Fourteenth
Biennial Status Report

March 2017 – February 2019

Contents

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), Anja Feldmann (2018–2033) heading the Department on Internet Architecture, 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 153 scientists, out of which 96 are doctoral students and 57 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, 27 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(s). The coordinating scientists are senior researchers or postdoctoral researchers with strong potential for becoming senior researchers.

[1] Zeynep Akata, Bjoern Andres, Klaus Berberich, Jasmin Blanchette, Karl Bringmann, Andreas Bulling, Renjie Chen, Mario Fritz, Olga Kalinina, Andreas Karrenbauer, Christoph Lenzen, Tobias Marschall, Pauli Miettinen, Karol Myszkowski (tenured), Nico Pfeifer, Gerard Pons-Moll, Simon Razniewski, Michael Sagraloff, Rishiraj Saha Roy, Marcel Schulz, Daria Stepanova, Jannik Strötgen, Thomas Sturm, Christian Theobalt (tenured), Jilles Vreeken, Christoph Weidenbach (tenured), Andrew Yates

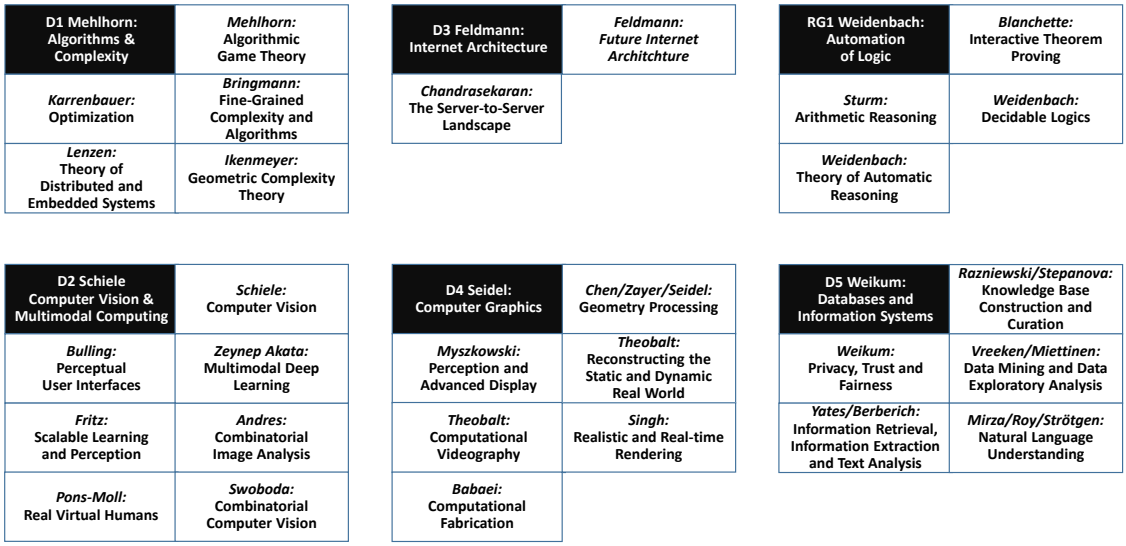

Figure 1.1: Research Areas of the Institute (as of March 2019)

Doctoral Students: As of March 2019, there are 96 doctoral students who are supervised by members of the institute, including 21 women (about 22%). 59 of the 96 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 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, adaptive networks). 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's group has been working on tracking multiple people for several years now. 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 won the Multi-Object Tracking Challenge twice (at ECCV'16 and at CVPR'17).
- Anja Feldmann's research vision is to obtain insights from Internet measurements as foundation for shaping the evolution of Internet by proposing optimizations and investigating alternative designs. As such her research group has and continue to address many challenges in *Internet measurement*, e.g., the use/abuse of Interdomain routing, as well as *cross-layer network optimization*, e.g., by highlighting the benefits of joint optimization of ISPs/CDNs or applications/network.
- Hans-Peter Seidel and his group stand out for their integrated view of 3D image analysis and synthesis. The group developed groundbreaking results on multiresolution modeling with special emphasis on new metaphors for editing and shape deformation, and novel data structures for high performance geometry processing. Recent research on markerless performance capture has pioneered methods for the reconstruction of detailed dynamic 3D models of humans from both multiview and monocular video and has led to entirely new ways of fusing and deeply integrating model-based and deep learning-based scene reconstruction. The group has been and continues to be 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 first-order calculi and their implementations, the group is now driving towards calculi for more expressive logics that also 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 2017–2019 and New Research Directions

Two important developments for the institute are the retirement of Thomas Lengauer (who became an Emeritus in 2018) and the upcoming retirement of the institute's founding director, Kurt Mehlhorn, later this year. Kurt Mehlhorn will move to Emeritus status at the end of August 2019. Similar to Thomas Lengauer he will continue as a researcher in the institute. In order to have some guidance for the upcoming search for directors we decided early on to (i) keep the algorithms and complexity slot and (ii) to make an effort to attract at least one outstanding woman as a director.

In the meantime we have been able to identify and attract Anja Feldmann. Anja joined the institute as a new director in 2018 and is heading the Internet Architecture Department. The appointment process for a new director in Algorithms and Complexity is on-going.

