Thirteenth Biennial Status Report

April 2015 – February 2017

Contents

1 The Institute **1**

1.1 Organization and Staff . . . 1
1.2 Scientific Vision and Strategic Goals . . . 2
1.3 Long-Term Achievements and Impact . . . 3
1.4 Highlights 2015–2017 and New Research Directions . . . 4
1.5 Career Mentoring . . . 5
1.6 Collaborations and Strategic Partnerships . . . 6
1.7 Results 2015–2017 . . . 7
1.8 Outlook Beyond 2017 . . . 9

2 The Research Units **10**

2.1 The Algorithms and Complexity Group (D1) . . . 10
2.2 The Computer Vision and Multimodal Computing Group (D2) . . . 15
2.3 The Computational Biology and Applied Algorithmics Group (D3) . . . 20
2.4 The Computer Graphics Group (D4) . . . 25
2.5 The Databases and Information Systems Group (D5) . . . 30
2.6 The Automation of Logic Group (RG1) . . . 35

1 The Institute

1.1 Organization and Staff

Directors and Departments: The Max Planck Institute for Informatics was established in 1990, with Kurt Mehlhorn as founding director. Harald Ganzinger was appointed shortly afterwards; he passed away in 2004. Currently, the institute has five directors (with appointment periods given in parentheses): Kurt Mehlhorn (1990–2019) heading the Department on Algorithms and Complexity (D1), Bernt Schiele (2010–2035) heading the Department on Computer Vision and Multimodal Computing (D2), Thomas Lengauer (2001–2018) heading the Department on Computational Biology and Applied Algorithmics, Hans-Peter Seidel (1999–2026) heading the Department on Computer Graphics (D4), Gerhard Weikum (2003-2023) heading the Department on Databases and Information Systems. In addition to these five departments, the institute has one permanent independent Research Group on Automation of Logic, headed by Christoph Weidenbach. In total, the institute currently has 177 scientists, out of which 117 are doctoral students and 60 have a doctoral degree.

Senior Researchers: The institute has four scientific ranks: directors, senior researchers with tenure, senior researchers, and researchers. Senior researchers with tenure and senior researchers roughly correspond to tenured associate professor and non-tenured assistant professor in the North-American system. The process for appointing senior researchers is similar to faculty appointment procedures, and involves reference letters from international top researchers. The appointment committee comprises the directors of the institute and a faculty member of Saarland University. Currently the institute has, in addition to its directors, 22 senior researchers.[1]

Senior researcher positions are not tenure-track, in general. We expect our senior researchers to become professors, ultimately full professors, or leading researchers in industry. A strong indicator for the success of the model is the high number of faculty positions that our alumni reached at universities all over the world (see Section 1.5).

Figure 1.1 shows the organizational structure and the research areas of the institute. Each department pursues a number of research areas and each area has its coordinator. The coordinating scientists are senior researchers or postdoctoral researchers with strong potential for becoming senior researchers.

Doctoral Students: As of March 2017, there are 117 doctoral students who are supervised by members of the institute, including 28 women (about 24%). 82 of the 117 students are non-German. The institute does not grant degrees. In this regard, we closely collaborate with Saarland University. Members of the institute teach courses at the University and

[1] Zeynep Akata, Bjoern Andres, Klaus Berberich, Jasmin Blanchette, Andreas Bulling, Piotr Didyk, Mario Fritz, Christian Ikenmeyer, Olga Kalinina, Andreas Karrenbauer, Christoph Lenzen, Tobias Marschall, Pauli Miettinen, Karol Myszkowski (tenured), Nico Pfeifer, Michael Sagraloff, Marcel Schulz, Jürgen Steimle, Thomas Sturm, Christian Theobalt (tenured), Jilles Vreeken, Christoph Weidenbach (tenured)

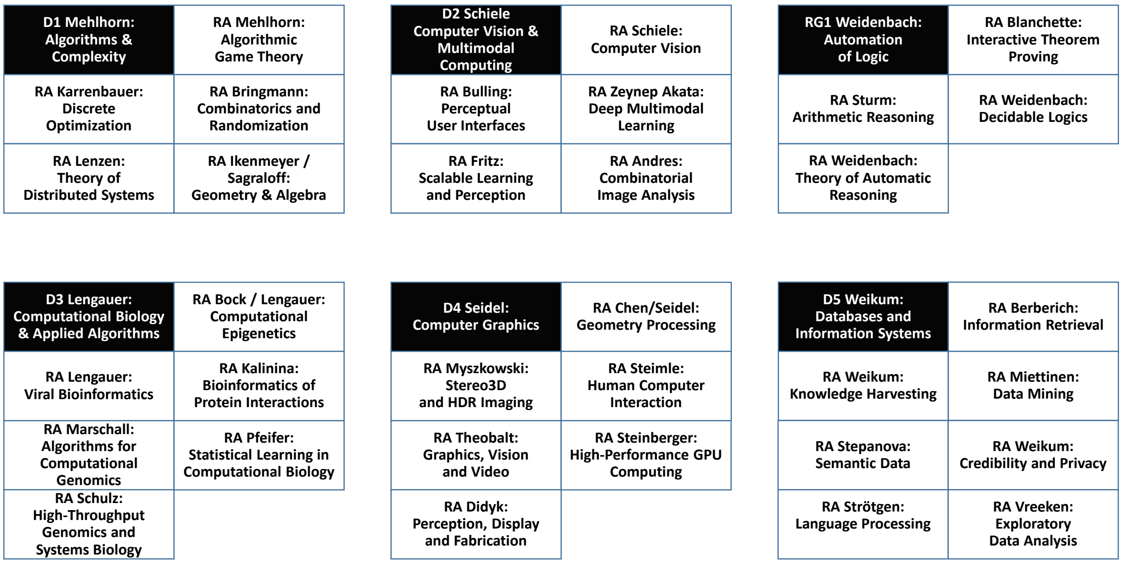

Figure 1.1: Research Areas of the Institute (as of March 2017)

supervise students at all levels. Upon appointment to Senior Researcher, the Computer Science Department of Saarland University decides whether to grant the Senior Researcher the right to supervise doctoral students. Currently, this right has been granted to all senior researchers of the institute.

Joint Administration and Technical Support: The institute shares the IT support group, administration, library, and facility management with the Max Planck Institute for Software Systems. The current head counts fore both institutes are 54 full-time employees including 20 IT support people.

1.2 Scientific Vision and Strategic Goals

Our central research theme is algorithms for multimodal computing. Algorithms have always been the main focus of the institute. They are the core of what makes computer systems useful and productive. They influence every aspect of our daily lives and are the basis for industrial change. Throughout the last decade, major parts of our research effort have focused on multimodal computing. The grand challenge is to understand, search, and organize large, distributed, noisy, incomplete, and diverse information in a robust, efficient, and intelligent manner. Our research ranges from foundations (algorithms and complexity, automation of logic) to a variety of multimodal domains (computer graphics and vision, geometric computation, intelligent information systems, computational biology). The overarching mission of the institute is to be one of the world's top players and strategic trendsetters on this theme of multimodal computing.

Most of the major advances in computer science have come through the combination of new theoretical insights and application-oriented experimental validation, driven by outstanding researchers. Our goal is, thus, to have impact through i) *research and publications*, ii) *software and services* enabled by our research, and iii) *people* alike. In the following, Section 1.3 presents our achievements regarding the first two dimensions, and Section 1.5 discusses our track record with respect to the third dimension.

1.3 Long-Term Achievements and Impact

Over the last twenty years, the institute has pursued a number of high-risk high-gain endeavors, starting with foundational science and ultimately making great practical impact. In the following we outline the highlights of the institute's scientific achievements.

- Kurt Mehlhorn and his group initiated and advanced the field of *Algorithm Engineering*: setting the trend and making ground-breaking contributions towards reconciling the rigorous design and complexity analysis of algorithms with the development of practically viable software libraries. Highlight results with huge impact are the software libraries LEDA and CGAL, where advanced algorithms have been coded with guarantees on their correctness and run-time properties. These libraries are widely used all over the world; LEDA was part of the software that Celera used for sequencing the human genome in the early 2000s.
- Bernt Schiele and his research group have been addressing the challenging problem of *human pose estimation* for about a decade. While initial work addressed pose estimation for single people, recent work has expanded to multi-person pose estimation in video sequences. Essential to the success of the proposed models are, on the one hand, strong part appearance models and, on the other hand, the representation of human pose priors. Our most advanced model is trained in an end-to-end fashion recovering the full human body configurations conditioned on the location of individual people. We introduced the MPI-INF Human Pose Dataset at CVPR'14 which has become the standard benchmark for human pose estimation.
- Thomas Lengauer was one of the founders of the field of computational biology in Europe in the early nineties. He and his group have performed seminal work in the fields of protein structure prediction, molecular docking, viral resistance analysis and computational epigenetics. The comprises methodical innovation, the development of widely used software and innovative analysis of biological data. The work has had significant impact beyond science, e. g., for the treatment of HIV patients.
- Hans-Peter Seidel and his group developed groundbreaking results on level-of-detail and multiresolution modeling with special emphasis on the development of powerful new metaphors for editing and shape deformation. Recent research on markerless performance capture has pioneered methods for the reconstruction of detailed dynamic 3D models of humans in arbitrary loose-fitting clothing, even for complex and rapid moves, in outdoor scenes, and with only two cameras. The group has been one of the pioneers and a driving force in perception-based graphics.
- Christoph Weidenbach and his group have been investigating foundations of automated reasoning since the early 90's. Starting from the development of sound and complete calculi

and their implementations, the group is now driving towards calculi that automatically discover and exploit structures.

- Gerhard Weikum and his team pioneered the theme of *Knowledge Harvesting*: automatically building comprehensive knowledge bases from Internet contents. This work provided the blueprint for industrial-strength knowledge graphs that are key assets for search engines, question answering and text analytics (at Google, Microsoft, etc.). The knowledge base YAGO was used by IBM Watson when it won the Jeopardy quiz show in 2011.

High-Risk Research: When the above long-term projects started, they were far from the mainstream in their scientific communities; most were considered elusive and some even characterized as useless. Needless to say, not all of our bold endeavors worked out. The sections on the five departments give examples. Nevertheless, most of this work resulted in novel insights (sometimes about what is, fundamentally or practically, non-viable) and often in influential publications.

1.4 Highlights 2015–2017 and New Research Directions

The following are selected highlights from the scientific results that the institute has achieved in the last two years.

- Kurt Mehlhorn has set himself the following goals for the next two years: a strongly polynomial algorithms for market equilibria in Arrow-Debreu markets, progress on the dynamic optimality conjecture, and a model for the network formation capabilities of the slime-mould Physarum.
- Bernt Schiele's group has been working on tracking multiple people for a number of years. In joint work with Bjoern Andres published at CVPR'15 and CVPR'17 multi person tracking has been formulated as a graph decomposition problem leveraging strong person detectors as unary costs and re-identification classifiers as pairwise costs. This work as recently won the Multi-Object Tracking Challenge at ECCV'16.
- Thomas Lengauer's group has contributed substantially to the results of the International Human Epigenome Consortium. On the topic of viral infections the group is pioneering a new approach to the interactive composition of drug combination therapies for HIV patients. Furthermore, there has been important progress on the analysis and use of broadly neutralizing antibodies. On a more fundamental side, the group has revisited the thermodynamical basis of molecular interactions and opened the path to a data-driven analysis of the molecular basis of resistance development.
- Hans-Peter Seidel's group developed powerful algorithms for shape interpolation. Markerless performance capture now succeeds in very general settings, with a low number of cameras, and the developed methods have opened up new applications towards real-time facial reenactment and real-time hand tracking for computer interaction. New perceptual models lead to significantly improved processing for stereoscopic displays. iSkin, a novel class of skin-worn sensors, opens up new interaction metaphors in mobile computing.
- Christoph Weidenbach's group has discovered a new criterion characterizing complexity for classes of decidable first-order logic fragments (LICS'2016, LICS'2017).

- In the context of the ERC Synergy Grant imPACT, Gerhard Weikum's group addressed the issue of trust in online information, and how it can be assessed even for evolving communities with anonymous posts (ICDM2015, CIKM2015, KDD2016). This was further extended by analyzing the credibility of textual statements in news and social media (CIKM2016, WWW2017). On the privacy dimension, we have devised a "non-standard" model for assessing a user's risk of being targeted by exposing sensitive information in online communities (SIGIR2016, in collaboration with MPI for Software Systems).

