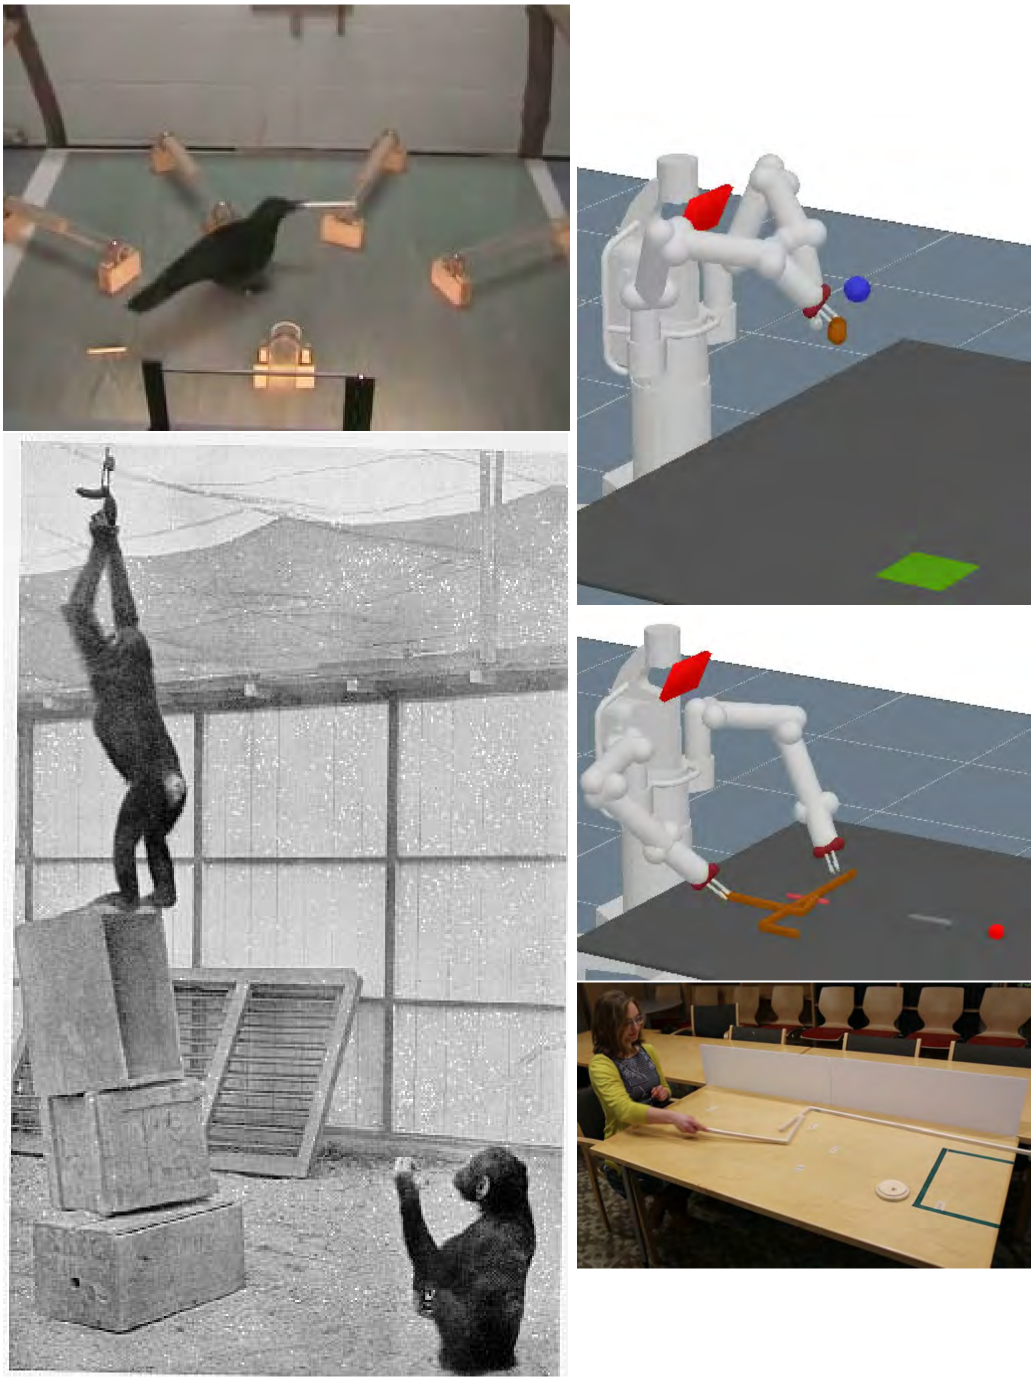

Figure 14.1: Figure 1: Left: A crow solving a sequential hook problem (Wimpenny et al., 2009); bottom: classical intelligence test with apes (Köhler, 1917). Right: preliminary work (submitted, Toussaint et al. (2018)) on using a stick to hit a flying blue ball to get it into the green area, and using a hook to get a second hook to reach for the red ball; bottom: a human subject on the same task.

The demonstration goal of this research program is to **reproduce the mentioned hallmarks of animal behavior on a real robotic system**, where only prior knowledge equivalent to a physics simulator is available to the system. This includes experimental setups such as Betty The Crow's, and a wide spectrum of variations of problems that involve sequencing stable and dynamic interactions with objects, e.g. hitting a ball to eventually get it into a target area. In collaboration with researchers from the Cognitive Sciences (planned with Josh Tenenbaum, continuing our joint prior work, see below) we will continuously consider to what degree our developed methods are models of animal and human intelligent behavior in similar tasks.

**Own prior work** We introduced[8] for the first time a non-linear programming (optimization) formulation of TAMP, which we extended and proposed a better solver[9] for. Using this we could plan cooperative manipulation sequences for a human and robot. In this prior work, the particular choice of how the categorical structure of the problem is represented (using logic) is still somewhat ad hoc. We aim to resolve this with the proposed research, deriving this structure purely from the fundamental problem of "inverting physics".

We currently cooperate with Josh Tenenbaum's group in investigating models of human manipulation strategies[10]. While the proposed project focuses on the development of a better fundamental understanding of the problem and novel algorithms for general purpose physical reasoning, we also aim to continue this cooperation to compare our results with human capabilities.

In our lab in Stuttgart we have extensive prior work on manipulation learning[11,12], learning from demonstration[13,14,15,16], reasoning and manipulation under uncertainty[17,18,19], and generally the application of probabilistic inference and constrained optimization methods in robotics. A recurring question raised by our methods that use a higher-level representation is the origin and grounding of the used abstractions. We previously invested in what could be called abstraction learning[20,21]. The proposed project fundamentally reconsiders what are appropriate abstractions for reasoning and aims to derive them from first principles of the underlying problem.

---

8 M. Toussaint. Logic-geometric programming: An optimization-based approach to combined task and motion planning. *Proc. of the Int. Joint Conf. on Artificial Intelligence*, 2015.

9 M. Toussaint, M. Lopes. Multi-bound tree search for logic-geometric programming in cooperative manipulation domains. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2017.

10 I. Yildirim, T. Gerstenberg, B. Saeed, M. Toussaint, J. Tenenbaum. Physical problem solving: Joint planning with symbolic, geometric, and dynamic constraints. *Thirty-Night Annual Conf. of the Cognitive Science Society*, 2017.

11 P. Englert, M. Toussaint. Learning manipulation skills from a single demonstration. *International Journal of Robotics Research*, 2017.

12 P. Englert, M. Toussaint. Combined optimization and reinforcement learning for manipulations skills. *Proc. of Robotics: Science and Systems*, 2016.

13 T. Munzer, M. Toussaint, M. Lopes. Preference learning on the execution of collaborative human-robot tasks. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2017.

14 M. Toussaint, T. Munzer, Y. Mollard, L. Wu, N. Vien, et al. Relational activity processes for modeling concurrent cooperation. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2016.

15 P. Englert, M. Toussaint. Inverse kkt - learning cost functions of manipulation tasks from demonstrations. *Proceedings of the International Symposium of Robotics Research*, 2015.

16 Y. Mollard, T. Munzer, A. Baisero, M. Toussaint, M. Lopes. Robot programming from demonstration, feedback and transfer. *Proc. of the Int. Conf. on Intelligent Robots and Systems*, 2015.

17 M. Toussaint, N. Ratliff, J. Bohg, L. Righetti, P. Englert, et al. Dual execution of optimized contact interaction trajectories. *Proc. of the Int. Conf. on Intelligent Robots and Systems*, 2014.

18 D. Driess, P. Englert, M. Toussaint. Active learning with query paths for tactile object shape exploration. *Proc. of the Int. Conf. on Intelligent Robots and Systems*, 2016.

19 D. Driess, P. Englert, M. Toussaint. Constrained bayesian optimization of combined interaction force/task space controllers for manipulations. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2017.

20 J. Kulick, M. Toussaint, T. Lang, M. Lopes. Active learning for teaching a robot grounded relational symbols. *In Proc. of the Int. Joint Conf. on Artificial Intelligence*, 2013.

21 N. Jetchev, T. Lang, M. Toussaint. Learning grounded relational symbols from continuous data for abstract reasoning. *ICRA Workshop on Autonomous Learning*, 2013.

## 14.2 Research group leader: Marc Toussaint

### Prof. Dr. rer. nat. Marc Toussaint

| | |
|---|---|
| 2003 | PhD at the Institute for Neuroinformatics, Ruhr-University Bochum (Prof. Werner von Seelen) |
| 1994 – 1999 | University of Cologne, Physics (Diploma 1999) and Mathematics (Pre-diploma) |

### Professional Career

| | |
|---|---|
| 2018 – present | Max Planck Fellow at the MPI for Intelligent Systems |
| 2012 – present | Professor (W3), head of the Machine Learning and Robotics Lab, University of Stuttgart |
| 2017 – 2018 | Research Scholar at MIT |
| 2011 – 2018 | Coordinator of the DFG Priority Programme (Schwerpunktprogram) SPP 1527 Autonomous Learning |
| 2010 – 2012 | Assistant Professor (W1) for Machine Learning and Robotics at FU Berlin |
| 2010 – 2012 | Co-Speaker of the DFG graduate school Sensory Computation in Neural Systems |
| 2007 – 2012 | Head of the independent research group Machine Learning and Robotics (Emmy Noether Programme) before 10/2010 at TU Berlin, later at FU Berlin |
| 2006 – 2007 | Honda Research Institute: Guest scientist at the robotics department of HRI Europe Offenbach |
| 2004 – 2006 | University of Edinburgh: Postdoc (Emmy-Noether stipendiary) at the Institute for Adaptive and Neural Computation (Prof. Chris Williams) und the Institute for Perception Action, and Behavior (Prof. Sethu Vijayakumar) |
| 2000 – 2004 | Ruhr-University Bochum: Research assistant |

# 15 PHYSICS FOR INFERENCE AND OPTIMIZATION

## 15.1 Research Overview

The research goal of the Physics for Inference and Optimization group is *understanding, optimizing and predicting* relations between the microscopic and macroscopic properties of complex *large-scale interacting systems.* We pursue this agenda by addressing application-oriented problems of *inference and optimization on networks* via developing models and algorithms derived from **statistical physics principles**.

The two main motivations behind this interest are the idea that there is a pressing need for theory to be grounded in concrete applications in order to solve relevant scientific problems in rigorous ways, improving both methodological and domain-specific knowledge.

In addition, in recent years statistical physics has been able to provide new insights and novel approaches to problems in computer science.

Our approach is to address this problem with two main research questions that tackle this issue under different angles and together should provide a cohesive and coherent analysis of the problem:

- How can we exploit the *distributed* character of complex interacting systems to perform network optimization?
- How can we understand the mechanism leading to *large-scale pattern emergence* and thus perform robust and scalable inference?

The objective is that by answering these questions one can provide a comprehensive analysis of the bigger problem of understanding large-interacting systems in their different aspects.

Our research reflects this interdisciplinary approach and considers:

- *Developing theoretical models* capable of describing the mechanism driving the optimization or inference problems, assessing their properties and limitations.
- *Building efficient and scalable algorithms* to apply them leveraging various statistical physics tools such as message-passing techniques, variational inference or matrix product factorization.
- *Addressing data-rich open problems* involving domain experts from other disciplines. This is stimulated by the ongoing interactions with more empirically-oriented scientists from various disciplines.

## 15.2 Selected Research Projects

## Dealing with missing and noisy data in inference on networks

Often real networks are not fully observed, either because of limited experimental capacity to collect the information or because of the magnitude of the system's size. Similarly, data collection can be noisy, as when edges are mistakenly reported or when their magnitude is not clear enough. This leads on one hand to the challenge of estimating hidden parameters, labels, community memberships or centrality measures from an observed dataset subsampled from some larger and noisy network. On the other hand, given the biasing effect of how the network information was collected, an open challenge is to develop theoretical ideas into feasible data collection algorithms that are capable of correcting for these effects.

More information: https://pio.is.mpg.de/project/dealing-with-missing-and-noisy-data-in-inference-on-networks

## Analyzing complex energy landscape and optimal trajectories of routing optimization

*Caterina De Bacco*

One of the main challenges in solving routing optimization problems is the complex energy landscape of the cost function which is rich in local optima. Finding how the solutions' configurations corresponding to the different local optima are related can give important insights on our limits/hardness to solve the optimization problem. A similar problem affects other combinatorial optimization problems or energy minima of disordered systems as spin-glasses and phase diagrams have been analyzed in those contexts. We investigate this problem by adopting an approach that combines recent insights from statistical physics and a novel methodology developed in optimal transport theory that maps the problem of solving a routing optimization problem into the easier one of solving a dynamical system of Monge-Kantorovich differential equations. This will allow to study the properties and topologies of optimal path trajectories both at equilibrium and far from it and thus find if networks corresponding to the same energy level display ultrametricity and how are networks corresponding to different energy levels related to each other.

More information: https://pio.is.mpg.de/project/analyzing-complex-energy-landscape-and-optimal-trajectories-of-routing-optimization

## Analyzing complex energy landscapes for inference on networks

*Caterina De Bacco*

A relevant problem when dealing with community detection of real-world data is how to test the partition discovered by a model. Generative models allow to define a likelihood for each parameters' configuration, and usually one selects the one that maximizes it. But what if two configurations have slightly different values of likelihood but are associated to very different partitions? It seems unreasonable to simply rule out one of the two. To tackle this problem it would be relevant to analyze the entire energy landscape and derive phase diagrams to quantitatively measure how certain or uncertain the model predictions are.

More information: https://pio.is.mpg.de/project/analyzing-complex-energy-landscapes-for-inference-on-networks

## Distributed routing optimization

*Caterina De Bacco*

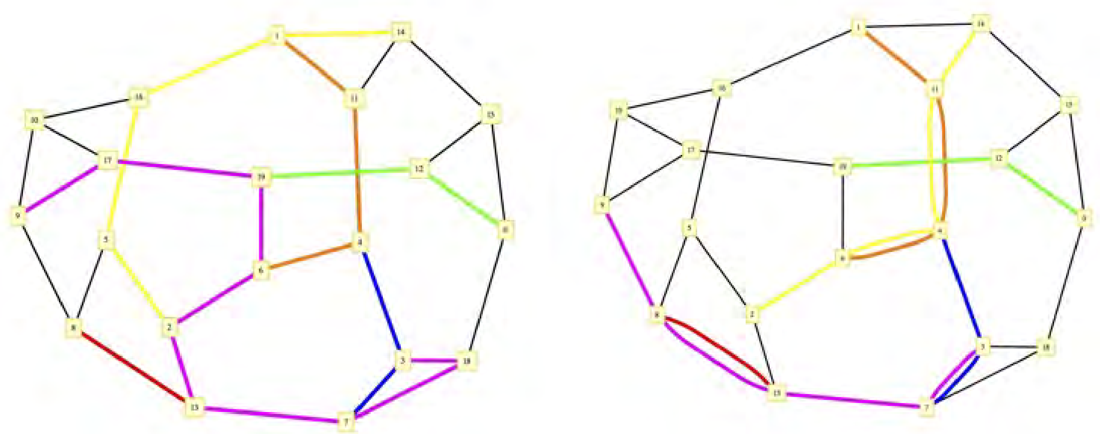

Figure 15.1: Example of an optimal (left) and non optimal (right) path configuration. The optimal is minimizing both traffic congestion and total path length; in the non optimal each user takes its shortest path, causing congestion on some edges.

Large interacting systems are often controlled by local interactions: the state of a variable depends directly only on the state of the surrounding neighborhood. The question is how to exploit this distributed character to optimize the network and in particular, we are interested in the problem of routing optimization. In communication networks, it is often the case that a network manager (e.g an Internet Service Provider) wants to optimize a global cost function where, ideally, all the network users' communication requests are not only satisfied but also with a guaranteed quality-of-service. This translates in a cost function that penalizes both the total traffic and transmission delay, summed over the users. To study this problem we deploy the distributed formalism of message-passing algorithms borrowing ideas from statistical physics of disordered systems. By exploiting the distributed character of this approach we are able to improve state-of-the-art results both in terms of finding a more optimal solution and in faster time.

More information: https://pio.is.mpg.de/project/distributed-routing-optimization

## Incorporating metadata information to perform inference on networks

*Caterina De Bacco*

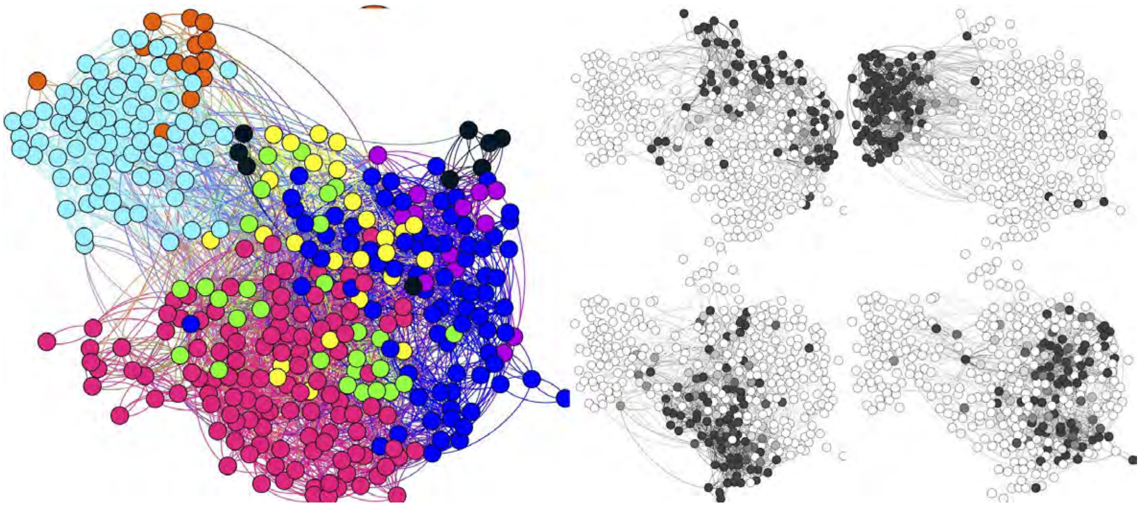

Figure 15.2: Division by caste membership for an Indian village (left). To the right we see the membership inferred in each of the four community, with color ranging from white for nodes not belonging to that group, to black, for highest membership in that group.

One of the main problems when working with real-world data is the absence of ground truth. The structure predicted by a model might not, in general, correspond to the structure suggested by additional information, such as metadata or covariates associated to nodes or edges of the network. Nevertheless, this additional information might bring in new insights to the problem, therefore it is very important to understand how to use this information, either to guide pattern prediction or to asses a posteriori how the results correlate with the metadata. We investigate modeling approaches capable of incorporating metadata for detecting communities on multilayer networks. In particular, we are interested to analyze how this additional information influences the resulting node partitions in real datasets. Knowing how the metadata contributes to clustering nodes can provide useful insights to understand the mechanism driving the observed topology. This can, for instance, be used to predict the connection of nodes for which we only know metadata but not their patterns of interactions. In addition, one can use the metadata to steer the detected partition towards a given type of division, by selecting which metadata to be incorporated based on the details of the problem considered.

More information: https://pio.is.mpg.de/project/incorporating-metadata-information-to-perform-inference-on-networks

## Investigating hierarchies in directed network structures

*Caterina De Bacco*

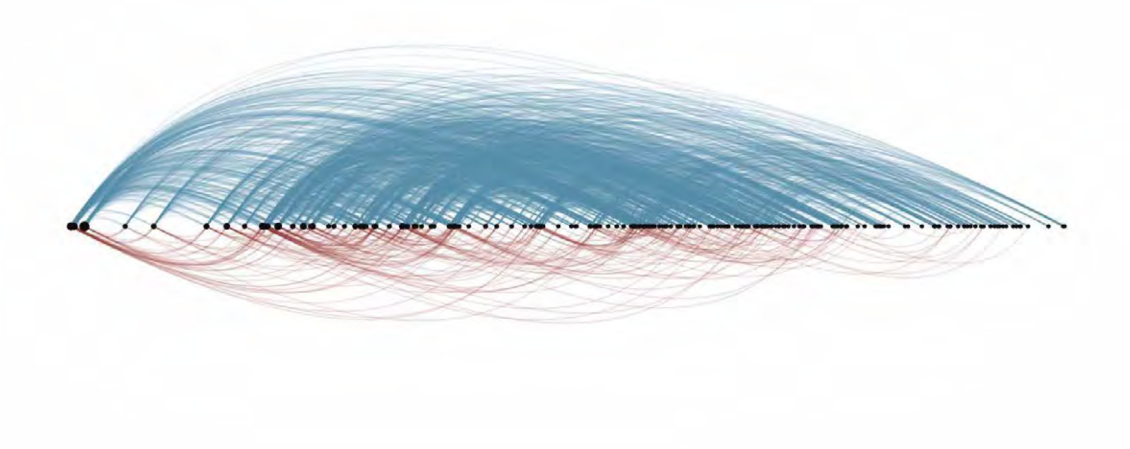

Figure 15.3: Linear hierarchy diagram with nodes embedded at their inferred SpringRank scores. Blue edges point down the hierarchy and red edges point up. Nodes on the left are higher ranked.

The elements of an interacting system are often affected by hidden hierarchies which might play a role in determining the interaction patterns that we observe between them. A relevant problem in this context is how is a global hierarchy of individuals being capable of reproducing the pairwise interaction patterns that we observe and how can we find it. But sometimes the hierarchy that one observes might be resulting from random chance rather than an actual underlying pattern induced by the observed topology, therefore one other relevant problem is to be able to distinguish between these two cases and find limits in inferring this structure. Finding an hierarchy structure and assessing its stability can allow to act on the network, for instance by redistributing resources as certain nodes become more or less important. Relevant examples are found in networks of influence in organizations, ecological networks of animals fighting each others in pairs, or hiring dynamics in job networks.

More information: https://pio.is.mpg.de/project/investigating-hierarchies-in-directed-network-structures

## 15.3 Research group leader: Caterina De Bacco

### Dr. Caterina De Bacco

| | |
|---|---|
| 2018 – persent | Max Planck Research Group Leader, Cyber Valley, MPI Tübingen. |
| 01/2018 – 06/2018 | Postdoc at Data Science Institute, Columbia University, New York. |
| 10/2015 – 12/2017 | Postdoc at Santa Fe Institute, Santa Fe, New Mexico. |
| 10/2012 – 09/2015 | PhD in Statistical Physics at Universitè Paris Sud 11, LPTMS, Paris. |
| 07/2015 | Master degree in Theoretical Physics, Università di Padova, Padova. |
| 07/2012 | Bachelor degree in Physics, Università di Padova, Padova. |
| 10/2010 – 06/2011 | Erasmus student at Imperial College, London. |

### Publications (selected)

C. De Bacco, D. B. Larremore, C. Moore. A physical model for efficient ranking in networks. *Science Advances* **4** (7), 2018

C. Brelsford, C. De Bacco. AreWater Smart Landscapes' Contagious? An epidemic approach on networks to study peer effects. *Networks and Spatial Economics (NETS)*: 1572–9427, 2018

T. Barthel, C. De Bacco, S. Franz. Matrix product algorithm for stochastic dynamics on networks applied to nonequilibrium Glauber dynamics. *Physical Review E* **97** (1): 010104, 2018

C. De Bacco, E. A. Power, D. B. Larremore, C. Moore. Community detection, link prediction, and layer interdependence in multilayer networks. *Physical Review E* **95** (4): 042317, 2017

T. Lesieur, C. De Bacco, J. Banks, F. Krzakala, C. Moore, et al. Phase transitions and optimal algorithms in high-dimensional Gaussian mixture clustering. In *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on*, pages 601–608, 2016

U. Bhat, C. De Bacco, S. Redner. Stochastic search with Poisson and deterministic resetting. *Journal of Statistical Mechanics: Theory and Experiment* **2016** (8): 083401, 2016

A. Berdahl, C. Brelsford, C. De Bacco, M. Dumas, V. Ferdinand, et al. Dynamics of beneficial epidemics. *arXiv preprint arXiv:1604.02096*, 2016

C. De Bacco, A. Guggiola, R. Kühn, P. Paga. Rare events statistics of random walks on networks: localisation and other dynamical phase transitions. *Journal of Physics A: Mathematical and Theoretical* **49** (18): 184003, 2016

# 16 PROBABILISTIC NUMERICS

## 16.1 Research Overview

### Computation is Inference

What does it mean to *compute* a number? For simple operations like divisions, floating point arithmetic gives a concise and formalized answer: A computer can compute such numbers in a fixed amount of time, to a specified precision. But the situation is more complicated when the quantity to be computed is described in terms of more advanced mathematical operations. There are no general closed-form solutions to tasks like finding the extremum of a function, the value of an integral, the solution to a linear system of equations, or to a differential equation. Methods designed to find approximate solutions to such jobs are known as *numerical algorithms*. Numerical methods are inherentely imperfect. They excel on some tasks and fail on others, and they do not always 'notice' the difference. Some numerical methods are expensive and slow, but work on a large class of problems. Others are fickle, specific to a small domain, on which they nevertheless work extremely well.

The computational cost of contemporary machine learning is dominated by such numerical tasks, and the properties of the algorithms used to solve them:

- Optimization methods (sgd, quasi-Newton, Frank-Wolfe,...) train and fit estimators,
- Integration algorithms (MCMC, free energy methods, quadrature,...) compute marginals, conditionals and expectations in probabilistic inference,
- Linear algebra tools (Gauss-Jordan, conjugate gradients, Cholesky decompositions,...) solve elementary tasks like least-squares estimation and do the heavy lifting in the innermost loop of many more high-level numerical methods
- Differential equations are solved to *simulate*, *predict* and *control* future states of the environment

As a young field, machine learning has been able to borrow many extremely well-designed numerical algorithms from other fields. Some of these (e.g. linear algebra subroutines) are so good that people have essentially stopped thinking of them as algorithms with properties and flaws, and blindly trust the black box. This confidence is not always justified, for machine learning and data science poses some new numerical challenges that are not addressed well by classic methods. An example is the strong computational noise caused by data sub-sampling in big-data applications, which can make advanced optimization methods (like quasi-Newton) essentially unusable in applications like deep learning.

The improve this situation, our group aims to develop and broaden the understanding of numerical algorithms for intelligent systems. The twist is that we do so *in the language and the setting of probabilistic machine learning*. Our core tenet is that machine learning not just poses new challenges for numerics, it can also offer its own contributions to their solution, because *numerical methods* ***are*** *elementary learning machines themselves.* A numerical method *estimates* an unknown, latent quantity (e.g. an integral), *given* the value of certain tractable, observable quantities (values of the integrand at discrete locations). So they solve an inference task. But numerical methods are not passive statistical estimators; they actively decide, often in a closed feedback loop, which numbers to compute. They thus really are elementary autonomous agents.

In our work, we phrase this active inference process probabilistically, as the actions of an agent equipped with a notion of uncertainty about its task, captured by a probability measure. Exact computations performed on a chip provide information about the analytically intractable task, yielding a *posterior* probability measure whose location and width should ideally provide a point estimate and meaningful surrounding uncertainty over the true solution. Such algorithms are known as **probabilistic numerical methods**. Over the past years, we have helped establish this young research field, together with international collaborators.

On the theoretical side, we have contributed in-depth analysis of classic numerical algorithms to show that probabilistic numerical algorithms can be as fast, reliable and flexible as the widely used classics - simply because the classics already *are* probabilistic numerical methods, they just are not usually presented in this way. Classic methods thus provide a point of reference, a solid foundation from which one can venture out to create new functionality for contemporary challenges. We also showed that in some cases (like optimization for deep learning), the probabilistic viewpoint can help uncover the cause for failures and problems of state-of-the-art methods, and suggest fixes.

For practitioners, our algorithms offer tangible improvements, some of which have helped our work enter into industrial use. Our Bayesian optimization framework *Entropy Search* has been used for advanced experimental design in areas like robotics and automated machine learning. For deep learning, we have contributed several tools that automate algorithmic hyperparameters in the inner loop, freeing users from tedious and wasteful tasks like step-size adaptation (some of our work can also be directly accessed as code packages on our github page[1]). As the theoretical foundations of probabilistic numerics continue to expand and improve, we now increasingly turn our attention to develop practical algorithms for and in machine learning.

Last but not least, we have also invested in the establishment of a dedicated research community. We created a community website[2] and organized and supported several international workshops and meetings.[3] Partly as a result of this, there is now a nascent by rapidly growing international community studying both the theory and practice of probabilistic numerics.

### History of the Group

Within the SAB reporting period, the group existed in two different funding setups; and it left the MPI IS before the end of the reporting period:

The Probabilistic Numerics group at the MPI for Intelligent Systems was initially founded and funded, in 2015, as an Emmy Noether Research group of the German Research Union (DFG). Prior to this time, Philipp Hennig was employed as a senior research scientist in the department of Empirical Inference. In late 2016, the Emmy Noether funding was replaced, when the group became one of the first independent Max Planck Research Groups of the MPI for Intelligent Systems. In 2017 our work was honoured with the award of an ERC Starting Grant, which in turn led to several offers from German universities to move the group to a more permanent home. In May 2018, Philipp Hennig was appointed as a full professor at the Eberhard Karls University of Tübingen, where the group now lives on in a new form, as the Chair for the Methods of Machine Learning.

We are grateful for our time at the MPI IS. To institute provided an ecosystem that allowed rapid growth and open collaboration. We look forward to continued close collaboration.

## 16.2 Selected Research Projects



[1] https://github.com/ProbabilisticNumerics
[2] http://probnum.org/
[3] http://probabilistic-numerics.org/meetings/

## Probabilistic Methods for Linear Algebra

*Philipp Hennig, Simon Bartels, Filip de Roos*

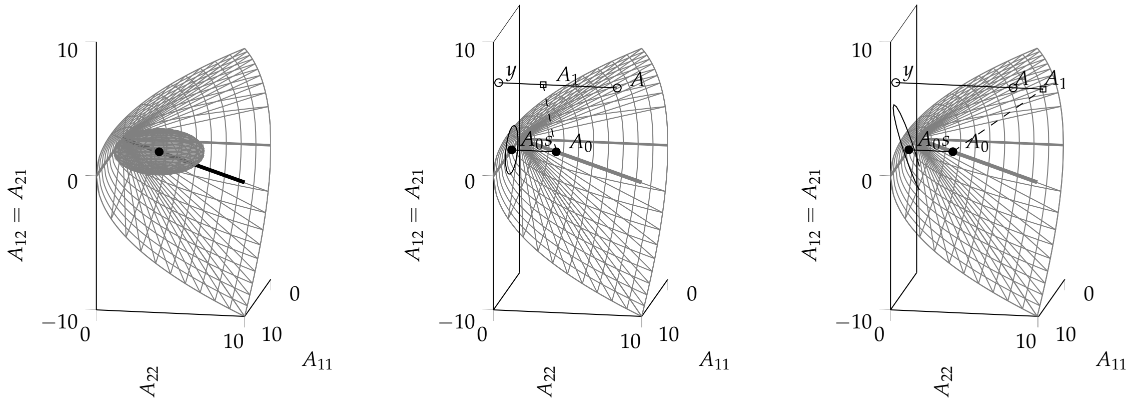

Figure 16.1: Probabilistic linear algebra methods assign matrix-valued probability measures to unknown matrix quantities; for example to elements of the positive definite cone.

Linear algebra methods form the basis for the majority of numerical computations. Because of this foundational, "inner-loop" role, they have to satisfy strict requirements on computational efficiency and numerical robustness.

Our work has added to a growing understanding that many widely used linear solvers can be interpreted as performing probabilistic inference on the elements of a matrix or a vector from observations linear projections of this latent object. In particular, this is true for such foundational algorithms as the method of conjugate gradients and other iterative algorithms in the *conjugate directions* and *projection method* classes.

Our ongoing research effort focusses on ways to use these insights in the large-scale linear algebra problems encountered in machine learning. There, the most frequent linear algebra task is the least-squares problem of solving

$$Kx = b$$

where $K$ is a very large symmetric positive definite matrix (e.g. a kernel Gram matrix, or the Hessian of a deep network loss function). A key challenge in the big data regime is that the matrix — defined as a function of a large data-set — can only be evaluated with strong stochastic noise caused by data sub-sampling. Classic iterative solvers, particularly those based on the Lanczos process, like conjugate gradients, are known to be unstable to such stochastic disturbances, which is part of the reason why second-order methods are not popular in deep learning. In recent work we have developed and tested extensions to classic solvers that remain stable [296] and tractable in this setting.

Work on probabilistic linear algebra methods is more challenging than one might perhaps think, precisely because linear algebra algorithms play such a foundational role across all computational science. Any new functionality added has to work with an extremely constrained computational budget, but also produce reliable and robust output. Our recent works (some of them still in review[4]) have not just produced new functionality, but also helped to shed light on the limitations imposed by computational complexity on the kinds of prior information that can feasibly be included in a prior over a matrix while remaining within the computational complexity class of existing solvers.

More information: https://pn.is.mpg.de/project/probabilistic-methods-for-linear-algebra

[4] e.g. Roos, F. D., Hennig, P. *Krylov Subspace Recycling for Fast Iterative Least-Squares in Machine Learning*, arXiv:1706.00241, and Bartels, Cockayne, Ipsen, Girolami and Hennig, *Probabilistic Linear Solvers: A Unifying View*, accepted conditional on minor revisions, Statistics and Computing.

## Probabilistic Methods for Nonlinear Optimization

*Philipp Hennig, Maren Mahsereci, Lukas Balles, Frank Schneider*

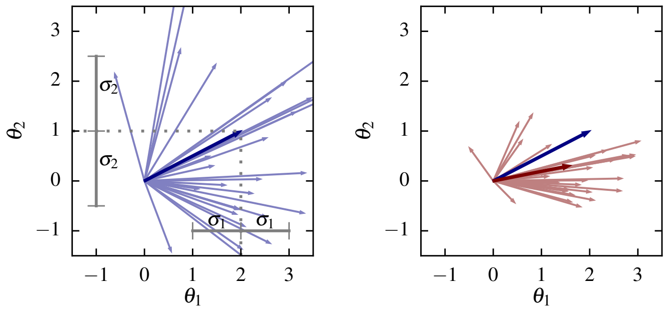


Figure 16.2: Raw batch gradients (left, blue) are approximately Gaussian distributed around their expectation, the population mean (solid blue). Their variance can be very large. But when it is known, search directions of reduced variance (right, red) can be constructed efficiently, at low overhead (image from Balles & Hennig, 2018)

Optimization problems arising in intelligent systems are similar to those studied in other fields (such as operations research, control, and computational physics), but they have some prominent features that set them apart, and which are not addressed by classic optimization methods. Numerical optimization is a domain where probabilistic numerical methods offer a particularly interesting theoretical contribution.

One key issue is computational noise. Big Data problems often have the property that computational precision can be traded off against computational cost. The most widely occuring task structure is empirical risk minimization: local optimization of a high-dimensional function $L$ that is the sum of many similar terms, each arising from an individual data point $y_i$

$$L(x) = \frac{1}{N}\sum_{i=1}^{N} \ell(y_i, x)$$

Examples of this problem include the training of neural networks, of logistic regressors, and many other linear and nonlinear regression/classification algorithms. If the dataset is very large or even infinite, it is impossible, or at least inefficient, to evaluate the entire sum. Instead, one draws $M \ll N$ (hopefully representative) *samples* $y_j$ from some distribution and computes the approximation

$$\hat{L}(x) = \frac{1}{M}\sum_{j=1}^{M} \ell(y_j, x) \approx L(x)$$

If the representers $y_j$ are drawn independently and identically from some distribution, then this approximation deviates, relative to the true $L(x)$, by an approximately Gaussian disturbance.

Classic optimizers like quasi-Newton methods are unstable to these disturbances, hence the popularity of first-order methods, like stochastic gradient descent (sgd), in deep learning. But even such simple methods become harder to use in the stochastic domain. In particular, sgd and its variants exhibit free parameters (e.g. step-sizes / learning rates) in the stochastic setting, even though such parameters can be easily tuned automatically in the noise-free case. Thus, even at the world's leading large AI companies, highly trained engineers spent their work time hand-tuning parameters by repeatedly running the same training routine on high-performance hardware. A very wasteful use of both human and hardware resources. A NeurIPS workshop organized by us in 2016 highlighted the urgency of this issue.

The probabilistic perspective cleanly captures the problem as caused by a non-trivial *likelihood* term: Instead of the point measure on $L(x)$ that describes the structure of classic optimizers, the batch setting yields a *Gaussian* distribution $p(\hat{L} \mid L) = \mathcal{N}(\hat{L}; L, \Sigma)$. This seemingly straightforward formulation immediately offers an analytical avenue to understand why existing optimizers fundamentally require hand-tuning: While a point measure only has a single parameter (the location), a Gaussian has *two* parameters: mean and (co-) variance. But the latter does not feature in classic analysis, and is simply unknown to the algorithm. It is possible to show [390] that this lack of information can make certain parameters (such as step sizes) fundamentally un-identified. Identifying them not just requires new algorithms, but also concretely *computing* a new object: In addition to batch gradients, also batch *square* gradients, to empirically estimate the variance. Doing so is not free, but it has low and bounded computational cost [390], because it can re-use the back-prop pass, the most expensive part of deep learning training steps.

Over years we have built a series of tools that use this additional quantity to tune various learning hyperparameters such as the learning rate [313, 683], batch size [390] and stopping criterion[5]. We have also contributed theoretical analysis for some of the most popular deep learning optimizers [686] and are now working towards a new code platform for the automation of deep learning in the inner loop, to free practitioners hands to build models, rather than tune algorithms.

More information: https://pn.is.mpg.de/project/probabilistic-methods-for-nonlinear-optimization

[5]Mahsereci, Balles, Lassner, Hennig. *Early Stopping Without a Validation Set*. In Rev. arXiv 1703.09580

## Probabilistic Solvers for Ordinary Differential Equations

*Philipp Hennig, Michael Schober, Hans Kersting*

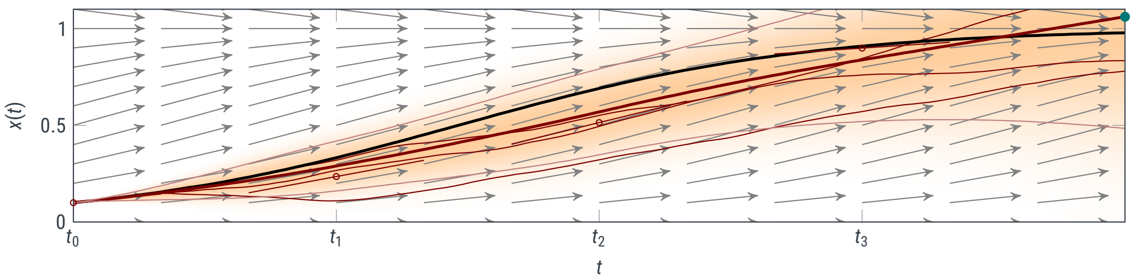


Figure 16.3: Probabilistic ODE solvers assign a probability measure (yellow shading, red samples, solid red mean prediction) to the unknown solution (black) of an ODE described an initial value (circle at left) and vector field (arrows in background). Our algorithms come with rate guarantees for speed, the error of the estimate, and the error of the estimated error (uncertainty).

Solvers for ordinary differential equations (ODEs) belong to the best-studied algorithms of numerical mathematics. An ODE is an implicit statement about the relationship of a curve $x : \mathbb{R}_{\geq 0} \to \mathbb{R}^N$ to its derivative, in the form $x'(t) = f(x(t), t)$, where $x'$ is the derivative of the curve, and $f$ is some function. To identify a unique solution, it is typically also necessary to provide additional constraints, such as the *initial value* $x(t_0) = x_0$. An ODE solver maps function and initial value $(f, x_0)$ to an estimate $x(t)$ for the solution curve. Good solvers have certain analytical guarantees about this estimate, such as the fact that its deviation from the true solution is of a high polynomial order in the step size used by the algorithms to discretize the ODE.

One of the main theoretical contributions of the group is the development of *probabilistic* versions of these solvers. In several works, we established a class of solvers for initial value problems that generalize classic solvers by taking as inputs Gaussian distributions $\mathcal{N}(x(t_0); x_0, \Psi)$, $\mathcal{GP}(f; \hat{f}, \Sigma)$ over the initial value and vector field, and return a Gaussian process posterior $\mathcal{GP}(x; m, k)$ over the solution. We were able to show that these methods[6]

- have the same (linear) computational complexity as classic methods [680] (they are Bayesian filters)
- can inherit the famous local and global polynomial convergence rates of classic solvers[7] (i.e. $\|m - x\| \leq Ch^q$ for $q \geq 1$)
- produce posterior variance estimates that are calibrated worst-case error estimates[7] (i.e. $\|m - x\|^2 \leq Ck$). In Short, they produce meaningful uncertainty
- are a generalization of certain famous classic ODE solvers (they reduce to Runge Kutta methods and multi-step Nordsieck methods in the opposing limits of uninformative prior and steady-state operation, respectively. In between, they offer a third, novel type of solver) [680]
- they can be generalized to produce non-Gaussian, nonparametric output while retaining many of the above properties [8].

Together, these results provide a rich and reliable new theory for probabilistic simulation that current ongoing research projects are seeking to leverage to speed up structured simulation problems inside of machine learning algorithms.

More information: https://pn.is.mpg.de/project/probabilistic-solvers-for-ordinary-differential-equations

[6]To ensure vetting by the expert community, we strife to publish some of our results in academic journals of applied mathematics, instead of computer science. An unfortunate downside of this is that, since journals in mathematics have review cycles up to two years, some of our most advanced theoretical work remains available only as pre-prints.
[7]Kersting, Sullivan, Hennig. *Convergence Rates of Gaussian ODE Filters*. In review at SIAM JUQ. arXiv:1807.09737
[8]Tronarp, Kersting, Särkkä, Hennig. *Probabilistic Solutions To Ordinary Differential Equations As Non-Linear Bayesian Filtering: A New Perspective*. In review at Statistics and Computing. arXiv::1810.03440

## Probabilistic Radiation Treatment Planning

*Philipp Hennig, Mark Bangert, Niklas Wahl, Hans-Peter Wieser*

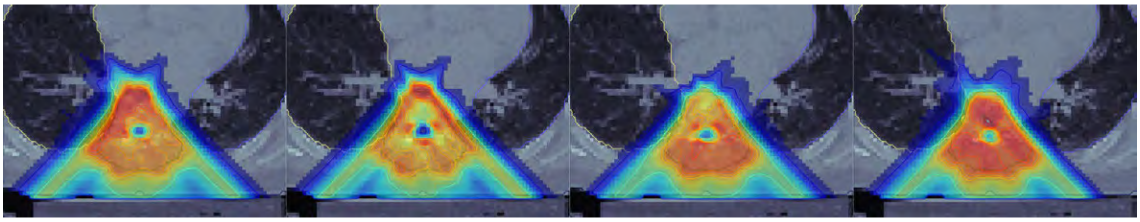

Figure 16.4: Several hypotheses for the dosis deposited in a spinal tumor under a traditional treatment plan. The samples are representative of an entire population simultaneously considered and optimized for reduced variance by the novel algorithm produced in this project.