Thomas Lengauer retired on May 31, 2018, after more than 15 years as Director of the Bioinformatics Department and dedicating his life to science. With a symposium on bioinformatics and beyond the institute and more than 200 participants celebrated the work and career of Thomas on May 25 2018. The list of scientific speakers ranged from collaborators to former students to colleagues from the institute: Gene Myers, Nico Pfeiffer, Yves Moreau, Kasia Bozek, Robert Tarjan, Jörg Rahnenführer, Bastian Beggel, and Anja Feldmann.

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

- During the reporting period, Kurt Mehlhorn's group worked mainly on fine-grained complexity, distributed computing, reliable computing, optimization, and algorithmic game theory.
- Bernt Schiele's group has been working on variety of essential problems that will be also followed in the future. An important example is how to learn without or very

little supervision. In joint work with Zeynep Akata we have been able to push the state-of-the-art in zero-shot as well as few-shot learning substantially e.g. by combining the strength of VAE and GANs within a joined framework. In joint work with Mario Fritz we have worked on various aspects at the intersection of privacy and security on the one side and computer vision and machine learning on the other side. E.g. we have contributed to the better understanding of privacy implications of visual data dissemination through and also showed, that compute vision model stealing is possible under weak assumptions, thereby raising concerns about the security of current machine learning techniques.

- Anja Feldmann's group has been working on understanding the capabilities of today's Internet Inter-Domain routing. In recent work they have been highlighting the good (CoNext 18), the bad (IMC 17), and the evil (IMC 18) aspects of today's collaborative system. This work has also been presented at RIPE as well as at IETF (via an IETF Applied Networking Research award).
- Hans-Peter Seidel's group developed a novel sparse matrix representation for unstructured grids and demonstrated its power in different geometry processing applications (EG'17, ICS'17, SC'18, SIGAsia'18) (joint with R. Zayer and M. Steinberger). C. Theobalt developed *VNect* (SIG'17), the first method for accurate and real time 3D skeletal motion capture from monocular RGB and *MoFA* (ICCV'17, PAMI'18, CVPR'18), the first unsupervised approach for the reconstruction of high-quality 3D faces from a single color image. *Deep Video Portraits* enables realistic re-animation of the entire posture and face expression in a portrait video using only an input color video (SIG'18, CommACM'19). With T. Ritschel and G. Singh we obtained far reaching results on the simulation of distribution effects and the generation of well-behaved sampling patterns in rendering (SIG'17, SIG'18, CGF'19). K. Myszkowski introduced an award winning advanced near-eye display design (TVCG'17), and we developed novel fundamental algorithms for perception-aware gaze contingent rendering (SIG'17), 2D-to-3D conversion (TVCG'18), and light field quality metrics (CVPR'17, TOG'18).
- Christoph Weidenbach's group has continued the discovery of new (un)decidable first-order fragments (CADE 2017, FROCOS 2017) and has established new techniques in solving linear arithmetic constraints over the integers (CADE 2017, IJCAR 2018, SMTCOMP 2018).
- In the context of the ERC Synergy Grant imPACT, we address the issue of trust in online information, particularly, aiming to analyze and assess the credibility of textual statements in news, discussion forums and other kinds of online communities. This research led to the doctoral thesis of Subhabrata Mukherjee which won an honorable mention for the KDD dissertation award 2018. More recent work (with publications in WWW 2017, WWW 2018 and EMNLP 2018) focused on incorporating web evidence for estimating the credibility of claims made in news or by influential people (politicians, bloggers, partisan media etc.). We devised a suite of methods, including pipelined classifiers, probabilistic graphical models for joint inference on source trustworthiness and content credibility, all the way to deep-learning networks with tailored attention mechanisms.

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, internet architecture and knowledge discovery in the coming years. Moreover, we will continue our collaboration with the Bioinformatics Center at Saarland University even after Thomas Lengauer's departure, and we remain committed to a strong research environment in this area.

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 313 doctoral students graduated. These include 56 women, and 130 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). Maximillian Dylla won the DBIS Dissertation Award in 2014 Marcus Rohrbach won the DAGM MVTec Dissertation Award in 2015. Karl Bringmann won the EATCS Distinguished Dissertation Award in 2015. Pablo Garrido and Petr Kellnhofer both won Eurographics Dissertation Awards in 2018.

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 76 young scientists held such positions. 52 of them have meanwhile left the institute. Out of these, 51 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], and other honors.

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

[3]Susanne Albers 2016, Jasmin Christian Blanchette 2016, Niko Beerenwinkel 2014, Christoph Bock 2015, Andreas Bulling 2018, Parinya Chalermsook 2017, Piotr Didyk 2018, 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, Pscal Schweitzer 2018, Jürgen Steimle 2016, Christian Theobalt 2013, 2017

Support for Women: The fraction of women in our institute currently is 22% for doctoral students and 16% 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 10 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 2019. 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

[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, Eunjin Oh 2018-2020, Qiuhong Ke 2018-2020, Sophie Tourret 208-2020

[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), Hang Zhou (École Polytechnique Paris, France), Anke van Zuylen (College of William & Mary, USA)

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 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).

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).

The Research Centre on Interactive Media, Smart Systems and Emerging Technologies (RISE): This centre is located in Cyprus and was founded in 2017. It is funded within the framework of Horizon 2020. MPI is a strategic international partner of RISE and represented on the Board by Hans-Peter Seidel.