New Research Directions: The overriding theme of multimodal computing will be continued. However, we expect a leap in research directions by the ongoing blending of virtual worlds with the real world. This trend involves several grand challenges on which we will embark. As digital contents and output from all kinds of sensors keeps exploding, we need to move from data to understanding situations and anticipating user behavior. We need to support immersive interactions across all modalities, considering visual signals like facial expressions, gestures and body language, in conjunction with language, contextual knowledge and social interactions. Finally, with the rapid advance of machine learning and data-driven algorithmic decision making, we need to better understand how to make computer behavior comprehensible. These challenges are the motivation for our foundational research on algorithms, visual computing and knowledge discovery in the coming years.

1.5 Career Mentoring

The institute has a strong track record on educating and mentoring young researchers, at both the doctoral student level and the postdoctoral level.

Doctoral Student Training: Since the institute was established in 1990, a total number of 278 doctoral students graduated. These include 46 women, and 110 non-Germans. Among these students, 13 have been awarded the Otto Hahn Medal of the Max Planck Society, Andrey Rybalchenko (2007) and Fabian Suchanek (2012) have received the Otto Hahn Award for the best dissertations in the Chemistry, Physics, Technology Section of the Max Planck Society. This award includes a 5-year scholarship for an independent research group, and only one award per year is given (since 2007). Karl Bringmann won the EATCS Distinguished Dissertation Award in 2015.

Young Scientist Career Advancement: A unique strength of the institute is its successful fostering of young scientists. These mentoring and career support efforts are most pronounced at the level of our senior researchers. Since the notion of senior researchers was explicitly introduced in 2007, a total of 68 young scientists held such positions. 51 of them have meanwhile left the institute. Out of these, 50 have accepted a tenured or tenure-track position offer at universities or university-like research organizations. Several alumni (i. e., people who graduated here or spent at least two years at the institute and left) and current senior researchers of the institute have won prestigious awards: Leibniz Prizes[2], ERC Grants[3],

[2] Susanne Albers 2008, Leif Kobbelt 2014, Peter Sanders 2012

[3] Susanne Albers 2016, Jasmin Christian Blanchette 2016, Niko Beerenwinkel 2014, Christoph Bock 2015, Jürgen Gall 2015, Bastian Goldlücke 2013, Leif Kobbelt 2013, Christoph Lenzen 2016, Marcus Magnor 2010, Thomas Neumann 2016, Rafal Mantiuk 2016, Antti Oulasvirta 2014, Bodo Rosenhahn 2011, Thomas Sauerwald 2015, Jürgen Steimle 2016, Christian Theobalt 2013

and other honors.

Support for Women: The fraction of women in our institute currently is 24% for doctoral students and 13% at the postdoctoral level (including senior researchers). To increase the representation of women in our field, we have established the Lise Meitner Fellowship for outstanding female scientists at the postdoctoral level. So far, these two-year fellowships have been awarded to 7 women.[4] A large number of female alumni of the institute have become professors.[5]

1.6 Collaborations and Strategic Partnerships

An overriding goal that our institute has been contributing to since its beginning in 1990, is to establish Saarbrücken as one of the world's premier sites in computer science. Indeed, over the last three decades the site as a whole has a unique track record, as exemplified by 4 ACM Fellows, 6 Leibniz Prizes (the highest scientific honor in Germany) and a total of 11 ERC Grants (at all levels).

Saarland Informatics Campus: Our most important partner is the Computer Science Department of Saarland University. We have a long-standing tradition of teamwork and joint engagement in both research and teaching. The are numerous collaborations with faculty members of the university, and with researchers of various institutes on campus, including colleagues from the Department for Computational Linguistics, the Department for Biology and the Medical School of the university.

Cluster of Excellence on Multimodal Computing and Interaction (MMCI): The MMCI cluster was established by the German Science Foundation (DFG) in the framework of the German Excellence Initiative in 2007 and successfully renewed in 2012, with funding until 2018. The institute is a key contributor to the success of the cluster through its scientific contributions. All directors of the institute are principal investigators of the cluster, and the scientific coordinator of the cluster is Hans-Peter Seidel.

Max Planck Center for Visual Computing and Communication (MPC-VCC): The center was established jointly by MPG and Stanford University with support from the German Ministry of Research and Education (BMBF) in 2003. Following a successful evaluation in 2012, MPC-VCC has been extended until 2020. Selected young scientists have the opportunity to work two years at Stanford as visiting assistant professor, and then continue

[4] Anke van Zuylen 2010-2012, Anna Adamaszek 2012-2014, Zeynep Akata 2014-2016, Daria Stepanova 2015-2017, Hang Zhou 2015-2017, Shida Beigpour 2016-2018, Qianru Sun 2016-2018

[5] Anna Adamaszek (Univ. Copenhagen, Denmark), Zeynep Akata (Univ. Amsterdam, Netherlands), Susanne Albers (TU Munich), Iris Antes (TU Munich), Hannah Bast (Univ. Freiburg), Carola Dörr (CNRS, France), Panagiota Fatourou (Univ. Ionnina, Greece), Lilia Georgieva (Heriot-Watt Univ., UK), Katja Hose (Aalborg Univ., Denmark), Georgiana Ifrim (UC Dublin, Ireland), Mouna Kacimi (Univ. Bozen-Bolzano, Italy), Ruxandra Lasowski (Univ. of Applied Sciences Furtwangen), Petra Mutzel (Univ. Dortmund), Alice McHardy (Univ. Düsseldorf), Nicole Megow (TU Munich), Ndapa Nakashole (UC San Diego, USA), Marina Papatriantafilou (Univ. Gothenburg, Sweden), Ruzica Piskac (Yale Univ., USA), Nicoleta Preda (Univ. Versailles, France), Maya Ramanath (IIT Delhi, India), Ina Schäfer (TU Braunschweig), Renate Schmidt (Univ. Manchester, UK), Viorica Sofronie-Stokkermans (Univ. Koblenz-Landau), Kavitha Telikepalli (Tata Institute, India), Yafang Wang (Shandong Univ., China), Nicola Wolpert (Univ. of Applied Sciences Stuttgart), Shanshan Zhang (Nanjing U of Science and Technology, China), Anke van Zuylen (College of William & Mary, USA)

as leaders of independent research groups at our institute. The center is jointly directed by Hans-Peter Seidel (MPI-INF) and Bernd Girod and Leo Guibas (Stanford).

Center for Bioinformatics: This center has been founded in 2000 through a grant by DFG. Its main partners are three faculties of Saarland University (Informatics, Biology, Medicine) and several research institutes on campus including our institute. The center is the hub for interdisciplinary research centered around computational biology. Thomas Lengauer is its speaker.

Indo Max Planck Center for Computer Science (IMPECS): This center fosters collaboration between top Indian universities, MPI-INF and MPI-SWS, and the universities in Kaiserslautern and Saarbrücken. It is funded by the BMBF, the Indian Ministry of Science and Technology and the MPG. The center is directed by Kurt Mehlhorn (MPI-INF) and Naveen Garg (IIT Delhi).

IMPRS for Computer Science: The International Max Planck Research School for Computer Science (IMPRS-CS) was established in 2001, as a joint program of the MPI for Informatics and Saarland University. Currently, IMPRS-CS supports 117 doctoral students, including 28 women (about 24%) and 82 non-Germans. In addition, IMPRS-CS currently supports 27 students with fellowships towards Master's degrees.

Saarbrücken Gradudate School for Computer Science: The school was established in 2007 and encompasses all doctoral training in computer science on campus. The school was largely modeled after the IMPRS-CS and adopted many of its elements. IMPRS-CS provides fellowships for doctoral students, within the structural framework of the Graduate School.

Further Projects: The most notable projects with competitive grants are the *ERC Synergy Grant imPACT* and the *German Epigenome Program DEEP*.

The ERC Synergy Grant has been awarded to Michael Backes (Saarland University), Peter Druschel (MPI for Software Systems), Rupak Majumdar (MPI for Software Systems) and Gerhard Weikum (MPI for Informatics) for joint research on the strategic research theme of "Privacy, Accountability, Compliance, and Trust for the Internet of Tomorrow", imPACT for short. The project is funded with a total of 10 Million Euros for the timeframe 2015–2020.

DEEP is a national flagship project and contributes to the International Human Genome Consortium (IHEC), whose goal is to sequence and annotate 100 reference epigenomes. Thomas Lengauer coordinates the data analysis. The project is funded by the German Ministery of Research and Education with a total of 20 Million Euros for the timeframe 2012–2017.

1.7 Results 2015–2017

Publications, Software, Startups: In the two-year timeframe 2015–2017, the institute published more than 800 papers in peer-refereed conferences and journals. Many of these appeared in top-tier venues, with competitive conferences typically accepting only 10 to 20 percent of their submissions. Several publications won best paper awards or best student paper awards (at CHI 15, CVPR 15, MICCAI 15, 2xUIST 15, SSS 16, GI 16, HVEI 16, ICASSP 16, IJCAR 16, PKDD 16, and SMI 16).

Two startups that spun off from our research in 2014 and 2012, Captury and Logic4Business, respectively, are gaining traction in their respective markets. A new startup, Ambiverse, was founded in 2016 with seed-funding from the EXIST program of the German Ministery for

Economy (BMWi). It aims to market the scalable software tools on entity linking and deep competence on knowledge-based language understanding for text analytics.

People: In the two-year timeframe 2015–2017, 47 of our doctoral students graduated. These include 7 women. In the same time period, 21 of our researchers left the institute to take a tenured or tenure-track faculty position[6]. These include 3 women.

Awards: Members of the institute won prestigious awards. The following are the most prominent examples; full lists are in the respective sections of the departments. Kurt Mehlhorn won the Erasmus Medal of Academia Europaea; Bernt Schiele was inducted as an IEEE Fellow; Thomas Lengauer won the Hector Science Prize; Hans-Peter Seidel was honored as an SMA Solid Modeling Pioneer; Gerhard Weikum won the ACM SIGMOD Edgar F. Cood Innovations Award. At the level of Senior Researchers, Christoph Lenzen and Jürgen Steimle won ERC Starting Grants, and Christoph Bock received an ERC Starting Grant and the Overton Award of the International Society for Computational Biology. At the student level, Karl Bringmann received the EATCS Distinguished Dissertation Award, and Marcus Rohrbach won the DAGM MVTec Dissertation Award.

Further honors with considerable visibility include the following. Kurt Mehlhorn was appointed to the Scientific Council of the European Research Council (ERC). Thomas Lengauer was appointed to the Executive Board of the Leopoldina National Academy of Sciences and became President-Elect of the International Society for Computational Biology (ISCB). Christoph Weidenbach has been elected President of the CADE steering committee. Bernt Schiele has been elected member of the Review Panel of the German Science Foundation (DFG) for the field of computer science. Hans-Peter Seidel was appointed to the ERC StG panel for the field of computer science. Gerhard Weikum gave the plenary talk at the annual meeting of the Max Planck Society in 2016. Christian Theobald was selected as one of three young scientists of the Max Planck Society for a round-table meeting with the German Chancellor, Angela Merkel, on Innovation from Foundational Research.

Gender Proportion and Diversity: Our search for new directors gives highest priority to female scientists (see below). At the level of senior researchers, we will expand and intensify our efforts to attract more women. At the postdoc level, the institute offers a distinguished fellowship for women, named Lise Meitner Scholarship, since 2013. The interest in this program has been strongly increasing in terms of both quantity and quality of applicants. In the last two years, we selected four young women as recipients of this fellowship. The best postdocs are often candidates for becoming senior researchers after two years.

Outreach and Visibility: To increase the international visibility of Saarbrücken as a world-class CS hub, we have established an agreement for joint branding: all CS players on campus now use the label “Saarland Informatics Campus” as part of their official addresses. The label will also be used in Google Scholar profiles, academic rating sites, and other PR

[6] Zeynep Akata (U. Amsterdam, Netherlands), Manu Basavaraju (NIT Karnataka, India), Jasmin Christian Blanchette (U. Amsterdam, Netherlands), Xiaohui Bei (Nanyang U., Singapore), Jasmin Christian Blanchette (O. Amsterdam, Netherlands), Parinya Chalermsook (Aalto U., Finland), Matthias FÃ¼gger (ENS Paris-Saclay, France), Jugal Garg (U. Illinois, USA), Martin Hoefer (U. Frankfurt), Arjun Jain (IIT Bombay, India), Thomas Kesselheim (TU Dortmund), Sebastian Krinninger (U. Salzburg, Austria), Erik Jan van Leeuwen (Utrecht U., Netherlands), Belen Masia (U. de Zaragoza, Spain), Christian Richardt (U. Bath, UK), Tobias Ritschel (U. College London, UK), Gaurav Sharma (IIT Kanpur, India), Vinay Setty (Aalborg U., Denmark), Jürgen Steimle (Saarland U., Germany), Thomas Sturm (CNRS Nancy, France), Shanshan Zhang (Nanjing U of Science and Technology, China)

efforts (see http://sic.saarland/).

Within the Max Planck Society, several institutes have set up a joint web site to increase visibility towards young talents (see http://www.cis.mpg.de/).