Software solutions for scientific and technical tasks usually do not consist of a single computational step, but rather are a *pipeline* of computations. An ongoing collaboration between the research group on probabilistic numerics and the Optimization Group at the German Cancer Research Centre in Heidelberg has offered an opportunity for us to test mathematical ideas for the *propagation of uncertainty* through such chains of computation while also producing tangible medical results.

The semi-automated production protocol of treatment plans is an everyday occurence in clinical practice: Tumor patients who are scheduled for radiation treatment by their oncologist come in for an imaging session in a CT or MRI scanner. The resulting 3D image is annotated by a physician, outlining both tumor tissue and surrounding organs at risk of unwanted radiation damage. This volumetric data provides the input to an optimization algorithm, which sets the parameters of a treatment system (angles, energies, and shape of treatment beams). In advanced treatment systems, in particular those using heavy ions as the treatment probe, this optimization problem regular has several thousand parameters. Then the patient comes in for a sequence of treatments.

This entire process is subject to a host of sources of uncertainty, from the imaging process, through human labeling, the optimization algorithm itself, the mechanical imperfections of patient placement during treatment, to complicated and correlated physical and biological sources of uncertainty about the reaction of each cell to the delivered radiation dose.

In a number of sub-projects with our collaborators in Heidelberg, we were able to construct a framework for the analytical and numerically efficient computation and propagation of such input uncertainties that can separate the effects of various sources of uncertainty [681], including highly nonlinear biological effects [682], and efficiently propagate them through the optimization process [684], to produce an improved treatment plan that is more robust to errors, and reduces the risk of complications for patients. From the point of view of research in probabilistic numerics, this strand of work provides examples both for the feasibility and concrete useability of uncertainty propagation in compartmental computations: By casting each step of the pipeline in terms of structured Gaussian distributions, uncertainty from various sources can be tracked, monitored, and controlled at feasible computational overhead. There is a deeper philosophical inside hidden inside of these practical tools: Because probability theory does not differentiate between different types of uncertainty but captures everything the the universal language of probability measures, there is no need to distinguish between numerical, physical, experimental or epistemological uncertainty in computational practice, either. The kinds of uncertainty caused by finite data, by finite computational budget, imperfect physical measurements and even quantum-mechanical aspects of the interaction between accelerated heavy ions and DNA molecules all fit into one joint Gaussian distribution.

More information: https://pn.is.mpg.de/project/probabilistic-radiation-treatment-planning

## Controller Learning using Bayesian Optimization

*Alonso Marco Valle , Sebastian Trimpe, Philipp Hennig, Alexander von Rohr, Jeannette Bohg, Stefan Schaal*

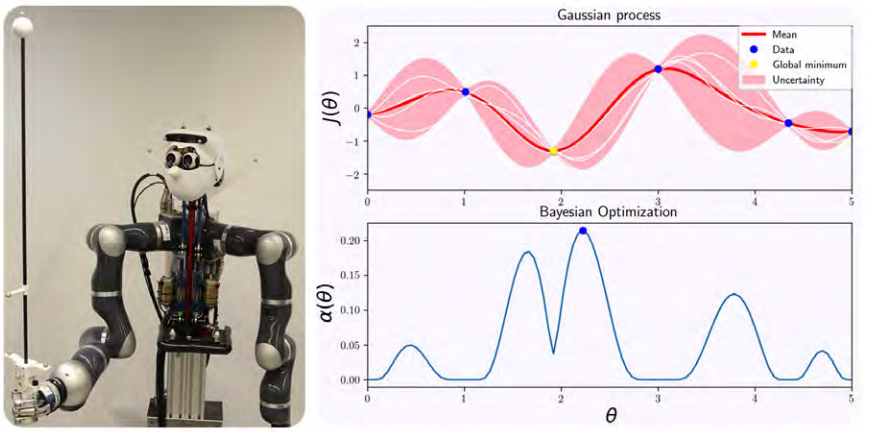


Figure 16.5: Left: Humanoid robot Apollo learning to balance an inverted pole using Bayesian optimization. Right: One-dimensional synthetic example of an unknown cost J($\theta$) modeled as a Gaussian process for controller parameter $\theta$, conditioned on observed data points. The next controller to evaluate is suggested by the Bayesian optimizer where the acquisition function $\alpha(\theta)$ finds its maximum.

Autonomous systems such as humanoid robots are characterized by a multitude of feedback control loops operating at different hierarchical levels and time-scales. Designing and tuning these controllers typically requires significant manual modeling and design effort and exhaustive experimental testing. For managing the ever greater complexity and striving for greater autonomy, it is desirable to tailor intelligent algorithms that allow autonomous systems to learn from experimental data. In our research, we leverage automatic control theory, machine learning, and optimization to develop automatic control design and tuning algorithms.

In [58], we propose a framework where an initial controller is automatically improved based on observed performance from a limited number of experiments. Entropy Search (ES)[9] serves as the underlying Bayesian optimizer for the auto-tuning method. It represents the latent control objective as a Gaussian process (GP) (see above figure) and sequentially suggests those controllers that are most informative about the location of the optimum. We validate the developed approaches on the experimental platforms at our institute (see figure).

We have extended this framework into different directions to further improve data efficiency. When auto-tuning real complex systems (like humanoid robots), simulations of the system dynamics are typically available. They provide less accurate information than real experiments, but at a cheaper cost. Under limited experimental cost budget (i.e., experiment total time), our work [31] extends ES to include the simulator as an additional information source and automatically trade off information vs. cost.

The aforementioned auto-tuning methods model the performance objective using standard GP models, typically agnostic to the control problem. In [24], the covariance function of the GP model is tailored to the control problem at hand by incorporating its mathematical structure into the kernel design. In this way, unforeseen observations of the objective are predicted more accurately. This ultimately speeds up the convergence of the Bayesian optimizer.

More information: https://pn.is.mpg.de/project/cont-learn-bayes-opt

[9] Hennig & Schuler, JMLR 13 (2012), 1809–1837

## 16.3 Research group leader: Philipp Hennig

### Prof. Dr. Philipp Hennig

From mid 2016 to early 2018, Philipp Hennig led the Probabilistic Numerics research group at the MPI for Intelligent Systems. He studied Physics in Heidelberg, Germany and at Imperial College, London, before moving to the University of Cambridge, UK, where he attained a PhD in the group of the late Sir David JC MacKay for research on approximate inference methods. Since this time, he is interested in connections between computation and inference. He began his postdoctoral career at the MPI IS in the department for Empirical Inference, before his group received independent funding through the Emmy Noether programme of the German Research Union (DFG) in 2015. An ERC Starting Grant awarded in 2017 triggered several offers for full professorships from German universities. In May 2018, Hennig accepted such an offer from the University of Tübingen, where he now holds the Chair for the Methods of Machine Learning. In November 2018, Capital Magazine listed Philipp Hennig as one of Germany's "40 under 40" influential young people in the category "Science and Society". Hennig retains an adjunct position at the MPI IS, in the Empirical Inference department.

#### Employment

| | |
|---|---|
| 2018 – present | Full Professor for the Methods of Machine Learning (W3), Eberhard Karls University of Tübingen, Germany |
| 2016 – 2018 | W2 Independent Max Planck Research Group Leader, MPI-IS, Germany |
| 2015 – 2016 | Independent Emmy Noether Group Leader, MPI-IS, Germany |
| 05/2015 – 01/2016 | Parental Leave |
| 2013 – 2015 | Senior Research Scientist, Empirical Inference Department, MPI-IS, Germany |
| 2011 – 2013 | Research Scientist, Empirical Inference Department, MPI-IS, Germany |

#### Education

| | |
|---|---|
| 2011 | PhD, University of Cambridge (Robinson College / Cavendish Lab.), UK. Supervisor: Sir D.J.C. MacKay |
| 2005 – 2016 | Erasmus studies at Imperial College, London (*Quantum Fields and Fundamental Forces* program) |
| 2007 | Diploma in Physics, Ruprecht Karls University Heidelberg, Germany |

#### Fellowships and Awards

| | |
|---|---|
| 2018 | Named as one of Germany's *40 under 40 in Science and Society* by Capital Magazine |
| 2017 | ERC Starting Grant, 5 year project PANAMA |
| 2015 | Emmy Noether Fellowship of the German Research Union (DFG) |

## Publications and Talks

| | |
|---|---|
| Publications | 14 Refereed Journal, 29 Refereed Conference, 1 Ed. Book, 1 US Patent<br>h-index: 22, i-10 index: 32, 1331 citations (Google Scholar 06/12/18) |

## Teaching/Supervision Activities

| | |
|---|---|
| 2012 – present | Supervision of 11 PhD students (3 graduated) |
| 2004 – present | Supervision of 4 Master Theses (1 current) |
| 2012 – present | Several Full Lecture Courses on Bayesian Machine Learning<br>Several Seminars, Invited Tutorials |
| 2017 | Lecturer at Dobbiaco Summer School on Probabilistic Numerics |
| 2012–13–15 | Co- and Principal Organizer of three Machine Learning Summer Schools (MLSS) |

## Professional Activities (selected)

| | |
|---|---|
| 2013 – present | Member of Editorial Board, Journal of Machine Learning |
| 2014 – present | Area Chair for Neural Information Processing Systems (3×), International Conference on Machine Learning (1×), Artificial Intelligence and Statistics (3×) |
| 2017 – 2018 | Member of the Exectuive Board for the International Max Planck Research School for Intelligent Systems (IMPRS-IS) |
| 2014 – 2018 | Steering Board, Max Planck-ETH Center for Learning Systems |

## (Co-) Organization of Scientific Workshops

| | |
|---|---|
| 2017 | ICERM workshop on Probabilistic Scientific Computing |
| 2016 | Dagstuhl Seminar: New Directions for Learning with Kernels and Gaussian Processes |
| 2016 | NIPS Workshop on Optimizing the Optimizers |
| 2015 | NIPS Workshop on Probabilistic Integration |
| 2015 | Workshop on Probabilistic Numerics for Differential Equations, Warwick |
| 2015 | Workshop on Probabilistic Numerics at DALI |
| 2014 | Roundtable on Probabilistic Numerics |
| 2012 | Inaugural NIPS workshop on Probabilistic Numerics |

## Third Party Funding

| | |
|---|---|
| 2018–2021 | Project ADIMEM (*automatic data-driven inference in mechanistic models*) by the German Federal Ministry of Research (BMBF 01 I S1 8052 B). with J. Macke (TUM), P. Berens (Tübingen), M. Oberländer (CAESAR Bonn) — € 160 000 for the sub-project |
| 2018–2021 | Sub-Project in the TUE.AI competence center of the German Federal Ministry of Research (BMBF ) — € 160 000 |
| 2018–2023 | ERC Starting Grant *Probabilistic Automated Numerical Analysis in Machine Learning and Artificial Intelligence*(757275 / PANAMA) — € 1 500 000 incl. overheads |
| 2015–2016 | Emmy Noether Grant of the German Research Union (DFG) — approx. € 1 000 000 including overheads, returned after 1.5 years, replaced by the Max Planck Research Group funding (with which it is formally incompatible) |
| 2015–2018 | DFG main-line grant *Analytical Probabilistic Radiation Therapy Treatment Planning* with Mark Bangert, DKFZ Heidelberg. Approx. € 300.000 |

# 17 RATIONALITY ENHANCEMENT

## 17.1 Research Overview

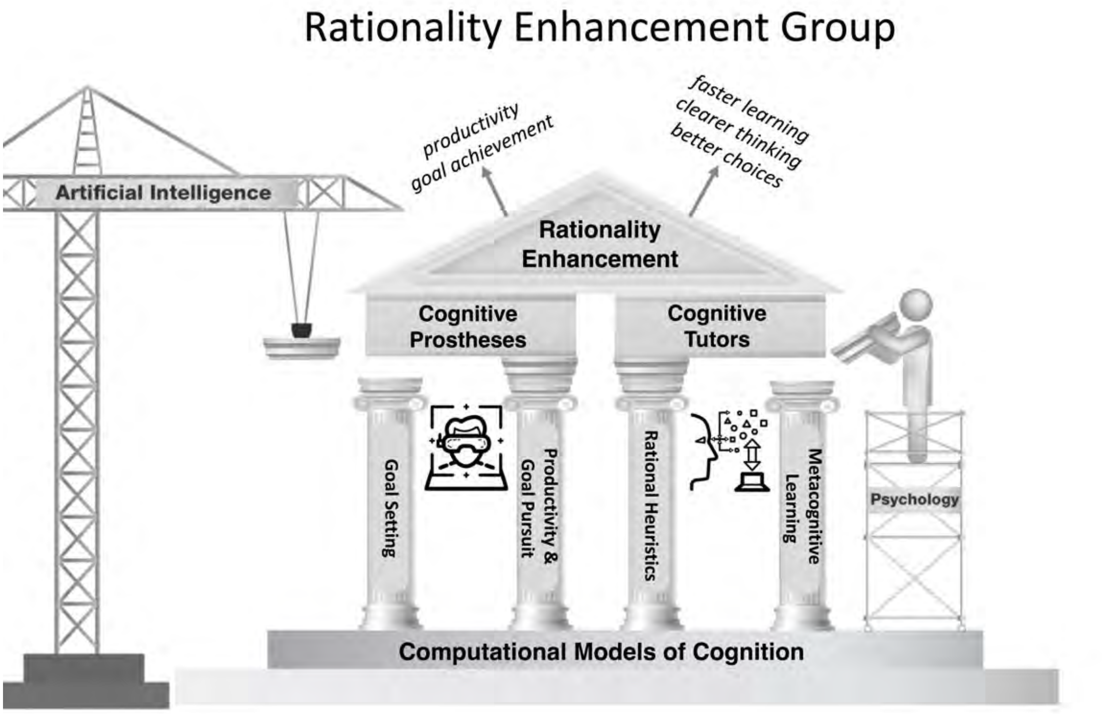


The scientific mission of the Max Planck Research Group for Rationality Enhancement is to lay the cognitive and technological foundations for helping people become more effective. Our two-pronged approach synergistically combines basic research on the computational mechanisms of human learning and decision-making with the development of intelligent systems for enhancing human rationality. Our basic research focuses on how people learn how to decide, decision strategies that make optimal use of limited time and bounded cognitive resources, cognitive control, planning, and goal-setting. Our applied research translates the resulting theories and computational models into cognitive tutors that teach people how to make better decisions and cognitive prostheses that augment people's limited cognitive capacities with artificial intelligence. We evaluate the models, cognitive tutors, and cognitive prostheses we develop in experiments conducted online, in the lab, and in the field. Ultimately, we aim to produce technologies that enhance human productivity, accelerate personal development, and assist people in goal setting and goal achievement.

## 17.2 Selected Research Projects

## Developing Cognitive Prostheses

*Falk Lieder*

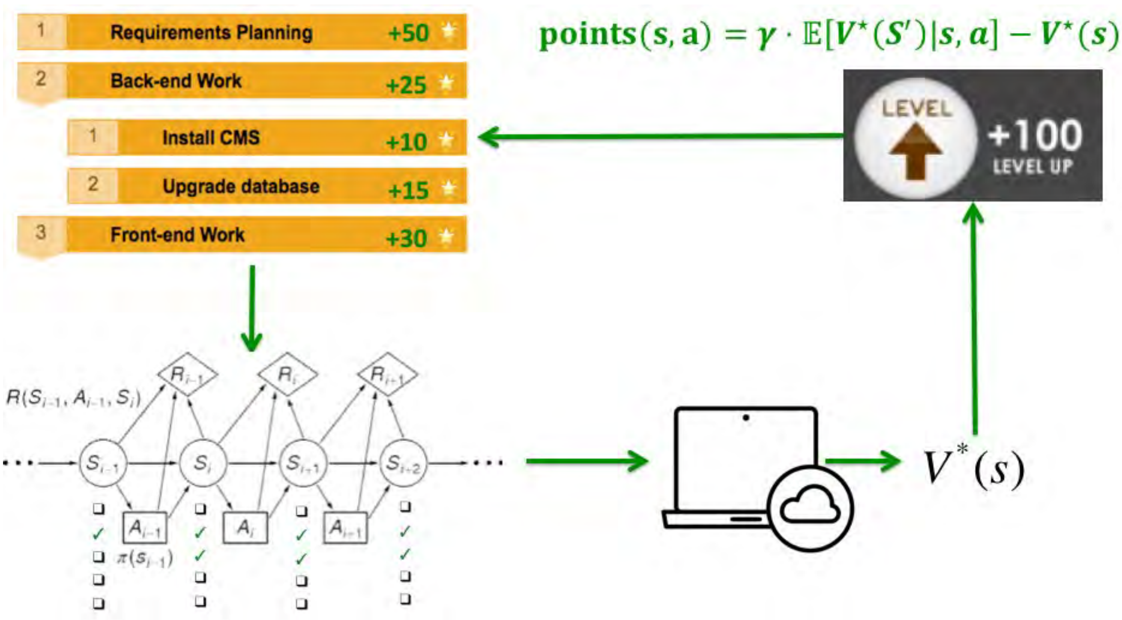


Today's rapid advances in artificial intelligence present an unprecedented opportunity to augment the human mind with the help of technology. Our goal is to leverage insights from cognitive science to develop technologies that help humanity overcome its cognitive limitations and enable people to make better decisions, learn faster, become more productive, and achieve their goals.

We are currently working on a cognitive prosthesis for helping people achieve their goals on time. The basic idea is to align each action's immediate reward with its long-term value so as to make good decisions easier[1]. We have developed a mathematical framework for designing incentive structures that can be used to guarantee that the rewards won't accidentally incentivize counter-productive behaviors and a computational method for computing incentives that make it as easy as possible for people to choose the course of action that is best in the long run. We are working on instantiating this approach in a to-do list gamification app. Initial results suggest that this is a promising approach to helping people overcome procrastination[2].

More information: https://re.is.mpg.de/project/developing-cognitive-prostheses

[1] F. Lieder, T. L. Griffiths. Helping people make better decisions using optimal gamification. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, 2016.

[2] F. Lieder, T. L. Griffiths. Optimal gamification can help people procrastinate less. In *Annual Meeting of the Society for Judgment and Decision Making*, 2017.

## Developing Cognitive Tutors

*Falk Lieder, Maria Wirzberger*

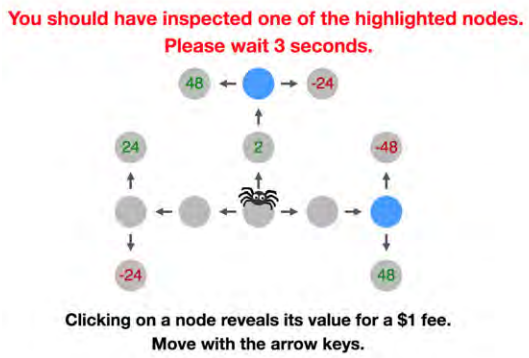


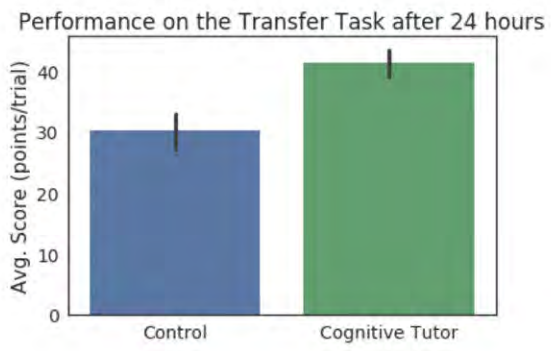


Building on insights into how people learn how to think and how to decide[3],[4],[5] and novel methods for discovering optimal cognitive strategies[6],[690], we develop intelligent systems that teach people how to make better decisions and accelerate the acquisition of cognitive skills. As a first step, we have developed an intelligent tutor that teaches people optimal planning strategies[7],[689]. The basic idea is to create reward structures that are conducive to human learning[8] and to provide metacognitive feedback so as to accelerate metacognitive reinforcement learning towards the optimal cognitive strategy [689],[9]. Future applications of this principle may include training executive functions, reducing impulsivity, and promoting learned industriousness.

More information: https://re.is.mpg.de/project/developing-cognitive-tutors

[3]F. Lieder, A. Shenhav, S. Musslick, T. L. Griffiths. Rational metareasoning and the plasticity of cognitive control. *PLoS Computational Biology* **14** (4): e1006043, 2018.

[4]P. M. Krueger, F. Lieder, T. L. Griffiths. Enhancing metacognitive reinforcement learning using reward structures and feedback. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*, 2017.

[5]F. Lieder, T. Griffiths. Strategy selection as rational metareasoning. *Psychological Review* **124**: 762–794, 2017.

[6]F. Lieder, P. M. Krueger, T. L. Griffiths. An automatic method for discovering rational heuristics for risky choice. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society. Austin, TX: Cognitive Science Society*, 2017.

[7]F. Lieder, P. M. Krueger, F. Callaway, T. L. Griffiths. A reward shaping method for promoting metacognitive learning. In *Proceedings of the Third Multidisciplinary Conference on Reinforcement Learning and Decision-Making*, 2017.

[8]P. M. Krueger, F. Lieder, T. L. Griffiths. Enhancing metacognitive reinforcement learning using reward structures and feedback. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*, 2017.

[9]F. Lieder, P. M. Krueger, F. Callaway, T. L. Griffiths. A reward shaping method for promoting metacognitive learning. In *Proceedings of the Third Multidisciplinary Conference on Reinforcement Learning and Decision-Making*, 2017.

## Reverse-Engineering the Computational Mechanisms of Metacognitive Learning

*Falk Lieder*

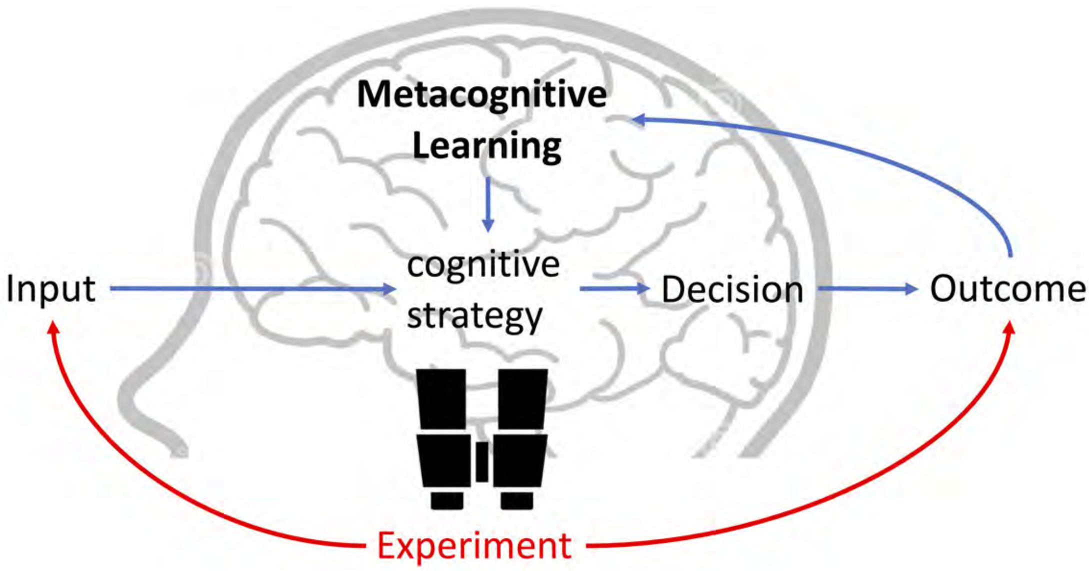


How does the human brain learn how to think and how to decide? What are the learning mechanisms that give rise to human intelligence and enable us to get better at what we do? How can this learning be promoted and accelerated? To answer these questions, we reverse-engineer how people learn when to use which cognitive strategy[10][11][12][13], how the brain learns to control its own information processing[14],[15], and how people discover and continuously refine their own cognitive strategies[16]. We translate what we have learned about metacognitive learning into robust and sample-efficient machine learning methods for discovering efficient decision-making strategies. This line of work has already led to a new, cognitively inspired sample-efficient learning algorithm called Bayesian Meta-Level Policy Search (BMPS) that outperforms the state-of-the-art methods for approximate metareasoning[17] and a method for automatically discovering optimal heuristics for human decision-making[18],[690]. The resulting psychological insights and technical advances lay a scientific foundation for leveraging technology to accelerate human learning and improve human decision-making.

More information: https://re.is.mpg.de/project/reverse-engineering-the-computational-mechanisms-of-metacognitive-learning

[10] F. Lieder, D. Plunkett, J. B. Hamrick, S. J. Russell, N. J. Hay, et al. Algorithm selection by rational metareasoning as a model of human strategy selection. In *Advances in Neural Information Processing Systems 27*, 2014.

[11] F. Lieder, T. L. Griffiths. When to use which heuristic: A rational solution to the strategy selection problem. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, 2015.

[12] F. Lieder, T. L. Griffiths. Model-based strategy selection learning, 2015.

[13] F. Lieder, T. Griffiths. Strategy selection as rational metareasoning. *Psychological Review* **124**: 762–794, 2017.

[14] F. Lieder, A. Shenhav, S. Musslick, T. L. Griffiths. Rational metareasoning and the plasticity of cognitive control. *PLoS Computational Biology* **14** (4): e1006043, 2018.

[15] A. Shenhav, S. Musslick, F. Lieder, W. Kool, T. Griffiths, et al. Toward a rational and mechanistic account of mental effort. *Annual Review of Neuroscience* **40**: 99–124, 2017.

[16] P. M. Krueger, F. Lieder, T. L. Griffiths. Enhancing metacognitive reinforcement learning using reward structures and feedback. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*, 2017.

[17] F. Callaway, S. Gul, P. Krueger, T. L. Griffiths, F. Lieder. Learning to select computations. In *Uncertainty in Artificial Intelligence: Proceedings of the Thirty-Fourth Conference*, 2018.

[18] F. Lieder, P. M. Krueger, T. L. Griffiths. An automatic method for discovering rational heuristics for risky choice. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society. Austin, TX: Cognitive Science Society*, 2017.

## 17.3 Research group leader: Falk Lieder

### Dr. Falk Lieder

Falk Lieder is a Max Planck Research Group Leader at the MPI for Intelligent Systems in Tübingen. His newly established Rationality Enhancement Group strives to lay the cognitive and technological foundations for helping people become more effective. He completed his Ph.D. in Tom Griffiths's Computational Cognitive Science Lab at the University of California, Berkeley. Prior to this, he worked as a research assistant in Klaas Stephan's Translational Neuromodeling Unit, received a masters degree in Neural Systems and Computation from ETH Zurich, and completed two simultaneous bachelor's degrees in Cognitive Science and Mathematics/Computer Science at the University of Osnabrück.
He has received the outstanding paper award of the NIPS workshop on cognitively inspired AI, the best poster award of the 14th Biannual Conference of the German Cognitive Science Society and received prestigious fellowships including the Berkeley fellowship for graduate study and two fellowships from the German national academic foundation as well as travel awards to the Annual meeting of the Cognitive Science society and the multidisciplinary conference on reinforcement learning and decision-making.

### Appointments

| | |
|---|---|
| 07/2018 – present | Max Planck Research Group Leader at the Max Planck Institute for Intelligent Systems |

### Awards and Honors (Selected)

| | |
|---|---|
| 09/2018 | Best poster award of the 14$^{th}$ Biannual Conference of the German Cognitive Science Society |
| 12/2017 | Outstanding Paper Award, NIPS workshop on Cognitively Informed Artificial Intelligence |
| 06/2017 | RLDM 2017 Student Travel Award |
| 2017 – 2018 | Golden Key International Honor Society |
| 08/2013 – 07/2015 | Berkeley Fellowship for graduate study offered to the top 4% of prospective graduate students |
| 07/2014 | German Society for Cognitive Science Student Travel Award |
| 08/2013 | German Society for Cognitive Science Student Travel Award |
| 2011 – 2012 | ETH Scholarship for International Students |
| 2008 – 2012 | Scholarship of the German National Academic Foundation awarded to the top 0.5% of German students |
| 2009 – 2011 | Study Abroad Scholarship of the German National Academic Foundation |

## Selected Organization and Community Service (2011-2015)

| | |
|---|---|
| 2013 – present | Reviewer for Cognitive Psychology, Cognition, Psychonomic Bulletin Review, NeuroImage, Brain and Behavior, Proceedings of the Annual Meeting of the Cognitive Science Society (12 times) |
| 2015 | Co-Organizer of the NIPS 2015 Workshop on Bounded Optimality and Rational Metareasoning |
| 2016 – 2017 | Co-Organizer of the Student Invited Speaker Series of the Helen Wills Neuroscience Institute |

## Invited Presentations

Int. Intuitive Computational Cognitive Science Spring School, Bernstein Center Freiburg, 2019
Psychology Guest Seminar Series, University of Warwick, 2019
Cognitive Science Colloquium, University of Tübingen, 2018
Workshop on Contemporary Cognitive Approaches to Decision-Making, Madison, 2018
Summer School on Cognitive Foundations of Economic Behavior, Vitznau, Switzerland, 2018
Computational Cognitive Science Lab, MIT, 2018
Economics Department, Harvard University, 2018
Cognitive Brownbag Seminar, UCSD, Department of Psychology, 2018
Google DeepMind, London, UK, 2017
Rational Process Models Symp., 50th Meet. of the Soc. for Math. Psychology, Warwick, 2017
Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, 2017
Shenhav Lab, Brown University, Providence, RI, USA, 2017
Neuroscape, University of California, San Francisco (UCSF), 2017
5th CSLI Workshop on Logic, Rationality Intelligent Interaction, Stanford, CA, USA, 2016
NIPS 2015 Workshop on Bounded Optimality and Rational Metareasoning, Montreal, 2015
Dagstuhl Seminar on Resource-Bounded Problem Solving, 2014
Bay Area Cognitive Science Meeting, Stanford, 2012

# 18 STATISTICAL LEARNING THEORY

## 18.1 Research Overview

*The Statistical Learning Theory Group has been founded in 2017. It is funded by the Max Planck Fellowship of Ulrike von Luxburg, whose main affiliation is the University of Tübingen. Compared to most of the other groups in the institute, our funding is more limited, the group only consists of two postdocs and, from time to time, an intern.*

The goal of statistical learning theory is to provide solid theoretical analysis of the behavior of machine learning algorithms. Under the assumption that the data has been sampled from some underlying, but unknown ground truth, we want to assess whether the results achieved by machine learning algorithms are trustworthy, whether the algorithms are well-behaved or erratic, or what is their complexity in terms of data required or computation time needed.

Some branches of statistical learning theory are well-studied and "more or less solved," while others are just beginning to be investigated. We would like to highlight the following two areas:

**Interactive and interpretable machine learning.** Here we ask how a fruitful interaction between machine learning algorithms and human users can be achieved. This is clearly a question of rising importance: machine learning systems get more and more complex and involved, which makes it hard to judge the meaning, implications, and trustworthiness of a machine's inference result. On the other hand, machine learning systems start to have serious impact on every-day life, hence being able to control their results gets more important. The question about interactive and interpretable machine learning clearly has aspects of human-computer interface, but also raises lots of algorithmic and also theoretical issues.

Up to now, our group has focused on one particular aspect, namely the question of how an algorithm's input can be provided by humans in a more natural way. We consider a setting where the input to a machine learning algorithm is not given in terms of similarity values ("On a scale from 0 to 1, the similarity between image A and image B is 0.8"), but rather in terms of distance comparisons ("Image A is more similar to image B than to image C"). Many studies in psychology give evidence that for human users such qualitative comparisons are much easier to provide than quantitative similarity scores. From a theoretical point of view, however, such an approach raises many questions: How can machine learning be performed on such qualitative input data, in which way do the approaches and results differ from the ones in the standard setting, and what kind of statistical guarantees can we give? We address these questions in our project on comparison-based learning below.

**Machine learning for scientific environments.** While machine learning methods have been used since more than a decade in some areas of science, for example in bioinformatics or the neurosciences, we currently observe a rising trend to use machine learning methods in many diverse scientific areas, ranging from social sciences over physics to geoscience. When machine learning methods are used in scientific contexts, it is of highest importance to have reliable statistical guarantees.

A prime example is the area of network science. For more than a decade, the community has focused on exploratory work, investigating properties of particular networks (social networks, brain networks, protein-networks, transport networks, etc). However, confirmatory statistical analysis is rare, even though it is of utter importance: due to the inherent randomness in networks it is not obvious how to distinguish "random artifacts" from "true structure". Hence, in one of our projects below we focus on constructing two-sample tests for populations of networks. For example, given brain networks of some Alzheimer patients and brain networks of a control group, how can we infer in a statistically solid way whether the two sets of brain networks behave similarly or not?

## 18.2 Selected Research Projects



### Comparison-Based Machine Learning

*Ulrike von Luxburg, Damien Garreau, Michael Perrot, Leena Chennuru Vankadara, Siavash Haghiri*

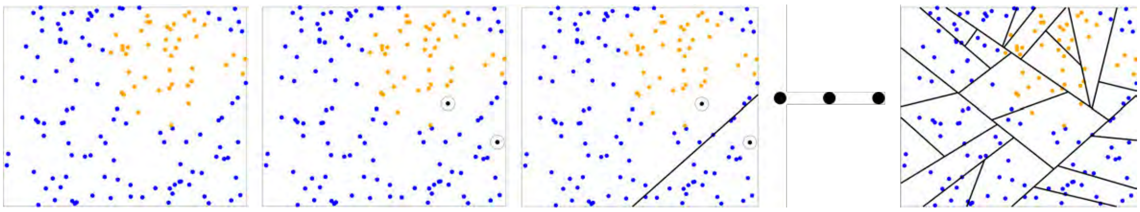

Figure 18.1: Construction of a comparison tree in the Euclidean setting: (i) current cell, (ii) pick random pivots with opposite labels, (iii) split the cell according to the ordinal comparisons involving the pivots, (iv) iterate.

In the typical machine learning setting we are given a data set $\mathcal{D}$ of objects and a distance function $\delta$ that quantifies how "close" the objects are to each other. In recent years, a whole new branch of the machine learning literature has emerged that relaxes this scenario. Instead of being able to evaluate the distance function $\delta$ itself, we only get to see the results of comparisons such as $\delta(A, B) < \delta(C, D)$, where $A, B, C, D \in \mathcal{D}$. We refer to any collection of answers to such comparisons as ordinal distance information. Our group investigates how machine learning algorithms can work with such data. We can learn a Euclidean representation that respects the comparisons but evaluating the quality of such an embedding is difficult [693]. As an alternative we proposed several algorithmic solutions that learn directly from the comparisons.

First, we developed kernels, allowing us to use kernel methods with ordinal distance comparisons [696]. In another line of work, we used ordinal information of the form "$A$ is most central among $A$, $B$, and $C$" to develop algorithms for medoid estimation, outlier identification, classification, and clustering [691]. Our approach consists in studying the intersections between disks (the "lens") and relative nearest neighbor graphs. We also proposed a method for large scale classification with generalization guarantees. The idea is to aggregate the ordinal comparisons into many weak classifiers and then to use boosting to learn a weighted combination with good predictive capabilities.

The search for nearest neighbors is another example of a classical machine learning task for which we proposed a comparison-based approach. The comparison tree [697] is a random tree constructed by recursively splitting the space by choosing random pivots and assigning data points to the nearest pivot. This structure allows for efficient search of the nearest neighbors of a query point, and we proved theoretical guarantees on the performance of this method. Since individual decision trees have a tendency to overfit on the training set, it is natural to aggregate many of these trees to construct a *random forest*. Albeit having access to very little information, these "comparison-based random forests" perform about as well as methods that have access to the true distances. We can also prove that they are statistically consistent [695]. In a related study [694], we investigated the theoretical properties of classic random forests. More precisely, we showed that the subsampling step in the construction of a random forest is important, in the sense that no subsampling or too severe subsampling can lead to inconsistency.

We also proposed clustering algorithms that are similar in spirit to agglomerative hierarchical clustering paradigms, but use only ordinal information. We theoretically proved that, for a planted hierarchical model, single linkage and complete linkage only recover the correct hierarchy for prohibitively large signal-to-noise ratios. Conversely, our average linkage based methods recover it even for small ratios.

More information: https://slt.is.mpg.de/project/comparison-based-ml

## Two-Sample Tests For Random Graphs and Networks

*Ulrike von Luxburg, Debarghya Ghoshdastidar*

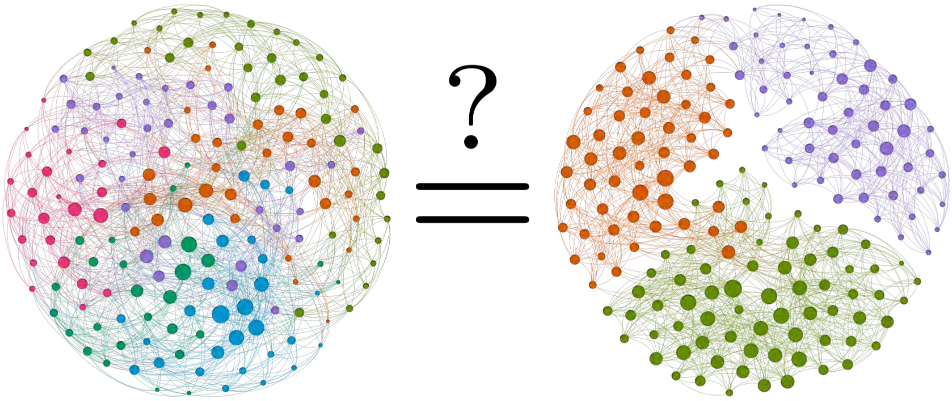

Network analysis is an integral part of many scientific disciplines ranging from bioinformatics to social sciences. A key question in this field is whether two different networks are likely to have been generated from the same underlying process or not.

This question is often formulated in terms of statistical hypothesis testing. There has been a tremendous amount of research on this topic, and empirical observations are often used to validate the hypotheses that real networks fit certain mathematical models or that two networks are dissimilar since their degree distributions or diameters are different. It is natural to wonder if there is a sound theoretical basis for these claims. Surprisingly, there has been little research on the theoretical aspects of network testing, and it is often not clear whether one can meaningfully test large networks by observing only a few replicates, which is a typical setting in applications of network analysis.

In this project, we focus on the problem of two-sample hypothesis testing of networks and provide a clear understanding of the statistical complexity of testing small populations of networks, each defined on a large number of nodes. We study the following formal problem. Given two populations of random graphs, $G_1, \ldots, G_m$ and $H_1, \ldots, H_m$, the problem is to test whether both populations are generated from the same model or not. We first consider the setting where the graphs are generated from an inhomogeneous Erdős-Rényi (IER) model defined on a common set of $n$ vertices, that is, $G_1, \ldots, G_m \sim_{\text{iid}} \text{IER}(P)$ and $H_1, \ldots, H_m \sim_{\text{iid}} \text{IER}(Q)$, where $P, Q \in \mathbb{R}^{n \times n}$ denote the model parameters. Under this setup, we test whether

$$\mathcal{H}_{\text{null}} : P = Q \text{ or } \mathcal{H}_{\text{alternative}} : d(P, Q) > \rho,$$

where $d$ is a metric, and $\rho \geq 0$ is a scalar. In the setting of fixed $m$ and $n \to \infty$, we show that popular $\chi^2$-type tests are not reliable since the associated formal testing problem is not solvable for $m \ll n$ under a minimax testing framework. We also identify two testing problems, based on the Frobenius norm and the spectral norm of $(P - Q)$, that are solvable in the small sample, large graph regime, and derive minimax-optimal theoretical hypothesis tests.

We put theory into practice and develop practical two-sample tests based on asymptotic distributions [692]. We show that these tests are both theoretically consistent in the small sample regime and exhibit good empirical performance for moderate sized networks.

We also study the more general problem of comparing two large graphs defined on different vertex sets. Standard tests in this setting are based on some network statistics. Using such testing principles, we provide a formal framework and demonstrate that the testing problem can be ill-posed, and network statistics may not distinguish highly dissimilar networks [698].

More information: https://slt.is.mpg.de/project/two-sample-tests

## 18.3 Research group leader: Ulrike von Luxburg

### Dr. Ulrike von Luxburg

Ulrike von Luxburg studied mathematics at the Universities of Konstanz, Grenoble and Tübingen. She received her PhD in Computer Science in 2004 (MPI for Biological Cybernetics / University of Berlin). After being a research group leader at Fraunhofer IPSI, she became a Minerva Research Group leader at the MPI for Intelligent Systems (2007-2012). From 2012, she was Heisenberg-Professor (W3, full professor) for Machine Learning at the University of Hamburg. In 2015 she got appointed as Full Professor for Machine Learning Theory at the University of Tübingen. In 2017 she also got appointed as a Max Planck Fellow at the MPI for Intelligent Systems.

In the period from 2003 – 2008, she received six best paper awards at leading machine learning conferences (3 COLT, 2 NIPS, 1 ALT). From 2008 – 2013 she was a member of the "Junge Akademie", the young scientists' branch of the Berlin-Brandenburg Academy of Sciences and Humanities (Berlin-Brandenburgische Akademie der Wissenschaften) and the German Academy of Natural Scientists Leopoldina (Deutsche Akademie der Naturforscher Leopoldina). More recently, she is one of the two spokespersons of the new Cluster of Excellence *Machine Learning: New Perspectives for Science*. Ulrike von Luxburg's research focus is on the theory of machine learning, in particular the statistical analysis of unsupervised learning algorithms.

#### Appointments

| | |
|---|---|
| 10/2015 – present | Full professor for Computer Science (Theory of Machine Learning) at the Department of Computer Science, University of Tübingen, Germany |
| 2012 – 2015 | Heisenberg-Professor (German W3 level) for Computer Science (Machine Learning) at the Department of Computer Science, University of Hamburg. |
| 2007 – 2012 | Research Group Leader (German W2 level), Max Planck Institute for Intelligent Systems, Tübingen |
| 2005 – 2006 | Head of the Data Mining Group at the Fraunhofer Institute for Integrated Publication and Information Systems (IPSI), Darmstadt |

#### Awards & Honors (Selected)

| | |
|---|---|
| 2008 – 2013 | Member of the *Junge Akademie*, the young scientists' branch of the Berlin-Brandenburg Academy of Sciences and Humanities and the German Academy of Natural Scientists Leopoldina. |
| 2008 | Co-winner of Best Student Paper Award at the 21rst Annual Conference on Neural Information Processing Systems (NIPS) 2008, Vancouver, Canada |
| 2007 | Co-winner of Best Student Paper Award at the International Conference on Algorithmic Learning Theory (ALT), Sendai, Japan |
| 2006 | Co-winner of Outstanding Student Paper Award at the International Conference on Learning Theory (COLT), Pittsburgh, USA |
| 2005 | Co-winner of Outstanding Student Paper Award at the International Conference on Learning Theory (COLT) 2005, Bertinoro, Italy |

| | |
|---|---|
| 2004 | Winner of Outstanding Student Paper Award at the Eighteenth Annual Conference on Neural Information Processing Systems (NIPS) 2004, Vancouver, Canada |
| 2003 | Winner of Best Student Paper Award at the International Conference on Learning Theory (COLT) 2003, Washington, USA |

### Selected Organization and Community Service (2011-2015)

| | |
|---|---|
| 2009 – 2018 | Action Editor for the Journal of Machine Learning Research |
| 2017 | General Chair for Neural Information Processing Systems (NIPS) |
| 2016 | Program Chair for Neural Information Processing Systems (NIPS) |
| 2011 | Program Chair, Conference on Learning Theory (COLT) |
| 2011–2018 | Organized 5 international workshops / spring schools in the field of learning theory and statistics |

### Board Memberships (2011 – 2015)

| | |
|---|---|
| Since 2017 | Executive Board of the NIPS conference |
| Since 2016 | Member of the steering committee of the Association for Computational Learning (runs the COLT conference) |
| 2015 – 2018 | Board member of the International Machine Learning Society (runs the ICML conference) |

# 19 CENTRAL SCIENTIFIC FACILITIES

The institute operates several “Central Scientific Facilities” (CSFs, or ZWEs in German), which are led by scientists and staffed by engineers and technicians. The objective of a CSF is to conduct research and provide high-level scientific and technical support to all departments and research groups in key research areas. Each CSF has a home campus but typically collaborates with scientists at both sites of our institute. The current CSFs in Stuttgart and in Tübingen are as follows:

**Materials (Stuttgart)** Was formed by combining a range of central scientific services in Stuttgart. The resulting new Materials CSF provides a range of thin film deposition and material science analysis techniques to support the fabrication of smart nanosystems. Current research includes investigating size effects of physical properties in nanodimensions.