ERC Synergy Grant imPACT: The ERC Synergy Grant has been awarded to Michael Backes (Helmholtz Center for Information Security), 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". The project is funded with a total of 10 Million Euros for the timeframe 2015–2020.

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 96 doctoral students, including 21 women (about 22%) and 59 non-Germans. In addition, IMPRS-CS currently supports 20 students with fellowships towards Master's degrees.

The Max Planck Graduate Center for Computer and Information Science: The center is a highly selective doctoral program that grants admitted students full financial support to pursue doctoral research in the broad area of computer and information science, with faculty at the Max Planck Institute for Informatics, the Max Planck Institute for Software Systems, the Max Planck Institute for Intelligent Systems, and some of the best German universities. Students normally start their graduate studies at the Saarland Informatics Campus in Saarbrücken. While taking courses, they have the opportunity to explore research in different areas as part of immersion labs at different MPIs and universities.The center started in 2018 and the first batch of admitted students started beginning of 2019.

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.

1.7 Results 2017–2019

Publications, Software, Startups: In the two-year timeframe 2017–2019, the institute published more than 600 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 CVPR 17, EMNLP 17, Eurographics Symposium 17, IEEE VR 17, IMWUT 17, ISWC 17, MUM 17, UIST 17, 3DV 18, ACM CHI 18, ETRA 18, Eurographics 18, IEEE ICDM 18, IMC 18, ISS 18, MobileHCI 18, SOS 18, WWW 18).

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 2017–2019, 39 of our doctoral students graduated. These include 10 women. In the same time period, 16 of our researchers left the institute to take a tenured or tenure-track faculty position[6]. These include 1 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.

Anja Feldmann received the Vodafone Innovation award 2018 and the Shelling award of the bavarian academy of sciences 2018; Bernt Schiele has been elected IEEE Fellow as well as IAPR Fellow; Hans-Peter Seidel received the Eurographics Medal 2017; Christoph Weidenbach won the Thoralf Skolem Award 2017; Gerhard Weikum's team received the Seoul WWW Test of Time Award at the Web Conference 2018 for the impact of the original Yago paper at WWW 2007, co-authored by Fabian Suchanek, Gjergji Kasneci and Gerhard Weikum.

At the level of Senior Researchers, Christian Theobalt won the Karl Heinz Beckurts-Prize 2017 and received the ERC Consolidator Grant 4DRepLy; Piotr Didyk received the ERC Starting Grant PERDY; Gerard Pons-Moll received an Emmy-Noether grant; Andreas Bulling received the ERC Starting Grant ANTICIPATE; Karol Myszkowski was recognized with the IEEE TVCG Best Associate Editor Award in 2017; Jilles Vreeken won the IEEE ICDM Tao Li Early Career Award 2018; Karl Bringmann won the EATCS Presburger Award 2019 and the Heinz Maier-Leibnitz-Prize 2019. Gerard Pons-Moll received a Google Faculty Research Award 2019.

At the student level, Franziska Müller received the Google PhD Fellowship in Machine Perception 2017 and the Women STEM Award 2017. Marvin Künnemann won the Otto-Hahn-

[6] Bastian Beggel (U. of Applied Sciences Kaiserslautern), Mario Boley (Monash U. Melbourne, Australia), Andreas Bulling (U. Stuttgart), Dan Casas (U. Madrid, Spain), Avishek Chatterjee (IIT Kharagpur, India), Renjie Chen (U. Science and Technology of China (USTC)), Piotr Didyk (USI Lugano, Switzerland), Mario Fritz (CISPA), Christian Ikenmeyer (University of Liverpool), Caigui Jiang (Xi'an Jiatong U.), Moti Medina (Ben-Gurion U., Israel), Pauli Miettinen (U. of Eastern Finland, Finland), Emanuele Natale (U. Côte d'Azur, France), Tim Oosterwijk (Maastricht U., Netherlands), Helge Rhodin (Univ. British Columbia, Vancouver, Canada), Michael Sagraloff (U. of Applied Sciences Landshut), Karteek Sreenivasaiah (IIT Hyderabad, India), Jilles Vreeken (CISPA), Hang Zhou (École Polytechnique Paris, France); Björn Andres Adjunct Professor (U. of Tübingen)

Medal 2017, Anna Rohrbach won the Otto-Hahn-Medal 2018, Petr Kellnhofer received the Dr.-Eduard-Martin Prize 2017, Mateusz Malinowski received the Dr.-Eduard-Martin Prize 2018. Mateusz Malinowksi won the DAGM MVTec Dissertation Award. The Guenter-Hotz Medal 2017 was awarded to Marc Habermann. David Stutz received the STEM Award IT 2018; Pablo Garrido and Petr Kellnhofer won Eurographics PhD Dissertation Awards 2018.

Further honors with considerable visibility include the following. Kurt Mehlhorn received a Honorary Doctorate Degree from the University of Patras and was appointed Honorary Member of the "Gesellschaft zur Förderung des Forschungstransfers". Bernt Schiele was elected Fellow of the IAPR 2018.

Gender Proportion and Diversity: We have been successful in attracting an outstanding woman as a director. 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 three 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 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 continues its role as a provider of for the BWINF, "Bundesweit Informatik Fördern", promoting young computer science talent in Germany.