The institute celebrated its 25th anniversary with a two-day scientific symposium on September 24–25, 2015 (see http://25years.mpi-inf.mpg.de/index.html). The program featured 12 top-notch speakers, including three Turing Award winners and chief scientists of leading high-tech companies.

The institute continues its role as a provider of for the BWINF, “Bundesweit Informatik Fördern”, promoting young computer science talent in Germany.

1.8 Outlook Beyond 2017

New External Member of the Institute: Klaus-Robert Müller (Professor at TU Berlin), an internationally leading scientist on machine learning and big data, has been elected by the Max Planck Society as an external member of our institute. This opens up various forms of collaboration.

Search for Directors: With the upcoming retirement of two directors, the search for new directors is a top priority. The instruments for identifying candidates include several calls for nominations, and a distinguished lecture series with two-day or three-day visits of the institute. The distinguished lecture series, over the last two years, allowed us to obtain an in-depth impression of 25 promising scientists including 15 women. Our search is thematically open, but gives priority to the field of algorithms and theory of computing. The institute is rooted in this area, and the area is central to the institute's research mission. We prioritize female scientists with world-class track record or potential for world-class leadership and young scientists, in the age range of 35 to 50.

Senior Researchers: Currently, three of our senior researchers have tenure. We use different models for appointing tenured senior researchers: tenure-track, promotion, hiring from outside. We expect most of our senior researchers to obtain faculty positions at universities or other research labs. This model has been very successful in fostering the next generation of academic leaders and has contributed considerably to the impact of the institute. We will continue to use a flexible mix of the different models.

2 The Research Units

2.1 The Algorithms and Complexity Group (D1)

History, Group Organization, and Development

The Algorithms and Complexity group (D1) was established in 1990 as one of the two initial groups of the institute. Kurt Mehlhorn leads the group since its foundation. Kurt Mehlhorn is scheduled to retire as director of the department on August 31st, 2019. The search for a new director in on-going. He will step down as director on the day the new director starts and will continue as a scientist after this date.

As of March 1st, 2017, the senior scientists and subgroup coordinators are Karl Bringmann (since March 2017), Christian Ikenmeyer (since March 2017), Andreas Karrenbauer (since January 2013), Christoph Lenzen (since July 2014), and Michael Sagraloff (since October 2005).

Section D1-Personnel in the full report lists the names of current and recent group members and the current positions of the group members that left during the report period. Our alumni continue to get very good positions.[1] Seven group members completed their PhD. Christoph Lenzen was awarded an ERC starting grant, Sebastian Krinninger received the dissertation award of the Austrian Computer Society, Hang Zhou received the Lise Meitner Award, Sandy Heydrich received a Google PhD Fellowship, and Kurt Mehlhorn became Foreign Associate of the US Academy of Sciences.

Vision and Research Areas

The overarching vision of the Saarland Informatics Campus is to turn the region into one of the top ten locations for research in informatics worldwide. The specific vision for D1 is to be known among the peers as one of the first class algorithm groups and as a trendsetter in algorithmics, and to have impact on the research community and society through research results, people, software, and scientific leadership.

[1] Fidaa Abed (Postdoc at TU Muenchen) Rob van Stee (Lecturer, University of Leicester), Ali Pourmiri (Postdoc at Institute for Research in Fundamental Sciences (IPM), Iran), He Sun (Lecturer, University of Bristol), Adrian Neumann (Siemens AG), Eric Berberich (Head of Software Infrastructure Team, Assembrix Ltd., Israel), Manu Basavaraju (Assistant Professor, IIT Bangalore), Kunal Dutta (Postdoc at Inria Sophia, France), Arijit Ghosh (Assistant Professor, IIT Chennai), Joel Rybicki (Postdoc at University of Helsinki), Fahimeh Ramezani (Assistant Professor at University of Isfahan, Iran), Jugal Garg (Assistant Professor, University of Illinois), Parinya Chalermsook (Assistant Professor at Aalto University), Mayank Goswami (Assistant Professor at New York University), Andreas Wiese (Assistant Professor at University Chile), Xiaohui Bei (Assistant Professor at Nanyang University), Marvin Künnemann (Postdoc at University of San Diego), Matthias Függer (CNRS Researcher at LSV, ENS Paris-Saclay) Syamantak Das (Postdoc at TU Muenchen), Sebastian Krinninger (Assistant Professor at Universität Salzburg), Martin Hoefer (Full Professor at Frankfurt University), Erik Jan van Leeuwen (Assistant Professor at University of Utrecht).

About 80% of our effort is theoretical work, and about 20% is experimental and software construction. Our research is organized into five areas. The area coordinators are shown in brackets (Coordinators as of March 1st, 2017, in italics).

- Combinatorics, Computing, and Randomness (*Karl Bringmann*, Erik Jan van Leeuwen, and *Kurt Mehlhorn*).
- Combinatorial Optimization (Parinya Chalermsook, *Andreas Karrenbauer*, and Andreas Wiese)
- Geometry and Algebra (*Christian Ikenmeyer* and *Michael Sagraloff*)
- Algorithmic Game Theory and Online Optimization (Martin Hoefer, Thomas Kesselheim, and *Kurt Mehlhorn*)
- Theory of Distributed and Embedded Systems (*Christoph Lenzen*).

We are involved in the Cluster of Excellence on *Multimodal Computing and Interaction* and in the *Indo Max Planck Center for Computer Science (IMPECS).*

High Impact Research

I will highlight some of our high impact results. Much of this work spans several reporting periods.

Determining the Roots of a Polynomial

Investigators: Michael Sagraloff, Ruben Becker, Eric Berberich, Cornelius Brand, Arno Eigenwillig, Pavel Emeliyanenko, Kurt Mehlhorn, Michael Kerber, Alexander Kobel, Pengming Wang in cooperation with Werner Krandick (Drexel), Sylvain Lazard (INRIA), Fabrice Rouillier (INRIA), Vikram Sharma (IMSc, Chennai).

Determining the roots of a polynomial is one of the key tasks in computer algebra that has applications, for example, in non-linear geometric computing, robotics, and control theory. I concentrate here on the problem of isolating the real roots of a real polynomial; the goal is to compute a set of pairwise disjoint real intervals each containing exactly one real root. Once the roots are isolated, standard methods can refine them to any desired precision.

We have developed an algorithm which on the one hand achieves the best complexity bounds and on the other hand is competitive with the best solvers on all instances. On difficult instances it beats the known solvers by far. The algorithm has recently been integrated into Maple and will soon replace the standard routine for computing roots of polynomials. The paper (Sagraloff, Mehlhorn, JSC ’16) introduces the new algorithm and the paper (Kobel, Rouillier, Sagraloff, ISSAC ’16) describes the implementation and gives experimental data.

We started working on the problem more than a decade ago. We needed efficient algorithms for root isolation for our work on nonlinear geometric computing such as computing arrangement of algebraic curves and surfaces and the existing algorithms and implementations did not satisfy our needs. In theory, the problem was solved at the time. An algorithm due to V. Pan (JSC ’02) was known to be near-optimal (assuming a certain computational model).

However, the method is quite involved and no implementation existed. There still is none. In practice, simpler methods are used whose worst-case behavior is far from optimal and whose excellent empirical behavior was not proven.

In our first papers, we (Eigenwillig, Kettner, Krandick, Mehlhorhn, Schmitt, Wolpert, CASC '05, Eigenwillig, PhD-Thesis '08) improved upon a classical subdivision approach based on Descrartes' Rule of Signs and showed how to make bisection work with approximate arithmetic, how to guarantee that the polynomial is large at bisection points, and how to extend the realm of inputs to polynomials whose coefficients are not integers, but are given through their binary expansion. The algorithm required randomization. In (Mehlhorn, Sagraloff, JSC '11), we presented a deterministic algorithm with the same time complexity. The subdivision trees of the bisection methods can have long non-branching paths (= path of nodes where one child does not branch further). Sagraloff (ISSAC '12) showed how to traverse such paths in time proportional to the logarithm of their length. Finally, Sagraloff and Mehlhorn (JSC '16) combined all the ingredients. The implementation paper (Kobel, Rouillier, Sagraloff, ISSAC '16) carefully evaluates the design space spanned by the theory paper. The analysis of the algorithm requires a sophisticated amortization argument (Krandick, Mehlhorn, JSC '06 and Eigenwillig, Sharma, Yap, ISSAC '06).

We obtained many additional results on the way, for example, algorithms for isolating all complex roots (Becker, Sagraloff, Sharma, Yap, JSC, to appear), algorithms for sparse polynomials (Sagraloff, ISSAC '14), and algorithms for solving polynomial systems (Brand, Sagraloff, ISSAC '16, Emeliyanenko, Sagraloff ISSAC '12).

This line of research certainly required a lot of stamina. It spans over almost a dozen years. Was it also high risk? We published very good results in very good venues continuously on the way and would have stopped the project if results would have dried up. In this sense, the project was not high risk. However, it was not clear until very recently, that the research would lead to a theoretical and practical breakthrough and not just excellent publications. In this sense, the project was high risk.

Fine-Grained Complexity and Algorithms

Investigators: Karl Bringmann and Marvin Künnemann

An emerging field is to study the complexity of problems in P via fine-grained reductions from certain hard core problems, such as SATISFIABILITY. This yields quantitative lower bounds, conditional on hypotheses for the core problems. In cases where such obtained lower bounds do not match the best complexity of a known algorithm, the failed proof attempt may suggest new algorithms. Over the past four years, Karl Bringmann and Marvin Künnemann have co-shaped this field. Most of their results are relative to the *Strong Exponential Time Hypothesis (SETH)*, which essentially states that that classic NP-hard SATISFIABILITY problem has no $O((2-\varepsilon)^n)$ algorithm (Impagliazzo, Paturi, Zane, JCSS '01).

Fréchet Distance: Karl (PhD-thesis '14, FOCS '14) proved one of the first conditional lower bounds for similarity measures of sequences, specifically for the Fréchet distance, a very popular measure of similarity of two curves. Matching its well-known quadratic time algorithm, he shows that the Fréchet distance cannot be computed in truly sub-quadratic time unless the Strong Exponential Time Hypothesis fails.

Already before this work, restricted classes of curves have been studied that attempt to capture realistic input curves and overcome the worst-case quadratic time barrier. The most popular such class is the class of c-packed curves, for which Driemel, Har-Peled, and Wenk (DCG '12) gave a $(1+\varepsilon)$-approximation in time $O(cn/\varepsilon + cn\log n)$. The conditional lower bound from Karl's thesis also extends to this setting; however, it does not give a matching lower bound. Karl and Marvin (ISAAC '15) resolved this issue positively, by giving a faster algorithm with time complexity $O(cn\log^2(1/\varepsilon)/\sqrt{\varepsilon} + cn\log n)$. This dependence on c, n, and ε is optimal in high dimensions apart from lower order factors, unless SETH fails. In fact, the new algorithm was obtained by examining and exploiting properties that prevented a stronger lower bound, thus providing an example showing that conditional lower bounds may also inspire algorithmic improvements.

Longest Common Subsequence: For longest common subsequence (LCS), Karl and Marvin (FOCS '15) prove that, even restricted to binary strings, there are no truly sub-quadratic time algorithms unless the Strong Exponential Time Hypothesis fails. They also generalize the result to edit distance for arbitrary fixed costs of the four operations (deletion in one of the two strings, matching, and substitution), by identifying trivial cases that can be solved in constant time, and proving quadratic-time hardness on binary strings for all other cost choices. This improves and generalizes the known hardness result for edit distance (Backurs, Indyk, STOC '15). Some of these results were independently obtained in (Abboud, Backurs, Vassilevska William, FOCS '15).

A Near-Linear Pseudopolynomial Time Algorithm for Subset Sum. Given a set Z of n positive integers and a target value t, the SUBSET SUM problem asks whether any subset of Z sums to t; it is one of the original NP-complete problems. A textbook pseudopolynomial time algorithm by Bellman from 1957 solves SUBSET SUM in time $O(n\,t)$. This algorithm has been taught in undergraduate algorithms courses for decades.

Karl (SODA '17) presented a simple randomized algorithm running in time $\tilde{O}(n+t)$, ignoring logarithmic factors. This improves upon a classic algorithm and is likely to be near-optimal.

Fine-grained complexity is an emerging field; it receives a lot of attention by many of the best researchers world-wide. And this is exactly where the risk lies. Would Karl and Marvin be fast enough to claim the important results? They were.