**Medical Systems (Stuttgart)** Conducts research and provides scientific services related to pre-clinical and medical imaging, instrumentation, clinical applications, and related technologies. Furthermore, the CSF runs a state-of-the-art flow cytometer and provides standard cellular and molecular biology technologies to strengthen and enable advanced research in the field of biomedical milli- and nano-robotics.

**Optics and Sensing Laboratory (Tübingen)** Provides support for camera-based 3D capturing systems and different types of optical equipment. Ongoing development of the CSF focuses on the design of optical systems and computational mechanisms for computational photography and intelligent sensing. Current research includes multi-sensor systems technology and psychophysiological signal analysis.

**Robotics (Stuttgart)** Works with the researchers of the institute to design and prototype novel robotic and mechatronic systems, including mechanical assemblies, micro-electronic layout and fabrication, programming of embedded processors, system integration, control, and validation.

**Scientific Computing (Tübingen)** Builds and maintains computing and storage infrastructure for the research in the institute while also collaborating with researchers to develop scalable algorithms and applications. In the future it is envisioned that this CSF will do research at the interface of high-performance computing and machine learning.

**Software Workshop (Tübingen)** Conducts research and offers support to the institute in the areas of software engineering and computer science. It trains researchers and helps them to optimize their software development projects, ensuring professional design, implementation, testing, documentation, maintenance and distribution. The major goals of the Software Workshop are to improve the effectiveness of the research software developments in our institute, and thereby to increase the impact of our scientific publications.

The following subsections outline the past and current work of each of these CSFs in more detail.

## 19.1 Materials

Physical properties of materials change if the dimension of the investigated devices decreases. This is especially true if one of the dimensions falls below a certain threshold; often this critical size is found to be on the order of $\sim$100 nm. However, the threshold depends strongly on the physical property which is investigated. One reason for the size effects is the changing ratio of bulk volume to surfaces, it is moot whether the surfaces are free (crystal)-surfaces or internal interfaces. Also, the influence of other defects, as dislocations or vacancies increases with decreasing dimensions.

As a consequence the entirety of surfaces, defects, grains and crystal structure forms the micro- or nanostructure (German: Gefüge) which defines the behavior of devices. Since individual nanostructures are anticipated to be the building blocks for future nanodevices in several key research areas as nanorobotics, biomedicine or sensor technology, knowledge on fabrication, handling and corresponding properties of zero- (e.g., nanoparticles), one- (e.g., nanowires) and two- (e.g., ultra-thin films) dimensional nanostructures has to be gathered. Quantitative results can only be achieved by fabrication and analysis of structures with tailored microstructure.

### 19.1.1 Service

The Materials ZWE is accessible to all departments and groups of the institute as a service facility. The services can be summarised as follows:

**Micro- and nanostructure fabrication:** Physical vapour deposition (PVD) is an established bottom-up fabrication for nano- and microstructure formation. The Materials ZWE of the MPI for Intelligent Systems has extensive knowledge gathered in tailoring microstructures since its establishment in 1998. Within the concept of PVD a well-defined energy landscape on a substrate surface is used to assemble smart nanosystems by agglomeration and self-organization starting from individual atoms.

The Materials ZWE houses six state of the art sputtering and two thermal evaporation systems providing low dimensional microstructures The materials available are metals, semiconductors or ceramics. A new 10 source magnetron sputtering system for alloys and multilayers will be installed in 2019. Additionally, a 3D printer provides a tool for rapid prototyping and fabrication of macroscopic devices. This service will eventually taken over by the ZWE Robotics. For basic 3D structuring a plasma induced reactive ion etching system is accessible for users.

**Micro- and nanostructure analysis:** Microstructures can be analysed by either diffraction, imaging or spectroscopy. The two former are used for structure analysis, the latter resolves the composition of the microstructure. Since most nanostructures are well beyond the Abbe diffraction limit for light microscopy, electron microscopy is a common tool for imaging. The Materials ZWE houses a scanning electron microscope (SEM) with a point-to-point resolution of 2 nm. It is capable to image conducting, non-conducting and biological systems by using either low voltage or variable pressure conditions. X-ray diffraction (XRD) is employed for investigating textures, grain sizes, orientation relationships or phase formation. Two XRD setups are employed for the different tasks arising from the institute's needs. Chemical analysis can be carried out by either x-ray induced photoelectron spectroscopy (XPS) or energy dispersive x-ray spectroscopy (EDXS). Additionally, transmission electron microscopy (TEM) is used extensively via the StEM of the neighbour institute.

**Service policy:** The extensive knowledge of the scientists and the technicians employed in the materials lab are open for the institute's scientists and research groups along the following guide lines: *i*) Discussing of the project and the needed/desired physical properties. *ii*) Advising on the possible microstructure suitable for the investigations. *iii*) Counselling and advising on existing standard fabrication protocols. *iv*) Development of new manufacture procedures for novel devices. *v*) Helping in construction of new experimental setups under vacuum conditions.

Depending on the amount of specimen and the complexity of the fabrication procedures, the fabrication is either carried out by the student/scientist themselves enabling learning-by-doing, or full service is provided by the technicians. The former is especially encouraged for students and longer term postdoctoral researchers interested in the fabrication of their own designs. The latter is used for either very complex or time consuming standard fabrication approaches.

### 19.1.2 Science of the Materials ZWE

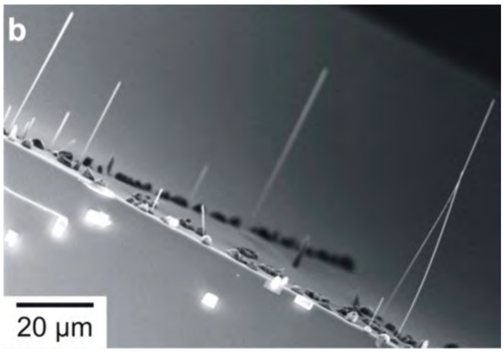


Figure 19.1: Typical SEM micrographs of Cu nanowires grown at elevated temperatures by PVD.

The Materials ZWE's own research focuses on the physical properties of unique **one-dimensional metal nanowhiskers**. One-dimensional nanostructures attracted in the past years a high interest due to their properties. Ceramics, semiconductor and carbon materials are easily synthesised in nanowire geometry with typical diameters of several dozen nanometers and length-diameter ratios of 1000:1 via methods like the vapour liquid solid process. However, only the metals as one of the oldest materials used in technology were difficult to fabricate in similar geometries. The Materials ZWE devised a unique and novel synthesis route, based on PVD, to grow similar sized metal nanowhiskers. Metals with face centered (e.g., Cu, Ag, Au, Pd), body centered (e.g., Fe), and hexagonal (e.g., Co) crystal structure were fabricated successfully.

**Microstructure** Electron Microscopy analysis on the nanowhiskers microstructure revealed that they are: *i*) **freestanding** and prismatic with no taper, *ii*) the surface facets are the **low energy surface** planes, *iii*) the **Wulff shape** dominates the geometrical shape, *iv*), **no flaws**, as dislocations or grain boundaries are observed. Overall it can be stated, that the nanowhiskers are **perfect single-crystals** (see 19.1). At the present, no other technique is able to yield such unique microstructures.

**Properties** Due to the unique microstructure the one-dimensional nanowires are perfect model systems to investigate, on a quantitative level, the structure-composition-properties relationship and to explore their size effects. The following fundamental physical properties were investigated: **mechanical strength**, **magnetic domain structure**, **conductivity**.

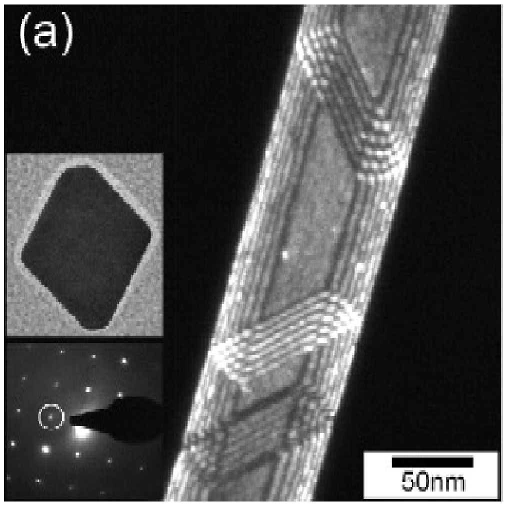


Figure 19.2: Dark field TEM image of individual stacking faults after mechanical testing.

**Mechanical properties** Mechanical reliability is dominated in bulk materials by the nucleation and mobility of dislocations. The onset of plasticity is often described irreversible dislocation interaction. Therefore, the transition from elastic to plastic behavior, the yield stress was measured for nanowires in tensile and bending configuration. The result for the yield stress ($\sigma_{Cu}$ = 6 GPa, $\sigma_{Au}$ = 1.5 GPa) reached the **theoretical strength**. It can be concluded from the experiments, that the first leading Shockley partial dislocation is nucleated at the surface of the nanowire. This leads to the formation of stacking faults as primary defects (see 19.2). Further dislocation nucleation on adjacent glide planes allow the growth of coherent mechanical twins. The nucleation event can be engineered by additional coating of the nanostructure surface. Overall, only a weak size effect is observed.

**Magnetic domain structure** Ferromagnetic behavior is often determined by the magnetic domain structure. The magnetic domain structure rearranges itself for a minimum energy configuration depending on internal and external magnetic fields. Different magnetic domains are separated by domain walls which require a certain volume. Therefore, if the structures are sufficiently small only one Weiss domain is formed. In equiaxed nanostructures superparamagnetism is observed where the thermal energy is sufficient to switch the magnetisation direction. Magnetocrystalline or shape anisotropy can stabilise the magnetizaton direction. Nanowhiskers grown from ferromagnetic metals form single magnetic domain structures with a magnetization direction parallel to its axis (see 19.3). No spontaneous reversal of the magnetic moment is observed under typical lab conditions.

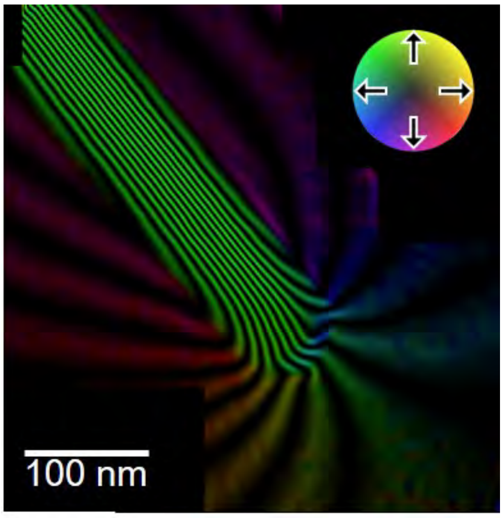


Figure 19.3: False colour electron holography map of individual Fe nanowire. Inset indicates the magnetic directions.

**Conductivity** A well known effect is the increase of resistivity by decreasing dimensions. On an atomistic level, the conductivity is determined by either electron-defect, -phonon, or -surface scattering. Since the density of defect increases with traditional fabrication methods, the conductivity decreases substantial with reduced dimensions for those structures. Since metallic nanowires are perfect single crystals without any scattering centres apart the surfaces the resistivity should be lower compared to other nanostructures. It is observed by 4-point probe measurements that perfect bulk like resistivity is achieved for diameters below 100 nm.

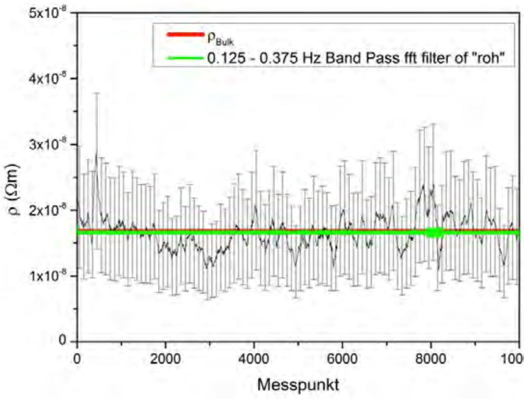


Figure 19.4: Resistivity of Cu nanowire measured by 4-point probe technique.

**Collaborations** Apart from internal institute collaborations, the Materials ZWE pursues collaborations with following research groups: Erik Bitzek (University Erlangen), Thomas Cornelius (University Marseille), Horatio Espinosa (Northwestern University), Joel Eymery (CNRS Grenoble), Dan S. Gianola (University of California, Santa Barbara), Daniel Kiener (University Leoben), Mark Legros (CEMES-CNRS Toulouse), Sang Ho Oh (POSTECH), Eugen Rabkin (Technion), Etienne Snoeck (CEMES-CNRS Toulouse), Olivier Thomas (University Marseille), and Cynthia Volkert (University Göttingen).

From the ZWE's own projects and the collaborations, 22 journal articles were published in the reporting period (see publication list).

## 19.2 Medical Systems

Figure 19.5: Installation of the seven Tesla preclinical magnetic resonance scanner on the MPI-IS campus in Stuttgart in January 2019.

The Medical Systems CSF was founded in 2016 and is still being formed. It aims to conduct research and provide scientific services related to preclinical and medical imaging, instrumentation, clinical applications, and related technologies. The Medical Systems CSF will be located in the near future in a small-animal research facility at the Max Planck Institute for Intelligent Systems in Stuttgart. In the meanwhile, instruments for preclinical imaging will be installed temporarily in the physics hall at the MPI for Intelligent Systems in Stuttgart (Figure 19.5).

The ultimate goal of the Medical Systems CSF is to improve human healthcare by developing new technologies by miniaturization of diagnostic, interventional, and therapeutic tools and agents. Non-invasive, small-animal imaging is a continuously developing field that opens new avenues for basic biomedical research. While small-animal imaging originates from classical clinical imaging disciplines of radiology and nuclear medicine, it has gained attention from nearly all medical and biological research fields over the last years. The potential of small-animal imaging goes far beyond detecting anatomical details or pathological changes in morphology using high-resolution computed tomography (CT) or magnetic resonance imaging (MRI), and it extends towards revealing complex biochemical pathways or quantitative measurements of receptor, transporter or gene expression.

Functional imaging applications rely on methodologies like positron emission tomography (PET), single-photon emission computed tomography (SPECT) or optical imaging (OI), providing excellent sensitivity to track biomolecules labeled with a radioactive isotope- or light-emitting marker. However, it is not only the optical or nuclear methods that are able to provide functional information; functional MRI, MR spectroscopy and chemical shift imaging have evolved to become promising tools for detecting changes in blood flow, tissue oxygenation or concentrations of endogenous molecules such as lactate or choline.

The manifold options dedicated to small-animal imaging modalities become even more powerful when they are combined as multimodality systems, allowing temporal and spatial correlation of the data with high accuracy. The strength of PET lies in its very high detection sensitivity in the picomolar range and accurate quantification, but PET lacks good spatial resolution and tissue contrast. MRI, however, enables high-resolution imaging of morphology with good soft-tissue contrast, detects endogenous metabolite distributions by spectroscopy and allows dynamic acquisition of tissue perfusion and additional functional parameters. Therefore, PET/MRI paves the way for non-invasive imaging to be used in a wide range of applications in biomedicine that have historically been reserved for classical *in vitro* and *ex vivo* molecular biology methods. Nowadays, PET/MRI systems are well established in preclinical imaging and are progressing into clinical applications to provide further insights into specific diseases, therapeutic assessments, and biological pathways.

### 19.2.1 Service

The Medical Systems CSF will use a seven Tesla dedicated small-animal MRI for several research projects within the MPI for Intelligent Systems. The device will enable a wide variety of research across the areas of medical robotics and medicine. Novel microrobotic and millirobotic technologies have the potential to revolutionize medicine, enabling high precision therapies directly at the site within the body that requires intervention. Millimeter-scale tethered and untethered robotic devices can deliver drugs, perform interventions and close the loop between the gathering of therapeutic application and diagnostic outcome. The small-animal MRI system will serve as a platform for real-time imaging, feedback and control of magnetic microswimmers for precise drug delivery. The gradient system of the MRI can be used to apply forces on magnetic biohybrid microrobots for real-time steering and custom software will be integrated into the MRI system to allow for operator-guided or fully automated steering and imaging of the biohybrid microrobots. After moving the MRI system to the final installation site in the planned small-animal research facility, the Medical Systems CSF will be able to use the combined and simultaneous PET/MRI system to provide localization of the microrobots when they are labeled with radioactive isotopes.

### 19.2.2 Technical Support Activities

Providing support in biological research activities is the other main goal of the Medical Systems CSF. The CSF provides and runs a state-of-the-art flow cytometer (Fluorescence-Activated Cell Sorting, FACS) to strengthen and enable cutting edge research in the field of biomedical milli- and nano-robotics. Not only *in vitro* cell culture experiments and *ex vivo* analyses can be performed using this device, but also the boundaries of science and technology in various research fields located at the MPI campus can be pushed using this technology. Standard cellular and molecular biology technologies complete the equipment for advanced research in the field of biomedical applications.

## 19.3 Optics and Sensing Laboratory

The activities of the Optics and Sensing Laboratory (OSLab), established in 2014, include cross-institute technical research support and original research work.

OSLab provides design, training and technical support for various motion capturing systems, the institute's telescope dome and the variety of camera-based equipment. The lab members are involved in research work in collaboration with departments and other ZWEs.

In particular, the laboratory focuses on original research in the fields of multi-sensor systems technology, augmented reality (AR) visualization, and psychophysiological signal analysis. At the same time, OSLab conducts experiments on self-organization of magnetic droplets.

### 19.3.1 Team, Facility, and Equipment

As for 2019, the OSLab team consists of two research engineers with expertise in multi-sensing technology and optics, two engineers, two research assistants, an intern, and a lab assistant. Furthermore, OSLab is in charge of two technical members of the PS department. In addition, several master's theses from the University of Tübingen were supervised by OSLab members.

OSLab facilities include video and AR within a capturing space focusing on motion capture sensor integration and AR experiments. This space offers a professional truss system, high-resolution, high-speed video cameras, synchronized camera and Inertial Measurement Unit (IMU)-based motion capturing systems, and controllable lighting.

OSLab also runs a dedicated optics and laser space, offering stable optical tables and the possibility to develop custom laser applications. The available lasers with precisely controlled characteristics (such as wavelength, output power, beam shape) are used for metrology (interferometry) and in various research projects (e.g., laser speckle for depth estimation). In addition to testing of existing optical equipment, OSLab possesses the expertise to simulate optical systems as well as design entirely new optics using the software Zemax.

The institute's experimental facilities for computation photography have been mainly set up and are operated by OSLab. They include a photography lab and a telescope dome.

The current sensing equipment of the laboratory consists of a portable biofeedback logger from Mindmedia, an IMU setup, audio recording equipment, high-speed and high-resolution cameras, and studio video fixtures and light measurement devices. In addition, we devised an experimental setup for AR interaction with a head-mounted display (Meta 2) and a portable computer backpack that allow us to execute highly dynamic AR experiments and integrate computationally demanding algorithms developed in our institute.

In 2018, the lab was granted a budget of 200k EUR for basic research equipment for the institute. Two systems were purchased: a high-resolution, high-speed laser scanner for digitization of indoor and outdoor environment, and a middle-IR-band camera to perform remote temperature analysis.

### 19.3.2 Research Support Activities

We would like to mention several projects that were conducted by OSLab:

**Capture facility** The OSLab is overseeing the institute's effort to establish the world leading multi-sensor capturing facility, which will allow the institute to acquire unique data to advance computer vision and machine learning research. In particular, we installed a 54-camera VICON MoCap system on a flexible truss system, and are in progess of integrating studio lights as well as 20 high-resolution, high-speed cameras. A new foot scanner from 3dMD Ltd. was installed and upgrades of lighting at body and face scanners were successfully executed. As of 2019, a new highly flexible multi-scale hand scanner will be co-designed with 3dMD Ltd and installed in the Tübingen laboratory of the Haptic Intelligence department.

**High-resolution face scanner** We developed a high-resolution face scanner solution base on 4 $\times$ 12M cameras and 4 $\times$ 5M cameras. Our system is using AgiSoft library to implement

3D reconstruction on the face, and custom algorithms for camera calibration. The scanning results are meeting the state-of-the-art accuracy. In addition, the project serves as a building block for 20 cameras capturing system that is currently under development in our laboratory.

**PSF panel** The OSLab participated in the design and construction of a novel optics test-panel, the so-called PSF-panel. It allows capturing the response of the optics to a single point source (point-spread function, PSF) over the entire field in a single shot. Members of OSLab used the PSF-panel to generate a ground-truth data set containing the PSFs of a large number of commercial DSLR lenses. Naturally, OSLab was strongly involved in the research project 'Automatic Estimation of Modulation Transfer Functions' based on this data set, which eventually led to a scientific publication [210]. Currently, the data set is being extended to different zoom settings and focusing distances of the lenses for a research project on PSF-estimation. In 2019, the data set will be further extended to a larger number of camera bodies.

**Telescope array** The OSLab was responsible for setting up a remotely operable observatory for the EI department in the telescope dome on the institute's roof. At heart of the observatory is a multi-telescope array, currently consisting of six telescopes, as shown in Fig. 19.6. The telescope array is expected to generate enormous amounts of data that will be used for superresolution and ptychography algorithms.

Figure 19.6: Telescope array of the EI department.

**Wave-front sensing** In addition to regular image sensors, a wave-front sensor (WFS) was devised and built by the OSLab in collaboration with the institute's mechatronics workshop. A setup consisting of an artificial star (large, collimated laser beam), a turbulence simulator, and the WFS is currently used in a research project, which aims at precise image-deblurring using wave-front data.

**Speckle for depth estimation** The OSLab is collaborating with the Autonomous Vision group on developing a high-resolution laser speckle. This laser pattern will be used for novel algorithms for depth estimation. As such algorithms require a high signal-to-noise ratio, a custom high-power diode laser is devised and implemented by the OSLab in order to reach resolutions of up to 500k dots.

These projects have a direct impact on the success of institute's research based on innovative data. As part of these projects, a highly qualified technical team was established and trained to maintain and support the future development of the facilities.

### 19.3.3 Research Activities

**Psychophysiological analysis** Human psychophysiological analysis is one of the prime research interest of the lab, in particular, detection and analysis of acute-stress by utilizing non-invasive, camera-based sensors. The ability to detect stress level is an essential technology for a variety of medical applications as well as more recent challenge of passenger stress analysis for autonomous vehicles.

Figure 19.7: Experimental setup for psychophysiological analysis of acute stress.

Our effort in this direction include the collection of a dataset during acute stress situations. It contains data from high-resolution cameras, an IR camera, biofeedback physiological data as well as biological ground-truth based endocrinological markers samples extracted from saliva (see Fig. 19.7). In addition, we continue our efforts in facial micro-expressions analysis for human emotional state prediction. We are developing a highly flexible and customizable GUI tool for facial expression ground-truth tagging that will be open for the research community.

In the future, we are going to combine the research effort in emotional human state prediction and analysis with visualization techniques based on AR technology. In collaboration with the PS department, we are looking into establishing a new synthetic emotional data set, which will be used for training deep networks for facial micro-expression detection and analysis.

**3D scanner accuracy estimation** In 2016, OSLab was granted a grassroots budget of 24k EUR for investigating the development of multi-scale scanner technology. As a result of this project, an automatic accuracy estimation methodology for multi-scale scanners was developed.

**Ultrasound fields** Together with the *Micro, Nano, and Molecular Systems Group*, the OSLab is investigating the potential of ultrasound fields. On the one hand, ultrasound can be used for haptic feedback, e.g., creating haptic perseption in mid-air. In the future, OSLab will integrate ultrasonic haptics into AR technology. On the other hand, ultrasound fields can be used to levitate and manipulate small objects. In this context, the OSLab investigates the rich self-assembly behavior of magnetic droplets.

## 19.4 Robotics

The Robotics Central Scientific Facility was established in 2017 to support the institute's research objectives via the design, rapid prototyping, and validation of novel robotic and mechatronic systems. Melding together the complementary disciplines of mechanical design, electrical engineering, embedded programming, and systems integration has made possible the vast majority of the technology that pervades our lives, from modern automobiles to smartphones and surgical robots. The majority of the work within the Robotics ZWE takes place in close collaboration with institute researchers. For those with completed designs, the Robotics ZWE also offers direct additive manufacturing and electrical circuit prototyping services.

### 19.4.1 Mechanical Design & Prototyping

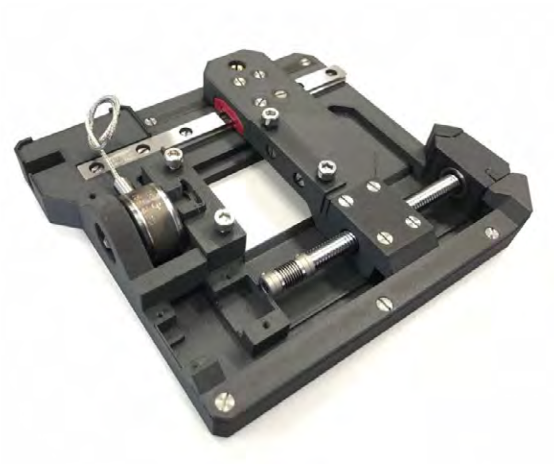

The design and rapid prototyping of mechanical structures can be accomplished using a variety of processes and techniques, including additive manufacturing, traditional subtractive techniques, casting, molding, and many more. When thinking about the integration of intelligent components (sensors, actuators, embedded electronics and processing), certain methods may be more or less viable for a given project.

**Computer-Aided Design** - Capturing geometric constraints, modeling sourced components, design of custom hardware, virtual assembly, motion, stress, and/or interference analysis, bill of materials generation, preparation of files or documents for prototyping and manufacturing

**Rapid Prototyping** - additive (3D printed) manufacturing of numerous materials, subtractive machining, laser cutting, composite layup, polymer casting

### 19.4.2 Electronics Design & Fabrication

The majority of embedded electronic systems require the integration of compact surface-mount components on either rigid or flexible substrates. The design of such circuitry often requires numerous iterations, and traditional methods can take weeks to produce. To accelerate the prototyping cycle, the Robotics ZWE has invested in the expertise and equipment to rapidly design and internally prototype fine-pitch, multi-layer, and flexible custom circuitry.

**Design** - includes the definition of requirements, selection of components, and specification of processes, taking into consideration what tools we have in house as well as our outsourcing vendor capabilities

**Capture** - schematic design and board layout

**Fabrication** - production of multi-layer rigid and/or flexible circuits, including fully printed three-dimensional electronics

**Assembly** - placement and reflow soldering of fine-pitch surface-mount components and hand-soldering of through-hole elements

**Test** - ensuring that the design and prototyping fit the requirements

### 19.4.3 Embedded Software Development

Many robotic and mechatronic systems require specific input, output, and computation functionality to be embedded within constrained, mobile, or remote environments. Yet other systems require wired and/or wireless interconnections between existing systems or various other computational elements.

**Microcontrollers** - experience with all major microcontroller architectures (AVR, ARM, TI, Freescale, Microchip, etc.)

**Programming Languages** - proficient in a variety of programming languages (C/C++, Java, Python, Matlab, etc.)

**Communication Protocols** - extensive experience with various communication protocols (I2C, SPI, USB, CAN, etc.)

### 19.4.4 Additive Manufacturing

The design of complex integrated mechanical systems usually benefits significantly from the ability to rapidly iterate on the design and specific features of mechanical structures. For this reason, the Robotics ZWE has established a suite of additive manufacturing tools spanning both the most common techniques as well as some of those on the cutting edge of technology.

Originating in the late 1980s, the field of additive manufacturing ("3D Printing") continues to evolve and grow at an astonishing rate. The fundamental principle is the layered deposition of material to fabricate given three-dimensional geometry. The most common techniques are:

**Fused-Deposition Modeling (FDM)** - a thin strand of plastic is heated to near it's melting point, and laid out by the planar movement of a print head over a flat platter. FDM requires supporting material for overhanging geometry.

**Stereolithography (SLA)** - lasers are used to selectively cure voxels within the top layer of a vat of UV-sensitive resin. SLA benefits from the uncured resin acting as a supporting material.

**Multi-Jet Modeling (MJM)** - a print head, similar to an inkjet printer, is used to deposit small droplets of UV-curing polymer that are then cured by a UV lamp. Some MJM systems enable the user to alter the material stiffness and other properties within a single part.

**Selective Laser Sintering (SLS)** - lasers are used to sinter adjoining particles of material. SLS is self supporting, and can be used for polymers as well as metals.

Choosing a particular process and machine for a given part depends upon the complexity of the geometry, the required bulk mechanical properties, cost, time, and a number of other factors. The Robotics ZWE is currently operating seven different 3D printers spanning the range of available technologies described above.

## 19.5 Scientific Computing

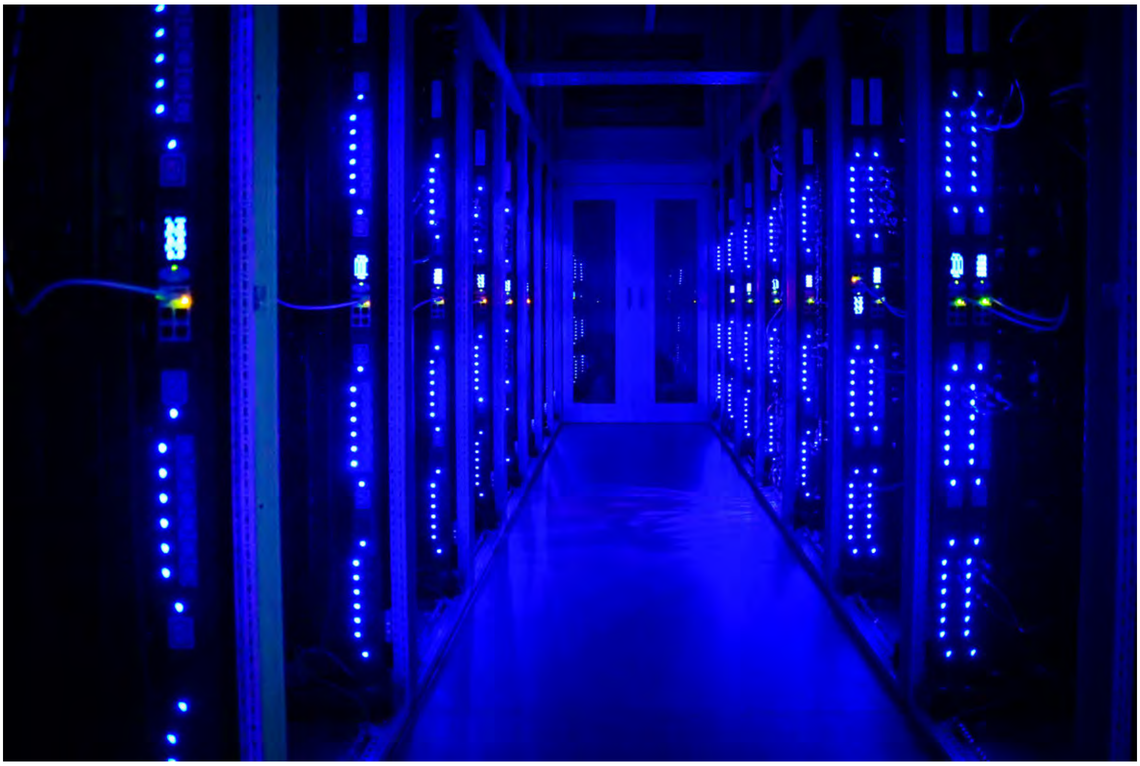

The primary responsibilities of the Scientific Computing Central Scientific Facility are to build and maintain computing and storage infrastructure for the research at the institute, and to collaborate with researchers in developing scalable algorithms and applications. It is also envisioned that the group will carry out independent research in the areas overlapping high performance computing and machine learning.

### 19.5.1 Service

The computing cluster, which started as part of the Department of Empirical Inference before being integrated and rebuilt as part of this Central Scientific Facility, was expanded in multiple stages to be able to support the growing needs from various departments and independent research groups. The cluster currently has 50 rack-mounted multi-processor nodes, with more than 2300 Intel and AMD 64-bit CPU cores, many of them with 1 TB RAM, as well as 190 modern Nvidia GPUs including the latest Tesla V100s. It also has 350 terabytes of fast, distributed storage space. The nodes are connected via 100 Gbps fiber-optic connections.

As a large number of the applications running on the cluster are data intensive, for which I/O performance is crucial, a parallel file system Lustre[1] is used. This is an open-source piece of software with good community support that is being used by many institutions around the world, and hence the possibility of vendor lock-in is avoided. HTCondor[2], again an open-source piece of software, is being used as the scheduler on the cluster, with a number of in-house improvements and adaptations to tailor it to the specific needs of the researchers. Based on the experience accumulated from the previous EI cluster, a new cluster banking system was developed taking into account the additional features that were needed. In setting up, maintaining and in some cases even building these tools, the Scientific Computing CSF has acquired a good amount of in-house expertise that allows it to keep adapting the software systems to the specific needs of the scientists.

The Scientific Computing CSF has taken ac-

[1] P. Schwan. Lustre: Building a File System for 1,000-node Clusters. In *Proceedings of the Linux Symposium*, pages 380–386, July 2003.

[2] D. Thain, T. Tannenbaum, M. Livny. Distributed computing in practice: the Condor experience. *Concurrency - Practice and Experience* **17** (2-4): 323–356, 2005.

tive interest in helping the users to use the cluster resources as efficiently as possible. Periodic tutorials were given on how to use the cluster effectively. The CSF also promotes the use of resources provided by Max Planck Computing and Data Facility at Garching, and works together with researchers to identify the best-suited use cases, and helps them port their applications so that they can be run there.

The CSF coordinated with the building contractors to ensure the successful completion of the new server room in the new building. The cluster was then successfully moved to and set up in the new server room at the beginning of the year 2018.

The group is continuously working on further expansion of the cluster to address the growing computing needs in the institute, closely working together with the departments, research groups, institute administration, and Max Planck Society administration to identify the needs, secure funding, and procure state-of-the-art hardware in a timely manner. Recently, a number of researchers from the Stuttgart campus of the institute has started using the cluster. It is likely that more users from Stuttgart start using the cluster in the near future, and the Scientific Computing CSF needs to support those users in a regular manner. Members of the group work closely with researchers to promote the use of cluster resources and collaborate in building scalable applications. The CSF also liaison with external agencies to support research collaborations; an example is the partnership with Nvidia through NVAIL[3].

## 19.5.2 Research

The computing cluster is growing very fast, with addition of very powerful GPU nodes. The energy consumption is also growing in a similar fashion. Hence, it is in our interest to ensure that the cluster is used in an energy efficient manner. Energy aware scheduling of jobs can potentially improve the energy usage by the cluster. The new server room is already equipped with facilities to measure energy usage by individual nodes, and we are currently building infrastructure to measure temperature at different parts of individual racks. Our cluster has an occupancy pattern with peaks and valleys that often depends on conference deadlines. It would be beneficial to shutdown part of the cluster when not used. For this, the scheduler should schedule jobs such that maximum number of nodes is kept free to be shut down, reducing fragmentation. Predicting run time of jobs for this purpose is an ongoing work. Another possible improvement is by scheduling jobs on nodes that are in areas with lower temperature, resulting in improved overall power consumption.

We usually buy the highest performing GPUs, while keeping a balance in the price-performance ratio. Because of this we usually have to pass on technologies such as NVLink, as their high price makes them unattractive. Some of these technologies could nonetheless offer perfomance wins for some of the applications that can use multi-GPUs, thus making them worthwhile overall. It would be interesting to understand the performance of machine learning/computer vision programs with respect to GPU platform architecture, since this would provide us with useful insights when facing this kind of decisions. To this aim, we have already initiated a project with external collaborators to investigate this aspect.

Large-scale machine learning is an area of research interest for the group. We intend to use high performance computing techniques to increase the computational performance of applications and algorithms dealing with large data sets (big data) and work in collaboration with departments and research groups.

[3] URL: https://www.nvidia.com/en-us/research/.

## 19.6 Software Workshop

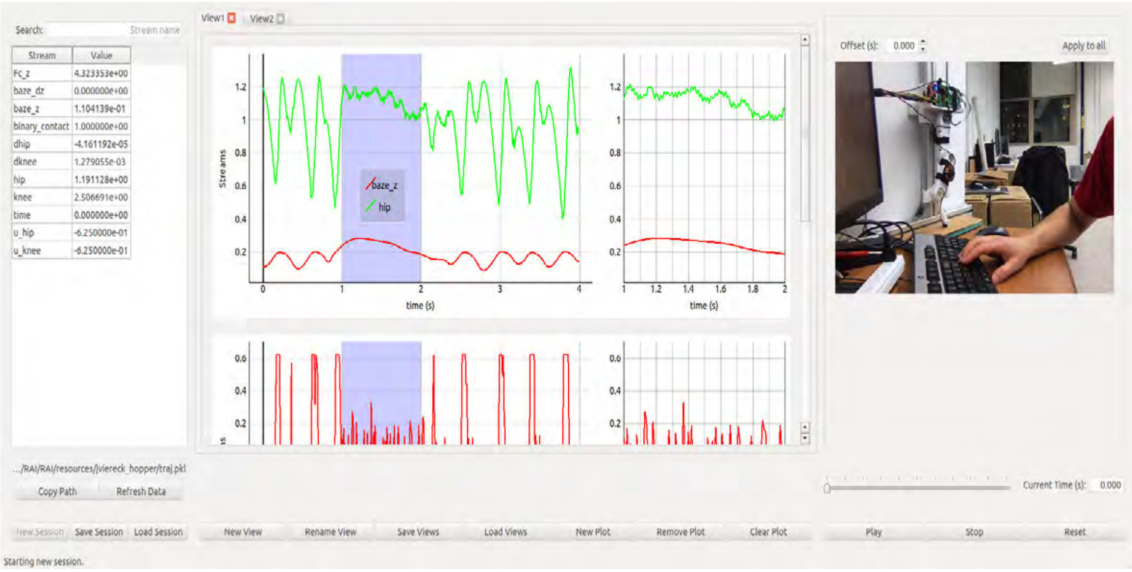


Figure 19.8: graphical user interface for analyzing data from robotic experiments.

The Software Workshop (SW), established in 2014, is a unique scientific facility at the intersection between software engineering and research. It conducts research in the area of computer science and support the departments and research groups of the institute in software development. Its main goal is to increase the quality and impact of the research at the institute, both internally and worldwide. Over the last few years, the standards and requirements for software development have significantly increased. As a consequence, a scientific publication benefits from providing the implementation in the following ways:

1) **Improving scientific quality**: the availability of the implementation enables the reproduction of scientific results, provides important details often not part of the publication, and sustains scientific credibility;

2) **Improving scientific visibility**: the availability of the code promotes a better scientific reference that can be used for comparison/benchmark, and promotes collaborations;

3) **Improving scientific impact**: releasing the scientific code enables an audience not primarily targeted (e.g., industries, open-source communities) to integrate scientific work published by the Max Planck Institute into their products.

Researchers need to integrate these aspects into their workflow. However designing, developing, and maintaining a high-quality software product requires time and specific expertise in computer science and software engineering, which researcher usually do not have. The SW is therefore meant to help them regarding these topics at the different phases of their research projects.

### 19.6.1 Team, Expertise, and Equipment

Between 2016 and 2018, three scientists, one software engineer, one post-doc, one Ph.D. candidate, and eight students have worked at the SW. As of January 2019, the SW is composed of two scientists and two research assistants (HiWi). The SW collaborates with all entities of the institute (departments, research groups, ZWEs) which involves covering a large variety of research topics. Consequently, the SW members must be versatile and have a general scientific background. The SW expertise includes several programming languages (C++, Python, Matlab, CMake, CUDA), technologies for software development (Django, Qt, OpenCV, Ansible, Boost, iOS), workflow management techniques (Agile, Scrum, Atlassian tools), and scientific fields (computer science, algorithmics, linear algebra, numerics, physics).

### 19.6.2 Service

This section describes the tasks of the SW as a service group.

**Training** The SW regularly organizes events to promote good software development practices and teach researchers about state-of-the-art technologies that could be useful to them. These events can take several forms: general cross-site presentations, hackathons, public code reviews, targeted tutorials, or individual meetings.

**Infrastructure for Software Development** The SW has set up, maintains, and develops the full infrastructure used by the institute for code development. It includes a wiki accessible from both sites (Confluence), secured applications for sharing data (Owncloud) and organizing meetings (Terminplaner), servers for hosting codes (GitLab, GitHub), communication tools (HipChat), workflow management engines (Jira), specific applications to monitor the quality of code developments over time (Bamboo, Crucible), and an internally-developed application for sharing released codes and their documentation (CodeDoc).

**Development of industry standard software** Most codes written and published by scientists are not matching the industry standards. The SW helps scientists to turn their working prototypes or "proofs of concept" into stable, well-documented, carefully tested, more efficient (e.g., through multi-threading) and maintainable software products that can be proudly shared with external collaborators and communities.

### 19.6.3 Research and Projects

This section highlights the different types of projects managed by the SW.

**Software Engineering Projects** The SW has developed an in-depth know-how for building infrastructures to handle scientific data. They allow researchers to collect, organize, analyze, and share their data among themselves efficiently and reliably. The SW has implemented several of these solutions for different departments, including a web application for collecting body scans and a graphical user interface to analyze simultaneously sensors data and videos from robotic experiments (Figure 19.8).

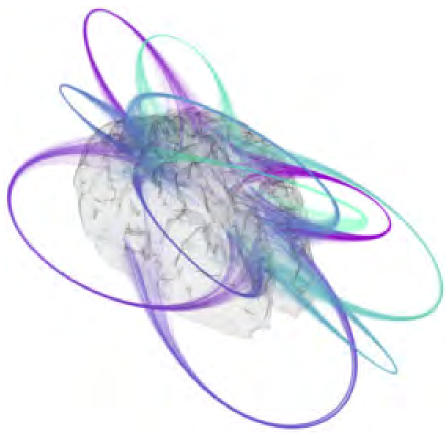

Figure 19.9: results of an algorithm used to cluster and visualize connectivities inside the brain.

**Research Projects** Several research projects were undertaken by the SW, either internally or in collaboration with other groups at the institute. Among others of these projects, the SW developed a Brain-Computer interface for collecting patients' data, implemented a fast, cross-platform, flexible $k$-means algorithm, and a visualization scheme for functional connectivity data in a human brain (Figure 19.9). The SW workshop is also a member of the *memmo*[4] consortium.

**External Projects** The SW is also involved in a few external projects usually connected to the background of some of its members. For instance, the SW has contributed to the open-source projects *Boost* and *CMake* The SW is also part of the *Enzo*[5] and the *Nugrid*[6] scientific collaborations for modeling and understanding astrophysical phenomena, and has provided numerical solutions to these efforts.

Within the last three years, more than 15 projects have been completed and one journal article has been published (see publication list).

[4] http://www.memmo-project.eu
[5] http://enzo-project.org
[6] https://nugrid.github.io

# 20 OUTREACH


## 20.1 Public Relations

In recent years, the field of artificial intelligence has attracted growing attention as AI has played an ever-greater role in the world. Accordingly, public interest in the research activities of the Max Planck Institute for Intelligent Systems (MPI-IS) has also risen. Especially since the Cyber Valley initiative gained momentum in 2017, media coverage of the institute and its research has increased significantly, both in Germany and around the world.