2 The Research Units

2.1 D1: Algorithms and Complexity

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 will retire as director of the department on August 31st, 2019. He will continue as a scientist after this date. The appointment process for a new director in on-going.

As of March 1st 2019, the senior scientists and subgroup coordinators are Alkimini Sgouritsa, Karl Bringmann, Andreas Karrenbauer, Christoph Lenzen, and Kurt Mehlhorn.

Section 35.1 of the Scientific 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, see page 228 of the Scientific Report. Nine group members completed their PhD, see page 228 of the Scientific Report. Group members received prestigious awards, see page 229 of the Scientific Report, e.g., Karl Bringmann received the EATCS Presburger Award 2019 and the Heinz Maier-Leibnitz-Prize 2019, Marvin Künnemann received an Otto Hahn Medal of the Max Planck Society in 2017, Emanuele Natale received an Award for Best Italian PhD Thesis in Theoretical Computer Science by the Italian Chapter of EATCS in 2017, Eunjiin Oh received the Lise Meitner Award from the MPI-INF in 2018, and Kevin Schewior received a Dissertation Award of the German Operations Research Society (GOR) in 2017. Some group members hold their own grants, see page 229 of the Scientific Report.

We have published extensively and in excellent venues, e.g., 12 papers in SODA 2018 and 2019, 4 papers in FOCS 2017 and 2018, 11 papers in ICALP 2017 and 2018, 5 papers in ESA 2017 and 2018, and 5 papers in SoCG 2017 and 2018.

Vision and Research Areas

The vision for D1 is to be a first class algorithm group and a trendsetter in algorithmics, and to have impact on the research community and society through research results, people, software, and scientific leadership.

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, 2019, in italics).

- Algorithmic Game Theory and Online Algorithms (*Kurt Mehlhorn* and *Alkimini Sgouritsa*)
- Combinatorics, Computing, and Randomness (*Karl Bringmann*).
- Combinatorial Optimization (*Andreas Karrenbauer*)

- Geometry and Algebra (Christian Ikenmeyer and Michael Sagraloff)
- Theory of Distributed and Embedded Systems (*Christoph Lenzen*).

Together, we span a large part of algorithmic research. The emphasis changes over the years as group members come and go. We hire postdocs mainly based on quality and less on thematic fit. There is considerable interaction and collaboration between the areas. The entire group meets twice a week to discuss science (Tuesday and Thursday noon seminar) and biweekly to discuss administrative matters.

We teach at all levels; see page 226 of the Scientific Report for details. Frequently group members pair for lectures.

We are involved in the *Indo Max Planck Center for Computer Science (IMPECS)*.

Main Results

I discuss some of the main results obtained in the reporting period under the headings *fine-grained complexity*, *algorithm engineering*, *geodesic Voronoi diagrams in simple polygons*, *dynamical systems*, *geometric complexity*, *distributed algorithms for network flow*, and *theory of distributed and embedding systems*. For all but the third heading, I cover groups of results.

Fine-Grained Complexity The classic theory of NP-hardness provides a coarse-grained classification of problems into efficient (= polynomial-time solvable) and intractable (= NP-hard). Modern *fine-grained complexity theory* yields quantitative running time lower bounds, and thus provides a more fine-grained classification of problems according to their time complexity. Naturally, such stronger lower bounds rely on stronger hypotheses than P≠NP, e. g., the *Strong Exponential Time Hypothesis* for Satisfiability or the *Orthogonal Vectors Conjecture*. Karl Bringmann, Marvin Künneman, Bhaskar Chaudhury, Andre Nusser, Philip Wellnitz and their network of co-workers develop fine-grained complexity theory by designing proof techniques for tight fine-grained lower bounds and applying them to a variety of problems. They also contributed to the foundations of fine-grained complexity theory, e. g., by studying the hypotheses that this theory relies on.

They obtained a multitude of results and so I mention only three referring to the Section of fine-grained complexity on page 131 of the Scientific Report and to Bringmann's section (see page 63 of the Scientific Report) for details.

- An algorithm for subset sum than runs in time $\widetilde{O}(n+t)$, where n is the number of input integers and t is the target and the corresponding conditional lower bound. The previous best algorithm dating back to the sixties ran in time $O(nt)$.
- Strong additional evidence for the orthogonal vectors conjecture is given, e.g., if the conjecture is false then there is a randomized a $O(n^{(1-\epsilon)k})$-time algorithm for the Min-Weight-k-Clique problem on d-hypergraphs with m vertices. The conjecture is a corner stone of fine-grained complexity. It states that given two sets A and B of n vectors from $\{0,1\}^d$, to decide whether there are vectors $a \in A$ and $b \in B$ such that a and b are orthogonal, requires essentially time $\Omega(n^2)$ provided that d is sufficiently large.
- An algorithm for the word-break problem that runs in $O(nm^{1/3}+m)$. Given a set D of strings of total length m and a text s of length n, is it possible to partition s into

Figure 2.1: Norm NF Z71-300, the new French keyboard layout.

substrings that all belong to D? Surprisingly, a matching conditional lower bound holds for "combinatorial algorithms" (essentially algorithms that are not allowed to use fast matrix multiplication).