Fault-tolerant Distributed Clock-Synchronization in Hardware

Investigators: Christoph Lenzen, Matthias Függer and Attila Kinali in cooperation with Florian Huemer (Technical University Vienna)

Moore's Law has led to an extreme increase of complexity in contemporary electronics. On the downside, the ever-shrinking sizes of transistors in modern chips have also led to a higher probability of failure. A fairly large fraction of these failures are transient in nature, i. e., they appear and disappear during regular operation; generally, they are probabilistic in nature. Even though it is theoretically possible to test for the probability of transient faults during production, doing so is prohibitively expensive and is only done for low-volume and safety critical applications. Furthermore, it significantly reduces the production yield.

Under adversarial conditions with a high probability of transient faults, like radiation in space applications, these faults cannot be avoided at all and need to be dealt with in other ways. The most common way to counter these faults is to triplicate the used circuits and use a, hopefully fault-free, majority vote on the results. This approach is not very resource efficient and different techniques for dealing with the high rate of faults are needed.

In search of novel methods to battle faults with minimal overhead, Christoph and co-workers have been investigating known results and algorithms from distributed computing, and how they can be adapted to implementation directly in hardware instead of using a software layer on top. Many fault-tolerant algorithms in distributed systems either require a common notion of time or a lockstep execution of the algorithm's rounds. In order to support this mode of operation, they have investigated fault-tolerant clock synchronization algorithms that are sufficiently simple to be implemented in hardware. They have chosen the algorithm by Welch and Lynch (I&C '88) and provided a more detailed analysis (Khanchandani, Lenzen, arXiv '16) with special focus on the limitations imposed by hardware implementation, showing that delay and measurement uncertainty, together with frequency uncertainty of the reference clocks of the nodes, are the main limiting factors for synchronization precision. They have also shown how to extend the algorithm to provide simultaneous phase and rate correction, which allows to reduce the rate at which clock corrections are performed without significant loss of precision. Additionally, they have provided a coupling algorithm that adds self-stabilization, relying on less precise, but self-stabilizing synchronization algorithms. This preserves the base algorithm's properties, essentially without regard for the precision of the self-stabilizing subroutine.

To illustrate the viability of the approach in practice, Kinali, Huemer, and Lenzen (ISVLSI '16) have built and measured a demonstrator of the system in FPGA and shown that the mathematical model faithfully predicts the performance of the implemented system. In fact, the performance exceeded expectations due to the fact that the worst-case model was overly pessimistic; however, one should note that they tested a single incarnation of the system only, while the worst-case bounds apply to all such incarnations that obey the system model.

This research is high risk for several reasons: (1) Christoph stepped out of his comfort zone (distributed algorithms and their realization in software) and started to investigate hardware implementations. (2) It was not clear, whether we could hire the missing expertise and find the suitable cooperation partners (we hired the missing expertise in Matthias Függer and Attila Kinali who also had the links to the suitable cooperation partners). (3) Would these results be appreciated by the hardware community? (4) Would it pay off for Christoph's career to commit substantial time to a topic for which it was not clear to what extent it would serve his career? The research has the potential of high impact. A first patent has been filed and Christoph Lenzen was awarded an ERC starting grant (ToRH: A theory of reliable hardware) to pursue this line of research.

2.2 The Computer Vision and Multimodal Computing Group (D2)

Group Overview

The Computer Vision and Multimodal Computing group (D2) was established in 2010 with the appointment of Bernt Schiele. At the time of writing, the group encompasses four research group leaders, three postdocs and 20 PhD students. Among those 22% are female (one research group leader, one postdoc, four PhD students). The research group leaders and senior scientists each have their own PhD-students to conduct research in their respective area: Mario Fritz (appointed 2011), Andreas Bulling (appointed 2013), Bjoern Andres (appointed 2013), and Zeynep Akata (appointed 2014).

Three group members completed their PhD during the reporting period and four more handed in their PhD documents recently. Our alumni get very good offers for faculty positions in academia[2] and research positions in industry[3].

Vision and Research Areas

Understanding sensor information is a fundamental problem in computer science. Scientific challenges cover the entire pipeline from single-sensor processing, over spatial and temporal fusion of multiple and divergent sensor modalities to the complete description of large-scale multimodal sensor streams. At the same time we observe a tremendous increase in both the quantity as well as the diversity of sensor information due to the increasing number of sensors (such as cameras, GPS, or inertial sensors) embedded in a wide variety of digital devices and environments as well as due to the increasing storage of multimodal sensor data (such as surveillance data, personal storage of digital information, multimedia databases, or simply the Internet). While storing and indexing large amounts of sensor data has made tremendous progress, understanding of this multimodal sensor data still lacks far behind. Therefore the long-term goal of D2 is to make progress on how to process, structure, access, and truly understand multi-sensory data both for online use as well as for large-scale databases.

The group currently focuses on two main areas of the broader field, namely computer vision and multimodal sensor processing. In the area of computer vision we address some of the most fundamental problems of image and video understanding such as object class recognition, people detection and tracking, and scene understanding. In the area of multimodal computing we currently focus on perceptual user interfaces. As a cross-cutting theme for both areas we also work in the area of machine learning. It is clear that only advanced machine learning techniques will enable to infer higher-level information from noisy sensor data and enable to deal with the large-scale nature of current and future multimodal databases and sensor-streams.

[2] Zeynep Akata (U Amsterdam, U Bonn, U Birmingham), Wei-Chen Chiu (National Chiao Tung U Taiwan), Mario Fritz (KTH Stockholm), Peter Gehler (U Würzburg), Gaurav Sharma (IIT Kanpur, IIT Delhi), Yusuke Sugano (U Osaka), Shanshan Zhang (Nanjing U)

[3] Mykhaylo Andriluka (Google), Rodrigo Benenson (Google), Maksim Lapin (Amazon), Mateusz Malinowski (DeepMind), Bojan Pepikj (Amazon), Marcus Rohrbach (Facebook), Leonid Pishchulin (Amazon)

Highlights of Longterm and Current Research

Given the fact that the group has been established relatively recently, we will focus on research highlights since around 2010 and also give examples of current research highlights also from the four research group leaders.

Human Pose Estimation *Investigators: Bernt Schiele, Bjoern Andres, Mykhaylo Andriluka, Eldar Insafutdinov*

Human pose estimation and articulated tracking have been in the focus of our work for nearly a decade. Our initial work in this area at CVPR'09 has been centered on part-based models that combine piecewise-trained body part appearance models with a kinematic model prior that places soft constraints on the set of feasible body configurations. The model allowed for efficient inference, but in retrospect has been too restrictive to adequately represent the space of valid body poses. Nonetheless the model has been sufficiently powerful to enable applications in articulated 3D pose tracking, 3D pose estimation, and activity recognition. At CVPR'13 we incorporated non-Gaussian kinematic models for the body-pose while still allowing the approach to be computationally efficient. The key insight has been to make the parameters of the kinematic model dependent on the image content. In ICCV'13 we extended this model with powerful appearance representations that captured appearance at different levels of granularity ranging from individual body joints and groups of body parts to the appearance of the whole body. Jointly these developments significantly boosted performance. To facilitate further progress in the area we introduced the MPII Human Pose Dataset at CVPR'14[4] that since has become the standard benchmark for human pose estimation.

Our recent work is focused on leveraging convolutional neural network (CNN) architectures and on addressing the more challenging scenario of multi-person pose estimation. In a collaboration with Bjoern Andres we proposed the so called Deep-Cut and Deeper-Cut models at ECCV'16 and CVPR'16 which rely on CNNs to generate body part proposals and assembles them into person configurations by solving an instance of a subgraph multi-cut problem. Similarly to our earlier work the pose assembly in DeeperCut relies on a body pose prior. However the model includes connections between all body parts in contrast to the tree-structured models. As in CVPR'13 the pose prior in DeeperCut is conditioned on the image, but relies on a more powerful appearance representation based on CNN features. In our most recent work at CVPR'17 we propose an approach that jointly recovers body poses and articulated motion of multiple people trained in an end-to-end fashion to recover the full human body configuration conditioned on the location of the person.

Multi Person Tracking *Investigators: Bernt Schiele, Mykhaylo Andriluka, Siyu Tang*

Tracking multiple people is another topic we have been investigating since nearly a decade. Our initial tracking-by-detection and detection-by-tracking paper at CVPR'08 and its extension to 3D at CVPR'10 where based on the intermediate representation of tracklets that allowed to detect and track people more robustly than previous approaches. In our IJCV'14 work we proposed an approach that leveraged the fact that occlusions result in very characteristic appearance patterns that help improving detection. In recent work published at CVPR'15 and

[4]http://human-pose.mpi-inf.mpg.de

CVPR'17 we formulate multi person tracking as a graph decomposition problem leveraging strong person detectors as unary costs and re-identification classifiers as pairwise costs. This work recently won the Multi-Object Tracking Challenge at ECCV'16.

Language and Vision *Investigators: Bernt Schiele, Marcus Rohrbach, Anna Rohrbach*

Since around 2010 we have been working at the intersection of visual processing and computational linguistic. At CVPR'10 we exploited semantic relatedness for knowledge transfer. At ICCV'13 we proposed an approach to generate language descriptions for activity videos probabilistically translating visual semantic representations into natural language descriptions. This problem is clearly challenging in general and we have been among the very first to propose an approach that is able to generate sensible descriptions for videos. At ECCV'16 we have proposed an approach for visual grounding where we aim to estimate the compatibility between a language query and multiple visual regions in order to find the most relevant region.

We strongly believe that the current work at the intersection of visual processing and computational linguistics is still in its infancy and that there is still high potential in this area. For example, extracting knowledge from language that can help computer vision tasks such as scene understanding has not yet shown full potential. At the same time, jointly exploiting video and linguistic resources should allow for better language models and understanding.

Scalable Learning and Perception *Investigator: Mario Fritz*

With the advance of new sensor technology and abundant data resources, machines can get a detailed "picture" of the real-world – unlike ever possible before. However, there is a big gap between these raw data sources and the semantic understanding a human may acquire by analyzing the same type of data. Our goal is to narrow and eventually close this gap between machine representations and the rich semantic understanding acquired by humans. Progress in this direction will facilitate a seamless interaction and exchange of information between machines and humans and enable applications in information retrieval, data mining, robotics, natural sciences, human machine collaboration and assisted living.

Therefore, we research novel tasks and evaluation metrics for scene and language comprehension that give us new handles to approach this complex topic. Along these lines, our groups has pioneered at NIPS'14 a *Visual Turing Challenge* that has received broad attention and is now an established research topic. Following up on this thread, we have presented the first Deep Learning approach to this problem at ICCV'15 that has been influential to many follow up approaches based on an end to end learning paradigm. Most recently, we have addressed scene understanding by building models that are predictive of the future fate of a scene. In this vein, we have proposed an approach that is learnt to predict the stability of a scene which in turn can be utilized for successful robotic manipulation (to appear at ICRA'17). In on going work, we have complemented this research with models that intend a much larger time horizon. Our models are the first to perform long-term video extrapolation via a focus on image boundaries and improved spatiotemporal modeling (ArXiv).

Perceptual User Interfaces *Investigator: Andreas Bulling*

Developing human-computer interfaces that fully exploit the information content available in natural human behavior is challenging, particularly in unconstrained daily life settings. The group works at the interface of human-computer interaction, computer vision, as well as wearable and ubiquitous computing. We develop novel ambient and on-body sensing systems as well as computational methods to analyze human behavior automatically. We specifically focus on visual and physical behavior as these modalities are most promising for developing interfaces that offer human-like interactive and social capabilities. We study these systems and methods in the context of specific application domains, most importantly pervasive eye-based human-computer interfaces and computational human behavior analysis.

Major achievements in the reporting period include the large-scale MPIIGaze dataset as well as end-to-end methods for person and device-independent appearance-based gaze estimation. These works represent the current state of the art in this area with the most recent full-face method significantly outperforming existing methods for both 2D and 3D gaze estimation with improvements of up to 14.3% on MPIIGaze and 27.7% on EYEDIAP. In a parallel line of work, we have presented the first fully morphable eye region model that allows us to rapidly synthesize large amounts of highly realistic and perfectly annotated eye region images for training these deep methods. We have also proposed novel learning-by-synthesis and analysis-by-synthesis methods that have further advanced the state of the art in appearance-based gaze estimation (emerging investigator award at ETRA 2016). In the area of human-computer interaction, we have presented novel interaction techniques for calibration-free interaction with smartwatches, eyewear computers, and ambient displays (best paper award at UIST'15) as well as methods and systems for estimating, analyzing, and predicting visual attention on public displays (best paper honorable mention awards at UIST'16 and CHI'16).

Combinatorial Image Analysis *Investigator: Bjoern Andres*

Along with spectacular progress in the field of computer vision toward applications, fundamental and persistent tasks crystallize. The research group *Combinatorial Image Analysis* contributes mathematical abstractions of such persistent tasks in the form of combinatorial optimization problems. We enable applications of these abstractions in the field of computer vision by defining and implementing practical algorithms for solving these problems exactly or approximatively. We analyze the feasible solutions output by these algorithms in terms of application-specific metrics defined by computer vision benchmarks.