As a publicly funded institution, MPI-IS has always been accountable to a range of stakeholders, from the general public to policymakers. The institute's public relations activities have thus focused on making the research of MPI-IS transparent and understandable to a broad audience.

Since 2016, the institute has further enhanced its press and public relations efforts in response to the growing need for external communication. The current approach is based on the four pillars described in the following subsections.

### 20.1.1 Press and Media Relations

MPI-IS has continued its traditional media relations activities with press releases and news items, both of which have recently increased in number.

Stories about the institute and its activities have appeared in regional media such as the Schwäbisches Tagesblatt, German national newspapers including the Süddeutsche Zeitung, Frankfurter Allgemeine, and Handelsblatt, as well as international dailies such as the Financial Times, Le Monde, and the Guardian. Cyber Valley was also featured on CNBC's The Edge, and the Millirobot appeared in the evening news on the German public broadcasters ZDF and SWR, as well as on Al Jazeera. MPI-IS was also the subject of radio reports on SWR and Deutschlandfunk in Germany, as well as on the BBC in the UK and CBC Radio in Canada.

### 20.1.2 Social Media

MPI-IS has increased its presence across social media channels since 2016, and the number of followers and subscribers has grown steadily. As of January 30, 2019, the institute's YouTube channel counted 4500 subscribers. It had over 1620 followers on Twitter, 946 on LinkedIn, and more than 370 on Instagram. On Facebook, the MPI-IS following has doubled since 2017, and stood at 2172 in January 2019.

### 20.1.3 Visitors' Groups and Delegations

Over the past two years, MPI-IS has welcomed many visitors' groups and delegations, to whom scientists and technicians have presented their research. Prominent visitors from politics, industry, and academia have included: Austria's President Dr. Alexander Van der Bellen, Baden-Württemberg's Minister President Winfried Kretschmann, Baden-Württemberg's Science Minister Theresia Bauer, the U.S. Consul General, Hungary's Minister of Technology and Innovation. Industry representatives included Robert Bosch GmbH, Daimler Trucks and Process Engineering, Hyundai Motor Company, IAV automotive engineering, Nvidia Corp., IBM Deutschland Research & Development GmbH, Samsung Electronics GmbH, Schwarz-Gruppe, BASF S. E., and Braun Melsungen AG.

Figure 20.1: In 2018, hundreds of visitors came to the Open House in Tübingen

### 20.1.4 Public Outreach

The institute regularly organizes a broad range of events for the general public, including interested laypeople, youth, and scientists.

**Open House:** MPI-IS opened its labs to the public twice in the period under review. About 1000 visitors attended the open house day in Tübingen in 2016. In 2018, the event was held simultaneously at both the Stuttgart and Tübingen sites, receiving 2500 and 1200 visitors, respectively.

Figure 20.2: Open House 2018 in Stuttgart: a young scientist shows a girl fascinating things.

**Max Planck Day 2018:** Max Planck Day was held on September 14, 2018, as part of a nationwide campaign that celebrated science. Both MPI-IS locations took part with the following events:

**AI and Society Symposium:** The Tübingen campus organized an AI and Society Symposium, which was followed by a panel discussion with Prof. Bernhard Schölkopf of MPI-IS, Prof. Thomas Hofmann of ETH Zürich, Dr. Michael Bolle of Robert Bosch GmbH, Dr. Ralf Herbrich of Amazon Development Center Germany, Prof. Sarah Spiekermann of Vienna University of Economics and Business, Dr. Sandra Wachter of Oxford Internet Institute, and Prof. Andreas L. Paulus of Georg-August-Universität Göttingen (who is also a member of the German Federal Constitutional Court). A total of 450 people attended this symposium.

**Science Slams:** On the Max Planck campuses in Tübingen and Stuttgart, Science Slams were held that were attended by 250-300 people at each site. At each event, scientists explained their research in an entertaining way that was easy to understand, and the audience chose the winner at the end.

**Summer Colloquium in Stuttgart:** The Summer Colloquium takes place every year in July and offers four public lectures on current scientific topics. The Günter Petzow Prize is also awarded on the occasion. The last three winners were as follows:

- 2016: Dr. Xing Ma: "Enzyme-Powered Mesoporous Silicia Micro/Nano-Motors: Towards Biocompatibility"
- 2017: Dr. Markus Weigand: "Exploring the Nanoscale with Soft X-rays"
- 2018: Dr. Wenqi Hu: "Smart Magnetic Soft Material and its Application in Miniature Robot" (see Fig. 20.3)

Figure 20.3: Wenqi Hu (second from left) accepting the 2018 Günter Petzow Prize, along with Günter Petzow, Metti Sitti, and Katherine J. Kuchenbecker.

**Max Planck Lecture in Stuttgart:** Every year, the institute's Stuttgart site organizes the Max Planck Lecture with top-class speakers and interesting topics on intelligent systems. Scientists and interested laypeople are invited to attend the event, which generally attracts 150-200 participants. The last three lectures were:

- 2016: Prof. Naomi Ehrich Leonard of Princeton University, USA: "On the Nonlinear Dynamics of Collective Decision-Making in Nature and Design"
- 2017: Prof. Amnon Shashua of Hebrew University and Mobileye Vision Technologies Ltd.: "The Three Pillars of Fully Autonomous Driving"
- 2018: Prof. Roland Siegwart of ETH Zürich: "Autonomous Robots that Walk and Fly"

**Inauguration ceremony of the new MPI-IS building in Tübingen on July 12, 2017:** 120 guests were invited to the event, among them the Minister-President of Baden-Württemberg Winfried Kretschmann, the state Science Minister Theresia Bauer, the President of the Max Planck Society Martin Stratmann, several members of the state parliament, and Cyber Valley partners.

**Girls' Day:** Each year, 40-50 girls are invited to the Girls' Day event at both MPI-IS sites to gain insights into science and technology professions.

**TÜFFF (Tübinger Fenster für Forschung):** The Tübinger Fenster für Forschung (Tübingen Window for Research) takes place every two years at the University of Tübingen and is similar to a Long Night of Science. In 2017, MPI-IS took part in the event with an information booth, where a dancing Nao and the Mosh camera app allowed interactive play.

## 20.2 Public Events

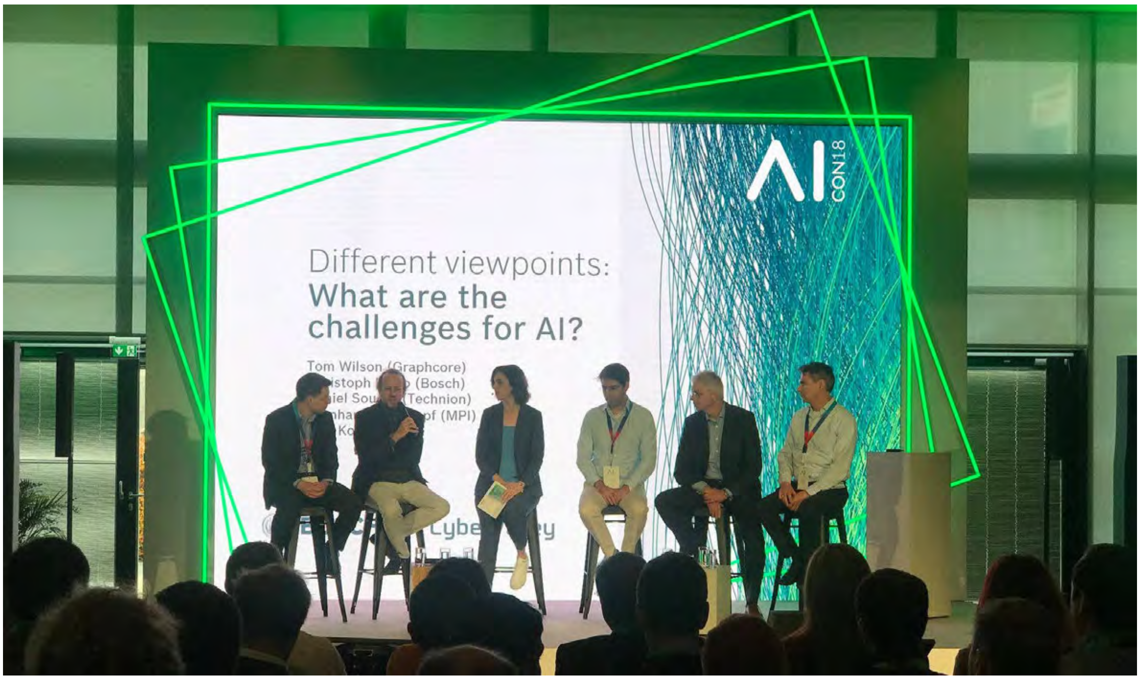


### 20.2.1 Public events hosted by the institute (selection)

| Date | Organizers | Description | Location |
|---|---|---|---|
| 2018-11-29 | Matthias Tröndle, Julia Braun | Visit by the President of Austria Dr. Alexander Van der Bellen and Baden-Württemberg's Minister-President Winfried Kretschmann | MPI-IS Tübingen |
| 2018-11-26 | Matthias Tröndle, Julia Braun | Visit of members of Baden-Württemberg's state parliament, Nico Weinmann and Daniel Karreis | MPI-IS Tübingen |
| 2018-10-26 | Matthias Tröndle, Linda Behringer | Visit by Baden-Württemberg's Science Minister Theresia Bauer | MPI-IS Tübingen |
| 2018-10-23 | Tamara Almeyda, Julia Braun | Visit of members of the German Bundestag, Andreas Steier and Mario Brandenburg | MPI-IS Tübingen |
| 2018-10-22 | Katherine J. Kuchenbecker, Barbara Kettemann | Max Planck Lecture: "Autonomous Robots that Walk and Fly", Prof. Roland Siegwart, ETH Zurich | MPI-IS Stuttgart |
| 2018-09-15 | Linda Behringer | Open House in Stuttgart: Erleben Sie Grundlagenforschung hautnah! | MPI-IS Stuttgart |
| 2018-09-15 | Claudia Däfler | Open House in Tübingen: Come and visit our labs and learn about our research! | MPI-IS Tübingen |
| 2018-09-14 | Linda Behringer | Science Slam for Max Planck Day in Stuttgart | MPI-IS Stuttgart |

| | | | |
|---|---|---|---|
| 2018-09-14 | Bernhard Schölkopf, Julia Braun, Matthias Tröndle | Max Planck Day in Tübingen: Symposium on Artificial Intelligence and Society | Kupferbau Tübingen |
| 2018-09-14 | Claudia Däfler | Max Planck Day in Tübingen (evening): Science Slam - Light Installation | Max Planck Campus Tübingen |
| 2018-07-13 | Katherine J. Kuchenbecker, Matthias Tröndle, Claudia Däfler | 2018 Intelligent Systems Summer Colloquium with the 2018 Günter Petzow Prize award ceremony and a celebration in honor of Prof. Manfred Rühle | MPI-IS Stuttgart |
| 2018-06-28 | Matthias Tröndle, Julia Braun | Visit of members of Baden-Württemberg's state parliament, Andreas Stoch and Gabi Rolland | MPI-IS Tübingen |
| 2018-06-06 | Linda Behringer | Visit of students from the Königin-Olga-Stift | MPI-IS Stuttgart |
| 2018-04-28 | Linda Behringer, Claudia Däfler | Girls' Day in Stuttgart and Tübingen | Both MPI-IS sites |
| 2018-04-27 | Tamara Almeyda | Visit of students from Stiftung Kinderland Baden-Württemberg | MPI-IS Tübingen |
| 2017-09-29 | Claudia Däfler | "Hausbesuch bei den Robotern" – Lab tour as part of of the "Robots" exhibition | MPI-IS Tübingen |
| 2017-09-18 | Annette Stumpf | Max Planck Lecture: "The Three Pillars of Fully Autonomous Driving", Prof. Amnon Shashua, Hebrew University and MobilEye | MPI-IS Stuttgart |
| 2017-07-12 | C. Däfler, M. Tröndle, J. Braun, D. Rebmann | Opening ceremony of the new institute building in Tübingen | MPI-IS Tübingen |
| 2017-07-07 | Metin Sitti, Annette Stumpf | 2017 Intelligent Systems Summer Colloquium with the 2017 Günter Petzow Prize award ceremony | MPI-IS Stuttgart |
| 2017-02-08 | Annette Stumpf | Visit of students from the Königin-Olga-Stift | MPI-IS Stuttgart |
| 2016-07-09 | Metin Sitti, Annette Stumpf | 2016 Günter Petzow Colloquium with the 2016 Günter Petzow Prize award ceremony | |
| 2016-06-18 | Claudia Däfler | Tübingen Open House | MPI-IS Tübingen |
| 2016-06-06 | Annette Stumpf | Max Planck Lecture: "On the Nonlinear Dynamics of Collective Decision-Making in Nature and Design", Prof. Naomi Ehrich Leonard, Princeton University | MPI-IS Stuttgart |
| 2016-04-28 | Annette Stumpf | Girls' Day | MPI-IS Stuttgart |

### 20.2.2 External public events (selection)

| Date | MPI-IS participants | Description | Location |
|---|---|---|---|
| 2018-11-19 | Bernhard Schölkopf, Linda Behringer, Tamara Almeyda | Information booth and talk by Bernhard Schölkopf at AI Con | Bosch Corporate Research, Renningen |
| 2018-11-09 | Bernhard Schölkopf | Public talk about artificial intelligence at the Falling Walls Conference | Berlin |
| 2018-11-08 | Linda Behringer, Tamara Almeyda | Information booth at the new.New festival | Schleyer-Halle Stuttgart |
| 2018-10-04 | Bernhard Schölkopf | Panel discussion about Cyber Valley at the Tübingen Art Gallery | Kunsthalle Tübingen |
| 2018-04-20 | Linda Behringer, Claudia Däfler | Information booth at the event "Schöne neue Welt?! Die Digitalisierung gestalten" held at the Baden-Württemberg state parliament | Stuttgart |
| 2017-05 until 2017-12 | Claudia Däfler | The exhibition at the Museum of Ancient Cultures focused on the most important steps of humankind. The institute supported the part "Origin of digital innovation" with a Nao robot and the Mosh Camera App | Hohentübingen Castle |
| 2017-07 until 2017-10 | Claudia Däfler | "Robots" exhibition The institute's research was presented at this exhibition in the Stadtmuseum Tübingen | Stadtmuseum Tübingen |
| 2017-04-28 | Claudia Däfler, Vincent Berenz | Information booth at TÜFFF – Tübinger Fenster für Forschung | University of Tübingen |
| 2017-03 | Georg Martius | Workshop for pupils: conversation with a scientist in regard to the film: "BAYMAX – Riesiges Robowabohu." This event was part of a film program for the "Arbeitswelten der Zukunft" science year | Backnang |
| 2016-12-15 | Bernhard Schölkopf, Michael Black | Cyber Valley Kick-off Meeting | Neues Schloss, Stuttgart |

## 20.3 Scientific Events


In addition to the selection of scientific events listed below, the institute frequently hosts talks by scientists from all over the world. From 2016 to 2018, MPI-IS hosted more than 230 individual scientific talks and colloquia; a complete list of talks can be found at https://is.mpg.de/talks.

| Date | Organizers (MPI-IS) | Description | Location |
|---|---|---|---|
| 2018-12-06 | Bernhard Schölkopf | Founding event of the European Laboratory for Learning and Intelligent Systems (ELLIS) at NeurIPS 2018 | Montreal, Canada |
| 2018-11-14 | Hasti Seifi, Farimah Fazlollahi, Gunhyuk Park, Katherine J. Kuchenbecker | AsiaHaptics Workshop: "Haptipedia: An Interactive Database for Selecting, Ideating, and Learning About Grounded Force-Feedback Devices" | Songdo, South Korea |
| 2018-09-18 until 2018-09-20 | Michael Black, Katherine J. Kuchenbecker, Bernhard Schölkopf, Metin Sitti | Special Symposium on Intelligent Systems | Both MPI-IS sites |
| 2018-06-13 | David Gueorguiev | EuroHaptics Workshop: "From Fingertip Mechanics to Tactile Sensation" | Pisa, Italy |
| 2018-04-03 until 2018-04-05 | Bernhard Schölkopf, Diana Rebmann | DALI 2018 - Data, Learning, and Inference | Lanzarote, Spain |
| 2018-02-08 | Sebastian Trimpe, Georg Martius, Melanie Zeilinger | Second Max Planck ETH Workshop on Learning Control | ETH Zürich, Switzerland |
| 2017-12-12 until 2017-12-15 | Co-Organizer: Sebastian Trimpe | Invited Session Series on Learning-based Control at CDC | Melbourne, Australia |

| | | | |
|---|---|---|---|
| 2017-10-29 until 2017-11-01 | Michael Black, Julia Braun | Scenes from Video III | Lago di Garda, Italy |
| 2017-09-20 until 2017-09-22 | Michael Black, Tamara Almeyda | Cyber Valley Symposium | Both MPI-IS sites |
| 2017-07-16 | Co-Organizer: Jeannette Bohg | Articulated Model Tracking - Workshop at RSS (Robotics: Science and Systems Conference) | Cambridge, USA |
| 2017-07-15 | Co-Organizer: Jeannette Bohg | Women in Robotics III - Workshop at RSS (Robotics: Science and Systems Conference) | Cambridge, USA |
| 2017-07-15 | Co-Organizer: Jeannette Bohg | Revisiting Contact - Turning a Problem into a Solution - Workshop at RSS (Robotics: Science and Systems Conference) | Cambridge, USA |
| 2017-07-03 until 2017-07-01 | Rohit Babbar, Fatma Güney | Max Planck ETH Center for Learning Systems: Pre-doc Summer School on Learning Systems | Zürich, Switzerland |
| 2017-06-19 until 2017-06-30 | Bernhard Schölkopf, Ruth Urner, Julia Braun, Diana Rebmann | Machine Learning Summer School | Tübingen |
| 2017-06-05 | Philipp Hennig | ICERM Seminar on Probabilistic Scientific Computing | Brown University, USA |
| 2017-06 | Philipp Hennig | Summer School on Probabilistic Numerics | Dobbiaco, Italy |
| 2017-04-18 until 2017-04-20 | Bernhard Schölkopf | DALI 2017 - Data, Learning, and Inference | Sestri Tenerife, Spain |
| 2017-03-27 | Jeannette Bohg | Interactive Multisensory Object Perception for Embodied Agents Symposium at the AAAI Spring Symposium Series in 2017 | Stanford University, USA |
| 2016-12-13 until 2016-12-16 | Michael Black, Katherine J. Kuchenbecker, Metin Sitti, Bernhard Schölkopf | Special Symposium on Intelligent Systems | Both MPI-IS sites |
| 2016-11 | Co-Organizers:, Philipp Hennig, Bernhard Schölkopf | Dagstuhl Seminar on the Future of Learning with Kernels and Gaussian Processes | Schloss Dagstuhl |
| 2016-09-25 until 2016-09-28 | Peer Fischer, Samuel Sanchez, Metin Sitti | Max Planck ETH Center for Learning Systems: “Biomedical Micro-/Nanosystems Engineering” workshop | Schloss Ringberg |

| | | | |
|---|---|---|---|
| 2016-08-26 | Philipp Hennig<br>Maren Mahsereci | NeurIPS: "Optimizing the Optimizer" workshop | Barcelona, Spain |
| 2016-07-11 | Peter Gehler | Max Planck ETH Center for Learning Systems: "Deep Learning" workshop | Donau-eschingen |
| 2016-06-09 until 2016-06-10 | Hakan Ceylan, Kirstin Petersen, Alexander Spröwitz | Max Planck ETH Center for Learning Systems: "Design and Coordination of Micro- to Macro-Scale Swarms" workshop | Radolfzell |
| 2016-06 | Peer Fischer, Samuel Sanchez | International Symposium on Micro- and Nanomachines | Hannover |
| 2016-03-30 until 2016-04-01 | Bernhard Schölkopf | DALI 2016 - Data, Learning, and Inference | Sestri Levante, Italy |
| 2016-03-15 until 2016-03-16 | Michael Black, Bernhard Schölkopf, Metin Sitti | Special Symposium on Intelligent Systems | Both MPI-IS sites |

# 21 PUBLICATIONS

## 21.1 Autonomous Motion Department

### 21.1.1 Journal Articles

#### 2018

[1] M. Anderson, S. Anderson, V. Berenz. A Value-Driven Eldercare Robot: Virtual and Physical Instantiations of a Case-Supported Principle-Based Behavior Paradigm. *Proceedings of the IEEE*: 1, 15, 2018.

[2] V. Berenz, S. Schaal. Playful: Reactive Programming for Orchestrating Robotic Behavior. *IEEE Robotics Automation Magazine* **25** (3): 49–60, 2018.

[3] M. Muehlebach, S. Trimpe. Distributed Event-Based State Estimation for Networked Systems: An LMI Approach. *IEEE Transactions on Automatic Control* **63** (1): 269–276, 2018 (cited on pages 198, 200).

[4] D. Kappler, F. Meier, J. Issac, J. Mainprice, C. Garcia Cifuentes, M. Wüthrich, V. Berenz, S. Schaal, N. Ratliff, J. Bohg. Real-time Perception meets Reactive Motion Generation. *IEEE Robotics and Automation Letters* **3** (3): 1864–1871, 2018.

#### 2017

[5] S. Trimpe. Event-based State Estimation: An Emulation-based Approach. *IET Control Theory & Applications* **11** (11): 1684–1693, 2017 (cited on pages 198, 200).

[6] C. Garcia Cifuentes, J. Issac, M. Wüthrich, S. Schaal, J. Bohg. Probabilistic Articulated Real-Time Tracking for Robot Manipulation. *IEEE Robotics and Automation Letters (RA-L)* **2** (2): 577–584, 2017.

[7] J. Bohg, K. Hausman, B. Sankaran, O. Brock, D. Kragic, S. Schaal, G. Sukhatme. Interactive Perception: Leveraging Action in Perception and Perception in Action. *IEEE Transactions on Robotics* **33**: 1273–1291, 2017.

[8] Z. Wang, A. Boularias, K. Mülling, B. Schölkopf, J. Peters. Anticipatory Action Selection for Human-Robot Table Tennis. *Artificial Intelligence* **247**. Special Issue on AI and Robotics: 399–414, 2017.

#### 2016

[9] A. Herzog, N. Rotella, S. Mason, F. Grimminger, S. Schaal, L. Righetti. Momentum Control with Hierarchical Inverse Dynamics on a Torque-Controlled Humanoid. *Autonomous Robots* **40** (3): 473–491, 2016.

[10] M. Wüthrich, S. Trimpe, C. Garcia Cifuentes, D. Kappler, S. Schaal. A New Perspective and Extension of the Gaussian Filter. *The International Journal of Robotics Research* **35** (14): 1731–1749, 2016 (cited on page 198).

[11] D. Laidig, S. Trimpe, T. Seel. Event-based Sampling for Reducing Communication Load in Realtime Human Motion Analysis by Wireless Inertial Sensor Networks. *Current Directions in Biomedical Engineering* **2** (1): 711–714, 2016.

[12] C. Daniel, H. van Hoof, J. Peters, G. Neumann. Probabilistic Inference for Determining Options in Reinforcement Learning. *Machine Learning, Special Issue* **104** (2). Best Student Paper Award of ECML-PKDD 2016: 337–357, 2016.

[13] N. Vitiello, A. J. Ijspeert, S. Schaal. Bioinspired Motor Control for Articulated Robots [From the Guest Editors]. *IEEE Robotics & Automation Magazine* **23** (1): 20–21, 2016.

### 21.1.2 Conference Papers

#### 2018

[14] N. Rotella, S. Schaal, L. Righetti. Unsupervised Contact Learning for Humanoid Estimation and Control. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 411–417, 2018.

[15] A. Gams, S. Mason, A. Ude, S. Schaal, L. Righetti. Learning Task-Specific Dynamics to Improve Whole-Body Control. In *Hua*, 2018.

[16] S. Mason, N. Rotella, S. Schaal, L. Righetti. An MPC Walking Framework With External Contact Forces. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1785–1790, 2018.

[17] B. Ponton, A. Herzog, A. Del Prete, S. Schaal, L. Righetti. On Time Optimization of Centroidal Momentum Dynamics. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5776–5782, 2018.

[18] G. Sutanto, Z. Su, S. Schaal, F. Meier. Learning Sensor Feedback Models from Demonstrations via Phase-Modulated Neural Networks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2018*, 2018.

[19] F. Meier, D. Kappler, S. Schaal. Online Learning of a Memory for Learning Rates. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2018*. accepted, 2018.

[20] A. Doerr, C. Daniel, M. Schiegg, D. Nguyen-Tuong, S. Schaal, M. Toussaint, S. Trimpe. Probabilistic Recurrent State-Space Models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018 (cited on pages 198, 201).

#### 2017

[21] K. Hausman, Y. Chebotar, S. Schaal, G. Sukhatme, J. Lim. Multi-Modal Imitation Learning from Unstructured Demonstrations using Generative Adversarial Nets. In *Proceedings from the conference "Neural Information Processing Systems 2017.* 2017.

[22] D. Kappler, F. Meier, N. Ratliff, S. Schaal. A New Data Source for Inverse Dynamics Learning. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.

[23] A. Doerr, C. Daniel, D. Nguyen-Tuong, A. Marco, S. Schaal, M. Toussaint, S. Trimpe. Optimizing Long-term Predictions for Model-based Policy Search. In *Proceedings of 1st Annual Conference on Robot Learning (CoRL)*. Vol. 78, pages 227–238, 2017 (cited on pages 198, 201).

[24] A. Marco, P. Hennig, S. Schaal, S. Trimpe. On the Design of LQR Kernels for Efficient Controller Learning. In *Proceedings of the 56th IEEE Annual Conference on Decision and Control (CDC)*, pages 5193–5200, 2017 (cited on pages 198, 199, 252).

[25] C. Rubert, D. Kappler, A. Morales, S. Schaal, J. Bohg. On the relevance of grasp metrics for predicting grasp success. In *Proceedings of the IEEE/RSJ International Conference of Intelligent Robots and Systems*, 2017.

[26] Y. Chebotar, K. Hausman, M. Zhang, G. Sukhatme, S. Schaal, S. Levine. Combining Model-Based and Model-Free Updates for Trajectory-Centric Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. Proceedings of Machine Learning Research, 2017.

[27] R. Akrour, D. Sorokin, J. Peters, G. Neumann. Local Bayesian Optimization of Motor Skills. In *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. Proceedings of Machine Learning Research, pages 41–50, 2017.

[28] L. Grossberger, M. R. Hohmann, J. Peters, M. Grosse-Wentrup. Investigating Music Imagery as a Cognitive Paradigm for Low-Cost Brain-Computer Interfaces. In *Proceedings of the 7th Graz Brain-Computer Interface Conference 2017 - From Vision to Reality*, pages 160–164, 2017.

[29] K.-H. Fiebig, V. Jayaram, T. Hesse, A. Blank, J. Peters, M. Grosse-Wentrup. Bayesian Regression for Artifact Correction in Electroencephalography. In *Proceedings of the 7th Graz Brain-Computer Interface Conference 2017 - From Vision to Reality*, pages 131–136, 2017.

[30] A. Rai, G. Sutanto, S. Schaal, F. Meier. Learning Feedback Terms for Reactive Planning and Control. In *Proceedings 2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[31] A. Marco, F. Berkenkamp, P. Hennig, A. P. Schoellig, A. Krause, S. Schaal, S. Trimpe. Virtual vs. Real: Trading Off Simulations and Physical Experiments in Reinforcement Learning with Bayesian Optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1557–1563, 2017 (cited on pages 198, 199, 252).

[32] A. Doerr, D. Nguyen-Tuong, A. Marco, S. Schaal, S. Trimpe. Model-Based Policy Search for Automatic Tuning of Multivariate PID Controllers. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 5295–5301, 2017 (cited on page 198).

### 2016

[33] D. Büchler, R. Calandra, J. Peters. Modeling Variability of Musculoskeletal Systems with Heteroscedastic Gaussian Processes. In *Workshop on Neurorobotics*, 2016.

[34] S. Mason, N. Rotella, S. Schaal, L. Righetti. Balancing and Walking Using Full Dynamics LQR Control With Contact Constraints. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 63–68, 2016.

[35] B. Ponton, A. Herzog, S. Schaal, L. Righetti. A Convex Model of Momentum Dynamics for Multi-Contact Motion Generation. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots Humanoids*, pages 842–849, 2016.

[36] B. Ponton, S. Schaal, L. Righetti. On the Effects of Measurement Uncertainty in Optimal Control of Contact Interactions. In *The 12th International Workshop on the Algorithmic Foundations of Robotics WAFR*, 2016.

[37] N. Rotella, S. Mason, S. Schaal, L. Righetti. Inertial Sensor-Based Humanoid Joint State Estimation. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1825–1831, 2016.

[38] A. Herzog, S. Schaal, L. Righetti. Structured contact force optimization for kino-dynamic motion generation. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2703–2710, 2016.

[39] M. Khadiv, S. Kleff, A. Herzog, S. A. A. Moosavian, S. Schaal, L. Righetti. Stepping Stabilization Using a Combination of DCM Tracking and Step Adjustment. In *2016 4th International Conference on Robotics and Mechatronics (ICROM)*, pages 130–135, 2016.

[40] The Role of Measurement Uncertainty in Optimal Control for Contact Interactions. In *Workshop on the Algorithmic Foundations of Robotics*, page 22, 2016.

[41] Y. Chebotar, K. Hausman, O. Kroemer, G. Sukhatme, S. Schaal. Generalizing Regrasping with Supervised Policy Learning. In *International Symposium on Experimental Robotics (ISER) 2016*, 2016.

[42] Y. Chebotar, K. Hausman, Z. Su, G. Sukhatme, S. Schaal. Self-Supervised Regrasping using Spatio-Temporal Tactile Features and Reinforcement Learning. In *International Conference on Intelligent Robots and Systems (IROS) 2016*, 2016.

[43] S. Trimpe. Predictive and Self Triggering for Event-based State Estimation. In *Proceedings of the 55th IEEE Conference on Decision and Control (CDC)*, pages 3098–3105, 2016 (cited on pages 198, 200).

[44] M. Anderson, S. L. Anderson, V. Berenz. Ensuring Ethical Behavior from Autonomous Systems. In *Artificial Intelligence Applied to Assistive Technologies and Smart Environments, Papers from the 2016 AAAI Workshop, Phoenix, Arizona, USA, February 12, 2016*, 2016.

[45] J. Bohg, D. Kappler, F. Meier, N. Ratliff, J. Mainprice, J. Issac, M. Wüthrich, C. Garcia Cifuentes, V. Berenz, S. Schaal. Interlocking Perception-Action Loops at Multiple Time Scales - A System Proposal for Manipulation in Uncertain and Dynamic Environments. In *International Workshop on Robotics in the 21st century: Challenges and Promises*, 2016.

[46] Y. Huang, D. Büchler, O. Koc, B. Schölkopf, J. Peters. Jointly Learning Trajectory Generation and Hitting Point Prediction in Robot Table Tennis. In *16th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 650–655, 2016.

[47] S. Gomez-Gonzalez, G. Neumann, B. Schölkopf, J. Peters. Using Probabilistic Movement Primitives for Striking Movements. In *16th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 502–508, 2016 (cited on page 54).

[48] C. Lassner, D. Kappler, M. Kiefel, P. Gehler. Barrista - Caffe Well-Served. In *ACM Multimedia Open Source Software Competition*, 2016.

[49] O. Koc, G. Maeda, J. Peters. A New Trajectory Generation Framework in Robotic Table Tennis. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 3750–3756, 2016 (cited on pages 36, 54).

[50] B. Sankaran, S. Ramalingam, Y. Taguchi. Parameter Learning for Improving Binary Descriptor Matching. In *International Conference on Intelligent Robots and Systems (IROS) 2016*, 2016.

[51] F. Meier, D. Kappler, N. Ratliff, S. Schaal. Towards Robust Online Inverse Dynamics Learning. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*, 2016.

[52] A. Kloss, D. Kappler, H. P. A. Lensch, M. V. Butz, S. Schaal, J. Bohg. Learning Where to Search Using Visual Attention. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*.

[53] R. Gadde, V. Jampani, M. Kiefel, D. Kappler, P. Gehler. Superpixel Convolutional Networks using Bilateral Inceptions. In *European Conference on Computer Vision (ECCV)*. Lecture Notes in Computer Science, 2016 (cited on page 88).

[54] S. Ebner, S. Trimpe. Communication Rate Analysis for Event-based State Estimation. In *Proceedings of the 13th International Workshop on Discrete Event Systems*, 2016.

[55] F. Meier, S. Schaal. Drifting Gaussian Processes with Varying Neighborhood Sizes for Online Model Learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2016*, 2016.

[56] D. Büchler, H. Ott, J. Peters. A Lightweight Robotic Arm with Pneumatic Muscles for Robot Learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 4086–4092, 2016 (cited on pages 36, 47).

[57] J. Bohg, D. Kappler, S. Schaal. Exemplar-based Prediction of Object Properties from Local Shape Similarity. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2016*, 2016.

[58] A. Marco, P. Hennig, J. Bohg, S. Schaal, S. Trimpe. Automatic LQR Tuning Based on Gaussian Process Global Optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 270–277, 2016 (cited on pages 198, 199, 252).

[59] J. Issac, M. Wüthrich, C. Garcia Cifuentes, J. Bohg, S. Trimpe, S. Schaal. Depth-based Object Tracking Using a Robust Gaussian Filter. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2016*, 2016.

[60] F. Widmaier, D. Kappler, S. Schaal, J. Bohg. Robot Arm Pose Estimation by Pixel-wise Regression of Joint Angles. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2016*, 2016.

[61] D. Kappler, S. Schaal, J. Bohg. Optimizing for what matters: the Top Grasp Hypothesis. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2016*, 2016.

[62] M. Wüthrich, C. Garcia Cifuentes, S. Trimpe, F. Meier, J. Bohg, J. Issac, S. Schaal. Robust Gaussian Filtering using a Pseudo Measurement. In *Proceedings of the American Control Conference (ACC)*, 2016 (cited on page 198).

### 21.1.3 Book Chapters

#### 2017

[63] J. Peters, D. Lee, J. Kober, D. Nguyen-Tuong, J. Bagnell, S. Schaal. Robot Learning. In. *Springer Handbook of Robotics*. 2nd. Springer International Publishing. Chap. 15, pages 357–394, 2017.

### 2016

[64] J.-A. Ting, F. Meier, S. Vijayakumar, S. Schaal. Locally Weighted Regression for Control. In. *Encyclopedia of Machine Learning and Data Mining*. Springer US, Boston, MA, pages 1–14, 2016.

[65] P. F. Dominey, T. J. Prescott, J. Bohg, A. K. Engel, S. Gallagher, T. Heed, M. Hoffmann, G. Knoblich, W. Prinz, A. Schwartz. Implications of Action-Oriented Paradigm Shifts in Cognitive Science. In *The Pragmatic Turn - Toward Action-Oriented Views in Cognitive Science*. Vol. 18. Strüngmann Forum Reports, vol. 18, J. Lupp, series editor. The MIT Press. Chap. 20, pages 333–356, 2016.

[66] J. Bohg, D. Kragic. Learning Action-Perception Cycles in Robotics: A Question of Representations and Embodiment. In *The Pragmatic Turn - Toward Action-Oriented Views in Cognitive Science*. Vol. 18. Strüngmann Forum Reports, vol. 18, J. Lupp, series editor. The MIT Press. Chap. 18, pages 309–320, 2016.

## 21.2 Empirical Inference Department

### 21.2.1 Books

#### 2017

[67] A. Gretton, P. Hennig, C. Rasmussen, B. Schölkopf. New Directions for Learning with Kernels and Gaussian Processes (Dagstuhl Seminar 16481). Vol. 6. 11, pages 142–167, 2017.

[68] J. Peters, D. Janzing, B. Schölkopf. Elements of Causal Inference - Foundations and Learning Algorithms. Adaptive Computation and Machine Learning Series. The MIT Press, Cambridge, MA, USA, 2017 (cited on pages 34, 40, 41).

### 21.2.2 Proceedings

#### 2016

[69] A. T. Ihler, D. Janzing, editors. Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence (UAI). AUAI Press, page 869, 2016.

### 21.2.3 Journal Articles

#### 2018

[70] O. Koc, J. Peters. Learning to serve: an experimental study for a new learning from demonstrations framework. *IEEE Robotics and Automation Letters (ICRA/RA-L)*, 2018 (cited on pages 36, 54).

[71] S. Gomez-Gonzalez, G. Neumann, B. Schölkopf, J. Peters. Adaptation and Robust Learning of Probabilistic Movement Primitives. *IEEE Transactions on Robotics*, 2018 (cited on pages 36, 54).

[72] O. Koc, G. Maeda, J. Peters. Optimizing Execution of Dynamic Goal-Directed Robot Movements with Learning Control. *IEEE Transactions on Robotics*, 2018 (cited on pages 35, 47, 54).

[73] J. A. Disselhorst, M. A. Krueger, S. M. M. Ud-Dean, I. Bezrukov, M. A. Jarboui, C. Trautwein, A. Traube, C. Spindler, J. M. Cotton, D. Leibfritz, B. J. Pichler. Linking imaging to omics utilizing image-guided tissue extraction. *Proceedings of the National Academy of Sciences* **115** (13): E2980–E2987, 2018.

[74] D. Janzing, P. Wocjan. Does universal controllability of physical systems prohibit thermodynamic cycles? *Open Systems and Information Dynamics*, 2018.

[75] E. J. Leehr, K. Schag, T. Dresler, M. Grosse-Wentrup, M. Hautzinger, A. J. Fallgatter, S. Zipfel, K. E. Giel, A.-C. Ehlis. Food specific inhibitory control under negative mood in binge-eating disorder: Evidence from a multimethod approach. *International Journal of Eating Disorders* **51** (2): 112–123, 2018.

[76] A. Vergari, N. Di Mauro, F. Esposito. Visualizing and understanding Sum-Product Networks. *Machine Learning*, 2018.

[77] D. Fioravanti*, Y. Giarratano*, V. Maggio*, C. Agostinelli, M. Chierici, G. Jurman, C. Furlanello. Phylogenetic convolutional neural networks in metagenomics. *BMC Bioinformatics* **19** (2). *equal contribution: 49 pages, 2018.

[78] F. J. R. Ruiz, I. Valera, L. Svensson, F. Perez-Cruz. Infinite Factorial Finite State Machine for Blind Multiuser Channel Estimation. *IEEE Transactions on Cognitive Communications and Networking* **4** (2): 177–191, 2018 (cited on page 52).

[79] M. B. Zafar, I. Valera, M. Gomez Rodriguez, K. Gummadi. A Flexible Approach for Fair Classification. *Journal of Machine Learning*, 2018 (cited on pages 35, 44).

[80] F. Veiga, J. Peters, T. Hermans. Grip Stabilization of Novel Objects using Slip Prediction. *IEEE Transactions on Haptics*, 2018.

[81] D. Tanneberg, J. Peters, E. Rueckert. Intrinsic Motivation and Mental Replay enable Efficient Online Adaptation in Stochastic Recurrent Networks. *Neural Networks* **109**: 67–80, 2018.

[82] A. Šošić, E. Rueckert, J. Peters, A. Zoubir, H. Koeppl. Inverse Reinforcement Learning via Nonparametric Spatio-Temporal Subgoal Modeling. *Journal of Machine Learning Research* **19** (69): 1–45, 2018.

[83] B. Tabibian, U. Upadhyay, A. De, A. Zarezade, B. Schölkopf, M. Gomez Rodriguez. Enhancing Human Learning via Spaced Repetition Optimization. *Proceedings of the National Academy of Sciences*, 2018 (cited on pages 48, 53).

[84] N. Shah*, B. Tabibian*, K. Muandet, I. Guyon, U. von Luxburg. Design and Analysis of the NIPS 2016 Review Process. *Journal of Machine Learning Research* **19** (49). *equal contribution: 1–34, 2018.

[85] M. R. Hohmann, T. Fomina, V. Jayaram, T. Emde, J. Just, M. Synofzik, B. Schölkopf, L. Schöls, M. Grosse-Wentrup. Case series: Slowing alpha rhythm in late-stage ALS patients. *Clinical Neurophysiology* **129** (2): 406–408, 2018 (cited on pages 36, 39).

[86] Z. Yi, Y. Zhang, J. Peters. Biomimetic Tactile Sensors and Signal Processing with Spike Trains: A Review. *Sensors and Actuators A: Physical* **269**: 41–52, 2018.

[87] J. Vinogradska, B. Bischoff, J. Peters. Approximate Value Iteration Based on Numerical Quadrature. *IEEE Robotics and Automation Letters* **3** (2): 1330–1337, 2018.

[88] A. Paraschos, E. Rueckert, J. Peters, G. Neumann. Probabilistic movement primitives under unknown system dynamics. *Advanced Robotics* **32** (6): 297–310, 2018.

[89] A. Paraschos, C. Daniel, J. Peters, G. Neumann. Using Probabilistic Movement Primitives in Robotics. *Autonomous Robots* **42** (3): 529–551, 2018.

[90] T. Osa, J. Pajarinen, G. Neumann, J. Bagnell, P. Abbeel, J. Peters. An Algorithmic Perspective on Imitation Learning. *Foundations and Trends in Robotics* **7** (1-2): 1–179, 2018.

[91] S. Manschitz, M. Gienger, J. Kober, J. Peters. Mixture of Attractors: A Novel Movement Primitive Representation for Learning Motor Skills From Demonstrations. *IEEE Robotics and Automation Letters* **3** (2): 926–933, 2018.

[92] O. Kroemer, S. Leischnig, S. Luettgen, J. Peters. A kernel-based approach to learning contact distributions for robot manipulation tasks. *Autonomous Robots* **42** (3): 581–600, 2018.

[93] V. Kapoor, M. Besserve, N. K. Logothetis, T. I. Panagiotaropoulos. Parallel and functionally segregated processing of task phase and conscious content in the prefrontal cortex. *Communications Biology* **1** (215): 1–12, 2018 (cited on pages 36, 49).

[94] V. Jayaram, A. Barachant. MOABB: Trustworthy algorithm benchmarking for BCIs. *Journal of Neural Engineering* **15** (6): 066011, 2018.

[95] J. F. Ramirez-Villegas, K. F. Willeke, N. K. Logothetis, M. Besserve. Dissecting the synapse- and frequency-dependent network mechanisms of in vivo hippocampal sharp wave-ripples. *Neuron*, 2018 (cited on pages 36, 49).

[96] M. Rojas-Carulla, B. Schölkopf, R. Turner, J. Peters. Invariant Models for Causal Transfer Learning. *Journal of Machine Learning Research* **19** (36): 1–34, 2018 (cited on pages 34, 41).

[97] D. Büchler, R. Calandra, B. Schölkopf, J. Peters. Control of Musculoskeletal Systems using Learned Dynamics Models. *IEEE Robotics and Automation Letters* **3** (4): 3161–3168, 2018.

[98] L. Xiao, F. Heide, W. Heidrich, B. Schölkopf, M. Hirsch. Discriminative Transfer Learning for General Image Restoration. *IEEE Transactions on Image Processing* **27** (8): 4091–4104, 2018.

[99] R. Babbar, M. Heni, A. Peter, M. Hrabě de Angelis, H.-U. Häring, A. Fritsche, H. Preissl, B. Schölkopf, R. Wagner. Prediction of Glucose Tolerance without an Oral Glucose Tolerance Test. *Frontiers in Endocrinology* **9**: 82, 2018.

[100] O. Koc, G. Maeda, J. Peters. Online optimal trajectory generation for robot table tennis. *Robotics and Autonomous Systems* **105**: 121–137, 2018 (cited on page 36).