Algorithm Engineering The algorithm ANewDsc developed by Kobel, Roullier (INRIA), and Sagraloff for isolating the roots of a univariate polynomial will become the default root finder in the next release of MAPLE; it is based on theoretical work by Mehlhorn and Sagraloff. On ill-conditioned instances, the implementation is considerably faster than previous methods, and on small or well-conditioned instances, it can compete with the fastest existing methods. Moreover, it applies not only to polynomials with integer coefficients but to polynomials with *arbitrary computable* coefficients. See page 193 of the Scientific Report for more details.

In 2015, the French Ministry of Culture published a document criticizing the lack of an official standard for a French keyboard layout. It pointed out that all layouts in use were lacking some special characters. The Ministry was concernced that this curbs proper use and development of the French language. AFNOR, the French national organization for standardization, was tasked to produce a standard that would support correct French and include the missing characters. AFNOR set up a committee to design the new standard. A. M. Feit (Aalto), M. Nancel (Inria), M. John, A. Karrenbauer, D. Weir (Aalto), A. Oulasvirta (Aalto) provided computational support to the committee. They modelled the problem as a quadratic assignment problem and optimized for performance, ergonomics, intuitiveness, and familiarity. Intermediate results were presented to the committee which would then typically further modified the objectives. The new standard, see Figure 2.1, was approved recently. The quadratic assignment problem is notoriously difficult. John and Karrenbauer contributed a novel column-generation method that exploits the sparsity of typical instances to allow for quicker assessments of solutions. I refer to Karrenbauer's section for more details, see page 75 of the Scientific Report.

Geodesic Voronoi Diagrams in Simple Polygons The *geodesic distance* between two points inside a simple polygon is the length of the shortest path connecting the two points and contained in the simple polygon. The *geodesic Voronoi diagram* of a set of m point sites in a

simple n-gon is defined as the subdivision of the simple polygon into cells, exactly one cell per site, such that every point in a cell has the same nearest site under the geodesic distance.

Clearly, $\Omega(n + m \log m)$ time is required for computing the geodesic Voronoi diagram and it was asked already 30 years ago, whether this time bound can be achieved. Since '89, progress was made in a sequence of papers. Eunjin Oh finally settled the problem and presented an algorithm with optimal running time, see page 186 of the Scientific Report for more details.

Dynamical Systems Systems that evolve over time play according to simple rules are ubiquitous and are a common theme for several group members.

An important question in Computational Economics is to provide algorithmic explanation why market equilibrium can be achieved in settings, where price updates on different goods in a market are made in a distributed and non-coordinated manner. In other words, price updates are *asynchronous*. Yun Kuen Cheung and Richard Cole (NYU) showed that *CES Fisher markets* converge toward market equilibrium amid general asynchronous *tatonnement* price updates.

Tatonnement price updates, already introduced by Léon Walras in his 1874 book "Elements of Pure Economics, or the theory of social wealth", can be briefly described as follows: if a good is over-demanded (i.e., its demand exceeds its supply), raise the price of the good; and vice versa. Synchronous tatonnement in a CES Fisher market was known to converge; each round of the process reduces a (carefully chosen) potential function by a certain fraction.

When tatonnement price updates are asynchronous, each update might be using outdated information, so the progresses will be countered by errors. The key question to ask is whether such errors can accumulate to nullify the progresses, and to bar the price updates from converging to equilibrium. Cheung and Cole show that this cannot happen, see page 118 of the Scientific Report for more details.

In a series of wet-lab experiments it was shown that the *slime mold Physarum polycephalum* possesses the ability to solve shortest path problems. The slime's adaption process can be modeled by a system of differential equations. In earlier work, we showed that this dynamical system converges for all graphs and all initial conditions.

Nature cannot be expected to follow the differential equation exactly. In this period, we (Karrenbauer, Kolev, Mehlhorn) showed convergence for a considerably more general dynamics. Let G be a graph with two designated nodes s_0 and s_1. Each edge e has a positive cost c_e and a time-varying positive capacity x_e. The capacity vector evolves according to the dynamics

$$\dot{x}_e = a_e(x,t)\left(|q_e| - x_e\right),$$

where q_e is the electrical flow through edge e in a network where edge e has resistance c_e/x_e and one unit of current is sent from s_0 to s_1, and $a_e(x,t)$ is any Lipschitz-continuous function that is bounded away from zero and from infinity, i.e., there is a positive constant C such that $1/C \leq a_e(x,t) \leq C$ for all x and t. In the original result, $a_e(x,t) = 1$ for all x and t. We refer to page 179 of the Scientific Report for more details.

Emanuele Natale and co-workers study *consensus dynamics* on the complete graph of n nodes. Initially, each node supports one from up to n opinions. Nodes randomly and in parallel sample the opinions of constantly many nodes. Based on these samples, they use an update

rule to change their own opinion. The goal is to reach consensus, a configuration where all nodes support the same opinion.

They compare two well-known dynamics: *2-Choices* and *3-Majority* dynamics. In the former, each node samples two nodes and adopts their opinion if they agree. In the latter, each node samples three nodes: If an opinion is supported by at least two samples the node adopts it, otherwise it randomly adopts one of the sampled opinions.

They show that the 3-Majority dynamics reaches consensus with high probability in $\widetilde{O}(n^{3/4})$ rounds, while the 2-Choices dynamics can need $\Omega(n/\log n)$ rounds. They thus get the first unconditional sublinear bound for 3-Majority dynamics and the first result separating the consensus time of these processes. Along the way, they develop a framework that allows a fine-grained comparison between different consensus processes, see page 217 of the Scientific Report for more details.