During the reporting period, our research has focused on the *minimum cost multicut* and *minimum cost arborescence* problems. We have established these problems and generalizations as suitable and often superior mathematical abstractions of diverse computer vision tasks. A major achievement of our methodological work is a new and improved algorithm for solving and tightening linear programming relaxations of the minimum cost multicut problem (CVPR'17). For instances of the problem that arise from computer vision applications, this algorithm outperforms the best branch-and-cut algorithms implemented in the framework of leading commercial solvers. A major achievement of our applied work are a new and improved model and algorithm for reconstructing vasular networks, a generalization of our PAMI'16 paper, which has won the Medical Image Analysis Best Paper Award 2015 by the MICCAI Society.

Multimodal Deep Learning *Investigator: Zeynep Akata*

As deep learning frameworks have been getting deeper and more expressive, learning representations to visualize and to explain what the network is learning and to generalize the network output to cover previously unseen concepts have become own research topics that are increasingly gaining popularity and interest the in computer vision, natural language processing and machine learning communities. The research group *Multimodal Deep Learning* contributes a ground for all these research areas to interact and feed each other. We cast deep learning as a phenomenon that is multimodal in nature, being effected by different kinds of sensory inputs, such as visual and textual, that complement each other. We study deep learning methods applied to various computer vision problems to shed light on understanding and improving these methods.

Major achievements in the reporting period have been our six CVPR papers on zero-shot learning which contributes a significant deal to this topic that is already well established and our leading role in establishing two research topics related to generative modeling. First, we proposed natural language conditional image generation as a task demonstrated by our ICML'16 and NIPS'16 papers published both as oral presentations (acceptance rates varying from 5% to 1%). Second, proposing a new model for generating visual explanations in the form of natural language demonstrated by our ECCV 2016 paper (acceptance rate 25%).

Publications, Cooperations and Awards

The Biennial Scientific Report contains a detailed report of the publications, cooperations and awards of the reporting period. From the journal publications 12 have been published or accepted at either IJCV or IEEE PAMI. From the conference publications, 54 have been published or accepted at one of the major computer vision conferences (CVPR, ICCV, ECCV). An additional 4 have been published at NIPS or ICML. Another 17 papers have been published at CHI, UIST, Ubicomp, and ETRA.

Marcus Rohrbach received the DAGM MVTec Dissertation Award and Qianru Sun the Lise Meitner Award. Several best paper awards have been awarded to members of the group including a MICCAI best paper award (Bjoern Andres) and a UIST best paper award (Andreas Bulling). Siyu Tang, Mykhayylo Andriluka, Bjoern Andres and Bernt Schiele are the winners of the the multiple object tracking challenge in 2016.

In the reporting period (2015–2017) the group has been cooperating with a wide range of research groups worldwide. Cooperations that have led to joint publications during the reporting period include: UC Berkeley, Stanford U, Harvard U, ETH Zurich, EPFL, KU Leuven, U Birmingham, Lancaster University, TU Munich, LMU Munich and Freiburg U. Cooperations at Saabrücken that have led to joint publications include the machine learning group and the computational linguistic department at Saarland University, as well as the MPI for Software Systems and the computer graphics (D4) and the databases and information systems (D5) groups at MPI for Informatics.

2.3 The Computational Biology and Applied Algorithmics Group (D3)

Group Overview

The department has been founded in August 2001 and has since been directed by Thomas Lengauer. Research in the department is focused exclusively on computational biology.

During the report period, the department has experienced some turnover, both at the PhD student and at the postdoc level. In preparation of the expected retirement of Thomas Lengauer and the consequential dissolution of the department at the end of May 2018 the department has been shrinking in size for about a year. As this text is written the department has 16 full-time scientists (three postdoc, 13 pre-doc) and three support people. Two additional scientists at the Group Leader level who have their main affiliation elsewhere (Nico Pfeifer, Univ. Tübingen, Christoph Bock, CeMM Vienna) are employed part-time.

Within the department, there are three research groups (i) on Computational Epigenetics, directed by Christoph Bock and Thomas Lengauer, (ii) on Statistical Learning in Computational Biology, directed by Senior Researcher Nico Pfeifer and (iii) on Structural Bioinformatics of Protein Interactions, directed by Senior Researcher Olga Kalinina. Two additional research groups are affiliated with the department, namely (iv) on High-throughput Genomics and Systems Biology, directed by Senior Researcher Marcel Schulz (Junior Group Leader in the Excellence Cluster Multimodal Computing and Interaction) and (v) on Algorithms for Computational Genomics, directed by Senior Researcher Tobias Marshall (Junior Professor at the Center for Bioinformatics Saar). This section includes the activities of the first three research groups, while the work of the affiliated groups Marschall and Schulz is reported in separate sections of the extended report. Furthermore, this section only reports on published results. Comments on ongoing unpublished work are provided in the extended part of the report.

The department is striving to have impact by engaging in computational biology research towards the following goals: (i) advancing the understanding of the molecular basis of disease, (ii) improving diagnosis, prognostication and therapy of diseases.

Towards these ends, we perform research which we mostly publish in international journals targeting the computational biology, biology and medicine communities and – to a lesser extent – in conferences or in computer science journals. Furthermore, for selected projects, we go beyond the proof-of-concept stage, providing and maintaining software that is used by a sizable user community worldwide. Examples include the geno2pheno server for the analysis of viral resistance and several epigenetics software packages (e. g. BiQ Analyzer, EpiExplorer, RnBeads).

We perform most of our research in interdisciplinary collaboration with biologists and medical experts, locally, nationally and internationally. Typically, the relevant research consortia are large and patterned around national and international third-party funding projects (e. g. BLUEPRINT, DEEP).

Research Areas and Achievements

Computational Epigenetics (dir. Christoph Bock and Thomas Lengauer)

In the area of epigenomics, our work is primarily characterized by our contributions to national and international epigenome mapping consortia. These are: the German Epigenomics Program DEEP (funded by the German Science Ministry from September 2012 to December 2017), the European BLUEPRINT consortium (funded by EU from October 2007 to September 2016), and the International Human Epigenome Consortium (IHEC), which provides the global umbrella for large-scale epigenome mapping projects. Whereas DEEP concentrates on metabolic and inflammatory diseases, BLUEPRINT concentrates on cells of the blood lineage. In BLUEPRINT, we are a partner contributing methods and data analyses, in DEEP we additionally coordinate the project-wide data analysis effort and supply the software pipelines for high-level data analysis. The reporting period has been characterized by major advances in data generation and data analysis in DEEP and by the successful completion of the BLUEPRINT project. The world-wide efforts culminated in a collective publication of 41 papers on epigenomes, out of which the Department contributed to seven. Our work on epigenomics is in close collaboration with the Chair of epigenomics at Saarland University held by Prof. Jörn Walter who is also Scientific Coordinator of the DEEP project.

We have added to our suite of epigenomics tools the DeepBlue software (Albrecht et al., Nucl Acids Res 2016, Albrecht et al., Bioinformatics 2017), a portal for epigenetic data that greatly simplifies handling large data sets such as generated in the IHEC context.

Furthermore, we provided computational data analysis to several epigenetic data analysis projects including (i) analyzing DNA methylation dynamics of human hematopoietic stem cell differentiation based on low-input and single-cell whole-genome bisulfite sequencing (Farlik et al., CellStemCell 2016), (ii) analyzing and manipulating DNA methylation patterns in leukemia (Amabile et al., Nat Commun 2015), (iii) analyzing the epigenomics changes involved in memory formation of human T cells (Durek et al., Immunity 2016), and (iv) dissecting the process of differentiation of monocytes to macrophages (Wallner et al., Epigenetics Chromatin 2016).

Viral Bioinformatics (dir. Thomas Lengauer)

Bioinformatics supporting diagnosis and therapy of viral infections has long been a major research activity in the Department. For well over a decade the department has been hosting the geno2pheno server which is freely available over the Internet and is used world-wide for interpreting genotypic data on viral drug resistance. The server offers over half a dozen types of analysis for viral resistance phenotypes of HIV, HBV and HCV. The analyses are based either on expert rules relating mutations in the viral genome to resistance phenotypes or on statistical models predicting from viral genotype that have been constructed on large sets of clinical and virologic training data. The server is broadly used in clinical practice, and some analyses have been recommended in the German-Austrian and European guidelines for treating HIV patients.

In the report period, the following enhancements have been provided on the server or are in development. (i) We developed an automatic method for calculating resistance-factor cutoffs – a problem that has evaded a systematic solution for a long time. The cutoffs for the latest version of geno2pheno[resistance] have been estimated using a recently published method (Pironti et al., J Clin Virol 2015). (ii) It has long been known that drug-resistance mutations selected under previous treatment regimens may be archived in several tissues in the body of the patient. We train statistical models for predicting both past drug exposure and current drug resistance from nucleotide sequences of HIV-1 circulating in the blood of the patient. These models are of interest in their own right, but our main purpose for developing them is to use them in a novel tool for prediction of the success of combination drug therapy that we will make available towards the middle of this year. We hope that this tool will significantly lower the bar for using therapy prediction engines rather than just virtual phenotypes in clinical routine. (iii) We have developed a statistical model for predicting the coreceptor usage of HIV-2 (Döring et al., Retrovirology 2016), addressing a topic of substantial clinical significance. (iv) The geno2pheno Web service is currently undergoing a major overhaul. The new system will offer higher integration between the different services for resistance prediction, better user guidance and substantially more functionality for user input to guide the configuration of highly promising drug combination therapies for patients. Finally we provided a thorough analysis of the sensitivity of the predictions of the especially widely used geno2pheno[coreceptor] tool to sequence variations that may be of both technical and biological nature. There is an ongoing discussion whether the testing HIV coreceptor usage with this tool should be done with three sample replicates or with just a single sample. Our analysis shows that a single sample is not only sufficient but is to be preferred, since it prevents false-positive predictions (Pironti et al., JAIDS 2017).

In a separate effort we have been performing basic research on inferring virus phylogenies from patients infected with viruses such as HIV. Such phylogenies are quite different from those inferred from data on traditional species. This is because viruses are typically sampled across evolutionary time and can therefore appear not only at the leaves of a phylogeny but also at interior vertices. We developed a new method that we call family joining (Kalaghatgi et al., Mol Biol Evol 2016) for constructing such trees on the basis of pairwise distances and validated it on synthetic data and on data from HIV transmission events.

Structural Bioinformatics of Protein Interactions (dir. Olga Kalinina)

The prime area of research in this group is the evolution of viral proteins on various scales, investigated by analyzing protein structure. The topics range from theoretical biophysical models of structural changes upon mutation to statistically grounded studies of viral intra-species and inter-species evolution. Additionally, we are interested in understanding the effect of genetic variation associated with human disease on protein interactions. We work on the development of novel methods and the application of existing tools.

On the basic research side, we have embarked on a fundamental reassessment of the thermodynamic energy balance in molecular systems. We have (i) put forth a theory that identifies exactly canceling terms of energy and entropy in molecular systems undergoing changes due to molecular interactions (Ahmad et al., J Chem Theor Comput 2015a), (ii) shown how these canceling terms mask the real contribution of the conformational changes to the binding energy when measured experimentally or computed conventionally (Ahmad et al., J Phys Chem 2016), (iii) introduced the perturbation-divergence formalism for decomposing free energy in a way that exactly accounts for the discriminant information needed to distinguish between the two distributions representing the initial and the final state of the molecular system (Ahmad et al. J Chem Theor Comput 2015b) and, (iv) shown that this formalism yields exact equations for the terms contributing to the binding free energy of the ligand receptor pair (Ahamd et al., J Chem Phys 2017). These theoretical findings do not only fundamentally revisit the way in which enthalpic and entropic contributions to the energy changes in molecular systems undergoing interactions should be handled but also open the door towards the selection of relevant features, i. e. particular atomic interactions or movements that are most informative of energy change in the multidimensional representations of such systems. Our motivation for investigating these issues is to shed more light on the molecular mechanics of resistance development in critical viral proteins. We are in the process of applying the new concepts to this question.

In another study we have investigated amino-acid conservation in interaction interfaces of HIV-1 proteins. Here we are clustering amino acids taking into account both their conservation and their proximity at the protein surface. We find that extremely conserved residues tend to form tight clusters in interaction interfaces and ligand-binding sites. More surprisingly, extremely variable residues have the same tendency, which we consider to be a reflection of the continual "arms race" between the pathogen and the host immune system (Voitenko et al., Bioinformatics 2016).