[101] A. Loktyushin, P. Ehses, B. Schölkopf, K. Scheffler. Autofocusing-based phase correction. *Magnetic Resonance in Medicine*. Epub ahead, 2018 (cited on page 42).

[102] K. Zhang, B. Schölkopf, P. Spirtes, C. Glymour. Learning Causality and Causality-Related Learning: Some Recent Progress. *National Science Review* **5** (1): 26–29, 2018.

[103] C. J. Simon-Gabriel, B. Schölkopf. Kernel Distribution Embeddings: Universal Kernels, Characteristic Kernels and Kernel Metrics on Distributions. *Journal of Machine Learning Research* **19** (44): 1–29, 2018 (cited on page 55).

## 2017

[8] Z. Wang, A. Boularias, K. Mülling, B. Schölkopf, J. Peters. Anticipatory Action Selection for Human-Robot Table Tennis. *Artificial Intelligence* **247**. Special Issue on AI and Robotics: 399–414, 2017.

[104] J. Chen, M. Hirsch, R. Heintzmann, B. Eberhardt, H. P. A. Lensch. A Phase-coded Aperture Camera with Programmable Optics. *Electronic Imaging* **2017** (17): 70–75, 2017.

[105] F. Seith, H. Schmidt, S. Gatidis, I. Bezrukov, C. Schraml, C. Pfannenberg, C. la Fougère, K. Nikolaou, N. Schwenzer. SUV-quantification of physiological lung tissue in an integrated PET/MR-system: Impact of lung density and bone tissue. *PLOS ONE* **12** (5): 1–13, 2017.

[106] L. O. M. Rothkegel, H. A. Trukenbrod, H. H. Schütt, F. A. Wichmann, R. Engbert. Temporal evolution of the central fixation bias in scene viewing. *Journal of Vision* **17** (13): 3, 2017.

[107] Z. Yi, Y. Zhang, J. Peters. Bioinspired tactile sensor for surface roughness discrimination. *Sensors and Actuators A: Physical* **255**: 46–53, 2017.

[108] J. Vinogradska, B. Bischoff, D. Nguyen-Tuong, J. Peters. Stability of Controllers for Gaussian Process Dynamics. *Journal of Machine Learning Research* **18** (100): 1–37, 2017.

[109] H. van Hoof, D. Tanneberg, J. Peters. Generalized exploration in policy search. *Machine Learning* **106** (9-10). Special Issue of the ECML PKDD 2017 Journal Track: 1705–1724, 2017.

[110] H. van Hoof, G. Neumann, J. Peters. Non-parametric Policy Search with Limited Information Loss. *Journal of Machine Learning Research* **18** (73): 1–46, 2017.

[111] S. Parisi, M. Pirotta, J. Peters. Manifold-based multi-objective policy search with sample reuse. *Neurocomputing* **263**. Special Issue on Multi-Objective Reinforcement Learning: 3–14, 2017.

[112] A. Paraschos, R. Lioutikov, J. Peters, G. Neumann. Probabilistic Prioritization of Movement Primitives. *Proceedings of the International Conference on Intelligent Robot Systems, and IEEE Robotics and Automation Letters (RA-L)* **2** (4): 2294–2301, 2017.

[113] V. Padois, S. Ivaldi, J. Babic, M. Mistry, J. Peters, F. Nori. Whole-body multi-contact motion in humans and humanoids: Advances of the CoDyCo European project. *Robotics and Autonomous Systems* **90**. Special Issue on New Research Frontiers for Intelligent Autonomous Systems: 97–117, 2017.

[114] T. Osa, A. M. Ghalamzan E., R. Stolkin, R. Lioutikov, J. Peters, G. Neumann. Guiding Trajectory Optimization by Demonstrated Distributions. *IEEE Robotics and Automation Letters* **2** (2): 819–826, 2017.

[115] G. Maeda, G. Neumann, M. Ewerton, R. Lioutikov, O. Kroemer, J. Peters. Probabilistic Movement Primitives for Coordination of Multiple Human-Robot Collaborative Tasks. *Autonomous Robots* **41** (3): 593–612, 2017.

[116] G. Maeda, M. Ewerton, G. Neumann, R. Lioutikov, J. Peters. Phase Estimation for Fast Action Recognition and Trajectory Generation in Human-Robot Collaboration. *International Journal of Robotics Research* **36** (13-14). Special Issue on the Seventeenth International Symposium on Robotics Research: 1579–1594, 2017.

[117] R. Lioutikov, G. Neumann, G. Maeda, J. Peters. Learning Movement Primitive Libraries through Probabilistic Segmentation. *International Journal of Robotics Research* **36** (8): 879–894, 2017.

[118] O. Kroemer, J. Peters. A Comparison of Autoregressive Hidden Markov Models for Multimodal Manipulations With Variable Masses. *IEEE Robotics and Automation Letters* **2** (2): 1101–1108, 2017.

[119] L. Gresele, M. Marsili. On Maximum Entropy and Inference. *Entropy* **19** (12): article no. 642, 2017.

[120] F. A. Wichmann, D. H. J. Janssen, R. Geirhos, G. Aguilar, H. H. Schütt, M. Maertens, M. Bethge. Methods and measurements to compare men against machines. *Electronic Imaging*: 36–45(10), 2017 (cited on pages 36, 53).

[121] T. S. A. Wallis, C. M. Funke, A. S. Ecker, L. A. Gatys, F. A. Wichmann, M. Bethge. A parametric texture model based on deep convolutional features closely matches texture appearance for humans. *Journal of Vision* **17** (12), 2017 (cited on pages 36, 53).

[122] H. H. Schütt, F. A. Wichmann. An image-computable psychophysical spatial vision model. *Journal of Vision* **17** (12), 2017 (cited on pages 36, 53).

[123] N. Pfister, P. Bühlmann, B. Schölkopf, J. Peters. Kernel-based tests for joint independence. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **80** (1): 5–31, 2017 (cited on page 41).

[124] D. Janzing, B. Schölkopf. Detecting Confounding in Multivariate Linear Models via Spectral Analysis. *Journal of Causal Inference* **6** (1), 2017.

[125] K. Muandet, K. Fukumizu, B. Sriperumbudur, B. Schölkopf. Kernel Mean Embedding of Distributions: A Review and Beyond. *Foundations and Trends in Machine Learning* **10** (1-2): 1–141, 2017 (cited on pages 33, 45).

[126] T. Fomina, S. Weichwald, M. Synofzik, J. Just, L. Schöls, B. Schölkopf, M. Grosse-Wentrup. Absence of EEG correlates of self-referential processing depth in ALS. *PLOS ONE* **12** (6): e0180136, 2017 (cited on pages 36, 39).

[127] V. Jayaram, M. Hohmann, J. Just, B. Schölkopf, M. Grosse-Wentrup. Task-induced frequency modulation features for brain-computer interfacing. *Journal of Neural Engineering* **14** (5): 056015, 2017 (cited on pages 36, 39).

[128] I. Tolstikhin, B. Sriperumbudur, K. Muandet. Minimax Estimation of Kernel Mean Embeddings. *Journal of Machine Learning Research* **18** (86): 1–47, 2017 (cited on pages 32, 55).

[129] G. Aguilar, F. A. Wichmann, M. Maertens. Comparing sensitivity estimates from MLDS and forced-choice methods in a slant-from-texture experiment. *Journal of Vision* **17** (1), 2017 (cited on pages 36, 53).

[130] T. S. A. Wallis, S. Tobias, M. Bethge, F. A. Wichmann. Detecting distortions of peripherally presented letter stimuli under crowded conditions. *Attention, Perception, & Psychophysics* **79** (3): 850–862, 2017.

[131] O. Ozdenizci, M. Yalcin, A. Erdogan, V. Patoglu, M. Grosse-Wentrup, M. Cetin. Electroencephalographic identifiers of motor adaptation learning. *Journal of Neural Engineering* **14** (4): 046027, 2017.

[132] P. Katiyar, M. R. Divine, U. Kohlhofer, L. Quintanilla-Martinez, B. Schölkopf, B. J. Pichler, J. A. Disselhorst. Spectral Clustering predicts tumor tissue heterogeneity using dynamic 18F-FDG PET: a complement to the standard compartmental modeling approach. *Journal of Nuclear Medicine* **58** (4): 651–657, 2017.

[133] P. Katiyar, M. R. Divine, U. Kohlhofer, L. Quintanilla-Martinez, B. Schölkopf, B. J. Pichler, J. A. Disselhorst. A Novel Unsupervised Segmentation Approach Quantifies Tumor Tissue Populations Using Multiparametric MRI: First Results with Histological Validation. *Molecular Imaging and Biology* **19** (3): 391–397, 2017.

[134] D. Grimm, D. Roqueiro, P. Salome, S. Kleeberger, B. Greshake, W. Zhu, C. Liu, C. Lippert, O. Stegle, B. Schölkopf, D. Weigel, K. Borgwardt. easyGWAS: A Cloud-based Platform for Comparing the Results of Genome-wide Association Studies. *The Plant Cell* **29** (1): 5–19, 2017.

[135] M. Khatami, T. Schmidt-Wilcke, P. C. Sundgren, A. Abbasloo, B. Schölkopf, T. Schultz. BundleMAP: Anatomically Localized Classification, Regression, and Hypothesis Testing in Diffusion MRI. *Pattern Recognition* **63**: 593–600, 2017.

[136] T. Huang, H. Peng, K. Zhang. Model Selection for Gaussian Mixture Models. *Statistica Sinica* **27** (1): 147–169, 2017.

[137] A. Kupcsik, M. Deisenroth, J. Peters, L. Ai Poh, V. Vadakkepat, G. Neumann. Model-based Contextual Policy Search for Data-Efficient Generalization of Robot Skills. *Artificial Intelligence* **247**. Special Issue on AI and Robotics: 415–439, 2017.

## 2016

[12] C. Daniel, H. van Hoof, J. Peters, G. Neumann. Probabilistic Inference for Determining Options in Reinforcement Learning. *Machine Learning, Special Issue* **104** (2). Best Student Paper Award of ECML-PKDD 2016: 337–357, 2016.

[138] F. Seith, S. Gatidis, H. Schmidt, I. Bezrukov, C. la Fougère, K. Nikolaou, C. Pfannenberg, N. Schwenzer. Comparison of Positron Emission Tomography Quantification Using Magnetic Resonance–and Computed Tomography–Based Attenuation Correction in Physiological Tissues and Lesions: A Whole-Body Positron Emission Tomography/Magnetic Resonance Study in 66 Patients. *Investigative Radiology* **51** (1): 66–71, 2016.

[139] E. Rueckert, D. Kappel, D. Tanneberg, D. Pecevski, J. Peters. Recurrent Spiking Networks Solve Planning Tasks. *Nature PG: Scientific Reports* **6** (Article number: 21142), 2016.

[140] E. Rueckert, J. Camernik, J. Peters, J. Babic. Probabilistic Movement Models Show that Postural Control Precedes and Predicts Volitional Motor Control. *Nature PG: Scientific Reports* **6** (Article number: 28455), 2016.

[141] G. Maeda, M. Ewerton, D. Koert, J. Peters. Acquiring and Generalizing the Embodiment Mapping from Human Observations to Robot Skills. *IEEE Robotics and Automation Letters* **1** (2): 784–791, 2016.

[142] C. Daniel, G. Neumann, O. Kroemer, J. Peters. Hierarchical Relative Entropy Policy Search. *Journal of Machine Learning Research* **17** (93): 1–50, 2016.

[143] K. Zhang, J. Li, E. Bareinboim, B. Schölkopf, J. Pearl. Preface to the ACM TIST Special Issue on Causal Discovery and Inference. *ACM Transactions on Intelligent Systems and Technologies* **7** (2): article no. 17, 2016.

[144] M. Gomez-Rodriguez, L. Song, N. Du, H. Zha, B. Schölkopf. Influence Estimation and Maximization in Continuous-Time Diffusion Networks. *ACM Transactions on Information Systems* **34** (2): 9:1–9:33, 2016.

[145] T. S. A. Wallis, M. Bethge, F. A. Wichmann. Testing models of peripheral encoding using metamerism in an oddity paradigm. *Journal of Vision* **16** (2), 2016 (cited on pages 36, 53).

[146] H. H. Schütt, S. Harmeling, J. H. Macke, F. A. Wichmann. Painfree and accurate Bayesian estimation of psychometric functions for (potentially) overdispersed data. *Vision Research* **122**: 105–123, 2016 (cited on page 53).

[147] L. O. M. Rothkegel, H. A. Trukenbrod, H. H. Schütt, F. A. Wichmann, R. Engbert. Influence of initial fixation position in scene viewing. *Vision Research* **129**: 33–49, 2016.

[148] F. Jäkel, M. Singh, F. A. Wichmann, M. H. Herzog. An overview of quantitative approaches in Gestalt perception. *Vision Research* **126**: 3–8, 2016.

[149] T. Fomina, G. Lohmann, M. Erb, T. Ethofer, B. Schölkopf, M. Grosse-Wentrup. Self-regulation of brain rhythms in the precuneus: a novel BCI paradigm for patients with ALS. *Journal of Neural Engineering* **13** (6): 066021, 2016 (cited on pages 36, 39).

[150] D. Foreman-Mackey, T. D. Morton, D. W. Hogg, E. Agol, B. Schölkopf. The population of long-period transiting exoplanets. *The Astronomical Journal* **152** (6): 206, 2016 (cited on page 38).

[151] D. Janzing, R. Chaves, B. Schölkopf. Algorithmic independence of initial condition and dynamical law in thermodynamics and causal inference. *New Journal of Phyiscs* **18** (9): 093052, 2016 (cited on pages 34, 41).

[152] M. Gomez Rodriguez, L. Song, H. Daneshmand, B. Schölkopf. Estimating Diffusion Networks: Recovery Conditions, Sample Complexity and Soft-thresholding Algorithm. *Journal of Machine Learning Research* **17** (90): 1–29, 2016.

[153] D. Wang, D. W. Hogg, D. Foreman-Mackey, B. Schölkopf. A Causal, Data-driven Approach to Modeling the Kepler Data. *Publications of the Astronomical Society of the Pacific* **128** (967): 094503, 2016 (cited on pages 34, 38, 40).

[154] P. Spirtes, K. Zhang. Causal discovery and inference: concepts and recent methodological advances. *Applied Informatics* **3** (3): 1–28, 2016.

[155] J. Townsend, N. Koep, S. Weichwald. Pymanopt: A Python Toolbox for Optimization on Manifolds using Automatic Differentiation. *Journal of Machine Learning Research* **17** (137): 1–5, 2016.

[156] M. R. Divine, P. Katiyar, U. Kohlhofer, L. Quintanilla-Martinez, J. A. Disselhorst, B. J. Pichler. A Population Based Gaussian Mixture Model Incorporating 18F-FDG-PET and DW-MRI Quantifies Tumor Tissue Classes. *Journal of Nuclear Medicine* **57** (3): 473–479, 2016.

[157] J. Peters, P. Bühlmann, N. Meinshausen. Causal inference using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* **78** (5). (with discussion): 947–1012, 2016 (cited on page 41).

[158] S. Weichwald, M. Grosse-Wentrup, A. Gretton. MERLiN: Mixture Effect Recovery in Linear Networks. *IEEE Journal of Selected Topics in Signal Processing* **10** (7): 1254–1266, 2016.

[159] V. Jayaram, M. Alamgir, Y. Altun, B. Schölkopf, M. Grosse-Wentrup. Transfer Learning in Brain-Computer Interfaces. *IEEE Computational Intelligence Magazine* **11** (1): 20–31, 2016 (cited on pages 36, 39).

[160] C. J. Schuler, M. Hirsch, S. Harmeling, B. Schölkopf. Learning to Deblur. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **38** (7): 1439–1451, 2016.

[161] E. D. Klenske, P. Hennig. Dual Control for Approximate Bayesian Reinforcement Learning. *Journal of Machine Learning Research* **17** (127): 1–30, 2016.

[162] M. Grosse-Wentrup, D. Janzing, M. Siegel, B. Schölkopf. Identification of causal relations in neuroimaging data with latent confounders: An instrumental variable approach. *NeuroImage* **125**: 825–833, 2016 (cited on page 34).

[163] R. Babbar, I. Partalas, E. Gaussier, M. Amini, C. Amblard. Learning Taxonomy Adaptation in Large-scale Classification. *Journal of Machine Learning Research* **17** (98): 1–37, 2016.

[164] K. Muandet, B. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf. Kernel Mean Shrinkage Estimators. *Journal of Machine Learning Research* **17** (48): 1–41, 2016 (cited on pages 32, 55).

[165] J. Mooij, J. Peters, D. Janzing, J. Zscheischler, B. Schölkopf. Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research* **17** (32): 1–102, 2016 (cited on page 41).

[166] B. Schölkopf, D. Hogg, D. Wang, D. Foreman-Mackey, D. Janzing, C. J. Simon-Gabriel, J. Peters. Modeling Confounding by Half-Sibling Regression. *Proceedings of the National Academy of Science* **113** (27): 7391–7398, 2016 (cited on pages 34, 38, 40).

[167] E. D. Klenske, M. Zeilinger, B. Schölkopf, P. Hennig. Gaussian Process-Based Predictive Control for Periodic Error Correction. *IEEE Transactions on Control Systems Technology* **24** (1): 110–121, 2016.

[168] K. Zhang, Z. Wang, J. Zhang, B. Schölkopf. On estimation of functional causal models: General results and application to post-nonlinear causal model. *ACM Transactions on Intelligent Systems and Technologies* **7** (2): article no. 13, 2016.

### 21.2.4 Conference Papers

#### 2018

[169] J. Kim, B. Tabibian, A. Oh, B. Schölkopf, M. Gomez Rodriguez. Leveraging the Crowd to Detect and Reduce the Spread of Fake News and Misinformation. In *Proceedings of the 11th ACM International Conference on Web Search and Data Mining (WSDM)*, pages 324–332, Feb. 2018 (cited on page 48).

[170] O. Nachum, S. Gu, H. Lee, S. Levine. Data-Efficient Hierarchical Reinforcement Learning. In *Advances in Neural Information Processing Systems 31*, pages 3307–3317, 2018 (cited on page 47).

[171] W. M. Harder F. Köhler J., P. M. DP-MAC: The Differentially Private Method of Auxiliary Coordinates for Deep Learning. In *Workshop on Privacy Preserving Machine Learning at the 32nd Conference on Neural Information Processing Systems*, 2018 (cited on page 50).

[172] J. Chen, M. Hirsch, B. Eberhardt, H. P. A. Lensch. A Computational Camera with Programmable Optics for Snapshot High Resolution Multispectral Imaging. In *Computer Vision - ACCV 2018 - 14th Asian Conference on Computer Vision*, 2018.

[173] R. Geirhos, C. R. M. Temme, J. Rauber, H. Schütt, M. Bethge, F. A. Wichmann. Generalisation in humans and deep neural networks. In *Advances in Neural Information Processing Systems 31*, pages 7549–7561, 2018 (cited on pages 36, 53).

[174] G. Maeda, O. Koc, J. Morimoto. Reinforcement Learning of Phase Oscillators for Fast Adaptation to Moving Targets. In *Proceedings of The 2nd Conference on Robot Learning (CoRL)*. Vol. 87, pages 630–640, 2018.

[175] J. Kapoor, A. Vergari, M. Gomez Rodriguez, I. Valera. Bayesian Nonparametric Hawkes Processes. In *Bayesian Nonparametrics workshop at the 32nd Conference on Neural Information Processing Systems*, 2018.

[176] Y. Mroueh, C.-L. Li*, T. Sercu*, A. Raj*, Y. Cheng. Sobolev GAN. In *6th International Conference on Learning Representations (ICLR)*. *equal contribution, 2018 (cited on page 46).

[177] P. K. Rubenstein, B. Schölkopf, I. Tolstikhin. Learning Disentangled Representations with Wasserstein Auto-Encoders. In *Workshop at the 6th International Conference on Learning Representations (ICLR)*, 2018 (cited on pages 33, 56).

[178] P. K. Rubenstein, B. Schölkopf, I. Tolstikhin. Wasserstein Auto-Encoders: Latent Dimensionality and Random Encoders. In *Workshop at the 6th International Conference on Learning Representations (ICLR)*, 2018 (cited on pages 33, 56).

[179] R. Pinsler, R. Akrour, T. Osa, J. Peters, G. Neumann. Sample and Feedback Efficient Hierarchical Reinforcement Learning from Human Preferences. In *IEEE International Conference on Robotics and Automation, (ICRA)*, pages 596–601, 2018.

[180] F. Muratore, F. Treede, M. Gienger, J. Peters. Domain Randomization for Simulation-Based Policy Optimization with Transferability Assessment. In *2nd Annual Conference on Robot Learning (CoRL)*. Vol. 87. Proceedings of Machine Learning Research, pages 700–713, 2018.

[181] R. Lioutikov, G. Maeda, F. Veiga, K. Kersting, J. Peters. Inducing Probabilistic Context-Free Grammars for the Sequencing of Movement Primitives. In *IEEE International Conference on Robotics and Automation, (ICRA)*, pages 1–8, 2018.

[182] D. Koert, G. Maeda, G. Neumann, J. Peters. Learning Coupled Forward-Inverse Models with Combined Prediction Errors. In *IEEE International Conference on Robotics and Automation, (ICRA)*, pages 2433–2439, 2018.

[183] R. Akrour, F. Veiga, J. Peters, G. Neuman. Regularizing Reinforcement Learning with State Abstraction. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.

[184] R. Akrour, J. Peters, G. Neuman. Constraint-Space Projection Direct Policy Search. In *14th European Workshop on Reinforcement Learning (EWRL)*, 2018.

[185] F. Groh*, P. Wieschollek*, H. P. A. Lensch. Flex-Convolution (Million-Scale Point-Cloud Learning Beyond Grid-Worlds). In *Computer Vision - ACCV 2018 - 14th Asian Conference on Computer Vision*. *equal contribution, 2018.

[186] P. Wieschollek, O. Gallo, J. Gu, J. Kautz. Separating Reflection and Transmission Images in the Wild. In *European Conference on Computer Vision (ECCV)*, 2018.

[187] M. van der Wilk, M. Bauer, S. T. John, J. Hensman. Learning Invariances using the Marginal Likelihood. In *Advances in Neural Information Processing Systems 31*, pages 9960–9970, 2018 (cited on pages 34, 51).

[188] F. Locatello*, G. Dresdner*, K. R., I. Valera, G. Rätsch. Boosting Black Box Variational Inference. In *Advances in Neural Information Processing Systems 31*. *equal contribution, pages 3405–3415, 2018 (cited on pages 34, 51).

[189] I. Valera, A. Singla, M. Gomez Rodriguez. Enhancing the Accuracy and Fairness of Human Decision Making. In *Advances in Neural Information Processing Systems 31*, pages 1774–1783, 2018 (cited on page 44).

[190] W. Jitkrittum, H. Kanagawa, P. Sangkloy, J. Hays, B. Schölkopf, A. Gretton. Informative Features for Model Comparison. In *Advances in Neural Information Processing Systems 31*, pages 816–827, 2018 (cited on pages 33, 45).

[191] A. Neitz, G. Parascandolo, S. Bauer, B. Schölkopf. Adaptive Skip Intervals: Temporal Abstraction for Recurrent Dynamical Models. In *Advances in Neural Information Processing Systems 31*, pages 9838–9848, 2018 (cited on pages 34, 41).

[192] M. W. Gondal, B. Schölkopf, M. Hirsch. The Unreasonable Effectiveness of Texture Transfer for Single Image Super-resolution. In *Workshop and Challenge on Perceptual Image Restoration and Manipulation (PIRM) at the 15th European Conference on Computer Vision (ECCV)*, 2018 (cited on page 42).

[193] M. S. M. Sajjadi, O. Bachem, M. Lucic, O. Bousquet, S. Gelly. Assessing Generative Models via Precision and Recall. In *Workshop on Theoretical Foundations and Applications of Deep Generative Models (TADGM) at the 35th International Conference on Machine Learning (ICML)*, 2018.

[194] M. Besserve, R. Sun, B. Schölkopf. Intrinsic disentanglement: an invariance view for deep generative models. In *Workshop on Theoretical Foundations and Applications of Deep Generative Models at ICML*, 2018.

[195] G. Tucker, S. Bhupatiraju, S. Gu, R. Turner, Z. Ghahramani, S. Levine. The Mirage of Action-Dependent Baselines in Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 5022–5031, 2018 (cited on page 47).

[196] A. Ścibior, O. Kammar, Z. Ghahramani. Functional Programming for Modular Bayesian Inference. In *Proceedings of the ACM on Functional Programming (ICFP)*. Vol. 2. Article No. 83, pages 1–29, 2018 (cited on page 52).

[197] A. Ścibior, O. Kammar, M. Vákár, S. Staton, H. Yang, Y. Cai, K. Ostermann, S. K. Moss, C. Heunen, Z. Ghahramani. Denotational Validation of Higher-order Bayesian Inference. In *Proceedings of the ACM on Principles of Programming Languages (POPL)*. Vol. 2. Article No. 60, pages 1–29, 2018 (cited on page 52).

[198] T. H. Kim, M. S. M. Sajjadi, M. Hirsch, B. Schölkopf. Spatio-temporal Transformer Network for Video Restoration. In *15th European Conference on Computer Vision (ECCV), Part III*. Vol. 11207. Lecture Notes in Computer Science, pages 111–127, 2018 (cited on pages 35, 42).

[199] M. S. M. Sajjadi, O. Bachem, M. Lucic, O. Bousquet, S. Gelly. Assessing Generative Models via Precision and Recall. In *Advances in Neural Information Processing Systems 31*, pages 5234–5243, 2018 (cited on page 46).

[200] A. Loktyushin, P. Ehses, B. Schölkopf, K. Scheffler. Learning-based solution to phase error correction in T2*-weighted GRE scans. In *1st International conference on Medical Imaging with Deep Learning (MIDL)*, 2018.

[201] M. Dehghani, A. Mehrjou, S. Gouws, J. Kamps, B. Schölkopf. Fidelity-Weighted Learning. In *6th International Conference on Learning Representations (ICLR)*, 2018.

[202] F. Solowjow, A. Mehrjou, B. Schölkopf, S. Trimpe. Efficient Encoding of Dynamical Systems through Local Approximations. In *Proceedings of the 57th IEEE International Conference on Decision and Control (CDC)*, 2018 (cited on page 198).

[203] G. Parascandolo, N. Kilbertus, M. Rojas-Carulla, B. Schölkopf. Learning Independent Causal Mechanisms. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 4033–4041, 2018 (cited on pages 34, 41, 46).

[204] D. Janzing, B. Schölkopf. Detecting non-causal artifacts in multivariate linear regression models. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 2250–2258, 2018.

[205] P. Parmas, C. Rasmussen, J. Peters, K. Doya. PIPPS: Flexible Model-Based Policy Search Robust to the Curse of Chaos. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 4065–4074, 2018.

[206] N. Kilbertus, A. Gascon, M. Kusner, M. Veale, K. Gummadi, A. Weller. Blind Justice: Fairness with Encrypted Sensitive Attributes. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 2635–2644, 2018 (cited on pages 43, 44).

[207] A. Yurtsever, O. Fercoq, F. Locatello, V. Cevher. A Conditional Gradient Framework for Composite Convex Minimization with Applications to Semidefinite Programming. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 5713–5722, 2018 (cited on page 46).

[208] F. Locatello, A. Raj, S. Praneeth Karimireddy, G. Rätsch, B. Schölkopf, S. U. Stich, M. Jaggi. On Matching Pursuit and Coordinate Descent. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 3204–3213, 2018 (cited on pages 33, 46).

[209] B. Huang, K. Zhang, Y. Lin, B. Schölkopf, G. C. Generalized Score Functions for Causal Discovery. In *Proceedings of the 24th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1551–1560, 2018 (cited on page 41).

[210] M. Bauer, V. Volchkov, M. Hirsch, B. Schölkopf. Automatic Estimation of Modulation Transfer Functions. In *IEEE International Conference on Computational Photography (ICCP)*, 2018 (cited on pages 42, 274).

[211] B. Eysenbach, S. Gu, J. Ibarz, S. Levine. Leave no Trace: Learning to Reset for Safe and Autonomous Reinforcement Learning. In *6th International Conference on Learning Representations (ICLR)*, 2018 (cited on page 47).

[212] V. Pong*, S. Gu*, M. Dalal, S. Levine. Temporal Difference Models: Model-Free Deep RL for Model-Based Control. In *6th International Conference on Learning Representations (ICLR)*. *equal contribution, 2018 (cited on page 47).

[213] I. Tolstikhin, O. Bousquet, S. Gelly, B. Schölkopf. Wasserstein Auto-Encoders. In *6th International Conference on Learning Representations (ICLR)*, 2018 (cited on pages 33, 56).

[214] M. S. M. Sajjadi, G. Parascandolo, A. Mehrjou, B. Schölkopf. Tempered Adversarial Networks. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 4448–4456, 2018 (cited on page 46).

[215] F. Locatello, R. Khanna, J. Ghosh, G. Rätsch. Boosting Variational Inference: an Optimization Perspective. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 84. Proceedings of Machine Learning Research, pages 464–472, 2018 (cited on pages 34, 46, 51).

[216] M. Besserve, N. Shajarisales, B. Schölkopf, D. Janzing. Group invariance principles for causal generative models. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 84. Proceedings of Machine Learning Research, pages 557–565, 2018 (cited on pages 34, 41, 49).

[217] P. Blöbaum, D. Janzing, T. Washio, S. Shimizu, B. Schölkopf. Cause-Effect Inference by Comparing Regression Errors. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 84. Proceedings of Machine Learning Research, pages 900–909, 2018 (cited on page 41).

[218] M. S. M. Sajjadi, R. Vemulapalli, M. Brown. Frame-Recurrent Video Super-Resolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 (cited on page 42).

[219] M. Balog, I. Tolstikhin, B. Schölkopf. Differentially Private Database Release via Kernel Mean Embeddings. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*. Vol. 80. Proceedings of Machine Learning Research, pages 423–431, 2018 (cited on pages 33, 45, 50, 56).

[220] P. K. Rubenstein, S. Bongers, B. Schölkopf, J. M. Mooij. From Deterministic ODEs to Dynamic Structural Causal Models. In *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2018 (cited on page 41).

## 2017

[27] R. Akrour, D. Sorokin, J. Peters, G. Neumann. Local Bayesian Optimization of Motor Skills. In *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. Proceedings of Machine Learning Research, pages 41–50, 2017.

[28] L. Grossberger, M. R. Hohmann, J. Peters, M. Grosse-Wentrup. Investigating Music Imagery as a Cognitive Paradigm for Low-Cost Brain-Computer Interfaces. In *Proceedings of the 7th Graz Brain-Computer Interface Conference 2017 - From Vision to Reality*, pages 160–164, 2017.

[29] K.-H. Fiebig, V. Jayaram, T. Hesse, A. Blank, J. Peters, M. Grosse-Wentrup. Bayesian Regression for Artifact Correction in Electroencephalography. In *Proceedings of the 7th Graz Brain-Computer Interface Conference 2017 - From Vision to Reality*, pages 131–136, 2017.

[221] J. Kim, B. Tabibian, A. Oh, B. Schölkopf, M. Gomez Rodriguez. Leveraging the Crowd to Detect and Reduce the Spread of Fake News and Misinformation. In *Workshop on Prioritising Online Content at the 31st Conference on Neural Information Processing Systems*, 2017.

[222] B. Tabibian, U. Upadhyay, A. De, A. Zarezade, B. Schölkopf, M. Gomez Rodriguez. Optimizing human learning. In *Workshop on Teaching Machines, Robots, and Humans at the 31st Conference on Neural Information Processing Systems*, 2017 (cited on page 48).

[223] M. B. Zafar, I. Valera, M. Gomez Rodriguez, K. Gummadi, A. Weller. From Parity to Preference-based Notions of Fairness in Classification. In *Advances in Neural Information Processing Systems 30*, pages 229–239, 2017 (cited on pages 35, 43).

[224] D. Wilbers, R. Lioutikov, J. Peters. Context-Driven Movement Primitive Adaptation. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3469–3475, 2017.

[225] V. Tangkaratt, H. van Hoof, S. Parisi, G. Neumann, J. Peters, M. Sugiyama. Policy Search with High-Dimensional Context Variables. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*, pages 2632–2638, 2017.

[226] E. Rueckert, M. Nakatenus, S. Tosatto, J. Peters. Learning inverse dynamics models in O(n) time with LSTM networks. In *IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 811–816, 2017.

[227] S. Parisi, S. Ramstedt, J. Peters. Goal-driven dimensionality reduction for reinforcement learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4634–4639, 2017.

[228] G. Maeda, M. Ewerton, T. Osa, B. Busch, J. Peters. Active Incremental Learning of Robot Movement Primitives. In *Proceedings of the 1st Annual Conference on Robot Learning (CoRL)*. Vol. 78. Proceedings of Machine Learning Research, pages 37–46, 2017.

[229] A. Gabriel, R. Akrour, J. Peters, G. Neumann. Empowered skills. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 6435–6441, 2017.

[230] F. B. Farraj, T. Osa, N. Pedemonte, J. Peters, G. Neumann, P. Giordano. A Learning-based Shared Control Architecture for Interactive Task Execution. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 329–335, 2017.

[231] F. End, R. Akrour, J. Peters, G. Neumann. Layered direct policy search for learning hierarchical skills. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 6442–6448, 2017.

[232] H. Abdulsamad, O. Arenz, J. Peters, G. Neumann. State-Regularized Policy Search for Linearized Dynamical Systems. In *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, (ICAPS)*, pages 419–424, 2017.

[233] M. Bauer*, M. Rojas-Carulla*, J. B. Świątkowski, B. Schölkopf, R. E. Turner. Discriminative k-shot learning using probabilistic models. In *Second Workshop on Bayesian Deep Learning at the 31st Conference on Neural Information Processing Systems*. *equal contribution, 2017.

[234] F. Locatello, M. Tschannen, G. Rätsch, M. Jaggi. Greedy Algorithms for Cone Constrained Optimization with Convergence Guarantees. In *Advances in Neural Information Processing Systems 30*, pages 773–784, 2017 (cited on page 46).

[235] T. Gebhard, N. Kilbertus, G. Parascandolo, I. Harry, B. Schölkopf. ConvWave: Searching for Gravitational Waves with Fully Convolutional Neural Nets. In *Workshop on Deep Learning for Physical Sciences (DLPS) at the 31st Conference on Neural Information Processing Systems*, 2017 (cited on page 38).

[236] S. U. Stich, A. Raj, M. Jaggi. Safe Adaptive Importance Sampling. In *Advances in Neural Information Processing Systems 30*, pages 4384–4394, 2017 (cited on pages 33, 45).

[237] N. Kilbertus, M. Rojas-Carulla, G. Parascandolo, M. Hardt, D. Janzing, B. Schölkopf. Avoiding Discrimination through Causal Reasoning. In *Advances in Neural Information Processing Systems 30*, pages 656–666, 2017 (cited on pages 34, 35, 44).

[238] S. Gu, T. Lillicrap, R. E. Turner, Z. Ghahramani, B. Schölkopf, S. Levine. Interpolated Policy Gradient: Merging On-Policy and Off-Policy Gradient Estimation for Deep Reinforcement Learning. In *Advances in Neural Information Processing Systems 30*, pages 3849–3858, 2017 (cited on page 47).

[239] K. Meding, A. Loktyushin, M. Hirsch. Automatic detection of motion artifacts in MR images using CNNs. In *42nd IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 811–815, 2017 (cited on page 42).

[240] M. Balog, N. Tripuraneni, Z. Ghahramani, A. Weller. Lost Relatives of the Gumbel Trick. In *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. Proceedings of Machine Learning Research, pages 371–379, 2017 (cited on page 52).

[241] S. Gu*, E. Holly*, T. Lillicrap, S. Levine. Deep Reinforcement Learning for Robotic Manipulation with Asynchronous Off-Policy Updates. In *Proceedings 2017 IEEE International Conference on Robotics and Automation (ICRA)*. *equal contribution, 2017 (cited on page 47).

[242] E. Jang, S. Gu, B. Poole. Categorical Reparametrization with Gumbel-Softmax. In *Proceedings International Conference on Learning Representations 2017*, 2017.

[243] S. Gu, T. Lillicrap, Z. Ghahramani, R. E. Turner, S. Levine. Q-Prop: Sample-Efficient Policy Gradient with An Off-Policy Critic. In *Proceedings International Conference on Learning Representations (ICLR)*, 2017 (cited on page 47).

[244] N. Jaques, S. Gu, D. Bahdanau, J. M. Hernández-Lobato, R. E. Turner, D. Eck. Sequence Tutor: Conservative fine-tuning of sequence generation models with KL-control. In *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. Proceedings of Machine Learning Research, pages 1645–1654, 2017.

[245] B. Huang, K. Zhang, J. Zhang, R. Sanchez-Romero, C. Glymour, B. Schölkopf. Behind Distribution Shift: Mining Driving Forces of Changes and Causal Arrows. In *IEEE 17th International Conference on Data Mining (ICDM)*, pages 913–918, 2017 (cited on pages 34, 41).

[246] A.-A. Mastakouri, S. Weichwald, O. Ozdenizci, T. Meyer, B. Schölkopf, M. Grosse-Wentrup. Personalized Brain-Computer Interface Models for Motor Rehabilitation. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 3024–3029, 2017.

[247] P. Wieschollek, M. Hirsch, B. Schölkopf, H. Lensch. Learning Blind Motion Deblurring. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, pages 231–240, 2017 (cited on pages 35, 42).

[248] T. H. Kim, K. M. Lee, B. Schölkopf, M. Hirsch. Online Video Deblurring via Dynamic Temporal Blending Network. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, pages 4038–4047, 2017 (cited on pages 35, 42).

[249] P. K. Rubenstein*, S. Weichwald*, S. Bongers, J. M. Mooij, D. Janzing, M. Grosse-Wentrup, B. Schölkopf. Causal Consistency of Structural Equation Models. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*. *equal contribution, 2017 (cited on page 41).

[250] M. Gong, K. Zhang, B. Schölkopf, C. Glymour, D. Tao. Causal Discovery from Temporally Aggregated Time Series. In *Proceedings Conference on Uncertainty in Artificial Intelligence (UAI) 2017*, page ID 269, 2017 (cited on page 41).

[251] K. Zhang, B. Huang, J. Zhang, C. Glymour, B. Schölkopf. Causal Discovery from Nonstationary/Heterogeneous Data: Skeleton Estimation and Orientation Determination. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1347–1353, 2017.

[252] S. Stich, A. Raj, M. Jaggi. Approximate Steepest Coordinate Descent. In *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. Proceedings of Machine Learning Research, pages 3251–3259, 2017 (cited on pages 33, 45).

[253] A. Raj, A. Kumar, Y. Mroueh, T. Fletcher, B. Schölkopf. Local Group Invariant Representations via Orbit Embeddings. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 54. Proceedings of Machine Learning Research, pages 1225–1235, 2017.

[254] M. Görner, B. Schölkopf, M. Grosse-Wentrup. Closing One's Eyes Affects Amplitude Modulation but Not Frequency Modulation in a Cognitive BCI. In *Proceedings of the 7th Graz Brain-Computer Interface Conference 2017 - From Vision to Reality*, pages 165–170, 2017.

[255] J. Moser, M. R. Hohmann, B. Schölkopf, M. Grosse-Wentrup. A Guided Task for Cognitive Brain-Computer Interfaces. In *Proceedings of the 7th Graz Brain-Computer Interface Conference 2017 - From Vision to Reality*, pages 326–331, 2017.

[256] M. W. Gondal, J. M. Köhler, R. Grzeszick, G. Fink, M. Hirsch. Weakly-Supervised Localization of Diabetic Retinopathy Lesions in Retinal Fundus Images. In *IEEE International Conference on Image Processing (ICIP)*, pages 2069–2073, 2017.

[257] M. Schober, A. Adam, O. Yair, S. Mazor, S. Nowozin. Dynamic Time-of-Flight. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 170–179, 2017.

[258] D. Lopez-Paz, R. Nishihara, S. Chintala, B. Schölkopf, L. Bottou. Discovering Causal Signals in Images. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 58–66, 2017 (cited on page 41).

[259] C. Lu, M. Hirsch, B. Schölkopf. Flexible Spatio-Temporal Networks for Video Prediction. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 2137–2145, 2017 (cited on page 42).

[260] V. Jayaram, B. Schölkopf, M. Grosse-Wentrup. Frequency Peak Features for Low-Channel Classification in Motor Imagery Paradigms. In *Proceedings of the 8th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 321–324, 2017.

[261] I. Tolstikhin, S. Gelly, O. Bousquet, C. J. Simon-Gabriel, B. Schölkopf. AdaGAN: Boosting Generative Models. In *Advances in Neural Information Processing Systems 30*, pages 5430–5439, 2017 (cited on pages 33, 46, 56).

[262] M. Balog, A. L. Gaunt, M. Brockschmidt, S. Nowozin, D. Tarlow. DeepCoder: Learning to Write Programs. In *Proceedings International Conference on Learning Representations 2017*, 2017 (cited on page 46).

[263] M. S. M. Sajjadi, B. Schölkopf, M. Hirsch. EnhanceNet: Single Image Super-Resolution through Automated Texture Synthesis. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, pages 4501–4510, 2017 (cited on pages 35, 42, 46).

[264] B. Tabibian, I. Valera, M. Farajtabar, L. Song, B. Schölkopf, M. Gomez Rodriguez. Distilling Information Reliability and Source Trustworthiness from Digital Traces. In *Proceedings of the 26th International Conference on World Wide Web (WWW)*, pages 847–855, 2017.

[265] R. Babbar, B. Schölkopf. DiSMEC – Distributed Sparse Machines for Extreme Multi-label Classification. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM)*, pages 721–729, 2017 (cited on page 46).

[266] P. Wieschollek, B. Schölkopf, H. P. A. Lensch, M. Hirsch. End-to-End Learning for Image Burst Deblurring. In *Computer Vision - ACCV 2016 - 13th Asian Conference on Computer Vision*. Vol. 10114. Image Processing, Computer Vision, Pattern Recognition, and Graphics, pages 35–51, 2017 (cited on page 42).

## 2016

[46] Y. Huang, D. Büchler, O. Koc, B. Schölkopf, J. Peters. Jointly Learning Trajectory Generation and Hitting Point Prediction in Robot Table Tennis. In *16th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 650–655, 2016.

[47] S. Gomez-Gonzalez, G. Neumann, B. Schölkopf, J. Peters. Using Probabilistic Movement Primitives for Striking Movements. In *16th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 502–508, 2016 (cited on page 54).

[49] O. Koc, G. Maeda, J. Peters. A New Trajectory Generation Framework in Robotic Table Tennis. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 3750–3756, 2016 (cited on pages 36, 54).

[56] D. Büchler, H. Ott, J. Peters. A Lightweight Robotic Arm with Pneumatic Muscles for Robot Learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 4086–4092, 2016 (cited on pages 36, 47).

[267] F. Leibfried, D. Braun. Bounded Rational Decision-Making in Feedforward Neural Networks. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 407–416, 2016.

[268] B. Tabibian, M. Farajtabar, I. Valera, L. Song, B. Schölkopf, M. Gomez Rodriguez. On the Reliability of Information and Trustworthiness of Web Sources in Wikipedia. In *Wikipedia workshop at the 10th International AAAI Conference on Web and Social Media (ICWSM)*, 2016.

[269] Z. Yi, R. Calandra, F. Veiga, H. van Hoof, T. Hermans, Y. Zhang, J. Peters. Active Tactile Object Exploration with Gaussian Processes. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4925–4930, 2016.

[270] H. van Hoof, N. Chen, M. Karl, P. van der Smagt, J. Peters. Stable Reinforcement Learning with Autoencoders for Tactile and Visual Data. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 3928–3934, 2016.

[271] V. Modugno, G. Neumann, E. Rueckert, G. Oriolo, J. Peters, S. Ivaldi. Learning soft task priorities for control of redundant robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 221–226, 2016.