Geometric Complexity Theory Geometric Complexity Theory was initiated by Mulmuley and Sohoni in order to resolve Valiant's famous determinant vs permanent conjecture (the permanent cannot be computed by a polynomial size algebraic circuit). Mulmuley and Sohoni laid out a line of attack by formulating a sequence of conjectures. Christian Ikenmeyer, Karl Bringmann and Gorav Jindal in cooperation with Markus Bläser, Vladimir Lysikov (Saarland University), Peter Bürgisser (TU Berlin), Fulvio Gesmundo (University of Copenhagen), Stefan Mengel (CRIL), Ketan Mulmuley (University of Chicago), Greta Panova (University of Southern California), Michael Walter (University of Amsterdam), and Jeroen Zuiddam (IAS Princeton) refuted some of these conjectures[1]; this is a severe blow to the program.

They also investigate the fundamental capabilities and limitations of algebraic proofs for complexity lower bounds. We provide a setting in which superpolynomial algebraic lower bounds proofs require superpolynomial circuit size under a reasonable Boolean complexity assumption. Moreover, we show how geometric complexity theory can in principle break this barrier by encoding hard functions by concise representation-theoretic labels. As a side result from our techniques we prove that there is a constant multiplicative error to which tensor rank is NP-hard to approximate.

We refer the reader to page 198 of the Scientific Report for details.

Distributed Algorithms for Networks Flows Shortest path problems and maximum flow problems are notoriously hard for distributed algorithms. The standard combinatorial algorithms for these problems are inherently sequential and do not seem to be a good starting point for distributed algorithms. This is different for newer algorithms based on interior point methods.

M. Ghaffari (ETH), A. Karrenbauer, F. Kuhn (Freiburg), C. Lenzen, and B. Patt-Shamir (Tel Aviv) designed a near-optimal distributed algorithm computing a $(1+o(1))$-approximation for the maximum flow in undirected weighted networks that runs in $O((D+n)n^{o(1)})$ rounds of communication in the Congest model; here D is the diameter of the network. This is the first improvement over the trivial bound of $O(n^2)$ and nearly matches the $\Omega(D+n)$ round complexity lower bound. The development of their algorithm entails two sub-results of

[1] This paragraph is more vague that the other paragraphs as KM has only superficial knowledge of these results

independent interest: (1) An $O((D+n)n^{o(1)})$ -round distributed construction of a spanning tree of average stretch $o(1)$ and (2) an $O((D+n)n^{o(1)})$-round distributed construction of an $o(1)$ -congestion approximator consisting of the cuts induced by $O(\log n)$ virtual trees.

Moreover, R. Becker, A. Karrenbauer, S. Krinninger, and C. Lenzen studied *undirected transshipment* and *single-source shortest paths* (SSSP) problems in distributed and streaming models of computation. They presented a tailored gradient descent algorithm that computes near-optimal solutions (up to a multiplicative error of $1+\varepsilon$) within $\varepsilon^{-O(1)}\log^{O(1)} n$ iterations, each comprising of a polylog-approximation of an auxiliary instance of a transshipment problem. Such solutions of poly-logarithmic error can be obtained by computing a solution on a sparse spanner of logarithmic stretch, which leads to interesting results in various distributed and streaming models of computation:

1. *Broadcast Congest model:* $(1+\varepsilon)$-approximate SSSP using $\widetilde{O}((\sqrt{n}+D)\varepsilon^{-O(1)})$ rounds.
2. *Broadcast Congested clique model:* $(1+\varepsilon)$-approximate shortest transshipment and SSSP using $\widetilde{O}(\varepsilon^{-O(1)})$ rounds.
3. *Multipass streaming model:* $(1+\varepsilon)$-approximate shortest transshipment and SSSP using $\widetilde{O}(n)$ space and $\widetilde{O}(\varepsilon^{-O(1)})$ passes.

I refer to page 173 of the Scientific Report for details.

Theory of Distributed and Embedded Systems The work on *Design of Fault-tolerant Hardware* (see page 203 of the Scientific Report) and *Hazard-freedom and Metastability-containment* (see page 207 of the Scientific Report) is to a large extent financed by Christoph Lenzen's ERC grant and so I give Christoph the priority to report about them; see page 77 of the Scientific Report.

The other topics in this research area are distributed algorithms and computational dynamics. I already reported about them in earlier paragraphs.

2.2 D2: Computer Vision and Multimodal Computing

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 three research group leaders, three associated research group leaders, three postdocs and 19 PhD students. Among those 22% are female (one research group leader, three postdocs, three PhD students).

The research group leaders and senior scientists each have their own PhD-students to conduct research in their respective area. The current three research group leaders are Zeynep Akata (appointed 2014), Gerard Pons-Moll (appointed 2017), and Paul Swoboda (appointed 2018). The three associated research group leaders have been appointed as senior researchers at the institute before and still supervise PhD students at the institute, but have accepted offers for faculty positions elsewhere in the meantime. These are: Andreas Bulling (faculty at Uni Stuttgart since November 2018), Mario Fritz (faculty at CISPA since July 2018), and Bjoern Andres (faculty at Uni Tübingen and group leader at Bosch BCAI Renningen since October 2017).