Finally, we have performed an analysis aimed at mining distant evolutionary relationships among virus proteins on the basis of the similarity of their protein structures. Using sequence- and structure-based analysis we could detect novel relations between proteins in distantly related or seemingly unrelated viruses (Caprari et al., Viruses 2015), including several horizontal gene transfer events between viruses and from their hosts. To address the long-standing problem of high computational cost of protein structure comparison, we are currently extending this approach methodically with a novel graph-based method for mining biologically relevant patterns in protein structures. This research involves a collaboration with Department 5.

On the topic of analyzing genetic variation, we have developed the web-based tool StructMAn (Gress et al., Nucl Acids Res 2016), which facilitates genome-wide analysis of the impact of single-nucleotide polymorphisms on protein structure. The tool analyzes contacts of the corresponding mutated amino acids in multiple related protein complexes, and thus provides insight into interactions of the affected site, even if the corresponding protein structure has not been resolved experimentally.

Statistical Learning in Computational Biology (dir. Nico Pfeifer)

Research in this group is targeting the following methodical issues: (i) integrative analysis of heterogeneous data sets, (ii) improving the interpretability of nonlinear estimators, (iii) efficient learning methods for large data sets. Regarding application areas we are focusing on (i) research related to vaccines for HIV, (ii) individualized treatment of HIV, (iii) cancer and (iv) computational epigenetics.

Regarding HIV vaccines we are focusing on broadly neutralizing antibodies, building on our previous research at Microsoft Research (Carlson et al., Nat Med 2016). In a larger collaboration to which we contributed essential statistical analyses we could show that therapy with a specific broadly neutralizing antibody (3BNC117) is able to enhance the host immune response against HIV-1 (Schoofs et al., Science 2016) and leads to significant delay of viral rebound after treatment interruption (Scheid et al., Nature 2016). In a separate study we could also show the therapeutic effectiveness of another antibody (Caskey et al., Nat Med 2017). Furthermore, we analyzed the evolution of neutralization resistance of HIV to different antibodies by applying our models to a very large database of HIV-1 sequences. Here, we could confirm that there is a global trend towards HIV-1 resistance over time, something that would be very expensive to analyze in an experimental setting.

Regarding integrative analysis we extended a multiple kernel learning framework that allows the use of different dimensionality reduction schemes with several input data matrices. We applied locality preserving projections and added the regularization constraint, thus affording results that are more stable with respect to small perturbations in the data while still being biologically meaningful (Speicher and Pfeifer, Bioinformatics 2015).

Prizes and Awards

Christoph Bock was awarded an ERC Starting Grant in 2016 and the Overton Award 2017 of the International Society for Computational Biology. Matthias Döring received the Best Poster Presentation Prize at the 14th European Meeting on HIV and Hepatitis 2016 in Rome, Italy. Marcel Schulz was nominated Fellow of the Sixth Exploratory Roundtable Conference 2015 in Shanghai, China. Ten travel fellowships were awarded to members of the group. Within the report period, Nico Pfeifer took the position of W3 Professor in the Department of Computer Science of the University of Tübingen.

Projects and Cooperations

Here we list projects with external funds amounting to at least one person-year. **EU:** BLUEPRINT (Epigenetics of the haematopoietic lineage, 2011–2016, Lengauer partner), PREDEMICS (Viral zoonoses, 2011–2017, Lengauer partner), EuResist (EU EEIG on viral resistance with changing funding, currently funding by Viiv (EUCOHIV project), Lengauer partner). **German Science Ministry:** DEEP (German Epigenome Program, 2012–2017, Lengauer partner and coordinator of data analysis), German Center for Infection Disease, Hepatitis Unit (2012–2019, Lengauer partner), i:DSem (Integrative Data Semantics in Systems Medicine, 2016-2021, Pfeifer partner). **German Health Ministry:** Master HIV-HEP (Project on viral resistance, 2013–2017, partner Lengauer). **DFG:** Excellence Cluster Multimodal Computing and Interaction (2007–2018, Lengauer PI, Schulz Junior Group Leader), Regulation in P. tetraurelia (2017–2019, PI Schulz). **Max Planck Society:** Max Planck Project of the Annual Meeting of 2016 (Elucidation of the association of viral tropism with patient parameters, 2016–2018, PI Pfeifer).

2.4 The Computer Graphics Group (D4)

Group Overview

The computer graphics group (D4) was established in 1999 with the appointment of H.-P. Seidel. At the time of writing, the group encompasses two senior researchers with tenure (K. Myszkowski (since 2000) and C. Theobalt (since 2009)), four group leaders and senior researchers without tenure (R. Chen (since 2015), P. Didyk (since 2014), J. Steimle (since 2013) and M. Steinberger (since 2015)), eight postdocs, and 17 PhD students.

Twelve group members completed their PhD during the reporting period and two more handed in their PhD documents recently. Several alumni received offers for postdoc positions[5] and faculty appointments[6]. Since its establishment in 1999 the group graduated more than 60 PhD students, and more than 40 former group members received offers for faculty positions.

Vision and Research Areas

During the last three decades computer graphics established itself as a core discipline within computer science. New and emerging technologies such as digital media, social networks, digital television, digital photography, telecommunication and telepresence, virtual and augmented reality further indicate its potential and pose new challenges in the years to come.

To address these challenges, and in particular to seamlessly blend real and synthetic footage, we have adopted a new and more integrated scientific view of computer graphics as *3D Image Analysis and Synthesis* that takes into account the whole image processing pipeline from scene acquisition to scene reconstruction to scene editing to scene rendering. We also take into account human perception on all levels of the pipeline, and we exploit the abundance of digital visual data to extract powerful priors that can assist us in the various tasks.

Our vision and long term goal are completely immersive visually rich environments with sophisticated scene representations and the highest visual quality, fused seamlessly with the real world. Standard 2D screens are being replaced with HDR displays, stereo and automultiscopic screens, portable and wearable displays. Imaging algorithms with embedded perceptual models ensure that the viewing comfort is maximized. Interaction is intuitive and lightweight.

To address these challenges, our research is currently organized into these research areas:

- Digital Geometry Processing (coordinated by R. Chen and H.-P. Seidel)
- GPU-Programming (coordinated by M. Steinberger)
- User-centric Computational Videography (coordinated by C. Theobalt)
- Reconstructing the Static and Dynamic Real World (coordinated by C. Theobalt)
- Rendering, Illumination, and Stereo3D (coordinated by P. Didyk, K. Myszkowski)
- HDR Imaging and Perception in Graphics (coordinated by K. Myszkowski, P. Didyk)
- Human Computer Interaction (coordinated by Jürgen Steimle)

[5] M. Bachynskyi (Aarhus), D. Casas (Madrid), O. Elek (Prague), A. Elhayek (DFKI), Y. Gryaditskaya (INRIA), P. Kellnhofer (MIT), C. Nguyen (Apple), B. Reinert (Google), H. Rhodin (EPFL), S. Sridhar (Stanford)

[6] A. Jain (IIT Bombay, India, 2016), B. Masia (Univ. Zaragoza, Spain, 2015), T. Ritschel (University College London, 2015), C. Richardt (Univ. Bath, UK, 2016), J. Steimle (Saarland University, 2016)

Contributions and Impact: Long-Term Results

Scene representation and interaction crucially depend on the efficient representation of geometry, and *Digital Geometry Processing* is therefore a core discipline within computer graphics. We have done groundbreaking work on level-of-detail and multiresolution modeling with special emphasis on the development of powerful new metaphors for editing and shape deformation. This includes work on interactive multi-resolution modeling on arbitrary meshes, on feature-sensitive extraction from volume data, on multi-level partition of unity implicits, and on Laplacian surface editing. Our work has shaped the field, and our results have had significant impact in the community (more than 3,000 citations to these four papers alone). Other influential work addressed mesh segmentation, bounded distortion mapping, correspondence analysis, symmetry detection and shape interpolation. Several young researchers (Kobbelt, Blanz, Theisel, Belyaev, Ohtake, Botsch, Wand, Hildebrandt) later received faculty offers from top universities and won prestigious awards. I just mention L. Kobbelt who later won an ERC Advaced Grant (2013), and the DFG Leibniz Prize (2014).

Of course, geometry is just one piece in the puzzle. In order to carry out our intended research program on 3D Image Analysis and Synthesis I felt that we had to build up additional expertise in computer vision and multi-view processing, and we successfully hired M. Magnor (IRG) and B. Rosenhahn as postdocs. Although I originally had some concerns about perhaps spreading out the group too much, in hindsight, this was clearly the right decision, and we were now well-positioned to address *3D Scene Reconstruction, Inverse Rendering, Multiview Stereo, and Optical Motion Capture* in sufficient depth. Based on our work on real-time image synthesis we first developed novel shading-based refinement techniques for the reconstruction of spatial appearance and geometric detail for static objects. J. Kautz, H. Lensch, and M. Goesele were among the first doctoral students to work in this area. With the dissertations of C. Theobalt and B. Goldlücke we then addressed free viewpoint video and eventually started our work on human motion capture. J. Gall and B. Rosenhahn added powerful filtering and optimization strategies, and A. Jain and T. Thormählen developed MovieReshape, a system for tracking and reshaping of humans in videos (1 Mio youtube clicks). Following his appointment to tenured associate professor (2009), C. Theobalt and his group took this research to the next level by developing some of the first approaches to capture dynamic shape models of humans in wide and general apparel. Recent advances enable marker-less skeleton capture in general environments, of humans and animals, in real-time, and with only a few possibly hand-held cameras. Key ingredients are again shading-based refinement and the combination of a fast generative pose estimation algorithm (based on a SoG representation) with a discriminative detection method. The technology is meanwhile commercialized by the spin-off the Captury that was founded by two former PhD students (J. Gall (2009) and N. Hasler (2010)) together with C. Theobalt. Again, several former students and postdocs who spent their formative years at MPI-INF later won faculty positions and prestigious awards: B. Rosenhahn (2009), C. Theobalt (2012), and J. Gall (2014) all won the German Pattern Recognition Award. M. Goesele (2008) and C. Theobalt (2009) both won the Eurographics Young Researcher Award. J. Kautz (2003), H. Lensch (2007), M. Goesele (2009), and I. Ihrke (2012) all won DFG Emmy Noether Fellowships. M. Magnor (2010), B. Rosenhahn (2011), B. Goldlücke (2013), and C. Theobalt (2013) all won ERC Starting Grants.

When we started our research effort on *Perception in Graphics* it was unclear how far this would lead, and there were several question marks: In particular, there were doubts on how to turn our findings on human perception into appropriate computational models of the human visual system, and on how to make the necessary computations efficient enough to be useful in practice. Meanwhile it has become absolutely clear that perceptual models are indispensible to meet the stringent real-time requirements in modern AR applications and at the same time ensure that the perceived quality and viewing comfort is maximized. Over the last decade, K. Myszkowski and his group have done groundbreaking work by developing computationally efficient perceptual models for a surprising range of problems including tone mapping, HDR image processing, HDR video processing (with R. Mantiuk), perceptually-motivated interactive rendering (with T. Ritschel), apparent display resolution enhancement (with P. Didyk) and disparity processing (with P. Kellnhofer, P. Didyk, T. Ritschel and E. Eisemann). Our survey on high dynamic range imaging (2010) has become a point of reference for the whole community (1500 citations), and certain aspects of our work on HDR compression (licensed by Dolby) are starting to appear in modern TV sets. T. Ritschel (2011) and E. Eisemann (2011) both received the Eurographics Young Researcher Award. R. Mantiuk (Cambridge U.) who graduated from the group in 2009 received an ERC Consolidator Grant (2016).

In recent years we also tried to build new bridges between computer graphics and *Human Computer Interaction*. This has led to several interesting joint results, e. g., on the digital fabrication of interactive shape-changing objects with foldable printed electronics, on the performance and ergonomics of touch surfaces (using biomechanical simulation), and on the use of real-time hand tracking in an HCI context. On the HCI side this work has been driven by A. Oulasvirta and J. Steimle. Both have meanwhile become trendsetters and established leaders in their field, and they both received ERC Starting Grants (in 2014 and 2016).