[272] S. Manschitz, M. Gienger, J. Kober, J. Peters. Probabilistic Decomposition of Sequential Force Interaction Tasks into Movement Primitives. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3920–3927, 2016.

[273] D. Koert, G. Maeda, R. Lioutikov, G. Neumann, J. Peters. Demonstration Based Trajectory Optimization for Generalizable Robot Motions. In *IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 351–358, 2016.

[274] M. Ewerton, G. Maeda, G. Kollegger, J. Wiemeyer, J. Peters. Incremental Imitation Learning of Context-Dependent Motor Skills. In *IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 351–358, 2016.

[275] M. Ewerton, G. Maeda, G. Neumann, V. Kisner, G. Kollegger, J. Wiemeyer, J. Peters. Movement Primitives with Multiple Phase Parameters. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 201–206, 2016.

[276] B. Belousov, G. Neumann, C. Rothkopf, J. Peters. Catching heuristics are optimal control policies. In *Advances in Neural Information Processing Systems 29*, pages 1426–1434, 2016.

[277] P. Wieschollek, O. Wang, A. Sorkine-Hornung, H. P. A. Lensch. Efficient Large-scale Approximate Nearest Neighbor Search on the GPU. In *29th IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2027–2035, 2016.

[278] B. Schölkopf, D. Janzing, D. Lopez-Paz. Causal and statistical learning. In *Oberwolfach Reports*. Vol. 13. 3, pages 1896–1899, 2016.

[279] S. Gu, S. Levine, I. Sutskever, A. Mnih. MuProp: Unbiased Backpropagation for Stochastic Neural Networks. In *4th International Conference on Learning Representations (ICLR)*, 2016.

[280] S. Gu, T. Lillicrap, I. Sutskever, S. Levine. Continuous Deep Q-Learning with Model-based Acceleration. In *Proceedings of the 33nd International Conference on Machine Learning (ICML)*. Vol. 48. JMLR Workshop and Conference Proceedings, pages 2829–2838, 2016 (cited on page 47).

[281] J. Borgström, A. D. Gordon, L. Ouyang, C. Russo, A. Ścibior, M. Szymczak. Fabular: Regression Formulas As Probabilistic Programming. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. POPL '16, pages 271–283, 2016.

[282] P. K. Rubenstein, K. P. Chwialkowski, A. Gretton. A Kernel Test for Three-Variable Interactions with Random Processes. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.

[283] M. Bauer, M. van der Wilk, C. E. Rasmussen. Understanding Probabilistic Sparse Gaussian Process Approximations. In *Advances in Neural Information Processing Systems 29*, pages 1533–1541, 2016 (cited on pages 34, 51).

[284] M. Balog, B. Lakshminarayanan, Z. Ghahramani, D. M. Roy, Y. W. Teh. The Mondrian Kernel. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016 (cited on pages 34, 51).

[285] L. Xiao, J. Wang, W. Heidrich, M. Hirsch. Learning High-Order Filters for Efficient Blind Deconvolution of Document Photographs. In *Computer Vision - ECCV 2016*. Vol. Lecture Notes in Computer Science, LNCS 9907, Part III, pages 734–749, 2016.

[286] E. D. Klenske, P. Hennig, B. Schölkopf, M. N. Zeilinger. Approximate dual control maintaining the value of information with an application to building control. In *European Control Conference (ECC)*, pages 800–806, 2016.

[287] I. Tolstikhin, B. K. Sriperumbudur, B. Schölkopf. Minimax Estimation of Maximum Mean Discrepancy with Radial Kernels. In *Advances in Neural Information Processing Systems 29*, pages 1930–1938, 2016 (cited on page 55).

[288] C. J. Simon-Gabriel*, A. Ścibior*, I. Tolstikhin, B. Schölkopf. Consistent Kernel Mean Estimation for Functions of Random Variables. In *Advances in Neural Information Processing Systems 29*. *joint first authors, pages 1732–1740, 2016 (cited on pages 35, 45, 52, 56).

[289] M. S. M. Sajjadi, R. Köhler, B. Schölkopf, M. Hirsch. Depth Estimation Through a Generative Model of Light Field Synthesis. In *Pattern Recognition - 38th German Conference (GCPR)*. Vol. 9796. Lecture Notes in Computer Science, pages 426–438, 2016.

[290] A. Kontorovich, S. Sabato, R. Urner. Active Nearest-Neighbor Learning in Metric Spaces. In *Advances in Neural Information Processing Systems 29*, pages 856–864, 2016 (cited on page 55).

[291] A. Pentina, R. Urner. Lifelong Learning with Weighted Majority Votes. In *Advances in Neural Information Processing Systems 29*, pages 3612–3620, 2016 (cited on page 55).

[292] S. Bauer, B. Schölkopf, J. Peters. The Arrow of Time in Multivariate Time Serie. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*. Vol. 48. JMLR Workshop and Conference Proceedings, pages 2043–2051, 2016 (cited on pages 34, 41, 53).

[293] J. González, Z. Dai, P. Hennig, N. Lawrence. Batch Bayesian Optimization via Local Penalization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 51. JMLR Workshop and Conference Proceedings, pages 648–657, 2016.

[294] H. Kersting, P. Hennig. Active Uncertainty Calibration in Bayesian ODE Solvers. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 309–318, 2016.

[295] S. Ben-David, R. Urner. On Version Space Compression. In *Algorithmic Learning Theory - 27th International Conference (ALT)*. Vol. 9925. Lecture Notes in Computer Science, pages 50–64, 2016 (cited on pages 32, 55).

[296] S. Bartels, P. Hennig. Probabilistic Approximate Least-Squares. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 51. JMLR Workshop and Conference Proceedings, pages 676–684, 2016 (cited on page 247).

[297] K. Zhang, J. Zhang, B. Huang, B. Schölkopf, C. Glymour. On the Identifiability and Estimation of Functional Causal Models in the Presence of Outcome-Dependent Selection. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 825–834, 2016 (cited on page 41).

[298] S. Etesami, N. Kiyavash, K. Zhang, K. Singhal. Learning Causal Interaction Network of Multivariate Hawkes Processes. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*. poster presentation, 2016 (cited on page 41).

[299] S. Weichwald, A. Gretton, B. Schölkopf, M. Grosse-Wentrup. Recovery of non-linear cause-effect relationships from linearly mixed neuroimaging data. In *Proceedings of the 6th International Workshop on Pattern Recognition in NeuroImaging (PRNI 2016)*, 2016.

[300] M. Gong, K. Zhang, T. Liu, D. Tao, C. Glymour, B. Schölkopf. Domain Adaptation with Conditional Transferable Components. In *Proceedings of the 33nd International Conference on Machine Learning (ICML)*. Vol. 48. JMLR Workshop and Conference Proceedings, pages 2839–2848, 2016 (cited on pages 34, 41).

[301] D. Lopez-Paz, B. Schölkopf, L. Bottou, V. Vapnik. Unifying distillation and privileged information. In *International Conference on Learning Representations (ICLR)*, 2016.

[302] R. Babbar, K. Muandet, B. Schölkopf. TerseSVM : A Scalable Approach for Learning Compact Models in Large-scale Classification. In *Proceedings of the 2016 SIAM International Conference on Data Mining (SDM)*, pages 234–242, 2016.

### 21.2.5 Book Chapters

#### 2018

[303] F. A. Wichmann, F. Jäkel. Methods in Psychophysics. In. *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*. 4th. Vol. 5 (Methodology). John Wiley & Sons, Inc. Chap. 7, 2018 (cited on pages 36, 53).

[304] V. Jayaram, K.-H. Fiebig, J. Peters, M. Grosse-Wentrup. Transfer Learning for BCIs. In *Brain–Computer Interfaces Handbook*. CRC Press. Chap. 22, pages 425–442, 2018.

#### 2017

[63] J. Peters, D. Lee, J. Kober, D. Nguyen-Tuong, J. Bagnell, S. Schaal. Robot Learning. In. *Springer Handbook of Robotics*. 2nd. Springer International Publishing. Chap. 15, pages 357–394, 2017.

[305] J. Peters, R. Tedrake, N. Roy, J. Morimoto. Robot Learning. In. *Encyclopedia of Machine Learning and Data Mining*. 2nd. Springer US, pages 1106–1109, 2017.

[306] J. Peters, J. Bagnell. Policy Gradient Methods. In. *Encyclopedia of Machine Learning and Data Mining*. 2nd. Springer US, pages 982–985, 2017.

#### 2016

[307] K. Zhang, A. Hyvärinen. Nonlinear functional causal models for distinguishing cause from effect. In. *Statistics and Causality: Methods for Applied Empirical Research*. 1st. John Wiley & Sons, Inc. Chap. 8, pages 185–201, 2016 (cited on page 41).

[308] M. Hohmann, T. Fomina, V. Jayaram, N. Widmann, C. Förster, J. Just, M. Synofzik, B. Schölkopf, L. Schöls, M. Grosse-Wentrup. A cognitive brain–computer interface for patients with amyotrophic lateral sclerosis. In *Brain-Computer Interfaces: Lab Experiments to Real-World Applications*. Vol. 228. Progress in Brain Research Supplement C. Elsevier. Chap. 8, pages 221–239, 2016 (cited on pages 36, 39).

### 21.2.6 Theses

#### PhD Theses

[309] C. J. Simon-Gabriel. Distribution-Dissimilarities in Machine Learning. PhD thesis. Eberhard Karls Universität Tübingen, Germany, 2018.

[310] V. Jayaram. A machine learning approach to taking EEG-based computer interfaces out of the lab. PhD thesis. Graduate Training Centre of Neuroscience, IMPRS, Eberhard Karls Universität Tübingen, Germany, 2018.

[311] O. Koc. Optimal Trajectory Generation and Learning Control for Robot Table Tennis. PhD thesis. Technical University Darmstadt, Germany, 2018.

[312] M. Schober. Probabilistic Ordinary Differential Equation Solvers — Theory and Applications. PhD thesis. Eberhard Karls Universität Tübingen, Germany, 2018.

[313] M. Mahsereci. Probabilistic Approaches to Stochastic Optimization. PhD thesis. Eberhard Karls Universität Tübingen, Germany, 2018 (cited on page 249).

[314] P. Geiger. Causal models for decision making via integrative inference. PhD thesis. University of Stuttgart, Germany, 2017.

[315] T. Fomina. Brain-Computer Interfaces for patients with Amyotrophic Lateral Sclerosis. PhD thesis. Eberhard Karls Universität Tübingen, Germany, 2017.

[316] E. D. Klenske. Nonparametric Disturbance Correction and Nonlinear Dual Control. PhD thesis. ETH Zurich, 2017.

[317] M. Kiefel. Tractable Structured Prediction using the Permutohedral Lattice. PhD thesis. ETH Zurich, 2016.

[318] R. Köhler. Advances in computational imaging: Benchmarking Deblurring Algorithms, Deep Neural Inpainting, Depth Estimation from Light Fields. PhD thesis. Eberhard Karls Universität Tübingen, Germany, 2016.

### Master Theses

[319] T. Lechner. Domain Adaptation Under Causal Assumptions. MA thesis. Eberhard Karls Universität Tübingen, Germany, 2018.

[320] R. Suter. A Causal Perspective on Deep Representation Learning. MA thesis. ETH Zurich, 2018.

[321] S. Zabel. Improving Tissue Differentiation based on Optical Emission Spectroscopy for Guided Electrosurgical Tumor Resection with Machine Learning. MA thesis. Eberhard Karls Universität Tübingen, Germany, 2018.

[322] T. Gebhard. On the Applicability of Machine Learning to Aid the Search for Gravitational Waves at the LIGO Experiment. MA thesis. Karlsruhe Institute of Technology, Germany, 2018.

[323] K. Sücker. Learning Optimal Configurations for Modeling Frowning by Transcranial Electrical Stimulation. MA thesis. Graduate School of Neural Information Processing, Eberhard Karls Universität Tübingen, Germany, 2017.

[324] T. Emde. Development and Evaluation of a Portable BCI System for Remote Data Acquisition. MA thesis. Graduate School of Neural Information Processing, Eberhard Karls Universität Tübingen, Germany, 2017.

## 21.3 Perceiving Systems Department

### 21.3.1 Journal Articles

#### 2019

[325] S. Pujades, B. Mohler, A. Thaler, J. Tesch, N. Mahmood, N. Hesse, H. H. Bülthoff, M. J. Black. The Virtual Caliper: Rapid Creation of Metrically Accurate Avatars from 3D Measurements. *IEEE Transactions on Visualization and Computer Graphics*, 2019.

#### 2018

[326] Y. Zhang, S. Tang, K. Muandet, C. Jarvers, H. Neumann. Local Temporal Bilinear Pooling for Fine-grained Action Parsing. *arXiv preprint arXiv:1812.01922*, 2018.

[327] A. Thaler, I. Piryankova, J. K. Stefanucci, S. Pujades, S. de la Rosa, S. Streuber, J. Romero, M. J. Black, B. J. Mohler. Visual Perception and Evaluation of Photo-Realistic Self-Avatars From 3D Body Scans in Males and Females. *Frontiers in ICT* **5**: 1–14, 2018.

[328] M. Keuper, S. Tang, B. Andres, T. Brox, B. Schiele. Motion Segmentation & Multiple Object Tracking by Correlation Co-Clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018 (cited on page 80).

[329] Y. Hu, C. J. Parde, M. Q. Hill, N. Mahmood, A. J. O'Toole. First Impressions of Personality Traits From Body Shapes. *Psychological Science* **29** (12): 1969–1983, 2018.

[330] Y. Huang, M. Kaufmann, E. Aksan, M. J. Black, O. Hilliges, G. Pons-Moll. Deep Inertial Poser: Learning to Reconstruct Human Pose from Sparse Inertial Measurements in Real Time. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)* **37**. Two first authors contributed equally: 185:1–185:15, 2018 (cited on pages 68, 83).

[331] M. A. Borno, L. Righetti, M. J. Black, S. L. Delp, E. Fiume, J. Romero. Robust Physics-based Motion Retargeting with Realistic Body Shapes. *Computer Graphics Forum* **37**: 6:1–12, 2018 (cited on pages 77, 84).

[332] E. Price, G. Lawless, R. Ludwig, I. Martinovic, H. H. Buelthoff, M. J. Black, A. Ahmad. Deep Neural Network-based Cooperative Visual Tracking through Multiple Micro Aerial Vehicles. *IEEE Robotics and Automation Letters* **3** (4). Also accepted and presented in the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).: 3193–3200, 2018 (cited on pages 68, 81).

[333] A. Ranjan, V. Jampani, K. Kim, D. Sun, J. Wulff, M. J. Black. Adversarial Collaboration: Joint Unsupervised Learning of Depth, Camera Motion, Optical Flow and Motion Segmentation, 2018 (cited on pages 66, 86, 90).

[334] A. Thaler, M. N. Geuss, S. C. Mölbert, K. E. Giel, S. Streuber, J. Romero, M. J. Black, B. J. Mohler. Body size estimation of self and others in females varying in BMI. *PLoS ONE* **13** (2), 2018 (cited on pages 67, 93).

[335] S. C. Mölbert, A. Thaler, B. J. Mohler, S. Streuber, J. Romero, M. J. Black, S. Zipfel, H.-O. Karnath, K. E. Giel. Assessing body image in anorexia nervosa using biometric self-avatars in virtual reality: Attitudinal components rather than visual body size estimation are distorted. *Psychological Medicine* **48** (4): 642–653, 2018 (cited on pages 67, 93).

#### 2017

[336] S. C. Mölbert, A. Thaler, S. Streuber, M. J. Black, H. Karnath, S. Zipfel, B. Mohler, K. E. Giel. Investigating Body Image Disturbance in Anorexia Nervosa Using Novel Biometric Figure Rating Scales: A Pilot Study. *European Eating Disorders Review* **25** (6): 607–612, 2017 (cited on pages 67, 93).

[337] J. Romero, D. Tzionas, M. J. Black. Embodied Hands: Modeling and Capturing Hands and Bodies Together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*. 245:1–245:17 **36** (6): 245:1–245:17, 2017 (cited on pages 64, 68, 73, 75).

[338] T. Li, T. Bolkart, M. J. Black, H. Li, J. Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics* **36** (6). Two first authors contributed equally: 194:1–194:17, 2017 (cited on pages 64, 68, 73, 74).

[339] A. Ahmad, G. Lawless, P. Lima. An Online Scalable Approach to Unified Multirobot Cooperative Localization and Object Tracking. *IEEE Transactions on Robotics (T-RO)* **33**: 1184–1199, 2017 (cited on pages 68, 82).

[340] M. Mahsereci, L. Balles, C. Lassner, P. Hennig. Early Stopping Without a Validation Set. *arXiv preprint arXiv:1703.09580*, 2017.

[341] G. Pons-Moll, S. Pujades, S. Hu, M. Black. ClothCap: Seamless 4D Clothing Capture and Retargeting. *ACM Transactions on Graphics, (Proc. SIGGRAPH)* **36** (4). Two first authors contributed equally: 73:1–73:15, 2017 (cited on pages 66, 73, 76).

[342] M. Kim, G. Pons-Moll, S. Pujades, S. Bang, J. Kim, M. J. Black, S.-H. Lee. Data-Driven Physics for Human Soft Tissue Animation. *ACM Transactions on Graphics, (Proc. SIGGRAPH)* **36** (4): 54:1–54:12, 2017 (cited on pages 65, 77).

[343] R. Gadde, V. Jampani, R. Marlet, P. Gehler. Efficient 2D and 3D Facade Segmentation using Auto-Context. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017 (cited on page 88).

[344] T. von Marcard, B. Rosenhahn, M. Black, G. Pons-Moll. Sparse Inertial Poser: Automatic 3D Human Pose Estimation from Sparse IMUs. *Computer Graphics Forum 36(2), Proceedings of the 38th Annual Conference of the European Association for Computer Graphics (Eurographics)*: 349–360, 2017 (cited on page 83).

### 2016

[345] M. Q. Hill, S. Streuber, C. A. Hahn, M. J. Black, A. J. O'Toole. Creating body shapes from verbal descriptions by linking similarity spaces. *Psychological Science* **27** (11): 1486–1497, 2016 (cited on page 93).

[346] S. Y. Yeo, J. Romero, M. Loper, J. Machann, M. Black. Shape estimation of subcutaneous adipose tissue using an articulated statistical shape model. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*: 1–8, 2016. eprint: http://dx.doi.org/10.1080/21681163.2016.1163508 (cited on pages 67, 94).

[347] S. Streuber, M. A. Quiros-Ramirez, M. Q. Hill, C. A. Hahn, S. Zuffi, A. O'Toole, M. J. Black. Body Talk: Crowdshaping Realistic 3D Avatars with Words. *ACM Trans. Graph. (Proc. SIGGRAPH)* **35** (4): 54:1–54:14, 2016 (cited on page 93).

[348] T. Feix, J. Romero, H.-B. Schmiedmayer, A. Dollar, D. Kragic. The GRASP Taxonomy of Human Grasp Types. *Human-Machine Systems, IEEE Transactions on* **46** (1): 66–77, 2016 (cited on page 75).

[349] D. Tzionas, L. Ballan, A. Srikantha, P. Aponte, M. Pollefeys, J. Gall. Capturing Hands in Action using Discriminative Salient Points and Physics Simulation. *International Journal of Computer Vision (IJCV)* **118** (2): 172–193, 2016 (cited on page 75).

[350] T. von Marcard, G. Pons-Moll, B. Rosenhahn. Human Pose Estimation from Video and IMUs. *Transactions on Pattern Analysis and Machine Intelligence PAMI* **38** (8): 1533–1547, Jan. 2016.

[351] M. A. Brubaker, A. Geiger, R. Urtasun. Map-Based Probabilistic Visual Self-Localization. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2016 (cited on page 179).

## 21.3.2 Conference Papers

### 2019

[352] P. Ghosh, A. Losalka, M. J. Black. Resisting Adversarial Attacks using Gaussian Mixture Variational Autoencoders. In *Proc. AAAI*, 2019.

## 2018

[353] L. Ma, S. Tang, M. J. Black, L. V. Gool. Customized Multi-Person Tracker. In *Computer Vision – ACCV 2018*, 2018 (cited on page 80).

[354] L. Sevilla-Lara, Y. Liao, F. Güney, V. Jampani, A. Geiger, M. J. Black. On the Integration of Optical Flow and Action Recognition. In *German Conference on Pattern Recognition (GCPR)*. Vol. LNCS 11269, pages 281–297, Oct. 2018 (cited on pages 86, 87, 179).

[355] J. Wulff, M. J. Black. Temporal Interpolation as an Unsupervised Pretraining Task for Optical Flow Estimation. In *German Conference on Pattern Recognition (GCPR)*. Vol. LNCS 11269, pages 567–582, Oct. 2018 (cited on page 86).

[356] M. Omran, C. Lassner, G. Pons-Moll, P. V. Gehler, B. Schiele. Neural Body Fitting: Unifying Deep Learning and Model-Based Human Pose and Shape Estimation. In *3DV*, 2018 (cited on pages 64, 79).

[357] Y. Zhang, S. Tang, H. Sun, H. Neumann. Human Motion Parsing by Hierarchical Dynamic Clustering. In *Proceedings of the British Machine Vision Conference (BMVC)*, page 269, 2018 (cited on page 85).

[358] T. von Marcard, R. Henschel, M. J. Black, B. Rosenhahn, G. Pons-Moll. Recovering Accurate 3D Human Pose in The Wild Using IMUs and a Moving Camera. In *European Conference on Computer Vision (ECCV)*. Vol. Lecture Notes in Computer Science, vol 11214, pages 614–631, 2018 (cited on pages 68, 83).

[359] J. Janai, F. Güney, A. Ranjan, M. J. Black, A. Geiger. Unsupervised Learning of Multi-Frame Optical Flow with Occlusions. In *European Conference on Computer Vision (ECCV)*. Vol. Lecture Notes in Computer Science, vol 11220, pages 713–731, 2018 (cited on pages 86, 179).

[360] A. Ranjan, T. Bolkart, S. Sanyal, M. J. Black. Generating 3D Faces using Convolutional Mesh Autoencoders. In *European Conference on Computer Vision (ECCV)*. Vol. Lecture Notes in Computer Science, vol 11207, pages 725–741, 2018 (cited on pages 68, 74).

[361] S. Prokudin, P. Gehler, S. Nowozin. Deep Directional Statistics: Pose Estimation with Uncertainty Quantification. In *European Conference on Computer Vision (ECCV)*, 2018.

[362] A. Ranjan, J. Romero, M. J. Black. Learning Human Optical Flow. In *29th British Machine Vision Conference*, 2018 (cited on pages 64, 87, 92).

[363] R. Tallamraju, S. Rajappa, M. J. Black, K. Karlapalem, A. Ahmad. Decentralized MPC based Obstacle Avoidance for Multi-Robot Target Tracking Scenarios. In *2018 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 1–8, 2018 (cited on page 82).

[364] N. Hesse, S. Pujades, J. Romero, M. J. Black, C. Bodensteiner, M. Arens, U. G. Hofmann, U. Tacke, M. Hadders-Algra, R. Weinberger, W. Muller-Felber, A. S. Schroeder. Learning an Infant Body Model from RGB-D Data for Accurate Full Body Motion Analysis. In *Int. Conf. on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2018 (cited on pages 67, 68, 94).

[365] Y. Suh, J. Wang, S. Tang, T. Mei, K. M. Lee. Part-Aligned Bilinear Representations for Person Re-identification. In *European Conference on Computer Vision (ECCV)*. Vol. 11218, pages 418–437, 2018 (cited on page 80).

[366] A. Kanazawa, M. J. Black, D. W. Jacobs, J. Malik. End-to-end Recovery of Human Shape and Pose. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 (cited on pages 64, 79).

[367] S. Zuffi, A. Kanazawa, M. J. Black. Lions and Tigers and Bears: Capturing Non-Rigid, 3D, Articulated Shape from Images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 (cited on page 68).

## 2017

[368] S. Kenny, N. Mahmood, C. Honda, M. J. Black, N. F. Troje. Effects of animation retargeting on perceived action outcomes. In *Proceedings of the ACM Symposium on Applied Perception (SAP'17)*, 2:1–2:7, 2017.

[369] Y. Huang, F. Bogo, C. Lassner, A. Kanazawa, P. V. Gehler, J. Romero, I. Akhter, M. J. Black. Towards Accurate Marker-less Human Shape and Pose Estimation over Time. In *International Conference on 3D Vision (3DV)*, pages 421–430, 2017 (cited on page 79).

[370] R. Gadde, V. Jampani, P. V. Gehler. Semantic Video CNNs through Representation Warping. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, 2017 (cited on page 91).

[371] S. Tang, M. Andriluka, B. Andres, B. Schiele. Multiple People Tracking by Lifted Multicut and Person Re-identification. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3701–3710, 2017 (cited on page 80).

[372] J. Wulff, L. Sevilla-Lara, M. J. Black. Optical Flow in Mostly Rigid Scenes. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 6911–6920, 2017 (cited on pages 66, 90).

[373] J. Martinez, M. J. Black, J. Romero. On human motion prediction using recurrent neural networks. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on page 84).

[374] J. Martinez, R. Hossain, J. Romero, J. J. Little. A simple yet effective baseline for 3d human pose estimation. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, Oct. 2017 (cited on page 84).

[375] A. Ranjan, M. Black. Optical Flow Estimation using a Spatial Pyramid Network. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on pages 86, 87).

[376] C. Lassner, G. Pons-Moll, P. V. Gehler. A Generative Model of People in Clothing. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, 2017 (cited on page 76).

[377] C. Lassner, J. Romero, M. Kiefel, F. Bogo, M. J. Black, P. V. Gehler. Unite the People: Closing the Loop Between 3D and 2D Human Representations. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on pages 64, 68, 79).

[378] G. Varol, J. Romero, X. Martin, N. Mahmood, M. J. Black, I. Laptev, C. Schmid. Learning from Synthetic Humans. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on pages 64, 68, 87, 92).

[379] J. Bütepage, M. Black, D. Kragic, H. Kjellström. Deep representation learning for human motion prediction and classification. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on page 84).

[380] J. Janai, F. Güney, J. Wulff, M. Black, A. Geiger. Slow Flow: Exploiting High-Speed Cameras for Accurate and Diverse Optical Flow Reference Data. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 1406–1416, 2017 (cited on pages 68, 179).

[381] F. Bogo, J. Romero, G. Pons-Moll, M. J. Black. Dynamic FAUST: Registering Human Bodies in Motion. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on pages 68, 73).

[382] A. O. Ulusoy, M. J. Black, A. Geiger. Semantic Multi-view Stereo: Jointly Estimating Objects and Voxels. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on pages 89, 178).

[383] S. Zuffi, A. Kanazawa, D. Jacobs, M. J. Black. 3D Menagerie: Modeling the 3D Shape and Pose of Animals. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 5524–5532, 2017 (cited on pages 68, 78).

[384] A. Rohrbach, M. Rohrbach, S. Tang, S. J. Oh, B. Schiele. Generating Descriptions with Grounded and Co-Referenced People. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4196–4206, 2017 (cited on page 85).

[385] E. Levinkov, J. Uhrig, S. Tang, M. Omran, E. Insafutdinov, A. Kirillov, C. Rother, T. Brox, B. Schiele, B. Andres. Joint Graph Decomposition and Node Labeling by Local Search. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1904–1912, 2017.

[386] T. Nestmeyer, P. V. Gehler. Reflectance Adaptive Filtering Improves Intrinsic Image Estimation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1771–1780, 2017.

[387] V. Jampani, R. Gadde, P. V. Gehler. Video Propagation Networks. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on page 91).

[388] C. Zhang, S. Pujades, M. Black, G. Pons-Moll. Detailed, accurate, human shape estimation from clothed 3D scan sequences. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Spotlight, 2017 (cited on pages 68, 73, 76).

[389] G. Riegler, O. Ulusoy, A. Geiger. OctNet: Learning Deep 3D Representations at High Resolutions. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on page 178).

[390] L. Balles, J. Romero, P. Hennig. Coupling Adaptive Batch Sizes with Learning Rates. In *Proceedings Conference on Uncertainty in Artificial Intelligence (UAI) 2017*, pages 410–419, 2017 (cited on page 249).

## 2016

[48] C. Lassner, D. Kappler, M. Kiefel, P. Gehler. Barrista - Caffe Well-Served. In *ACM Multimedia Open Source Software Competition*, 2016.

[53] R. Gadde, V. Jampani, M. Kiefel, D. Kappler, P. Gehler. Superpixel Convolutional Networks using Bilateral Inceptions. In *European Conference on Computer Vision (ECCV)*. Lecture Notes in Computer Science, 2016 (cited on page 88).

[391] F. Güney, A. Geiger. Deep Discrete Flow. In *Asian Conference on Computer Vision (ACCV)*, 2016 (cited on page 179).

[392] D. Tzionas, J. Gall. Reconstructing Articulated Rigged Models from RGB-D Videos. In *European Conference on Computer Vision Workshops 2016 (ECCVW'16) - Workshop on Recovering 6D Object Pose (R6D'16)*, pages 620–633, 2016 (cited on page 75).

[393] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, M. J. Black. Keep it SMPL: Automatic Estimation of 3D Human Pose and Shape from a Single Image. In *Computer Vision – ECCV 2016*. Lecture Notes in Computer Science, pages 561–578, 2016 (cited on pages 64, 68, 79).

[394] J. Xie, M. Kiefel, M.-T. Sun, A. Geiger. Semantic Instance Annotation of Street Scenes by 3D to 2D Label Transfer. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016 (cited on pages 88, 179).

[395] A. O. Ulusoy, M. J. Black, A. Geiger. Patches, Planes and Probabilities: A Non-local Prior for Volumetric 3D Reconstruction. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016 (cited on pages 89, 178).

[396] Y.-H. Tsai, M.-H. Yang, M. J. Black. Video segmentation via object flow. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016 (cited on page 91).

[397] H. Fu, C. Wang, D. Tao, M. J. Black. Occlusion boundary detection via deep exploration of context. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016 (cited on page 86).

[398] L. Pishchulin, E. Insafutdinov, S. Tang, B. Andres, M. Andriluka, P. Gehler, B. Schiele. DeepCut: Joint Subset Partition and Labeling for Multi Person Pose Estimation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4929–4937, 2016 (cited on page 80).

[399] V. Jampani, M. Kiefel, P. V. Gehler. Learning Sparse High Dimensional Filters: Image Filtering, Dense CRFs and Bilateral Neural Networks. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 4452–4461, 2016 (cited on page 88).

[400] L. Sevilla-Lara, D. Sun, V. Jampani, M. J. Black. Optical Flow with Semantic Segmentation and Localized Layers. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3889–3898, 2016 (cited on pages 66, 90).

[401] R. Fleming, B. Mohler, J. Romero, M. J. Black, M. Breidt. Appealing female avatars from 3D body scans: Perceptual effects of stylization. In *11th Int. Conf. on Computer Graphics Theory and Applications (GRAPP)*, 2016.

### 21.3.3 Patents

#### 2018

[402] M. Black, D. Hirshberg, M. Loper, E. Rachlin, A. Weiss. Co-Registration – Simultaneous Alignment and Modeling of Articulated 3D Shapes. U.S. Patent 9,898,848. Feb. 2018.

[403] M. J. Black, A. Balan, A. Weiss, L. Sigal, M. Loper, T. St Clair. Method and Apparatus for Estimating Body Shape. U.S. Patent 10,002,460. June 2018.

#### 2017

[404] M. Loper, N. Mahmood, M. Black. Method for providing a three dimensional body model. U.S. Patent 9,710,964 B2. July 2017.

[405] M. J. Black, P. Guan. System and method for simulating realistic clothing. U.S. Patent 9,679,409 B2. June 2017.

[406] M. J. Black, O. Freifeld, A. Weiss, M. Loper, P. Guan. Parameterized Model of 2D Articulated Human Shape. U.S. Patent 9,761,060. Sept. 2017.

### 21.3.4 Book Chapters

#### 2017

[407] T. Nestmeyer, P. Robuffo Giordano, H. H. Bülthoff, A. Franchi. Decentralized Simultaneous Multi-target Exploration using a Connected Network of Multiple Robots. In. Autonomous Robots, pages 989–1011, 2017.

[408] R. Fleming, B. J. Mohler, J. Romero, M. J. Black, M. Breidt. Appealing Avatars from 3D Body Scans: Perceptual Effects of Stylization. In. *Computer Vision, Imaging and Computer Graphics Theory and Applications: 11th International Joint Conference, VISIGRAPP 2016, Rome, Italy, February 27 – 29, 2016, Revised Selected Papers*. Springer International Publishing, pages 175–196, 2017.

[409] S. Prokudin, D. Kappler, S. Nowozin, P. Gehler. Learning to Filter Object Detections. In. *Pattern Recognition: 39th German Conference, GCPR 2017, Basel, Switzerland, September 12–15, 2017, Proceedings*. Springer International Publishing, Cham, pages 52–62, 2017.

### 21.3.5 Theses

#### PhD Theses

[410] J. Wulff. Model-based Optical Flow: Layers, Learning, and Geometry. PhD thesis. Tuebingen University, 2018.

[411] V. Jampani. Learning Inference Models for Computer Vision. PhD thesis. MPI for Intelligent Systems and University of Tübingen, 2017.

[412] M. M. Loper. Human Shape Estimation using Statistical Body Models. University of Tübingen, 2017.

[413] D. Tzionas. Capturing Hand-Object Interaction and Reconstruction of Manipulated Objects. PhD thesis. University of Bonn, 2017 (cited on page 75).

[414] A. Lehrmann. Non-parametric Models for Structured Data and Applications to Human Bodies and Natural Scenes. PhD thesis. ETH Zurich, 2016.

## 21.4 Physical Intelligence Department

### 21.4.1 Books

#### 2017

[415] M. Sitti. Mobile Microrobotics. The MIT Press, Cambridge, MA, 2017.

### 21.4.2 Journal Articles

#### 2018

[416] H. Ceylan, I. C. Yasa, O. Yasa, A. F. Tabak, J. Giltinan, M. Sitti. 3D-Printed Biodegradable Microswimmer for Theranostic Cargo Delivery and Release. *ACS Nano*, in press.

[417] V. Singh, V. Kumar, S. Kashyap, A. V. Singh, V. Kishore, M. Sitti, P. S. Saxena, A. Srivastava. Graphene oxide synergistically enhances antibiotic efficacy in Vancomycin resistance Staphylococcus aureus. *ACS Applied Bio Materials*, in press.

[418] P. Erkoc, I. C. Yasa, H. Ceylan, O. Yasa, Y. Alapan, M. Sitti. Mobile microrobots for active therapeutic delivery. *Advanced Therapeutics*, in press.

[419] Y. Alapan, O. Yasa, B. Yigit, I. C. Yasa, P. Erkoc, M. Sitti. Microrobotics and Microorganisms: Biohybrid Autonomous Cellular Robots. *Annual Review of Control, Robotics, and Autonomous Systems*, in press.

[420] A. V. Singh, P. Laux, A. Luch, C. Sudrik, S. Wiehr, A.-M. Wild, G. Santamauro, J. Bill, M. Sitti. Review of emerging concepts in nanotoxicology: opportunities and challenges for safer nanomaterial design. *Toxicology Mechanisms and Methods*, in press.

[421] K. V. Sreekanth, S. Sreejith, Y. Alapan, M. Sitti, C. T. Lim, R. Singh. Microfluidics Integrated Lithography-Free Nanophotonic Biosensor for the Detection of Small Molecules. *Advanced Optical Materials*, in press.

[422] J. Giltinan, M. Sitti. Simultaneous Six-Degree-of-Freedom Control of a Single-Body Magnetic Microrobot. *IEEE Robotics and Automation Letters*, in press.

[423] L. Koens, W. Wang, M. Sitti, E. Lauga. The near and far of a pair of magnetic capillary disks. *Soft Matter*, in press.

[424] A. V. Singh, T. Jahnke, S. Wang, Y. Xiao, Y. Alapan, S. Kharratian, M. C. Onbasli, K. Kozielski, H. David, G. Richter, J. Bill, P. Laux, A. Luch, M. Sitti. Anisotropic Gold Nanostructures: Optimization via in Silico Modeling for Hyperthermia. *ACS Applied Nano Materials* **1** (11): 6205–6216, 2018.

[425] G. Santomauro, A. Singh, B. W. Park, M. Mohammadrahimi, P. Erkoc, E. Goering, G. Schütz, M. Sitti, J. Bill. Incorporation of Terbium into a Microalga Leads to Magnetotactic Swimmers. *Advanced Biosystems* **2** (12): 1800039, 2018.

[426] I. S. M. Khalil, A. F. Tabak, M. A. Seif, A. Klingner, M. Sitti. Controllable switching between planar and helical flagellar swimming of a soft robotic sperm. *PloS One* **13** (11): e0206456, 2018.

[427] D. Ipparthi, T. A. G. Hageman, N. Cambier, M. Sitti, M. Dorigo, L. Abelmann, M. Mastrangeli. Kinetics of orbitally shaken particles constrained to two dimensions. *Physical Review E* **98** (4): 042137, 2018.

[428] U. T. Sanli, H. Ceylan, I. Bykova, M. Weigand, M. Sitti, G. Schütz, K. Keskinbora. 3D Nanoprinted Plastic Kinoform X-Ray Optics. *Advanced Materials* **30** (36): 1802503, 2018.

[429] A. V. Singh, T. Gharat, M. Batuwangala, B.-W. Park, T. Endlein, M. Sitti. Three-dimensional patterning in biomedicine: Importance and applications in neuropharmacology. *Journal of Biomedical Materials Research Part B: Applied Biomaterials* **106** (3): 1369–1382, 2018.

[430] A. V. Singh, Y. Alapan, T. Jahnke, P. Laux, A. Luch, A. Aghakhani, S. Kharratian, M. C. Onbasli, J. Bill, M. Sitti. Seed-mediated synthesis of plasmonic gold nanoribbons using cancer cells for hyperthermia applications. *Journal of Materials Chemistry B* **6** (46): 7573–7581, 2018.

[431] F. Munoz, G. Alici, H. Zhou, W. Li, M. Sitti. Analysis of Magnetic Interaction in Remotely Controlled Magnetic Devices and Its Application to a Capsule Robot for Drug Delivery. *IEEE Transactions on Magnetics* **23** (1): 298–310, 2018.

[432] I. S. M. Khalil, A. F. Tabak, Y. Hamed, M. E. Mitwally, M. Tawakol, A. Klingner, M. Sitti. Swimming Back and Forth Using Planar Flagellar Propulsion at Low Reynolds Numbers. *Advanced Science* **5** (2): 1700461, 2018. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/advs.201700461.

[433] D. Son, X. Dong, M. Sitti. A Simultaneous Calibration Method for Magnetic Robot Localization and Actuation Systems. *IEEE Transactions on Robotics*, 2018, in press.

[434] M. A. Woodward, M. Sitti. Morphological intelligence counters foot slipping in the desert locust and dynamic robots. *Proceedings of the National Academy of Sciences* **115**: E8358–E8367, 2018.

[435] P. Erkoc, I. C. Yasa, H. Ceylan, O. Yasa, Y. Alapan, M. Sitti. Mobile Microrobots for Active Therapeutic Delivery. *Advanced Therapeutics*: 1800064, 2018.

[436] U. Bozuyuk, O. Yasa, I. C. Yasa, H. Ceylan, S. Kizilel, M. Sitti. Light-Triggered Drug Release from 3D-Printed Magnetic Chitosan Microswimmers. *ACS Nano* **12**: 9617–9625, 2018.

[437] O. Yasa, P. Erkoc, Y. Alapan, M. Sitti. Microalga-Powered Microswimmers toward Active Cargo Delivery. *Advanced Materials*: 1804130, 2018.

[438] B. Yigit, Y. Alapan, M. Sitti. Programmable collective behavior in dynamically self-assembled mobile microrobotic swarms. *Advance Science*, 2018, in press.

[439] W. Wang, J. V. I. Timonen, A. Carlson, D.-M. Drotlef, C. T. Zhang, S. Kolle, A. Grinthal, T.-S. Wong, B. Hatton, S. H. Kang, S. Kennedy, J. Chi, R. T. Blough, M. Sitti, L. Mahadevan, J. Aizenberg. Multifunctional ferrofluid-infused surfaces with reconfigurable multiscale topography. *Nature* **559** (7712): 77–82, 2018.

[440] O. Schauer, B. Mostaghaci, R. Colin, D. Hürtgen, D. Kraus, M. Sitti, V. Sourjik. Motility and chemotaxis of bacteria-driven microswimmers fabricated using antigen 43-mediated biotin display. *Scientific Reports* **8** (1): 9801, 2018.

[441] T. Endlein, M. Sitti. Innate turning preference of leaf-cutting ants in the absence of external orientation cues. *Journal of Experimental Biology*, 2018. eprint: http://jeb.biologists.org/content/early/2018/06/01/jeb.177006.full.pdf.

[442] M. Amjadi, M. Sitti. Self-Sensing Paper Actuators Based on Graphite–Carbon Nanotube Hybrid Films. *Advanced Science*: 1800239, 2018. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/advs.201800239.

[443] M. Sitti. Miniature soft robots – road to the clinic. *Nature Reviews Materials* **3**: 74–75, 2018.

[444] Y. Alapan, O. Yasa, O. Schauer, J. Giltinan, A. F. Tabak, V. Sourjik, M. Sitti. Soft erythrocyte-based bacterial microswimmers for cargo delivery. *Science Robotics* **3** (17): eaar4423, 2018. eprint: http://robotics.sciencemag.org/content/3/17/eaar4423.full.pdf.

[445] I. S. M. Khalil, D. Mahdy, A. E. Sharkawy, R. R. Moustafa, A. F. Tabak, M. E. Mitwally, S. Hesham, N. Hamdi, A. Klingner, A. Mohamed, M. Sitti. Mechanical Rubbing of Blood Clots Using Helical Robots Under Ultrasound Guidance. *IEEE Robotics and Automation Letters* **3** (2): 1112–1119, 2018.

[446] L. Vannozzi, I. C. Yasa, H. Ceylan, A. Menciassi, L. Ricotti, M. Sitti. Self-Folded Hydrogel Tubes for Implantable Muscular Tissue Scaffolds. *Macromolecular Bioscience* **18** (4): 1700377, 2018. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/mabi.201700377.

[447] M. Yunusa, G. J. Amador, D.-M. Drotlef, M. Sitti. Wrinkling Instability and Adhesion of a Highly Bendable Gallium Oxide Nanofilm Encapsulating a Liquid-Gallium Droplet. *Nano Letters* **18** (4): 2498–2504, 2018. eprint: https://doi.org/10.1021/acs.nanolett.8b00164.

[448] A. V. Singh, T. Jahnke, V. Kishore, B.-W. Park, M. Batuwangala, J. Bill, M. Sitti. Cancer cells biomineralize ionic gold into nanoparticles-microplates via secreting defense proteins with specific gold-binding peptides. *Acta Biomaterialia* **71**: 61–71, 2018.

[449] S. Sheykhansari, K. Kozielski, J. Bill, M. Sitti, D. Gemmati, P. Zamboni, A. V. Singh. Redox metals homeostasis in multiple sclerosis and amyotrophic lateral sclerosis: a review. *Cell Death & Disease* **9** (3): 348, 2018.

[450] I. S. M. Khalil, A. F. Tabak, Y. Hamed, M. Tawakol, A. Klingner, N. E. Gohary, B. Mizaikoff, M. Sitti. Independent Actuation of Two-Tailed Microrobots. *IEEE Robotics and Automation Letters* **3** (3): 1703–1710, 2018.

[451] W. Hu, G. Z. Lum, M. Mastrangeli, M. Sitti. Small-scale soft-bodied robot with multimodal locomotion. *Nature* **554**: 81–85, 2018.