Nine group members completed their PhD during the reporting period and one more will defend his PhD in early April. Three of them obtained prizes for their dissertation[2]. Our researchers get very good offers for faculty positions in academia[3], postdoc positions in academia[4], and research positions in industry[5].

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 areas, namely computer vision and multimodal sensor processing. In the area of computer vision we address some of the most fundamental problems

[2] Anna Rohrbach obtained the Otto-Hahn Medal of the Max-Planck Society, Siyu Tang and Mateusz Malinowski both received the DAGM MVTec Dissertation Award 2018, and Mateusz Malinowski also the Dr. Eduard-Martin Award

[3] Björn Andres (U Tübingen), Zeynep Akata (U Würzburg), Andreas Bulling (U Stuttgart, U Freiburg), Mario Fritz (CISPA, U Marburg)

[4] Anna Rohrbach (UC Berkeley), Siyu Tang (MPI Intelligent Systems), Xucong Zhang (ETH Zurich)

[5] Jan Hosang (Google, Zurich), Anna Khoreva (Bosch, BCAI Renningen), Maksim Lapin (Amazon, Berlin), Wenbin Li (Continental, Regensburg), Mateusz Malinowski (Google Deepmind, UK), Seong Joon Oh (LINE Plus Clova ML, Korea)

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.

Highlights of Current Research

In the last report we discussed long-term research results e.g. in the areas of human pose estimation and tracking as well as object recognition. In this report we mainly focus on recent research highlights from the current reporting period.

Zero and Few-Shot Learning *Investigators: Zeynep Akata and Bernt Schiele*

Many approaches in generalized zero- and few-shot learning rely on cross-modal mapping between the image feature space and the class embedding space. As labeled images are rare, one direction is to augment the dataset by generating either images or image features. However, the former misses fine-grained details and the latter requires learning a mapping associated with class embeddings.

An important achievement in this area during the reporting period are CVPR 2017 and TPAMI 2018 papers that establish a benchmark on zero-shot learning which has since then become one of the major resource for benchmarking (their cumulative citation count is over 200 since June 2017). In addition, we proposed connections between zero-shot learning and human gaze (CVPR 2017, spotlight). Furthermore, we improved zero-shot learning by generating images of unseen classes (ICML and NIPS 2016) and image features of unseen classes (CVPR 2018) which established the new state of the art on six benchmark datasets.

Furthermore, in a recently accepted CVPR'19 paper, we propose to develop a conditional generative model that combines the strength of VAE and GANs and in addition, via an unconditional discriminator, and learns the marginal feature distribution of unlabeled images to learn highly discriminative CNN features generalizable across datasets. We have also demonstrated that our learned features are interpretable: we visualize them by inverting them back to the pixel space and we explain them by generating textual arguments of why they are associated with a certain label. This way we aim to connect low-shot learning with explainability.

Privacy and Security *Investigators: Mario Fritz and Bernt Schiele*

While recent advances in computer vision and machine learning bring many benefits to society, we equally have to be aware of emerging security and privacy risks.

Adersarial Perturbation are small changes to the input data, that can mislead state of the art classifiers. We make significant contributions to the understanding of this phenomena and investigate the impact on more complex vision/language tasks and deep learning policies of agents, as well as, investigate the interplay between attacks and defenses in a game theoretic framework. This line of work was published in top tier venues: 1 ICCV, 2 CVPR

Attacks and Defenses for Machine Learning have gained relevance as more and more such approaches are deployed in various application scenarios. We have shown for the first time that details such as architecture parameters can be reverse engineered by observing only the output of a neural network. Beyond this, we show that modeling stealing is possible under weak assumptions, thereby raising new concerns about the security of these techniques and how intellectual property can be protected. We complement these new attacks with new defenses to counter such strategies. The work is published in top venues: 1 ICLR, 1 CVPR.

Enforcing Privacy has received increasing attention, as the latest machine learning techniques can extract private information from multi-modal data with high accuracy. We have pioneered work on visual privacy where we take a systematic approach to categorize and sanitize a broad range of private information in visual data. Our work extends to natural language and also decentralized learning. The work is published in top venues: 2 CVPR, 1 ECCV, 1 USENIX Security, 1 AISTATS.

Understanding Privacy is important in order to assess the amount of private information in the first place. We focus on visual data in social media type data, where we evaluate and advance person recogntion as well as demonstrate recognition of social relations in a domain based approach. We have presented first steps towards the overall vision of a Visual Privacy Advisor that can support users to enforce their privacy policies on social media. The work is published in top venues: 1 ICCV, 1 CVPR.

Scalable Learning and Perception *Investigator: Mario Fritz*

With the advances of deep learning techniques, recent success has spurred the hope to address more holistic challenges that encompass not only scene understanding but also include light, surfaces, physics and multiple agents. We have made significant progress towards connecting the success of Deep Learning in vision to computer graphics, simple physics in robotics as well as in learning of agent behavior. In particular, we have proposed advanced approaches that use machine learning to extract surface and light information, synthesize new appearance or automatically edit visual content. We have proposed a first, fully data-driven approach ("intuitive physics") to learn simple stability events and demonstrated for the first time how such models can be leveraged for robotic manipulation. This line of work is published in top tier venues: 1 ICCV, 2 TPAMI, 1 ICRA, 1 NIPS, 1 CVPR.