Contributions and Impact: Major Results 2015-2017

We publish our results in top-tier conferences and journals. In the two-year timeframe 2015-2017 the group had 19 papers in ACM TOG (including ACM SIGGRAPH and ACM SIGGRAPH Asia), 22 papers in Computer Graphics Forum (including Eurographics, SGP, EG Rendering Symposium) and IEEE TVCG, 20 papers in major computer vision venues (ICCV, ECCV, CVPR, 3DV, PAMI, IJCV), and 7 papers in major HCI conferences (ACM CHI, UIST). Details are provided in the D4 section of the full report. Here we just mention some highlights:
Shape Modeling We continued our successful work on shape interpolation. We introduced an efficient approximation algorithm for nonlinear physically-based shape interpolation using model reduction (TOG2015) and developed a novel shape interpolation method that blends C^∞ planar harmonic mappings represented in closed-form. The intermediate mappings in the blending are guaranteed to be locally injective C^∞ harmonic mappings, with conformal and isometric distortion bounded by that of the input mappings (SIG2016). We also investigated generalizations of barycentric coordinates (SGP2016) and applications to warping (EG2016).
GPU-Programming We revisited sparse matrices and linear algebra primitives for high performance mesh processing on the GPU (HPEC2016, EG2017) and demonstrated that the approach offers significant potential for speedup as well as unified code generation. We introduced Operator Graphs (SIGAsia2016) to schedule procedural generation on the GPU and achieved significant speed-ups (8-30) over previous work. We also demonstrated how to

control the output of procedural modeling systems using genetic algorithms (EG2017).
Inverse Rendering and Intrinsic Video Space-time coherent estimates of scene illumination and reflectance are a precondition for the fusion of real and synthetic footage (e. g. virtual mirror). Our method for live intrinsic video (SIG2016) decomposes every pixel in a video in its albedo and shading component, without requiring a geometry model of the scene.
Static Scene Reconstruction with Lightweight Sensors Large-scale high-quality static 3D scene reconstruction with a moving depth scanner, such as an RGB-D camera, is the basis for many applications in mixed and augmented reality. We developed BundleFusion (TOG2017) for globally-consistend online reconstruction of large scale scenes with an RGB-D camera. We also demonstrated how to lift the very low resolution of depth scans to the much higher resolution of the RGB camera in an RGB-D setup (SIG2015).
Marker-less Skeletal Motion Capture We advanced our methods such that they succeed in much more general settings, and with a minimal number of cameras. In collaboration with B. Schiele's group we developed one of the first methods that combines body part detection with a CNN and model-based generative pose fitting in a combined pose optimization algorithm (PAMI2017). We proposed a new image formation model for generative model-based 4D reconstruction (ICCV2015, ECCV2016). We developed one of the first approaches for marker-less full body motion capture with egocentrically mounted video cameras on the body (SIGAsia2016). At its core is a new generative pose optimization scheme for fish-eye optics that fits a skeletal model to image cues and body part detections found with a CNN.
Facial Performance Capture and Facial Reenactment We developed a new method for automatic creation of a personalized high-quality 3D face rig (TOG2016). In our Face2Face approach (SIGAsia2015, SIG2016Demo, CVPR2017) we demonstrated robust facial reenactment at real-time frame rates based on just monocular video. Fully automatic reconstruction of expressive lip shapes from just monocular RGB video is presented in (SIGAsia2016).
Novel View Synthesis We improved novel-view synthesis by making use of the correlations observed in 3D models and applying them to new image instances (PAMI2016), and we proposed an intuitive control space for predictable editing of captured BRDF data (SIGAsia2016). Our results on efficient multi-image correspondences for online light field video processing (CGF16) improve novel view synthesis in light fields.
Depth Reconstruction and Processing for Steroscopic Displays Binocular disparity is the main depth cue that makes stereoscopic images appear 3D. However, the range of depth that can be reproduced by this cue is severely limited due to the vergence-accomodation conflict as well as constraints imposed by current displays. We (i) investigated motion parallax and developed a joint disparity-parallax computational model that predicts apparent depth resulting from both cues (SIGAsia2016), and (ii) proposed a new method for stereoscopic depth adjustment that utilizes eye tracking or other gaze prediction information (SIG2016).
Image and Video Enhancement and HDR Imaging We developed a new framework (PAMI2015) that solves a variety of image and video enhancement operations based on task-specific learning from example images. The visual quality of a motion picture is significantly influenced by the choice of the presentation frame rate. We developed a novel filtering technique which enables simulating the whole spectrum of presentation frame rates on a single-frame-rate display (SIG2016). We continued our work on HDR Imaging (EGSR15), and tone mapping (CVMP2015, ICCP2015) and summarized the state of the art in perception-inspired HDR video coding and compression in a survey (CHIPS2020).

Perception-aware Techniques for Digital Fabrication To produce objects with a particular haptic feedback using 3D printing, we addressed the problem of mapping a real-world material to its nearest 3D printable counterpart by constructing a perceptual model for the compliance of nonlinearly elastic objects (SIG2016). We also demonstrated how such models can be applied to complex geometries in an interaction-aware way.
Digital Fabrication of Customized Interactive Surfaces and On-Skin Interaction We developed new approaches for the digital fabrication of touch sensors and active displays of custom shape (UIST2015). These involve printed thin-film electronics (UIST2015) and electronic components that are embedded inside 3D printed objects (CHI2016). We significantly extended our pioneering work on iSkin (CHI2015) and developed conformal skin-worn electronics (SkinMarks) (CHI2017) that enable interactions on challenging body locations.
Hand Gesture-based Computer Input and Tablet Interaction We developed methods for real-time hand tracking from a single depth camera (CVPR2015) and empirical models of hand movement (CHI2015), as well as methods that can be used for interaction with emerging devices like smartwatches for interaction on the move (CHI2017). We also did a comparative study on the ergonomics of touch surfaces using biomechanical simulation (CHI2015).

Awards

H.-P. Seidel received an SMA Solid Modeling Pioneer Recognition for “his technical contributions to the theory of B-splines and his groundbreaking work on level-of-detail and multiresolution modeling with special emphasis on the development of powerful new metaphors for editing and shape deformation” (2016). C. Theobalt was selected as one of three young scientists of the Max Planck Society for a round-table meeting with German Chancellor, Angela Merkel (2017). Christian also received an Oculus Faculty Research Award (2016), and he was elected into the ‘top 40 under 40’ innovators in Germany by the business magazine Capital (2015). J. Steimle received an ERC Starting Grant (2016) and a Google Faculty Research Award (2015). M. Steinberger won the Heinz Zemanek Prize of the Austrian Computer Society (2016). F. Müller received a MINT-Award Informatik 2016. P. Kellnhofer (SPIE’15 and ACM SAP’15), T. Leimkühler (Graphics Interface’16), S. Oberding (ACM UIST’15), J. Steimle (ACM CHI’15, ACM CHI’16), M. Steinberger and R. Zayer (IEEE HPEC’16) all won best paper awards at major international conferences.

Cooperations

The group has been cooperating with a wide range of research groups worldwide. Cooperations that have led to joint publications during the reporting period include Stanford, MIT, Harvard, Brown, NYU, UNC Chapel Hill, Bar-Ilan U., U. Bath, Cambridge U., TU Delft, U. Erlangen, Fraunhofer IIS, TU Graz, Technion Haifa, KU Leuven, UCL London, Weizmann Institute, ETH Zürich. In addition, we also collaborated with some leading industrial research labs, including Google, Microsoft, Oculus, Intel, Nvidia, Disney, Adobe and Technicolor. Formal cooperations exist with Stanford (Max Planck Center for Visual Computing and Communication (MPC-VCC)) and Fraunhofer IIS (Perceptually-aware light field capture, processing and display). There are also several collaborations within the institute and with other groups on campus. For brevity we just mention the Cluster of Excellence MMCI and the Intel VCI.

2.5 The Databases and Information Systems Group (D5)

D5 has been established in 2003. It is headed by Gerhard Weikum and, as of March 2017, consists of 20 doctoral students, 3 senior researchers (Klaus Berberich, Pauli Miettinen and Jilles Vreeken) and 6 other post-doctoral researchers.

Scientific Vision and Research Areas

The group's general objective is to develop methods for *knowledge discovery* in a broad sense: extracting, organizing, searching and exploring various kinds of knowledge from structured, semistructured, textual and multimodal information sources. Our approach towards this goal combines concepts, models, and algorithms from several fields, including database systems, information retrieval, natural language processing, web science and data mining.

Our research has been driven by the overarching vision of automatically constructing, growing and curating large-scale and high-quality knowledge bases from Internet sources. We have spearheaded this research avenue and refer to it as *knowledge harvesting*. To boost search, data analysis and language understanding, machines need to be equipped with comprehensive knowledge about the world's entities, their semantic properties and their relationships among each other. In addition to such encyclopedic facts, machines should also have commonsense knowledge about properties of everyday objects (e. g., coffee being black, liquid and hot) and human activities (e. g., pouring milk in coffee or the visit of an evaluation panel), and should even capture socio-cultural contexts of propositions (e. g., on (who believes) who invented the computer).

For illustration, envision a next-generation social network with "cyber-human" agents as additional participants and augmented-reality facilities for human users. The agents should behave similarly to their human peers. They will understand discourse context, situated language, facial expressions, gestures, emotions and actions of users, and harness rich world knowledge to infer the users' intentions and anticipate their behavior. With these abilities, a software agent could be the host in a talk show with human-like behavior, or join in on a chat about movies. The agent will be able to answer questions about murders, suspects and motives in a crime movie, or discuss public protests and their underlying political controversies when watching a news clip.

Our research is currently organized into seven technical areas:

- Knowledge Harvesting, coordinated by Gerhard Weikum
- Semantic Data, coordinated by Daria Stepanova
- Language Processing, coordinated by Jannik Strötgen
- Information Retrieval, coordinated by Klaus Berberich
- Data Mining, coordinated by Pauli Miettinen
- Credibility and Privacy, coordinated by Gerhard Weikum
- Exploratory Data Analysis, coordinated by Jilles Vreeken

Contributions and Impact: Ten Year Results

Enhancing computers with "machine knowledge" that can power intelligent applications has been a long-standing goal of computer science (going back to AI pioneers like Feigenbaum and Lenat). Major advances in *knowledge harvesting*, with our group as a trendsetter, have made this formerly elusive vision practically viable today.

Our work on knowledge harvesting was motivated by the objective of semantic search, starting in 2004. Later it became the Yago-Naga project, with the first release of the Yago knowledge base (`yago-knowledge.org`) in February 2007. The unique strength of Yago is its high-quality type system with hundred thousands of classes. When IBM Watson won the Jeopardy quiz show, it leveraged Yago's knowledge of fine-grained entity types for semantic type checking.

Impact: Knowledge harvesting has been adopted at big industrial stakeholders, and knowledge bases have become a key asset in semantic search (for queries about entities), question answering, analytics (e. g., aggregating by entities), recommendations and data integration (i. e., to combine heterogeneous datasets). Examples are the knowledge graphs for search engines (e. g., Google, Bing, Baidu) and social networks (e. g., Facebook) as well as domain-specific knowledge bases (e. g., Bloomberg, Mayo Clinic, Walmart). In addition, knowledge bases have found wide use as a distant supervision source for a variety of tasks in natural language processing. Our Yago-Naga project has served as a blueprint for many of these follow-up endeavors. The original Yago paper at WWW 2007 has over 2000 citations, and the Yago2 paper from 2013 already has more than 500 citations.

High-Risk Research: The Yago-Naga project has been a high-risk (and high-gain) endeavor. In the first few years, hardly any of our colleagues believed that large knowledge bases would become viable and make impact. Some dimensions of this research theme did not work out, though. We started with the goal of developing a search engine that understands semantic concepts for computing precise and concise answers to sophisticated queries. This branch of our research led to insights and novel methods, influential publications and advanced prototypes like Bingo!, TopX and Naga, but did not succeed in building a full-scale system that could be deployed for Internet search. We also had the ambition to build the envisioned search engine in a completely decentralized manner as a peer-to-peer system. On this theme as well, we were very successful in terms of insight and publications, but our advanced prototype system, Minerva, did not make the practical impact that we had aimed for.

Further Highlights: The Yago-Naga theme also spun off side projects which have been very influential. One of these is the RDF-3X database engine for efficient storage and querying of subject-predicate-object triples, primarily developed by Thomas Neumann who is now a professor at TU Munich and has received the VLDB Early Career Award 2014 for this work. The four main papers on this research together have more than 1200 citations, and the RDF-3X system is widely used in the Semantic Web community. Another contribution with high impact is the AIDA method and software for Named Entity Recognition and Disambiguation. Here, the knowledge base is leveraged as a background asset for better language understanding and text analytics. The EMNLP 2011 paper has more than 400 citations. This work has spawned a startup called Ambiverse.

Contributions and Impact: Major Results 2015–2017

We publish our results in top-tier conferences in several communities: Web science (Web), data mining (DM), database systems (DB), information retrieval (IR) and natural language processing (NLP). In the two-year timeframe 2015–2017, the group had 11 full papers in first-rate Web venues (WWW, WSDM, ISWC, AAAI), 20 in DM (KDD, ICDM, SDM, ECML/PKDD), 5 in DB (SIGMOD, VLDB, TODS), 12 in IR (SIGIR, CIKM, ECIR), and 10 in NLP (ACL, TACL, EMNLP, COLING). In addition, we are successful in building prototype systems and publishing demo papers at top venues like ACL, VLDB, SIGIR, WWW, etc.