[452] M. Turan, Y. Y. Pilavci, I. Ganiyusufoglu, H. Araujo, E. Konukoglu, M. Sitti. Sparse-then-dense alignment-based 3D map reconstruction method for endoscopic capsule robots. *Machine Vision and Applications* **29** (2): 345–359, 2018.

[453] M. Turan, Y. Almalioglu, H. Araujo, E. Konukoglu, M. Sitti. Deep EndoVO: A recurrent convolutional neural network (RCNN) based visual odometry approach for endoscopic capsule robots. *Neurocomputing* **275**: 1861–1870, 2018.

[454] M. Amjadi, S. Sheykhansari, B. J. Nelson, M. Sitti. Recent Advances in Wearable Transdermal Delivery Systems. *Advanced Materials* **30** (7): 1704530, 2018. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.201704530.

[455] V. Sridhar, B. Park, M. Sitti. Light-Driven Janus Hollow Mesoporous TiO2–Au Microswimmers. *Advanced Functional Materials* **28** (5): 1704902, 2018. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adfm.201704902.

[456] M. A. Woodward, M. Sitti. Universal Custom Complex Magnetic Spring Design Methodology. *IEEE Transactions on Magnetics* **54** (1): 1–13, 2018.

## 2017

[457] L. Hines, K. Petersen, G. Z. Lum, M. Sitti. Soft Actuators for Small-Scale Robotics. *Advanced Materials* **29** (13): 1603483, 2017.

[458] H. Ceylan, I. C. Yasa, M. Sitti. 3D Chemical Patterning of Micromaterials for Encoded Functionality. *Advanced Materials* **29** (9): 1605072, 2017.

[459] T. Endlein, A. Ji, S. Yuan, I. Hill, H. Wang, W. J. P. Barnes, Z. Dai, M. Sitti. The use of clamping grips and friction pads by tree frogs for climbing curved surfaces. In *Proc. R. Soc. B*. Vol. 284. 1849, page 20162867, 2017.

[460] O. Yasa, O. Uysal, M. S. Ekiz, M. O. Guler, A. B. Tekinay. Presentation of functional groups on self-assembled supramolecular peptide nanofibers mimicking glycosaminoglycans for directed mesenchymal stem cell differentiation. *J. Mater. Chem. B* **5**: 4890–4900, 2017.

[461] M. J. Raymond, P. Ray, G. Kaur, M. Fredericks, A. V. Singh, L. Q. Wan. Multiaxial Polarity Determines Individual Cellular and Nuclear Chirality. *Cellular and Molecular Bioengineering* **10** (1): 63–74, 2017.

[462] M. Mastrangeli, Q. Zhou, V. Sariola, P. Lambert. Surface tension-driven self-alignment. *Soft Matter* **13**: 304–327, 2017.

[463] M. Kim, Y. Alapan, A. Adhikari, J. A. Little, U. A. Gurkan. Hypoxia-enhanced adhesion of red blood cells in microscale flow. *Microcirculation* **24** (5): e12374, 2017. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/micc.12374.

[464] D. Ipparthi, M. Mastrangeli, A. Winslow. Dipole codes attractively encode glue functions. *Theoretical Computer Science* **671**. Computational Self-Assembly: 19–25, 2017.

[465] D. Ipparthi, A. Winslow, M. Sitti, M. Dorigo, M. Mastrangeli. Yield prediction in parallel homogeneous assembly. *Soft Matter* **13**: 7595–7608, 2017.

[466] V. Flauraud, M. Mastrangeli, G. Bernasconi, J. Butet, D. Alexander, E. Shahrabi, O. Martin, J. Brugger. Nanoscale topographical control of capillary assembly of nanoparticles. *Scientific Reports* **12**: 73–80, 2017.

[467] V. Flauraud, G. D. Bernasconi, J. Butet, M. Mastrangeli, D. T. L. Alexander, O. J. F. Martin, J. Brugger. Mode Evolution in Strongly Coupled Plasmonic Dolmens Fabricated by Templated Assembly. *ACS Photonics* **4** (7): 1661–1668, 2017. eprint: https://doi.org/10.1021/acsphotonics.6b01026.

[468] G. J. Amador, D. L. Hu. Sticky Solution Provides Grip for the First Robotic Pollinator. *Chem* **2** (2): 162–164, 2017.

[469] M. Turan, J. Shabbir, H. Araujo, E. Konukoglu, M. Sitti. A deep learning based fusion of RGB camera information and magnetic localization information for endoscopic capsule robots. *International Journal of Intelligent Robotics and Applications* **1** (4): 442–450, 2017.

[470] L. Ricotti, B. Trimmer, A. W. Feinberg, R. Raman, K. K. Parker, R. Bashir, M. Sitti, S. Martel, P. Dario, A. Menciassi. Biohybrid actuators for robotics: A review of devices actuated by living cells. *Science Robotics* **2** (12), 2017. eprint: http://robotics.sciencemag.org/content/2/12/eaaq0495.full.pdf.

[471] B.-W. Park, J. Zhuang, O. Yasa, M. Sitti. Multifunctional Bacteria-Driven Microswimmers for Targeted Active Drug Delivery. *ACS Nano* **11** (9). PMID: 28873304: 8910–8923, 2017. eprint: http://dx.doi.org/10.1021/acsnano.7b03207.

[472] A. V. Singh, Z. Hosseinidoust, B.-W. Park, O. Yasa, M. Sitti. Microemulsion-Based Soft Bacteria-Driven Microswimmers for Active Cargo Delivery. *ACS Nano*. PMID: 28858477: null, 2017. eprint: http://dx.doi.org/10.1021/acsnano.7b02082.

[473] M. M. Stanton, B.-W. Park, D. Vilela, K. Bente, D. Faivre, M. Sitti, S. Sánchez. Magnetotactic Bacteria Powered Biohybrids Target E. coli Biofilms. *ACS Nano*. PMID: 28933815: null, 2017. eprint: http://dx.doi.org/10.1021/acsnano.7b04128.

[474] H. Koorehdavoudi, P. Bogdan, G. Wei, R. Marculescu, J. Zhuang, R. W. Carlsen, M. Sitti. Multi-fractal characterization of bacterial swimming dynamics: a case study on real and simulated Serratia marcescens. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **473** (2203), 2017. eprint: http://rspa.royalsocietypublishing.org/content/473/2203/20170154.full.pdf.

[475] G. Amador, T. Endlein, M. Sitti. Soiled adhesive pads shear clean by slipping: a robust self-cleaning mechanism in climbing beetles. *Journal of The Royal Society Interface* **14** (131): 20170134, 2017.

[476] D.-M. Drotlef, M. Amjadi, M. Yunusa, M. Sitti. Bioinspired Composite Microfibers for Skin Adhesion and Signal Amplification of Wearable Sensors. *Advanced Materials* **29** (28). Back Cover: 1701353, 2017.

[477] I. S. Khalil, A. F. Tabak, K. Sadek, D. Mahdy, N. Hamdi, M. Sitti. Rubbing Against Blood Clots Using Helical Robots: Modeling and In Vitro Experimental Validation. *IEEE Robotics and Automation Letters* **2** (2): 927–934, 2017.

[478] H. Ceylan, J. Giltinan, K. Kozielski, M. Sitti. Mobile microrobots for bioengineering applications. *Lab on a Chip* **17** (10): 1705–1724, 2017.

[479] B. Mostaghaci, O. Yasa, J. Zhuang, M. Sitti. Bioadhesive Bacterial Microswimmers for Targeted Drug Delivery in the Urinary and Gastrointestinal Tracts. *Advanced Science* **4** (6): 1700058, 2017.

[480] J. Zhuang, B.-W. Park, M. Sitti. Propulsion and Chemotaxis in Bacteria-Driven Microswimmers. *Advanced Science* **4** (9): 1700109, 2017.

[481] H. Kim, M. Sitti, T. Seo. Tail-Assisted Mobility and Stability Enhancement in Yaw/Pitch Motions of a Water-Running Robot. *IEEE/ASME Transactions on Mechatronics* **22** (3): 1207–1217, 2017.

[482] A. V. Singh, S. Baylan, B.-W. Park, G. Richter, M. Sitti. Hydrophobic pinning with copper nanowhiskers leads to bactericidal properties. *PloS One* **12** (4): e0175428, 2017.

[483] M. M. Stanton, B.-W. Park, A. Miguel-López, X. Ma, M. Sitti, S. Sánchez. Biohybrid microtube swimmers driven by single captured bacteria. *Small* **13** (19), 2017.

[484] W. Wang, J. Giltinan, S. Zakharchenko, M. Sitti. Dynamic and programmable self-assembly of micro-rafts at the air-water interface. *Science Advances* **3** (5): e1602522, 2017.

[485] K. Kozielski, M. Sitti. Gene delivery particle engineering strategies for shape-dependent targeting of cells and tissues. *Current Gene Therapy* **17**, 2017.

[486] M. Amjadi, B. Mostaghaci, M. Sitti. Recent Advances in Skin Penetration Enhancers for Transdermal Gene and Drug Delivery. *Current Gene Therapy* **17**, 2017.

[487] S. Song, D.-M. Drotlef, C. Majidi, M. Sitti. Controllable load sharing for soft adhesive interfaces on three-dimensional surfaces. *Proceedings of the National Academy of Sciences* **114** (22): E4344–E4353, 2017.

### 2016

[488] Z. Ye, G. Z. Lum, S. Song, S. Rich, M. Sitti. Gallium Adhesion: Phase Change of Gallium Enables Highly Reversible and Switchable Adhesion (Adv. Mater. 25/2016). *Advanced Materials* **28** (25): 5087–5087, 2016.

[489] I. S. Khalil, A. Fatih Tabak, A. Klingner, M. Sitti. Magnetic propulsion of robotic sperms at low-Reynolds number. *Applied Physics Letters* **109** (3): 033701, 2016.

[490] A. V. Singh, M. Sitti. Patterned and Specific Attachment of Bacteria on Biohybrid Bacteria-Driven Microswimmers. *Advanced Healthcare Materials* **5** (18): 2325–2331, 2016.

[491] G. Z. Lum, Z. Ye, X. Dong, H. Marvi, O. Erin, W. Hu, M. Sitti. Shape-programmable magnetic soft matter. *Proceedings of the National Academy of Sciences* **113** (41): E6007–E6015, 2016.

[492] M. Amjadi, M. Sitti. High-Performance Multiresponsive Paper Actuators. *ACS Nano* **10** (11): 10202–10210, 2016.

[493] H. Wu, V. Sariola, J. Zhao, H. Ding, M. Sitti, C. J. Bettinger. Composition-dependent underwater adhesion of catechol-bearing hydrogels. *Polymer International* **65** (11): 1355–1359, 2016.

[494] Z. Hosseinidoust, B. Mostaghaci, O. Yasa, B.-W. Park, A. V. Singh, M. Sitti. Bioengineered and biohybrid bacteria-based systems for drug delivery. *Advanced Drug Delivery Reviews* **106**: 27–44, 2016.

[495] E. Diller, J. Giltinan, G. Z. Lum, Z. Ye, M. Sitti. Six-degree-of-freedom magnetic actuation for wireless microrobotics. *The International Journal of Robotics Research* **35** (1-3): 114–128, 2016.

[496] A. Y. Stark, M. R. Klittich, M. Sitti, P. H. Niewiarowski, A. Dhinojwala. The effect of temperature and humidity on adhesion of a gecko-inspired adhesive: implications for the natural system. *Scientific Reports* **6**: 30936, 2016.

[497] J. Zhuang, M. Sitti. Chemotaxis of bio-hybrid multiple bacteria-driven microswimmers. *Scientific Reports* **6**: 32135, 2016.

[498] J. Giltinan, E. Diller, M. Sitti. Programmable assembly of heterogeneous microparts by an untethered mobile capillary microgripper. *Lab on a Chip* **16** (22): 4445–4457, 2016.

[499] M. Amjadi, K.-U. Kyung, I. Park, M. Sitti. Stretchable, Skin-Mountable, and Wearable Strain Sensors and Their Potential Applications: A Review. *Advanced Functional Materials* **26**: 1678–1698, 2016.

[500] M. Amjadi, M. Turan, C. P. Clementson, M. Sitti. Parallel microcracks-based ultrasensitive and highly stretchable strain sensors. *ACS Applied Materials & Interfaces* **8** (8): 5618–5626, 2016.

[501] A. V. Singh, M. Sitti. Targeted drug delivery and imaging using mobile milli/microrobots: A promising future towards theranostic pharmaceutical design. *Current Pharmaceutical Design* **22** (11): 1418–1428, 2016.

[502] D. Son, S. Yim, M. Sitti. A 5-D localization method for a magnetically manipulated untethered robot using a 2-D array of Hall-effect sensors. *IEEE/ASME Transactions on Mechatronics* **21** (2): 708–716, 2016.

[503] L. Hines, K. Petersen, M. Sitti. Inflated soft actuators with reversible stable deformations. *Advanced Materials* **28** (19): 3690–3696, 2016.

[504] F. Munoz, G. Alici, W. Li, M. Sitti. Size optimization of a magnetic system for drug delivery with capsule robots. *IEEE Transactions on Magnetics* **52** (5): 1–11, 2016.

## 21.4.3 Conference Papers

### 2018

[505] Z. Davidson, H. Shahsavan, Y. Guo, L. Hines, Y. Xia, S. Yang, M. Sitti. Electromechanical actuation of dielectric liquid crystal elastomers for soft robotics. *Bulletin of the American Physical Society*, in press.

[506] U. Sanli, M. Baluktsian, H. Ceylan, M. Sitti, M. Weigand, G. Schuetz, K. Keskinbora. X-ray Optics Fabrication Using Unorthodox Approaches. *Bulletin of the American Physical Society*, in press.

[507] H. Kim, M. A. Woodward, M. Sitti. Enhanced Non-Steady Gliding Performance of the MultiMo-Bat through Optimal Airfoil Configuration and Control Strategy. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pages 1382–1388, 2018.

[508] W. Wang, V. Kishore, L. Koens, E. Lauga, M. Sitti. Collectives of Spinning Mobile Microrobots for Navigation and Object Manipulation at the Air-Water Interface. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pages 1–9, 2018.

[509] M. Turan, Y. Almalioglu, H. B. Gilbert, U. Sari Alp Eren andSoylu, M. Sitti. Endo-VMFuseNet: A Deep Visual-Magnetic Sensor Fusion Approach for Endoscopic Capsule Robots. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pages 1–7, 2018.

[510] M. Turan, Y. Almalioglu, H. Gilbert, H. Araujo, T. Cemgil, M. Sitti. Endosensorfusion: Particle filtering-based multi-sensory data fusion with switching state-space model for endoscopic capsule robots. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pages 1–8, 2018.

## 2017

[511] M. Mastrangeli. Self-assembly of micro/nanosystems across scales and interfaces. In *2017 19th International Conference on Solid-State Sensors, Actuators and Microsystems (TRANSDUCERS)*, pages 676–681, 2017.

[512] I. S. M. Khalil, A. F. Tabak, M. A. Seif, A. Klingner, B. Adel, M. Sitti. Swimming in low reynolds numbers using planar and helical flagellar waves. In *International Conference on Intelligent Robots and Systems (IROS) 2017*, pages 1907–1912, 2017.

[513] X. Dong, M. Sitti. Planning spin-walking locomotion for automatic grasping of microobjects by an untethered magnetic microgripper. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6612–6618, 2017.

[514] O. Erin, J. Giltinan, L. Tsai, M. Sitti. Design and actuation of a magnetic millirobot under a constant unidirectional magnetic field. In *Proceedings 2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3404–3410, 2017.

[515] D. Son, M. D. Dogan, M. Sitti. Magnetically actuated soft capsule endoscope for fine-needle aspiration biopsy. In *Proceedings 2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1132–1139, 2017.

[516] H. Kim, Y. Liu, K. Jeong, M. Sitti, T. Seo. Dynamic analysis on hexapedal water-running robot with compliant joints. In *2017 14th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI)*, pages 250–251, 2017.

[517] I. S. M. Khalil, A. Alfar, A. F. Tabak, A. Klingner, S. Stramigioli, M. Sitti. Positioning of drug carriers using permanent magnet-based robotic system in three-dimensional space. In *2017 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 1117–1122, 2017.

[518] G. Z. Lum, M. T. Pham, T. J. Teo, G. Yang, S. H. Yeo, M. Sitti. An XY Θz flexure mechanism with optimal stiffness properties. In *2017 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 1103–1110, 2017.

## 2016

[519] F. Munoz, G. Alici, H. Zhou, W. Li, M. Sitti. Analysis of the magnetic torque on a tilted permanent magnet for drug delivery in capsule robots. In *Advanced Intelligent Mechatronics (AIM), 2016 IEEE International Conference on*, pages 1386–1391, 2016.

[520] H. Kim, K. Jeong, M. Sitti, T. Seo. Steering control of a water-running robot using an active tail. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 4945–4950, 2016.

[521] I. S. Khalil, A. F. Tabak, A. Hosney, A. Mohamed, A. Klingner, M. Ghoneima, M. Sitti. Sperm-shaped magnetic microrobots: Fabrication using electrospinning, modeling, and characterization. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pages 1939–1944, 2016.

[522] I. S. Khalil, A. F. Tabak, A. Hosney, A. Klingner, M. Shalaby, R. M. Abdel-Kader, M. Serry, M. Sitti. Targeting of cell mockups using sperm-shaped microrobots in vitro. In *Biomedical Robotics and Biomechatronics (BioRob), 2016 6th IEEE International Conference on*, pages 495–501, 2016.

### 21.4.4 Theses

#### PhD Theses

[523] D. Son. Manipulation of Magnetically Actuated Capsule Endoscopes. *Carnegie Mellon University, USA*, 2018.

[524] J. Giltinan. Actuation and Addressability for the Manipulation of Objects by Magnetic Microrobotic Capillary Grippers. *Carnegie Mellon University, USA*, 2018.

[525] M. Amjadi. Functional Nanomaterial Composites for Soft Sensing and Actuation. *ETH Zurich, Switzerland*, 2018.

[526] M. Turan. Intelligent Localization and Mapping Methods for Endoscopic Capsule Robots. *ETH Zurich, Switzerland*, 2018.

[527] M. Woodward. MultiMo-Bat: Biologically Inspired Integrated Multi-Modal Locomotion. *Carnegie Mellon University, USA*, 2017.

[528] D. Colmenares. Design and Control of a Resonant, Flapping Wing Micro Air Vehicle Capable of Sustained Flight. *Carnegie Mellon University, USA*, 2017.

[529] S. Song. Soft Robotic Grippers using Gecko-Inspired Fibrillar Adhesives for Three-Dimensional Surface Grasping. *Carnegie Mellon University, USA*, 2017.

[530] J. Zhuang. Taxis-Based Motion Control of Biohybrid Microrobots. *Carnegie Mellon University, USA*, 2017.

[531] G. Z. Lum. Design methodology for multi-degrees-of-freedom compliant mechanisms. *Carnegie Mellon University, USA*, 2016.

## 21.5 Haptic Intelligence Department

### 21.5.1 Journal Articles

#### 2019

[532] N. T. Fitter, K. J. Kuchenbecker. How Does It Feel to Clap Hands with a Robot? *International Journal of Social Robotics*, 2019.

#### 2018

[533] C. Pacchierotti, E. M. Young, K. J. Kuchenbecker. Task-Driven PCA-Based Design Optimization of Wearable Cutaneous Devices. *IEEE Robotics and Automation Letters* **3** (3). Presented at ICRA 2018: 2214–2221, 2018.

[534] A. E. Block, K. J. Kuchenbecker. Softness, Warmth, and Responsiveness Improve Robot Hugs. *International Journal of Social Robotics*, 2018 (cited on pages 163, 164).

[535] N. T. Fitter, K. J. Kuchenbecker. Teaching a Robot Bimanual Hand-Clapping Games via Wrist-Worn IMUs. *Frontiers in Robotics and Artificial Intelligence* **5** (85), 2018.

[536] E. Ambron, A. Miller, K. J. Kuchenbecker, L. J. Buxbaum, H. B. Coslett. Immersive Low-Cost Virtual Reality Treatment for Phantom Limb Pain: Evidence from Two Cases. *Frontiers in Neurology* **9** (67): 1–7, 2018 (cited on page 153).

[537] Y. A. Oquendo, E. W. Riddle, D. Hiller, T. A. Blinman, K. J. Kuchenbecker. Automatically Rating Trainee Skill at a Pediatric Laparoscopic Suturing Task. *Surgical Endoscopy* **32** (4): 1840–1857, 2018.

#### 2017

[538] A. M. Bur, E. D. Gomez, J. G. Newman, G. S. Weinstein, J. Bert W. O'Malley, C. H. Rassekh, K. J. Kuchenbecker. Evaluation of High-Fidelity Simulation as a Training Tool in Transoral Robotic Surgery. *Laryngoscope* **127** (12): 2790–2795, 2017.

[539] J. D. Brown, C. E. O'Brien, S. C. Leung, K. R. Dumon, D. I. Lee, K. J. Kuchenbecker. Using Contact Forces and Robot Arm Accelerations to Automatically Rate Surgeon Skill at Peg Transfer. *IEEE Transactions on Biomedical Engineering* **64** (9): 2263–2275, 2017.

[540] H. Culbertson, K. J. Kuchenbecker. Ungrounded Haptic Augmented Reality System for Displaying Texture and Friction. *IEEE/ASME Transactions on Mechatronics* **22** (4): 1839–1849, 2017.

[541] N. Gurari, A. M. Okamura, K. J. Kuchenbecker. Perception of Force and Stiffness in the Presence of Low-Frequency Haptic Noise. *PLoS ONE* **12** (6): e0178605, 2017.

[542] K. J. Kuchenbecker, R. Parajon, M. P. Maggio. Evaluation of a Vibrotactile Simulator for Dental Caries Detection. *Simulation in Healthcare* **12** (3): 148–156, 2017.

[543] R. P. Khurshid, N. T. Fitter, E. A. Fedalei, K. J. Kuchenbecker. Effects of Grip-Force, Contact, and Acceleration Feedback on a Teleoperated Pick-and-Place Task. *IEEE Transactions on Haptics* **10** (1): 40–53, 2017.

[544] H. Culbertson, K. J. Kuchenbecker. Importance of Matching Physical Friction, Hardness, and Texture in Creating Realistic Haptic Virtual Surfaces. *IEEE Transactions on Haptics* **10** (1): 63–74, 2017.

### 21.5.2 Conference Papers

#### 2019

[545] E. Cao, S. Machaca, T. Bernard, B. Wolfinger, Z. Patterson, A. Chi, G. L. Adrales, K. J. Kuchenbecker, J. D. Brown. Bimanual Wrist-Squeezing Haptic Feedback Changes Speed-Force Tradeoff in Robotic Surgery Training. Extended abstract presented as an ePoster at the Annual Meeting of the Society of American Gastrointestinal and Endoscopic Surgeons (SAGES). Baltimore, USA, 2019.

[546] M. Mohan, K. J. Kuchenbecker. A Design Tool for Therapeutic Social-Physical Human-Robot Interactions. Workshop paper (3 pages) presented at the HRI Pioneers Workshop. Daegu, South Korea, 2019 (cited on page 160).

[547] H. Seifi, F. Fazlollahi, M. Oppermann, J. A. Sastrillo, J. Ip, A. Agrawal, G. Park, K. J. Kuchenbecker, K. E. MacLean. Haptipedia: Accelerating Haptic Device Discovery to Support Interaction & Engineering Design. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2019 (cited on page 150).

## 2018

[548] G. Park, K. J. Kuchenbecker. Reducing 3D Vibrations to 1D in Real Time. Hands-on demonstration (4 pages) presented at AsiaHaptics. Incheon, South Korea, 2018 (cited on page 158).

[549] K. J. Kuchenbecker. Arm-Worn Tactile Displays. Cross-Cutting Challenge Interactive Discussion presented at the IEEE Haptics Symposium. San Francisco, USA, 2018.

[550] A. Burka, K. J. Kuchenbecker. Can Humans Infer Haptic Surface Properties from Images? Work-in-progress paper (3 pages) presented at the IEEE Haptics Symposium. San Francisco, USA, 2018 (cited on page 141).

[551] E. Young, K. J. Kuchenbecker. Delivering 6-DOF Fingertip Tactile Cues. Work-in-progress paper (5 pages) presented at EuroHaptics. Pisa, Italy, 2018 (cited on page 151).

[552] H. Lee, K. Park, J. Kim, K. J. Kuchenbecker. A Large-Scale Fabric-Based Tactile Sensor Using Electrical Resistance Tomography. Hands-on demonstration (3 pages) presented at AsiaHaptics. Incheon, South Korea, 2018 (cited on page 166).

[553] H. Seifi, F. Fazlollahi, G. Park, K. J. Kuchenbecker, K. E. MacLean. Haptipedia: Exploring Haptic Device Design Through Interactive Visualizations. Hands-on demonstration presented at EuroHaptics. Pisa, Italy, 2018 (cited on page 150).

[554] H. Seifi, K. E. MacLean, K. J. Kuchenbecker, G. Park. Haptipedia: An Expert-Sourced Interactive Device Visualization for Haptic Designers. Work-in-progress paper (3 pages) presented at the IEEE Haptics Symposium. San Francisco, USA, 2018 (cited on page 150).

[555] N. T. Fitter, M. Mohan, K. J. Kuchenbecker, M. J. Johnson. Exercising with Baxter: Design and Evaluation of Assistive Social-Physical Human-Robot Interaction. Workshop paper (6 pages) presented at the HRI Workshop on Personal Robots for Exercising and Coaching. Chicago, USA, 2018 (cited on page 162).

[556] D. Gueorguiev, D. Tzionas, C. Pacchierotti, M. J. Black, K. J. Kuchenbecker. Statistical Modelling of Fingertip Deformations and Contact Forces during Tactile Interaction. Extended abstract presented at the Hand, Brain and Technology conference (HBT). Ascona, Switzerland, Aug. 2018 (cited on pages 75, 149).

[557] D. Gueorguiev, D. Tzionas, C. Pacchierotti, M. J. Black, K. J. Kuchenbecker. Towards a Statistical Model of Fingertip Contact Deformations from 4D Data. Work-in-progress paper (3 pages) presented at the IEEE Haptics Symposium. San Francisco, USA, Mar. 2018 (cited on pages 75, 149).

[558] H. Lee, J. Kim, K. J. Kuchenbecker. Soft Multi-Axis Boundary-Electrode Tactile Sensors for Whole-Body Robotic Skin. Workshop paper (2 pages) presented at the RSS Pioneers Workshop. Pittsburgh, USA, 2018 (cited on page 166).

[559] A. E. Block, K. J. Kuchenbecker. Emotionally Supporting Humans Through Robot Hugs. Workshop paper (2 pages) presented at the HRI Pioneers Workshop. Chicago, USA, 2018 (cited on page 163).

[560] R. Burns, K. J. Kuchenbecker. Designing a Haptic Empathetic Robot Animal for Children with Autism. Workshop paper (4 pages) presented at the RSS Workshop on Robot-Mediated Autism Intervention: Hardware, Software and Curriculum. Pittsburgh, USA, 2018 (cited on page 167).

[561] G. Park, K. J. Kuchenbecker. Reducing 3D Vibrations to 1D in Real Time. Hands-on demonstration presented at EuroHaptics. Pisa, Italy, 2018 (cited on page 158).

## 2017

[562] A. E. Block, K. J. Kuchenbecker. Physical and Behavioral Factors Improve Robot Hug Quality. Workshop Paper (2 pages) presented at the RO-MAN Workshop on Social Interaction and Multimodal Expression for Socially Intelligent Robots. Lisbon, Portugal, 2017 (cited on page 163).

[563] A. Burka, K. J. Kuchenbecker. How Much Haptic Surface Data is Enough? Workshop paper (5 pages) presented at the AAAI Spring Symposium on Interactive Multi-Sensory Object Perception for Embodied Agents. Stanford, USA, 2017 (cited on page 141).

[564] Y. A. Oquendo, E. W. Riddle, D. Hiller, T. A. Blinman, K. J. Kuchenbecker. Automatic OSATS Rating of Trainee Skill at a Pediatric Laparoscopic Suturing Task. Extended abstract presented as a podium presentation at the Annual Meeting of the Society of American Gastrointestinal and Endoscopic Surgeons (SAGES). Houston, USA, 2017.

[565] N. T. Fitter, K. J. Kuchenbecker. Hand-Clapping Games with a Baxter Robot. Hands-on demonstration presented at ACM/IEEE International Conference on Human-Robot Interaction (HRI). Vienna, Austria, 2017 (cited on page 162).

[566] A. Burka, K. J. Kuchenbecker. Proton Pack: Visuo-Haptic Surface Data Recording. Hands-on demonstration presented at the IEEE World Haptics Conference (WHC). Munich, Germany, 2017 (cited on page 141).

[567] J. Carlson, K. J. Kuchenbecker. An Interactive Augmented-Reality Video Training Platform for the da Vinci Surgical System. Workshop paper (3 pages) presented at the ICRA Workshop on C4 Surgical Robots. Singapore, 2017.

[568] N. T. Fitter, K. J. Kuchenbecker. Physically Interactive Exercise Games with a Baxter Robot. Hands-on demonstration presented at the IEEE World Haptics Conference (WHC). Munich, Germany, 2017 (cited on page 162).

[569] S. Hu, K. J. Kuchenbecker. Teaching a Robot to Collaborate with a Human Via Haptic Teleoperation. Work-in-progress paper (2 pages) presented at the IEEE World Haptics Conference (WHC). Munich, Germany, 2017 (cited on page 161).

[570] A. E. Block, K. J. Kuchenbecker. How Should Robots Hug? Work-in-progress paper (2 pages) presented at the IEEE World Haptics Conference (WHC). Munich, Germany, 2017 (cited on page 163).

[571] C. Ng, K. Zareinia, Q. Sun, K. J. Kuchenbecker. Stiffness Perception during Pinching and Dissection with Teleoperated Haptic Forceps. In *Proceedings of the International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 456–463, 2017.

[572] N. T. Fitter, K. J. Kuchenbecker. Synchronicity Trumps Mischief in Rhythmic Human-Robot Social-Physical Interaction. In *Proceedings of the International Symposium on Robotics Research (ISRR)*, 2017.

[573] E. M. Young, K. J. Kuchenbecker. Design of a Parallel Continuum Manipulator for 6-DOF Fingertip Haptic Display. In *Proceedings of the IEEE World Haptics Conference (WHC)*. Finalist for best poster paper, pages 599–604, 2017 (cited on page 151).

[574] S. Hu, K. J. Kuchenbecker. High Magnitude Unidirectional Haptic Force Display Using a Motor/Brake Pair and a Cable. In *Proceedings of the IEEE World Haptics Conference (WHC)*, pages 394–399, 2017 (cited on page 152).

[575] J. D. Brown, J. N. Fernandez, S. P. Cohen, K. J. Kuchenbecker. A Wrist-Squeezing Force-Feedback System for Robotic Surgery Training. In *Proceedings of the IEEE World Haptics Conference (WHC)*, pages 107–112, 2017.

[576] A. Burka, A. Rajvanshi, S. Allen, K. J. Kuchenbecker. Proton 2: Increasing the Sensitivity and Portability of a Visuo-haptic Surface Interaction Recorder. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 439–445, 2017 (cited on page 141).

[577] A. Burka, K. J. Kuchenbecker. Handling Scan-Time Parameters in Haptic Surface Classification. In *Proceedings of the IEEE World Haptics Conference (WHC)*, pages 424–429, 2017 (cited on page 141).

### 21.5.3 Book Chapters

#### 2018

[578] K. J. Kuchenbecker. Haptics and Haptic Interfaces. In *Encyclopedia of Robotics*. Springer, 2018.

## 21.5.4 Theses

### PhD Theses

[579] A. L. Burka. Instrumentation, Data, and Algorithms for Visually Understanding Haptic Surface Properties. Department of Electrical and Systems Engineering. PhD thesis. Philadelphia, USA: University of Pennsylvania, 2018 (cited on page 141).

### Master's Theses

[580] M. P. Forte. Robust Visual Augmented Reality in Robot-Assisted Surgery. Department of Electronic, Information, and Biomedical Engineering. MA thesis. Milan, Italy: Politecnico di Milano, 2018 (cited on page 159).

## 21.6 Autonomous Learning

### 21.6.1 Journal Articles

#### 2018

[581] V. Botella-Soler, S. Deny, G. Martius, O. Marre, G. Tkačik. Nonlinear decoding of a complex movie from the mammalian retina. *PLOS Computational Biology* **14** (5): 1–27, 2018.

[582] S. Blaes, M. Vlastelica, J.-J. Zhu, G. Martius. Control What You Can: Intrinsically Motivated Hierarchical Reinforcement Learner. In *Deep Reinforcement Learning Workshop at NeurIPS 2018, Montreal*, 2018 (cited on page 171).

#### 2017

[583] R. Der, G. Martius. Self-Organized Behavior Generation for Musculoskeletal Robots. *Frontiers in Neurorobotics* **11**: 8, 2017.

### 21.6.2 Conference Papers

#### 2018

[584] H. Sun, G. Martius. Robust Affordable 3D Haptic Sensation via Learning Deformation Patterns. In *Proceedings International Conference on Humanoid Robots (IEEE Humanoids)*. Oral Presentation, pages 846–853, 2018 (cited on page 175).

[585] M. Rolinek, G. Martius. L4: Practical loss-based stepsize adaptation for deep learning. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, pages 6434–6444, 2018 (cited on page 174).

[586] D. Baumann, J.-J. Zhu, G. Martius, S. Trimpe. Deep Reinforcement Learning for Event-Triggered Control. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 943–950, 2018 (cited on pages 174, 198, 200).

[587] S. S. Sahoo, C. H. Lampert, G. Martius. Learning equations for extrapolation and control. In *Proc. 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, 2018*. Vol. 80, pages 4442–4450, 2018 (cited on page 173).

[588] C. Pinneri, G. Martius. Systematic self-exploration of behaviors for robots in a dynamical systems framework. In *Proc. Artificial Life XI*, pages 319–326, 2018 (cited on page 172).

## 21.7 Autonomous Vision

### 21.7.1 Journal Articles

#### 2018

[589] H. Alhaija, S. Mustikovela, L. Mescheder, A. Geiger, C. Rother. Augmented Reality Meets Computer Vision: Efficient Data Generation for Urban Driving Scenes. *International Journal of Computer Vision (IJCV), 2018*, 2018 (cited on page 180).

#### 2016

[351] M. A. Brubaker, A. Geiger, R. Urtasun. Map-Based Probabilistic Visual Self-Localization. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 2016 (cited on page 179).

### 21.7.2 Conference Papers

#### 2018

[354] L. Sevilla-Lara, Y. Liao, F. Güney, V. Jampani, A. Geiger, M. J. Black. On the Integration of Optical Flow and Action Recognition. In *German Conference on Pattern Recognition (GCPR)*. Vol. LNCS 11269, pages 281–297, Oct. 2018 (cited on pages 86, 87, 179).

[359] J. Janai, F. Güney, A. Ranjan, M. J. Black, A. Geiger. Unsupervised Learning of Multi-Frame Optical Flow with Occlusions. In *European Conference on Computer Vision (ECCV)*. Vol. Lecture Notes in Computer Science, vol 11220, pages 713–731, 2018 (cited on pages 86, 179).

[590] Conditional Affordance Learning for Driving in Urban Environments. In *Conference on Robot Learning 2018*, 2018 (cited on page 179).

[591] Towards Robust Visual Odometry with a Multi-Camera System. In *International Conference on Intelligent Robots and Systems (IROS) 2018*, 2018.

[592] I. Cherabier, J. Schönberger, M. Oswald, M. Pollefeys, A. Geiger. Learning Priors for Semantic 3D Reconstruction. In *European Conference on Computer Vision (ECCV)*, 2018 (cited on page 178).

[593] B. Coors, A. P. Condurache, A. Geiger. SphereNet: Learning Spherical Representations for Detection and Classification in Omnidirectional Images. In *European Conference on Computer Vision (ECCV)*, 2018 (cited on page 179).

[594] L. Mescheder, A. Geiger, S. Nowozin. Which Training Methods for GANs do actually Converge? In *International Conference on Machine learning (ICML)*, 2018 (cited on page 180).

[595] I. A. Barsan, P. Liu, M. Pollefeys, A. Geiger. Robust Dense Mapping for Large-Scale Dynamic Environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2018*, 2018.

[596] J. Schönberger, M. Pollefeys, A. Geiger, T. Sattler. Semantic Visual Localization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 (cited on page 179).

[597] D. Stutz, A. Geiger. Learning 3D Shape Completion from Laser Scan Data with Weak Supervision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 (cited on page 178).

[598] Y. Liao, S. Donne, A. Geiger. Deep Marching Cubes: Learning Explicit Surface Representations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 (cited on page 178).

[599] D. Paschalidou, A. O. Ulusoy, C. Schmitt, L. Gool, A. Geiger. RayNet: Learning Volumetric 3D Reconstruction with Ray Potentials. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018 (cited on page 178).

[600] B. Coors, A. Condurache, A. Mertins, A. Geiger. Learning Transformation Invariant Representations with Weak Supervision. In *International Conference on Computer Vision Theory and Applications*, 2018.

[601] M. Menze, C. Heipke, A. Geiger. Object Scene Flow. *ISPRS Journal of Photogrammetry and Remote Sensing (JPRS)*, 2018 (cited on page 179).

## 2017

[380] J. Janai, F. Güney, J. Wulff, M. Black, A. Geiger. Slow Flow: Exploiting High-Speed Cameras for Accurate and Diverse Optical Flow Reference Data. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 1406–1416, 2017 (cited on pages 68, 179).

[382] A. O. Ulusoy, M. J. Black, A. Geiger. Semantic Multi-view Stereo: Jointly Estimating Objects and Voxels. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on pages 89, 178).

[389] G. Riegler, O. Ulusoy, A. Geiger. OctNet: Learning Deep 3D Representations at High Resolutions. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on page 178).

[602] L. Mescheder, S. Nowozin, A. Geiger. The Numerics of GANs. In *Proceedings from the conference "Neural Information Processing Systems 2017.* 2017 (cited on page 180).

[603] P. Liu, L. Heng, T. Sattler, A. Geiger, M. Pollefeys. Direct Visual Odometry for a Fisheye-Stereo Camera. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.

[604] G. Riegler, A. O. Ulusoy, H. Bischof, A. Geiger. OctNetFusion: Learning Depth Fusion from Data. In *International Conference on 3D Vision (3DV) 2017*, 2017 (cited on page 178).

[605] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, A. Geiger. Sparsity Invariant CNNs. In *International Conference on 3D Vision (3DV) 2017*, 2017 (cited on page 178).

[606] A. Behl, O. H. Jafari, S. K. Mustikovela, H. A. Alhaija, C. Rother, A. Geiger. Bounding Boxes, Segmentations and Object Coordinates: How Important is Recognition for 3D Scene Flow Estimation in Autonomous Driving Scenarios? In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, 2017 (cited on page 179).

[607] H. A. Alhaija, S. K. Mustikovela, L. Mescheder, A. Geiger, C. Rother. Augmented Reality Meets Deep Learning for Car Instance Segmentation in Urban Scenes. In *Proceedings of the British Machine Vision Conference 2017*, 2017 (cited on page 180).

[608] H. Alhaija, S. Mustikovela, A. Geiger, C. Rothern. Geometric Image Synthesis. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 2018 (cited on page 180).

[609] F. Camposeco, T. Sattler, A. Cohen, A. Geiger, M. Pollefeys. Toroidal Constraints for Two Point Localization Under High Outlier Ratios. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on page 179).

[610] T. Schöps, J. L. Schönberger, S. Galliani, T. Sattler, K. Schindler, M. Pollefeys, A. Geiger. A Multi-View Stereo Benchmark with High-Resolution Images and Multi-Camera Videos. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, 2017 (cited on page 178).

[611] L. Mescheder, S. Nowozin, A. Geiger. Adversarial Variational Bayes: Unifying Variational Autoencoders and Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*. Vol. 70. Proceedings of Machine Learning Research, 2017 (cited on page 180).

## 2016

[391] F. Güney, A. Geiger. Deep Discrete Flow. In *Asian Conference on Computer Vision (ACCV)*, 2016 (cited on page 179).

[394] J. Xie, M. Kiefel, M.-T. Sun, A. Geiger. Semantic Instance Annotation of Street Scenes by 3D to 2D Label Transfer. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016 (cited on pages 88, 179).

[395] A. O. Ulusoy, M. J. Black, A. Geiger. Patches, Planes and Probabilities: A Non-local Prior for Volumetric 3D Reconstruction. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016 (cited on pages 89, 178).

## 21.8 Dynamic Locomotion

### 21.8.1 Journal Articles

#### 2018

[612] A. Spröwitz, A. Tuleu, M. Ajallooeian, M. Vespignani, R. Moeckel, P. Eckert, M. D'Haene, J. Degrave, A. Nordmann, B. Schrauwen, J. Steil, A. J. Ijspeert. Oncilla robot: a versatile open-source quadruped research robot with compliant pantograph legs. *Frontiers in Robotics and AI* **5** (67). https://www.frontiersin.org/articles/10.3389/frobt.2018.00067/full, 2018 (cited on page 184).

#### 2017

[613] S. Pouya, M. Khodabakhsh, A. Spröwitz, A. Ijspeert. Spinal joint compliance and actuation in a simulated bounding quadruped robot. *Autonomous Robots*: 437–452, 2017 (cited on page 184).

#### 2016

[614] C. Hubicki, J. Grimes, M. Jones, D. Renjewski, A. Spröwitz, A. Abate, J. Hurst. ATRIAS: Design and validation of a tether-free 3D-capable spring-mass bipedal robot. *The International Journal of Robotics Research* **35** (12): 1497–1521, 2016 (cited on page 186).

### 21.8.2 Conference Papers

#### 2018

[615] S. Heim, A. Spröwitz. Learning from Outside the Viability Kernel: Why we Should Build Robots that can Fail with Grace. In *Proceedings of SIMPAR 2018*, pages 55–61, 2018 (cited on page 185).

[616] S. Heim, F. Ruppert, A. Sarvestani, A. Spröwitz. Shaping in Practice: Training Wheels to Learn Fast Hopping Directly in Hardware. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2018*, pages 5076–5081, 2018 (cited on page 185).

#### 2017

[617] S. Heim, A. Spröwitz. Is Growing Good for Learning? In *Proceedings of the 8th International Symposium on Adaptive Motion of Animals and Machines AMAM2017*, 2017.

[618] S. Heim, F. Grimminger, D. Özge, A. Spröwitz. Linking Mechanics and Learning. In *Proceedings of Dynamic Walking 2017*, 2017.

[619] A. Spröwitz, C. Göttler, A. Sinha, C. Caer, M. U. Öztekin, K. Petersen, M. Sitti. Scalable Pneumatic and Tendon Driven Robotic Joint Inspired by Jumping Spiders. In *Proceedings 2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 64–70, 2017 (cited on page 184).

### 21.8.3 Theses

#### 2018

[620] J. Richter. Untersuchung und Charakterisierung von Teilelementen der Modifikation im Lumbosacralbereich von Vögeln. BSc thesis. DLG MPI IS Stuttgart and Hochschule Harz, July 2018.

#### 2017

[621] B. Györfi. Evaluation of the passive dynamics of compliant legs with inertia. MA thesis. University of Applied Science Pforzheim, Germany, 2017 (cited on page 184).

## 21.9 Embodied Vision

### 21.9.1 Journal Articles

#### 2018

[622] J. Vinogradska, B. Bischoff, J. Achterhold, T. Koller, J. Peters. Numerical Quadrature for Probabilistic Policy Search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*: 1–1, 2018.

## 21.10 Intelligent Control Systems

### 21.10.1 Journal Articles

#### 2019

[623] M. Neumann-Brosig, A. Marco, D. Schwarzmann, S. Trimpe. Data-efficient Auto-tuning with Bayesian Optimization: An Industrial Control Study. *IEEE Transactions on Control Systems Technology*, 2019 (cited on pages 198, 199).

#### 2018

[3] M. Muehlebach, S. Trimpe. Distributed Event-Based State Estimation for Networked Systems: An LMI Approach. *IEEE Transactions on Automatic Control* **63** (1): 269–276, 2018 (cited on pages 198, 200).

[624] M. Hertneck, J. Koehler, S. Trimpe, F. Allgöwer. Learning an Approximate Model Predictive Controller with Guarantees. *IEEE Control Systems Letters* **2** (3): 543–548, 2018 (cited on page 198).

#### 2017

[5] S. Trimpe. Event-based State Estimation: An Emulation-based Approach. *IET Control Theory & Applications* **11** (11): 1684–1693, 2017 (cited on pages 198, 200).

#### 2016

[10] M. Wüthrich, S. Trimpe, C. Garcia Cifuentes, D. Kappler, S. Schaal. A New Perspective and Extension of the Gaussian Filter. *The International Journal of Robotics Research* **35** (14): 1731–1749, 2016 (cited on page 198).

[11] D. Laidig, S. Trimpe, T. Seel. Event-based Sampling for Reducing Communication Load in Realtime Human Motion Analysis by Wireless Inertial Sensor Networks. *Current Directions in Biomedical Engineering* **2** (1): 711–714, 2016.

### 21.10.2 Conference Papers

#### 2018

[20] A. Doerr, C. Daniel, M. Schiegg, D. Nguyen-Tuong, S. Schaal, M. Toussaint, S. Trimpe. Probabilistic Recurrent State-Space Models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018 (cited on pages 198, 201).

[202] F. Solowjow, A. Mehrjou, B. Schölkopf, S. Trimpe. Efficient Encoding of Dynamical Systems through Local Approximations. In *Proceedings of the 57th IEEE International Conference on Decision and Control (CDC)*, 2018 (cited on page 198).

[586] D. Baumann, J.-J. Zhu, G. Martius, S. Trimpe. Deep Reinforcement Learning for Event-Triggered Control. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 943–950, 2018 (cited on pages 174, 198, 200).

[625] G. S. Lima, W. M. Bessa, S. Trimpe. Depth Control of Underwater Robots using Sliding Modes and Gaussian Process Regression. In *Proceeding of the 15th Latin American Robotics Symposium*, 2018.

[626] R. Soloperto, M. A. Müller, S. Trimpe, F. Allgöwer. Learning-Based Robust Model Predictive Control with State-Dependent Uncertainty. In *Proceedings of the IFAC Conference on Nonlinear Model Predictive Control (NMPC)*, 2018 (cited on page 198).

[627] A. von Rohr, S. Trimpe, A. Marco, P. Fischer, S. Palagi. Gait learning for soft microrobots controlled by light fields. In *Proceeding of the International Conference on Intelligent Robots and Systems (IROS)*, 2018 (cited on pages 198, 199).

[628] D. Baumann, F. Mager, H. Singh, M. Zimmerling, S. Trimpe. Evaluating Low-Power Wireless Cyber-Physical Systems. In *1st Workshop on Benchmarking Cyber-Physical Networks and Systems (CPSBench)*, pages 13–18, 2018 (cited on pages 198, 200).

[629] F. Solowjow, D. Baumann, J. Garcke, S. Trimpe. Event-triggered Learning for Resource-efficient Networked Control. In *Proceedings of the American Control Conference (ACC)*, pages 6506–6512, 2018 (cited on pages 198, 200).

## 2017

[23] A. Doerr, C. Daniel, D. Nguyen-Tuong, A. Marco, S. Schaal, M. Toussaint, S. Trimpe. Optimizing Long-term Predictions for Model-based Policy Search. In *Proceedings of 1st Annual Conference on Robot Learning (CoRL)*. Vol. 78, pages 227–238, 2017 (cited on pages 198, 201).

[24] A. Marco, P. Hennig, S. Schaal, S. Trimpe. On the Design of LQR Kernels for Efficient Controller Learning. In *Proceedings of the 56th IEEE Annual Conference on Decision and Control (CDC)*, pages 5193–5200, 2017 (cited on pages 198, 199, 252).

[31] A. Marco, F. Berkenkamp, P. Hennig, A. P. Schoellig, A. Krause, S. Schaal, S. Trimpe. Virtual vs. Real: Trading Off Simulations and Physical Experiments in Reinforcement Learning with Bayesian Optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1557–1563, 2017 (cited on pages 198, 199, 252).

[32] A. Doerr, D. Nguyen-Tuong, A. Marco, S. Schaal, S. Trimpe. Model-Based Policy Search for Automatic Tuning of Multivariate PID Controllers. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 5295–5301, 2017 (cited on page 198).

## 2016

[43] S. Trimpe. Predictive and Self Triggering for Event-based State Estimation. In *Proceedings of the 55th IEEE Conference on Decision and Control (CDC)*, pages 3098–3105, 2016 (cited on pages 198, 200).

[54] S. Ebner, S. Trimpe. Communication Rate Analysis for Event-based State Estimation. In *Proceedings of the 13th International Workshop on Discrete Event Systems*, 2016.

[58] A. Marco, P. Hennig, J. Bohg, S. Schaal, S. Trimpe. Automatic LQR Tuning Based on Gaussian Process Global Optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 270–277, 2016 (cited on pages 198, 199, 252).

[59] J. Issac, M. Wüthrich, C. Garcia Cifuentes, J. Bohg, S. Trimpe, S. Schaal. Depth-based Object Tracking Using a Robust Gaussian Filter. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) 2016*, 2016.

[62] M. Wüthrich, C. Garcia Cifuentes, S. Trimpe, F. Meier, J. Bohg, J. Issac, S. Schaal. Robust Gaussian Filtering using a Pseudo Measurement. In *Proceedings of the American Control Conference (ACC)*, 2016 (cited on page 198).

## 21.11 Locomotion in Biorobotic and Somatic Systems

### 21.11.1 Journal Articles

#### 2018

[630] J. Nirody, J. Jinn, T. Libby, T. Lee, A. Jusufi, D. Hu, R. Full. Geckos Race across Water using Multiple Mechanisms. *Current Biology*, 2018.

## 21.12 Movement Generation and Control

### 21.12.1 Journal Articles

#### 2018

[331] M. A. Borno, L. Righetti, M. J. Black, S. L. Delp, E. Fiume, J. Romero. Robust Physics-based Motion Retargeting with Realistic Body Shapes. *Computer Graphics Forum* **37**: 6:1–12, 2018 (cited on pages 77, 84).

[631] L. Righetti, Q. Pham, R. Madhavan, R. Chatila. Lethal Autonomous Weapon Systems [Ethical, Legal, and Societal Issues]. *IEEE Robotics & Automation Magazine* **25** (1): 123–126, 2018.

[632] Q. Pham, R. Madhavan, L. Righetti, W. Smart, R. Chatila. The Impact of Robotics and Automation on Working Conditions and Employment [Ethical, Legal, and Societal Issues]. *IEEE Robotics and Automation Magazine* **25** (2): 126–128, 2018.

[633] J. Viereck, J. Kozolinsky, A. Herzog, L. Righetti. Learning a Structured Neural Network Policy for a Hopping Task. *IEEE Robotics and Automation Letters* **3** (4): 4092–4099, 2018.

#### 2016

[9] A. Herzog, N. Rotella, S. Mason, F. Grimminger, S. Schaal, L. Righetti. Momentum Control with Hierarchical Inverse Dynamics on a Torque-Controlled Humanoid. *Autonomous Robots* **40** (3): 473–491, 2016.

### 21.12.2 Conference Papers

#### 2018

[14] N. Rotella, S. Schaal, L. Righetti. Unsupervised Contact Learning for Humanoid Estimation and Control. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 411–417, 2018.

[15] A. Gams, S. Mason, A. Ude, S. Schaal, L. Righetti. Learning Task-Specific Dynamics to Improve Whole-Body Control. In *Hua*, 2018.

[16] S. Mason, N. Rotella, S. Schaal, L. Righetti. An MPC Walking Framework With External Contact Forces. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1785–1790, 2018.

[17] B. Ponton, A. Herzog, A. Del Prete, S. Schaal, L. Righetti. On Time Optimization of Centroidal Momentum Dynamics. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5776–5782, 2018.

#### 2017

[634] M. Khadiv, S. A. A. Moosavian, A. Herzog, L. Righetti. Pattern Generation for Walking on Slippery Terrains. In *2017 5th International Conference on Robotics and Mechatronics (ICROM)*, 2017.

#### 2016

[34] S. Mason, N. Rotella, S. Schaal, L. Righetti. Balancing and Walking Using Full Dynamics LQR Control With Contact Constraints. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 63–68, 2016.

[35] B. Ponton, A. Herzog, S. Schaal, L. Righetti. A Convex Model of Momentum Dynamics for Multi-Contact Motion Generation. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots Humanoids*, pages 842–849, 2016.

[36] B. Ponton, S. Schaal, L. Righetti. On the Effects of Measurement Uncertainty in Optimal Control of Contact Interactions. In *The 12th International Workshop on the Algorithmic Foundations of Robotics WAFR*, 2016.

[37] N. Rotella, S. Mason, S. Schaal, L. Righetti. Inertial Sensor-Based Humanoid Joint State Estimation. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1825–1831, 2016.

[38] A. Herzog, S. Schaal, L. Righetti. Structured contact force optimization for kino-dynamic motion generation. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2703–2710, 2016.

[39] M. Khadiv, S. Kleff, A. Herzog, S. A. A. Moosavian, S. Schaal, L. Righetti. Stepping Stabilization Using a Combination of DCM Tracking and Step Adjustment. In *2016 4th International Conference on Robotics and Mechatronics (ICROM)*, pages 130–135, 2016.

[635] M. Khadiv, A. Herzog, S. A. A. Moosavian, L. Righetti. Step Timing Adjustement: a Step toward Generating Robust Gaits. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 35–42, 2016.

### 21.12.3 Book Chapters

#### 2017

[636] L. Righetti, A. Herzog. Momentum-Centered Control of Contact Interactions. In *Geometric and Numerical Foundations of Movements*. Vol. 117. Springer Tracts in Advanced Robotics. Springer, Cham, pages 339–359, 2017.

## 21.13 Micro, Nano, and Molecular Systems

### 21.13.1 Journal Articles

#### 2018

[637] J. Sachs, K. I. Morozov, O. Kenneth, T. Qiu, N. Segreto, P. Fischer, A. M. Leshansky. Role of symmetry in driven propulsion at low Reynolds number. *Phys. Rev. E* **98**: 063105, 6 Dec. 2018 (cited on pages 227, 228).

[638] C. M. Maier, M. A. Huergo, S. Milosevic, C. Pernpeintner, M. Li, D. P. Singh, D. Walker, P. Fischer, J. Feldmann, T. Lohmüller. Optical and Thermophoretic Control of Janus Nanopen Injection into Living Cells. *Nano Letters*, 2018. eprint: https://doi.org/10.1021/acs.nanolett.8b03885.

[639] Z. Wu, J. Troll, H. H. Jeong, Q. Wei, M. Stang, F. Ziemssen, Z. Wang, M. Dong, S. Schnichels, T. Qiu, P. Fischer. A swarm of slippery micropropellers penetrates the vitreous body of the eye. *Science Advances* **4** (11): eaat4388, 2018. eprint: http://advances.sciencemag.org/content/4/11/eaat4388.full.pdf (cited on pages 220, 226).

[640] T. Yu, P. Chuphal, S. Thakur, S. Y. Reigh, D. P. Singh, P. Fischer. Chemical micromotors self-assemble and self-propel by spontaneous symmetry breaking. *Chem. Comm.* **54**: 11933–11936, 2018.

[641] K. Melde, T. Qiu, P. Fischer. Fast spatial scanning of 3D ultrasound fields via thermography. *Applied Physics Letters* **113** (13): 133503, 2018. eprint: https://doi.org/10.1063/1.5046834 (cited on page 222).

[642] J.-P. Günther, M. Börsch, P. Fischer. Diffusion Measurements of Swimming Enzymes with Fluorescence Correlation Spectroscopy. *Accounts of Chemical Research* **51** (9): 1911–1920, 2018. eprint: https://doi.org/10.1021/acs.accounts.8b00276 (cited on pages 220, 229).

[643] E. Suraniti, P. Merzeau, J. Roche, S. Gounel, A. G. Mark, P. Fischer, N. Mano, A. Kuhn. Uphill production of dihydrogen by enzymatic oxidation of glucose without an external energy source. *Nature Communications* **9** (1): 3229, 2018 (cited on page 220).

[644] M. N. Popescu, W. E. Uspal, C. Bechinger, P. Fischer. Chemotaxis of Active Janus Nanoparticles. *Nano Letters* **18** (9): 5345–5349, 2018. eprint: https://doi.org/10.1021/acs.nanolett.8b02572 (cited on pages 220, 229).

[645] D. Paria, H. H. Jeong, V. Vadakkumbatt, P. Deshpande, P. Fischer, A. Ghosh, A. Ghosh. Graphene-silver hybrid devices for sensitive photodetection in the ultraviolet. *Nanoscale* **10**: 7685–7693, 2018.

[646] P. Fischer. Nanoparticles on the move for medicine. *Physics World Focus on Nanotechnology*: 26028, 2018.

[647] S. Palagi, P. Fischer. Bioinspired microrobots. *Nature Reviews Materials* **3**: 113–124, 2018.

[648] D. P. Singh, W. E. Uspal, M. N. Popescu, L. G. Wilson, P. Fischer. Photogravitactic Microswimmers. *Adv. Func. Mat.* **28**: 1706660, 2018. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adfm.201706660 (cited on pages 220, 229).

[649] M. Matuschek, D. P. Singh, J. Hyeon-Ho, M. Nesterov, T. Weiss, P. Fischer, F. Neubrech, L. Na Liu. Chiral Plasmonic Hydrogen Sensors. *Small* **14** (7): 1702990, 2018. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/smll.201702990.

[650] K. Melde, E. Choi, Z. Wu, S. Palagi, T. Qiu, P. Fischer. Acoustic Fabrication via the Assembly and Fusion of Particles. *Advanced Materials* **30** (3): 1704507, 2018 (cited on pages 220, 222).

#### 2017

[651] T. Qiu, F. Adams, S. Palagi, K. Melde, A. G. Mark, U. Wetterauer, A. Miernik, P. Fischer. Wireless Acoustic-Surface Actuators for Miniaturized Endoscopes. *ACS Applied Materials & Interfaces* **9** (49): 42536–42543, 2017. eprint: http://dx.doi.org/10.1021/acsami.7b12755 (cited on pages 220, 224).

[652] U. Choudhury, A. Straube, P. Fischer, J. Gibbs, F. Höfling. Active colloidal propulsion over a crystalline surface. *New Journal of Physics* **19**: 125010, 2017.

[653] H. H. Jeong, M. Alarcon-Correa, A. G. Mark, K. Son, T.-C. Lee, P. Fischer. Corrosion-Protected Hybrid Nanoparticles. *Advanced Science* **4** (12): 1700234, 2017.

[654] T. Qiu, S. Palagi, A. G. Mark, K. Melde, F. Adams, P. Fischer. Active Acoustic Surfaces Enable the Propulsion of a Wireless Robot. *Advanced Materials Interfaces* **4** (21): 1700933, 2017 (cited on page 224).

[655] D. P. Singh, U. Choudhury, P. Fischer, A. G. Mark. Non-Equilibrium Assembly of Light-Activated Colloidal Mixtures. *Advanced Materials* **29**. 32: 1701328, 2017 (cited on pages 220, 229).

[656] A. G. Mark, E. Suraniti, J. Roche, H. Richter, A. Kuhn, N. Mano, P. Fischer. On-chip enzymatic microbiofuel cell-powered integrated circuits. *Lab on a Chip* **17** (10). Recent HOT Article: 1761–1768, 2017.

[657] D. C. Hooper, A. G. Mark, C. Kuppe, J. T. Collins, P. Fischer, V. K. Valev. Strong Rotational Anisotropies Affect Nonlinear Chiral Metamaterials. *Advanced Materials* **29** (13): 1605110, 2017.

[658] P. J. Vach, D. (Walker) Schamel, P. Fischer, P. Fratzl, D. Faivre. Pattern formation and collective effects in populations of magnetic microswimmers. *J. of Phys. D: Appl. Phys.* **50** (11): 11LT03, 2017.

[659] F. Adams, T. Qiu, A. Mark, B. Fritz, L. Kramer, D. Schlager, U. Wetterauer, A. Miernik, P. Fischer. Soft 3D-Printed Phantom of the Human Kidney with Collecting System. *Ann. of Biomed. Eng.* **45** (4): 963–972, 2017 (cited on page 224).

[660] J. T. Kim, U. Choudhury, J. Hyeon-Ho, P. Fischer. Nanodiamonds That Swim. *Advanced Materials* **29** (30). Back Cover: 1701024, 2017.

## 2016

[661] M. Alarcon-Correa, D. Walker (Schamel), T. Qiu, P. Fischer. Nanomotors. *Eur. Phys. J.-Special Topics* **225** (11-12): 2241–2254, 2016.

[662] T. Qiu, S. Palagi, A. G. Mark, K. Melde, F. Adams, P. Fischer. Wireless actuation with functional acoustic surfaces. *Appl. Phys. Lett.* **109** (19). APL Editor's pick. APL News.: 191602, 2016 (cited on page 224).

[663] K. Melde, A. G. Mark, T. Qiu, P. Fischer. Holograms for acoustics. *Nature* **537**. Max Planck press release, Nature News & Views, Nature Video.: 518–522, 2016 (cited on pages 220, 221).

[664] H. H. Jeong, A. G. Mark, P. Fischer. Magnesium plasmonics for UV applications and chiral sensing. *Chem. Comm.* **52** (82): 12179–12182, 2016.

[665] P. Garbacz, P. Fischer, S. Kraemer. A loop-gap resonator for chirality-sensitive nuclear magneto-electric resonance (NMER). *J. Chem. Phys.* **145** (10): 104201, 2016.

[666] D. Walker (Schamel), D. P. Singh, P. Fischer. Capture of 2D Microparticle Arrays via a UV-Triggered Thiol-yne "Click" Reaction. *Advanced Materials* **28** (44): 9846–9850, 2016.

[667] H. H. Jeong, A. G. Mark, T.-C. Lee, M. Alarcon-Correa, S. Eslami, T. Qiu, J. G. Gibbs, P. Fischer. Active Nanorheology with Plasmonics. *Nano Letters* **16** (8): 4887–4894, 2016.

[668] H. H. Jeong, A. G. Mark, M. Alarcon-Correa, I. Kim, P. Oswald, T. C. Lee, P. Fischer. Dispersion and shape engineered plasmonic nanosensors. *Nature Communications* **7**: 11331, 2016 (cited on page 220).

[669] S. Palagi, A. G. Mark, S. Y. Reigh, K. Melde, T. Qiu, H. Zeng, C. Parmeggiani, D. Martella, A. Sanchez-Castillo, N. Kapernaum, F. Giesselmann, D. S. Wiersma, E. Lauga, P. Fischer. Structured light enables biomimetic swimming and versatile locomotion of photoresponsive soft microrobots. *Nature Materials* **15** (6). Max Planck press release, Nature News & Views.: 647–653, 2016 (cited on page 220).

[670] A. M. Maier, C. Weig, P. Oswald, E. Frey, P. Fischer, T. Liedl. Magnetic Propulsion of Microswimmers with DNA-Based Flagellar Bundles. *Nano Letters* **16** (2): 906–910, 2016.

## 21.13.2 Conference Papers

### 2018

[627] A. von Rohr, S. Trimpe, A. Marco, P. Fischer, S. Palagi. Gait learning for soft microrobots controlled by light fields. In *Proceeding of the International Conference on Intelligent Robots and Systems (IROS)*, 2018 (cited on pages 198, 199).

[671] T. Qiu, S. Palagi, J. Sachs, P. Fischer. Soft Miniaturized Linear Actuators Wirelessly Powered by Rotating Permanent Magnets. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3595–3600, 2018.

### 2017

[672] F. Adams, T. Qiu, A. G. Mark, K. Melde, S. Palagi, A. Miernik, P. Fischer. Wireless micro-robots for endoscopic applications in urology. In *Eur Urol Suppl*. Vol. 16. 3, e1914, 2017.

[673] S. Palagi, A. G. Mark, K. Melde, T. Qiu, H. Zeng, C. Parmeggiani, D. Martella, D. S. Wiersma, P. Fischer. Locomotion of light-driven soft microrobots through a hydrogel via local melting. In *2017 International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS)*, pages 1–5, 2017.

### 2016

[674] G. cerretti, D. Martella, H. Zeng, C. Parmeggiani, S. Palagi, A. G. Mark, K. Melde, T. Qiu, P. Fischer, D. Wiersma. Towards Photo-Induced Swimming: Actuation of Liquid Crystalline Elastomer in Water. In *Proc. of SPIE 9738*, Laser 3D Manufacturing III, 97380T, 2016.

[675] T. Qiu, S. Palagi, A. G. Mark, K. Melde, P. Fischer. Wireless actuator based on ultrasonic bubble streaming. In *2016 International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS)*, pages 1–5, 2016.

[676] S. Palagi, A. G. Mark, K. Melde, H. Zeng, C. Parmeggiani, D. Martella, D. S. Wiersma, P. Fischer. Soft continuous microrobots with multiple intrinsic degrees of freedom. In *2016 International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS)*, pages 1–5, 2016.

[677] A. G. Mark, S. Palagi, T. Qiu, P. Fischer. Auxetic Metamaterial Simplifies Soft Robot Design. In *2016 IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 4951–4956, 2016.

## 21.13.3 Book Chapters

### 2018

[678] Nanoscale robotic agents in biological fluids and tissues. In. *The Encyclopedia of Medical Robotics*. Vol. 2. World Scientific. Chap. 2, pages 19–42, 2018.

### 2017

[679] S. Palagi, D. (Walker) Schamel, T. Qiu, P. Fischer. Chapter 8 - Micro- and nanorobots in Newtonian and biological viscoelastic fluids. In *Microbiorobotics*. Second edition. Micro and Nano Technologies. Elsevier, Boston. Chap. 8, pages 133–162, 2017.

## 21.14 Probabilistic Numerics

### 21.14.1 Books

#### 2017

[67] A. Gretton, P. Hennig, C. Rasmussen, B. Schölkopf. New Directions for Learning with Kernels and Gaussian Processes (Dagstuhl Seminar 16481). Vol. 6. 11, pages 142–167, 2017.

### 21.14.2 Journal Articles

#### 2018

[680] M. Schober, S. Särkkä, P. Hennig. A probabilistic model for the numerical solution of initial value problems. *Statistics and Computing*, 2018 (cited on page 250).

[681] N. Wahl, P. Hennig, H.-P. Wieser, M. Bangert. Analytical incorporation of fractionation effects in probabilistic treatment planning for intensity-modulated proton therapy. *Medical Physics*, 2018 (cited on page 251).

#### 2017

[682] H.-P. Wieser, P. Hennig, N. Wahl, M. Bangert. Analytical probabilistic modeling of RBE-weighted dose for ion therapy. *Physics in Medicine and Biology (PMB)* **62** (23): 8959–8982, 2017 (cited on page 251).

[683] M. Mahsereci, P. Hennig. Probabilistic Line Searches for Stochastic Optimization. *Journal of Machine Learning Research* **18** (119): 1–59, 2017. eprint: jmlrpaper17-049 (cited on page 249).

[684] N. Wahl, P. Hennig, H. P. Wieser, M. Bangert. Efficiency of analytical and sampling-based uncertainty propagation in intensity-modulated proton therapy. *Physics in Medicine & Biology* **62** (14): 5790–5807, 2017 (cited on page 251).

#### 2016

[161] E. D. Klenske, P. Hennig. Dual Control for Approximate Bayesian Reinforcement Learning. *Journal of Machine Learning Research* **17** (127): 1–30, 2016.

[167] E. D. Klenske, M. Zeilinger, B. Schölkopf, P. Hennig. Gaussian Process-Based Predictive Control for Periodic Error Correction. *IEEE Transactions on Control Systems Technology* **24** (1): 110–121, 2016.

### 21.14.3 Conference Papers

#### 2018

[685] T. Kajihara, M. Kanagawa, K. Yamazaki, K. Fukumizu. Kernel Recursive ABC: Point Estimation with Intractable Likelihood. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2405–2414, 2018.

[686] L. Balles, P. Hennig. Dissecting Adam: The Sign, Magnitude and Variance of Stochastic Gradients. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018 (cited on page 249).

#### 2017

[24] A. Marco, P. Hennig, S. Schaal, S. Trimpe. On the Design of LQR Kernels for Efficient Controller Learning. In *Proceedings of the 56th IEEE Annual Conference on Decision and Control (CDC)*, pages 5193–5200, 2017 (cited on pages 198, 199, 252).

[31] A. Marco, F. Berkenkamp, P. Hennig, A. P. Schoellig, A. Krause, S. Schaal, S. Trimpe. Virtual vs. Real: Trading Off Simulations and Physical Experiments in Reinforcement Learning with Bayesian Optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1557–1563, 2017 (cited on pages 198, 199, 252).

[257] M. Schober, A. Adam, O. Yair, S. Mazor, S. Nowozin. Dynamic Time-of-Flight. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 170–179, 2017.

[390] L. Balles, J. Romero, P. Hennig. Coupling Adaptive Batch Sizes with Learning Rates. In *Proceedings Conference on Uncertainty in Artificial Intelligence (UAI) 2017*, pages 410–419, 2017 (cited on page 249).

[687] A. Klein, S. Falkner, S. Bartels, P. Hennig, F. Hutter. Fast Bayesian Optimization of Machine Learning Hyperparameters on Large Datasets. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS 2017)*. Vol. 54. Proceedings of Machine Learning Research, pages 528–536, 2017.

### 2016

[58] A. Marco, P. Hennig, J. Bohg, S. Schaal, S. Trimpe. Automatic LQR Tuning Based on Gaussian Process Global Optimization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 270–277, 2016 (cited on pages 198, 199, 252).

[286] E. D. Klenske, P. Hennig, B. Schölkopf, M. N. Zeilinger. Approximate dual control maintaining the value of information with an application to building control. In *European Control Conference (ECC)*, pages 800–806, 2016.

[293] J. González, Z. Dai, P. Hennig, N. Lawrence. Batch Bayesian Optimization via Local Penalization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 51. JMLR Workshop and Conference Proceedings, pages 648–657, 2016.

[294] H. Kersting, P. Hennig. Active Uncertainty Calibration in Bayesian ODE Solvers. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 309–318, 2016.

[296] S. Bartels, P. Hennig. Probabilistic Approximate Least-Squares. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 51. JMLR Workshop and Conference Proceedings, pages 676–684, 2016 (cited on page 247).

## 21.14.4 Theses

### PhD Theses

[312] M. Schober. Probabilistic Ordinary Differential Equation Solvers — Theory and Applications. PhD thesis. Eberhard Karls Universität Tübingen, Germany, 2018.

[313] M. Mahsereci. Probabilistic Approaches to Stochastic Optimization. PhD thesis. Eberhard Karls Universität Tübingen, Germany, 2018 (cited on page 249).

[316] E. D. Klenske. Nonparametric Disturbance Correction and Nonlinear Dual Control. PhD thesis. ETH Zurich, 2017.

## 21.15 Rationality Enhancement

### 21.15.1 Conference Papers

#### 2018

[688] F. Callaway, S. Gul, P. Krueger, T. L. Griffiths, F. Lieder. Learning to select computations. In *Uncertainty in Artificial Intelligence: Proceedings of the Thirty-Fourth Conference*, 2018 (cited on page 258).

[689] F. Lieder, F. Callaway, P. M. Krueger, P. Das, T. L. Griffiths, S. Gul. Discovering and Teaching Optimal Planning Strategies. In *The 14th biannual conference of the German Society for Cognitive Science, GK*, 2018 (cited on page 257).

[690] S. Gul, P. M. Krueger, F. Callaway, T. L. Griffiths, F. Lieder. Discovering Rational Heuristics for Risky Choice. In *The 14th biannual conference of the German Society for Cognitive Science, GK*, 2018 (cited on pages 257, 258).

## 21.16 Statistical Learning Theory

### 21.16.1 Journal Articles

#### 2018

[84] N. Shah*, B. Tabibian*, K. Muandet, I. Guyon, U. von Luxburg. Design and Analysis of the NIPS 2016 Review Process. *Journal of Machine Learning Research* **19** (49). *equal contribution: 1–34, 2018.

#### 2017

[691] M. Kleindessner, U. von Luxburg. Lens Depth Function and k-Relative Neighborhood Graph: Versatile Tools for Ordinal Data Analysis. *Journal of Machine Learning Research (JMLR)* **18**, 2017 (cited on page 262).

### 21.16.2 Conference Papers

#### 2018

[692] D. Ghoshdastidar, U. von Luxburg. Practical Methods for Graph Two-Sample Testing. In *Proceedings Neural Information Processing Systems*, 2018 (cited on page 263).

[693] L. Vankadara, U. von Luxburg. Measures of distortion for machine learning. In *Proceedings Neural Information Processing Systems*, 2018 (cited on page 262).

[694] C. Tang, D. Garreau, U. von Luxburg. When do random forests fail? In *Proceedings Neural Information Processing Systems*, 2018 (cited on page 262).

[695] S. Haghiri, D. Garreau, U. von Luxburg. Comparison-Based Random Forests. In *International Conference on Machine learning (ICML)*, 2018 (cited on page 262).

#### 2017

[696] M. Kleindessner, U. von Luxburg. Kernel functions based on triplet comparisons. In *Proceedings Neural Information Processing Systems*, 2017 (cited on page 262).

[697] S. Haghiri, D. Ghoshdastidar, U. von Luxburg. Comparison-based nearest neighbor search. In *Artificial Intelligence and Statistics*, 2017 (cited on page 262).

[698] D. Ghoshdastidar, M. Gutzeit, A. Carpentier, U. von Luxburg. Two-sample tests for large random graphs using network statistics. In *Conference on Computational Learning Theory (COLT)*, 2017 (cited on page 263).

# 21.17 Materials CSF

## 21.17.1 Journal Articles

### 2018

[428] U. T. Sanli, H. Ceylan, I. Bykova, M. Weigand, M. Sitti, G. Schütz, K. Keskinbora. 3D Nanoprinted Plastic Kinoform X-Ray Optics. *Advanced Materials* **30** (36): 1802503, 2018.

[699] A. Davydok, T. W. Cornelius, Z. Ren, C. Leclere, G. Chahine, T. Schülli, F. Lauraux, G. Richter, O. Thomas. *Quantum Beam Science* **2** (4), 2018.

[700] U.-T. Sanli, H. Ceylan, C. Jiao, M. Baluktsian, C. Grevent, K. Hahn, Y. Wang, V. Srot, G. Richter, I. Bykova, M. Weigand, M. Sitti, G. Schütz, K. Keskinbora. New concepts for 3d optics in x-ray microscopy. eng. *Microscopy and Microanalysis* **24** (Suppl 2): 288–289, Aug. 2018.

[701] S. Bagheri, N. Strohfeldt, M. Ubl, A. Berrier, M. Merker, G. Richter, M. Siegel, H. Giessen. Niobium as Alternative Material for Refractory and Active Plasmonics. *Acs Photonics* **5** (8): 3298–3304, 2018.

[702] Z. Ren, T. W. Cornelius, C. Leclere, A. Davydok, J. S. Micha, O. Robach, G. Richter, O. Thomas. Three-point bending behavior of a Au nanowire studied by in-situ Laue micro-diffraction. *Journal of Applied Physics* **124** (18), 2018.

[703] V. Kozhevnikov, A. M. Valente-Feliciano, P. J. Curran, G. Richter, A. Volodin, A. Suter, S. J. Bending, C. Van Haesendonck. Equilibrium Properties of the Mixed State in Superconducting Niobium in a Transverse Magnetic Field: Experiment and Theoretical Model. *Journal of Superconductivity and Novel Magnetism* **31** (11): 3433–3444, 2018.

[704] J. Shin, T. W. Cornelius, S. Labat, F. Lauraux, M. I. Richard, G. Richter, N. P. Blanchard, D. S. Gianola, O. Thomas. In situ Bragg coherent X-ray diffraction during tensile testing of an individual Au nanowire. *Journal of Applied Crystallography* **51**: 781–788, 2018.

[705] Z. Ren, T. W. Cornelius, C. Leclere, A. Davydok, J. S. Mieha, O. Robach, G. Richter, O. Thomas. Plasticity in inhomogeneously strained Au nanowires studied by Laue microdiffraction. *Mrs Advances* **3** (39): 2331–2339, 2018.

### 2017

[482] A. V. Singh, S. Baylan, B.-W. Park, G. Richter, M. Sitti. Hydrophobic pinning with copper nanowhiskers leads to bactericidal properties. *PloS One* **12** (4): e0175428, 2017.

[706] G. Cheng, S. Yin, T.-H. Chang, G. Richter, H. Gao, Y. Zhu. Anomalous Tensile Detwinning in Twinned Nanowires. eng. *Physical Review Letters* **119** (25), 256101, Dec. 22, 2017.

[707] B. N. Jaya, S. Goto, G. Richter, C. Kirchlechner, G. Dehm. Fracture behavior of nanostructured heavily cold drawn pearlitic steel wires before and after annealing. eng. *Materials Science and Engineering A: Structural Materials Properties Microstructure and Processing* **707**: 164–171, Nov. 7, 2017.

[708] D. Appy, M. Wallingford, D. Jing, R. Ott, M. C. Tringides, G. Richter, P. A. Thiel. Thermally activated diffusion of copper into amorphous carbon. eng. *Journal of Vacuum Science and Technology. A* **35** (6), 061401, Nov. 2017.

[709] R. Ramachandramoorthy, Y. Wang, A. Aghaei, G. Richter, W. Cai, H. D. Espinosa. Reliability of Single Crystal Silver Nanowire-Based Systems: Stress Assisted Instabilities. eng. *ACS Nano* **11** (5): 4768–4776, May 23, 2017.

[710] D. Song, A. H. Tavabi, Z.-A. Li, A. Kovacs, J. Rusz, W. Huang, G. Richter, R. E. Dunin-Borkowsk, J. Zhu. An in-plane magnetic chiral dichroism approach for measurement of intrinsic magnetic signals using transmitted electrons. eng. *Nature Communications* **8**, 15348, 2017.

[711] V. Kozhevnikov, A.-M. Valente-Feliciano, P. J. Curran, A. Suter, A. H. Liu, G. Richter, E. Morenzoni, S. J. Bending, C. Van Haesendonck. Equilibrium properties of superconducting niobium at high magnetic fields: A possible existence of a filamentary state in type-II superconductors. eng. *Physical Review B* **95** (17), 174509, May 1, 2017.

[712] J. Lipecka, J. Janczak-Rusch, M. Lewandowska, M. Andrzejcuk, G. Richter, L. P. Jeurgens. Thermal stability of Al-Si12 nano-alloys confined between AlN layers in a nanomultilayer configuration. *Scripta Materialia* **130**: 210–213, Mar. 15, 2017.

### 2016

[713] C. Leclere, T. Cornelius, Z. Ren, O. Robach, J.-S. Micha, A. Davydok, O. Ulrich, G. Richter, O. Thomas. KB scanning of X-ray beam for Laue microdiffraction on accelero-phobic samples: application to in situ mechanically loaded nanowires. eng. *Journal of Synchrotron Radiation* **23**: 1395–1400, 2016.

[714] N. Gazit, L. Klinger, G. Richter, E. Rabkin. Formation of hollow gold-silver nanoparticles through the surface diffusion induced bulk intermixing. eng. *Acta Materialia* **117**: 188–196, Sept. 15, 2016.

[715] C. Ensslen, C. Brandl, G. Richter, R. Schwaiger, O. Kraft. Notch insensitive strength and ductility in gold nanowires. eng. *Acta Materialia* **108**: 317–324, Apr. 15, 2016.

[716] S. Guo, N. Talebi, W. Sigle, R. Vogelgesang, G. Richter, M. Esmann, S. F. Becker, C. Lienau, P. A. van Aken. Reflection and Phase Matching in Plasmonic Gold Tapers. eng. *Nano Letters* **16** (10): 6137–6144, 2016.

[717] J. P. Liebig, M. Goken, G. Richter, M. Mackovic, T. Przybilla, E. Spiecker, O. N. Pierron, B. Merle. A flexible method for the preparation of thin film samples for in situ TEM characterization combining shadow-FIB milling and electron beam-assisted etching. *Ultramicroscopy* **171**: 82–88, 2016.

[718] S. Fernandez, M. Richard, D. Flötotto, G. Richter, D. Mandula, M. Elzo Aizarna, V. Favre-Nicolin, D. Burghammer, T. Schülli, O. Thomas. X-ray nanodiffraction in forward scattering and Bragg geometry of a single isolated Ag-Au nanowire. eng. *Thin Solid Films* **617**: 9–13, 2016.

## 21.18 Scientific Computing CSF

### 21.18.1 Journal Articles

### 2018

[719] M. Cóndor, T. Rüberg, C. Borau, J. Piles, J. García-Aznar. A web-based application for automated quantification of chemical gradients induced in microfluidic devices. *Computers in Biology and Medicine* **95**: 118–128, 2018.

## 21.19 Software Workshop CSF

### 21.19.1 Journal Articles

### 2018

[720] R. Iaconi, O. De Marco, J.-C. Passy, J. Staff. The effect of binding energy and resolution in simulations of the common envelope binary interaction. *Monthly Notices of the Royal Astronomical Society* **477**: 2349–2365, June 2018.

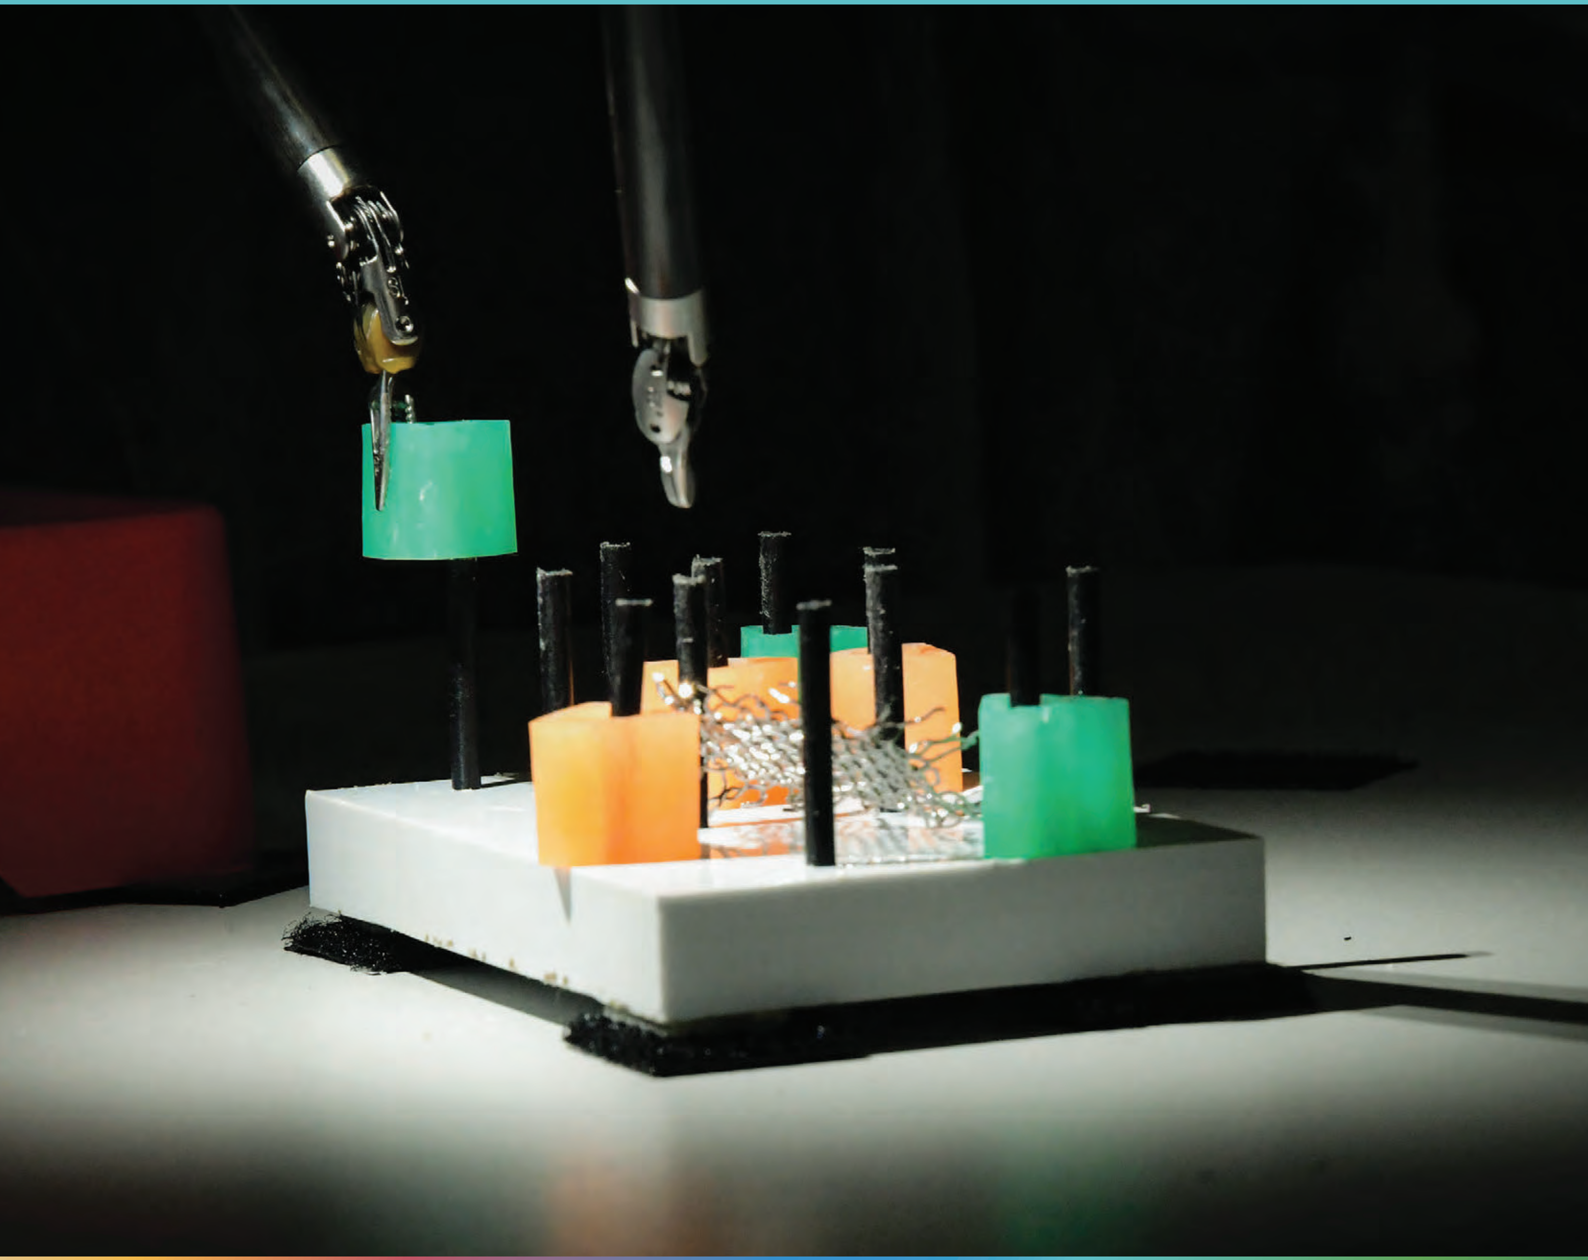

Max Planck Institute for Intelligent Systems

Heisenbergstrasse 3
70569 Stuttgart

Max-Planck-Ring 4
72076 Tübingen

https://is.mpg.de