Commonsense and Activity Knowledge: We developed methods to extract properties of general objects from Web sources, text and image tags. The resulting commonsense knowledge base, WebChild, features the largest and cleanest public collection of part-whole relations (AAAI2016). In addition, we developed methods to build large repositories of human activites and how-to tasks, along with semantic attributes like participants, typical time and location, steps involved, etc. (WWW2015, CIKM2015, WWW2017).

Named Entities and Relation Phrases: We further advanced our strong standing on discovering and linking named entities in text and tables (2xCIKM2016, TACL2016, WWW2016), including co-reference resolution (EMNLP2015, TACL2015), and on discovering, semantically typing and organizing text phrases that denote relations (EMNLP2015, EMNLP2016). Several group members founded a startup company, Ambiverse (`ambiverse.com`), which has built a product based on this competence.

Trust and Privacy: In the context of the ERC Synergy Grant, we addressed the issue of trust in online information, and how it can be assessed even for evolving communities with anonymous posts (ICDM2015, CIKM2015, KDD2016). We further extended this line by analyzing the credibility of textual statements in news and social media (CIKM2016, WWW2017). On the privacy dimension, we have devised a “non-standard” model for assessing a user’s risk of being targeted by exposing sensitive information in online communities (SIGIR2016).

Rule Mining: We successfully launched a new research direction on inferring intensional knowledge – rules – from extensional assertions (VLDBJ2015). Our unique contribution is that we obtain exception-aware rules, based on novel knowledge revision techniques (ISWC2016, ILP2016).

Event-Oriented Information Retrieval: Our prior work on temporal IR has been extended and enhanced by new ways of extracting and canonicalizing events from text (WWW2016, WSDM2017), learning statistical language models for events (CIKM2016), and computing event digests and rankings from diverse sources (WWW2015, SPIRE2016, SIGIR2016, ECIR2016).

Non-Standard Matrix Factorization: This line of work has specifically addressed the factorization of matrices and tensors in non-standard algebras, like the Boolean algebra or the (sub-) tropical algebra (with max and sum (or product) as operators). This has not only led to novel algorithms but also yields better interpretable decompositions of datasets into patterns (SDM2015, ECML/PKDD2016, SDM2016, ICDM2016). Moreover, it provides insight into how continuous factorizations relate to discrete ones, over different algebras (ICDM2016).

Exploratory Analysis of Correlations and Causality: Using information-theoretic principles, we developed methods to discover interesting correlations in high-dimensional data and analyze which of these constitute causal relations (SDM2015, SDM2016, ICDM2016, SDM2017). By treating pattern discovery as a form of data compression, more interpretable insights into data can be obtained (KDD2016, SDM2017). A related contribution is on subgroup discovery, with applications in material science (New J. of Physics 2017).

Young Researchers

Our group has a strong track record in promoting young researchers in their careers. The academic offspring of D5 includes internationally visible scientists such as Gerard de Melo (at Rutgers U, USA), Rainer Gemulla (at U Mannheim, Germany), Katja Hose (at U Aalborg, Denmark), Georgiana Ifrim (at UC Dublin, Ireland), Ndapa Nakashole (at UC San Diego, USA), Thomas Neumann (at TU Munich, Germany), Maya Ramanath (at IIT Delhi, India), Fabian Suchanek (at U Telecom ParisTech, France), Martin Theobald (at U Luxemburg), Yafang Wang (at Shandong U, China), and further alumni who have obtained faculty positions. A good fraction of our graduates joined industrial research labs, for example, Maximilian Dylla (Google), Niket Tandon (Allen Institute for Artificial Intelligence) and Mohamed Yahya (Bloomberg). In the two-year timeframe 2015–2017, we had 9 doctoral students graduating.

Awards

Gerhard Weikum received the ACM SIGMOD Edgar F. Codd Innovations Award 2016 for "pioneering work on the automatic construction of large knowledge bases". His lifetime achievements were also honored with the Robert-Piloty-Prize 2016. Sanjar Karaev (and his co-author and advisor Pauli Miettinen) won the Best Student Paper Award at the ECML/PKDD 2016 Conference. The startup Ambiverse, spun off from our research on language understanding, has won the 1-2-3-GO Award and the Digital Innovation Award for newly founded innovative companies.

Cooperations

Institute and Max Planck Society: Within the institute, we collaborate with D1 on efficient algorithms for large graphs and matrices, with D2 on commonsense knowledge for understanding visual contents, and with D3 on mining protein/gene patterns. The collaboration with D2 has led to jointly authored papers in CVPR and AAAI. Within the Max Planck Society, we collaborate with Peter Druschel, Krishna Gummadi and Rupak Majumdar, all at the Max Planck Institute for Software Systems in Saarbrücken and Kaiserslautern. This has resulted in joint papers at KDD and SIGIR. In addition, our group is a member of the MaxNet research network on Big Data Driven Materials Science, which involves 9 Max Planck Institutes. In this context, members of our group published joint work with the Fritz Haber Institute in the New Journal of Physics.

External Partners and Competitive Grants: D5 participates in the DFG-funded Excellence Cluster on Multimodal Computing and Interaction. Gerhard Weikum is one of the principal investigators and responsible for the research area on knowledge management. The Excellence Cluster is a framework for extensive cooperation on the Saarland University campus. We have collaborated with other groups in computational linguistics and in human-computer interfaces, with joint publications in premier venues. We also participate in the DFG-funded Collaborative Research Center (Sonderforschungsbereich) on Methods and Tools for Understanding and Controlling Privacy, where we collaborate with Michael Backes and Jens Dittrich, both professors at Saarland University. Within the IMPECS framework (Indo-German Max Planck Center for Computer Science), we have collaborated with IIT Delhi and IIT Kharagpur. There are various individual collaborations with researchers at other universities and research labs across the world (including AI2, CMU, U Glasgow, L3S Hannover, U Melbourne, TU Munich, U ParisTech, U Sheffield, Tsinghua U, etc.).

The single most important collaboration is the ERC Synergy Grant 610150 on Privacy, Accountability, Compliance and Trust for Tomorrow's Internet (imPACT), with Michael Backes, Peter Druschel, Rupak Majumdar and Gerhard Weikum as principal investigators. This ERC grant is one of the highest reputed scientific prizes in Europe; it provides the four PIs with a total budget of 10 Million Euros for the timeframe 2015–2020. Our research agenda aims at fundamental insight on reconciling the tensions between the four PACT properties (privacy, accountability, compliance, trust). This includes the goal of developing game-changing methods and tools that assist users in analyzing, understanding and managing their privacy risks in social media and other kinds of digital traces.

Future Research Focus

New Directions in Digital Knowledge: Current knowledge bases focus on facts about entities like prominent people, places and products, and still lack other knowledge dimensions like properties of everyday objects, human activities and socio-cultural contexts. These shortcomings are exemplified by the brittle and often poor behavior of prominent chatbots and question-answering agents. Our next goal is to equip computers with commonsense world knowledge to enable context-aware understanding of language, multimodal scenes and conversational behavior in human-machine interactions.

Trust and Privacy: In the age of overwhelming online information as well as biased or manipulative misinformation (aka. "post-truth" or "alternative truth"), a strategic goal is to analyze and assess the factuality or credibility of statements and claims, and explain them in terms of trustworthy sources and relevant context. We envision a polygraph for online media. To determine the "truth" and generate user-comprehensible explanations, we aim for methods to de-construct the dependencies among sources of evidence and track information back to its origin. These methods need to discover diverse perspectives on claims, like alternative viewpoints by different socio-cultural groups. As many sources may stem from anonymized users (with pseudonyms), respecting the privacy of these users is a major challenge.

2.6 The Automation of Logic Group (RG1)

Group Overview

The Automation of Logic Group has been established in September 2005 and is headed by Christoph Weidenbach. The group covers the complete pipeline from basic research on (new) logics and their automation up to applications in research and industry.

There are currently 7 researchers and 8 PhD students in the group. Noran Azmy, Willem Hagemann, Marek Košta, Manuel Lamotte-Schubert, and Martin Suda have finished their PhD theses during the reporting period and left the group. Viorica Sofronie-Stokkermans, has eventually finished her association status with the end of the AVACS project. Thomas Sturm has accepted a permanent position at CNRS but remains associated with the group. Jasmin Blanchette was working half time with the group as part of the VeriDis project but has meanwhile accepted a tenure track assistant professor position at VU Amsterdam. He also remains associated with the group and co-supervises Mathias Fleury. Matthias Horbach accepted an industry position.

Out of the current 8 PhD students, 4 PhD students are in the final phase of their PhD and are expected to submit their theses in 2017.

Thomas Sturm, working in computer algebra, Jasmin Blanchette, working in interactive theorem proving and its mechanization are, together with Christoph Weidenbach, the senior researchers of the group. Uwe Waldmann has meanwhile moved his research topics towards higher-order logic automation.

Research Areas and Achievements

The vision of the group is to increase the productivity of formal analysis/verification/problem solving technology through a higher degree of automation of the underlying logics. The following challenges motivate main parts of our work: (i) understand complete and sound reasoning for the combination of theories, (ii) drive the development of first-order and propositional reasoning calculi, (iii) understand the mechanics of reasoning with respect to concrete models, (iv) build specific reasoning procedures for arithmetic theories, (v) show applicability of our methods to reasoning challenges from other areas, and (vi) scale the applicability of our methods to the size of real world industrial applications.

About half of our work is of theoretical nature and the other half is experimental, in particular on the basis of developed tools. Implementation of our methods is important to check their automation potential in practice, to increase the impact of our theoretical results by providing our software to other people and eventually to detect further challenges in theory development. For example, the Spass workbench is used worldwide by individuals and groups working in research and industry. The computer algebra system Redlog has been one of the leading systems for more than a decade. In 2012 we founded the company “Logic for Business” (L4B) providing consulting and software for the overall lifecycle management of complex products (http://www.logic4business.com/).

We have structured our research along the underlying logics: (i) (First-Order) Arithmetic, (ii) Free First-Order Logic, (iii) Theory Combinations, (iv) Higher-Order Logic, (v) Non-Classical Logic, and (vi) Software. The logic structure does not impose a structure on the group. In fact, most of us contribute to several areas and all of us to the development of software.

One syntactic indicator on the success of our work in automated reasoning can be seen from the CADE/IJCAR/FroCos conferences. In addition to our scientific contributions to these conferences, group members serve regularly on the respective program committees. Christoph Weidenbach was elected CADE president in 2016. Jasmin Blanchette, Mathias Fleury, and Christoph Weidenbach won the best paper award at IJCAR 2016 for their contribution "A verified SAT solver framework with learn, forget, restart, and incrementality". Jasmin Blanchette won the ERC starting grant "Matryoshka" on higher-order logic automation for fast verification in 2016. Uwe Waldmann won the best teaching award at the computer science department at Saarland University for his lecture "Automated Reasoning II".

A selection of scientific results for logics (i), (ii), and (iv), respectively, are:

First-Order Logic Separated Fragments (LICS 2016, LCC 2016, LICS2017). Existentially and universally quantified variables are separated, if they do not occur together in atoms. The first-order separated fragment is decidable. It unifies existing fragments, such as the monadic and Bernays-Schönfinkel fragment. Exploiting the concept of separation yields extended decidability results for known fragments, such as the monadic, the Bernays-Schönfinkel, and the Ackermann fragment. Certain standard classes in computational complexity can be characterized by subfragments of separated formulas of different degrees of interaction between variables.

Fast Decision Procedures for LIA Constraint Solving (CADE 2015, IJCAR 2016, SMT 2016). Linear integer arithmetic is an integral part of many applications of automated reasoning. With CUTSAT++ we established a terminating procedure on LIA constraints without exploring a priori bounds. It is based on partial model construction and quantifier elimination techniques. In addition, we developed several efficient satisfiability tests based on the following observation: an LIA constraint is satisfiable if it contains a cube of sufficient size. These tests can be efficiently and naturally integrated in LIA constraint solving. Eventually, they also support the discovery of implicit equations in LIA constraints.

Formalizing Automated Reasoning Calculi in Logic (IJCAR 2016, FoSSaCS 2017). Specifying and proving metatheorems of automated reasoning calculi by pen and paper is error-prone and partly tedious. We develop libraries and methodologies in Isabelle/HOL supporting the formal specification and metatheorem development of automated reasoning calculi. So far we have successfully established CDCL for propositional logic, resolution for first-order logic, and orderings (RPO, TKBO) for higher-order logics.

Projects and Cooperations

Together with Stephan Merz's group (Inria Nancy) we are the first Inria project with a joint French-German leadership: VeriDis (`http://www.inria.fr/en/teams/veridis`), where we investigate automated reasoning support for the verification of distributed algorithms. The project is accompanied by the ANR/DFG project SMArT on satisfiability modulo arithmetic and theories, and the EU FET Open action SC-Square on bringing together automated reasoning and computer algebra research, both led by Thomas Sturm.