The second focus, in cooperation with Bernt Schiele, is on modeling uncertainty in Deep Learning approaches. While we have seen a steady increase in accuracy – in many application scenarios this is insufficient. In particular, we focus on anticipating on future events in a scene, which is of high relevance to autonomous driving. We have advanced modeling of uncertainty in Deep Learning approaches in particular when multi-modal posteriors are concerned and also demonstrated such techniques in relevant application scenarios. Also this line of work is published in top tier venues: 2 CVPR, 1 ICLR, 1 AAAI.

Real Virtual Humans *Investigator: Gerard Pons-Moll*

The faithful digitization of human beings has the potential to redefine the way we think and communicate (with other humans and with machines), and it is necessary for many applications. While digital human models entail many components of artificial intelligence, we are focusing on modelling and capturing the 3D human appearance and motion of people.

Over the last years, we have made significant progress, and have produced statistical body models of pose, shape, soft-tissue and clothing (Dyna-Siggraph'15, SMPL-SiggraphAsia'15, CVPR'17, ClothCap-Siggraph'17) which are widely used in industry and academia, and are regarded as state of the art. While useful, to train these models we required high quality 3D scans which makes it difficult to scale. In addition, no easy to use models of people in clothing exist.

Our main challenge for the future is to faithfully reconstruct and model 3D people in clothing from single images and video, for which data is abundant. Our solutions will have to effectively incorporate 3D model knowledge within deep learning frameworks. In essence, we want to train machines to learn about the 3D world by looking at images and video. We have done first steps in this direction and introduced algorithms to reconstruct people in clothing from video for the first time (CVPR'18), and to reconstruct pose and shape from a single image (3DV'18 best paper). We also recently made publicly available the first dataset, recorded with a moving camera, of people in natural scenes with ground truth 3D poses–for which we devised a fusion algorithm that combines Inertial Sensors and video data (ECCV'18). In the future, we plan to learn models of clothing geometry, topology and appearance, as well as algorithms to reconstruct and model people interactions.

In summary, our research aims at answering the following two inter-related research questions: *How do we efficiently digitize humans without losing the detail that make us real?* and *How can we train machines to perceive such rich representations from visual data?*

The work has been published at the major international computer vision and computer graphics conferences and journals including (10 papers at CVPR, ICCV, ECCV), (6 papers at Siggraph, Siggraph Asia, TOG), BMVC(Best Paper Award 2013), Eurographics(Best Paper Award 2017), 3 papers at 3DV(2018 Best Student Paper Award), IJCV and TPAMI. Gerard Pons-Moll has received two starting grants, the Emmy Noether Programme 1.6Mi in December 2018, and a grant from Zentrum Digitalisierung Bayern 1.25Mi in January 2017 (declined), a Google Faculty Research Award (2019), a Facebook research cooperation agreement (2018), as well as 3 best paper awards at vision and graphics conferences.

Combinatorial Image Analysis *Investigator: Bjoern Andres*

Algorithms for image analysis are beginning to have major impact on our society, e.g., in the area of computer vision for autonomous driving, and in the areas of image and data analysis for biomedical research and applications. The research group Combinatorial Image Analysis studies mathematical abstractions of image analysis tasks in the form of combinatorial optimization problems. It develops algorithms for solving these problems exactly, partially or locally. It analyzes these algorithms and the solutions they output in terms of application-specific metrics and benchmarks. The goal of this research is to enable practical applications of expressive combinatorial optimization problems in the fields of computer vision and biomedical image and data analysis.

Our main achievement during the reporting period consists in the understanding that seemingly unrelated computer vision tasks, including instance-separating semantic image segmentation, multiple object tracking and human body pose estimation, are fundamentally linked by a single combinatorial optimization problem, a generalization of the well-known correlation clustering problem. A publication at CVPR 2017 can be seen as the culmination

of this work that spans a series of papers published over a period of three years at ICML, CVPR, ICCV and ECCV. Also during the reporting period, Bjoern Andres has accepted an offer of an Honorary Professorship by the University of Tübingen.

Vision and Language *Investigator: Zeynep Akata*

Deep models are frequently seen as opaque and are unable to explain their decisions. In contrast, humans can justify their decisions with natural language and point to the evidence in the visual world which led to their decisions. However, existing models which generate textual explanations enforce task relevance through a discriminative term loss function, but such mechanisms only weakly constrain mentioned object parts to actually be present in the image.

Our major achievements in this field during the reporting period are CVPR 2018 (spotlight) paper that proposed two datasets for both activity and VQA pointing and explanation tasks. In addition, we proposed improvements in ECCV 2018 over our generated textual explanations proposed in ECCV 2016 and extended it to videos in the context of self driving cars in ECCV 2018. These works were important in establishing a benchmarks and state of the art in various settings through well curated datasets, and well designed models which are generalizable across domains. Our efforts have been awarded with a DARPA grant with the title Deeply Explainable Artificial Intelligence to a team composed of UC Berkeley, University of Boston and University of Amsterdam. At the University of Amsterdam this funding supports 30% of an assistant professor and a PhD student.

In the short term, we propose to generate explanations by utilizing localized grounding of constituent phrases in generated explanations to ensure image relevance. From the perceptual input point of view, this grounding can have different forms, e.g. attention maps, bounding boxes, etc. where the important regions of the image are determined in a post-hoc manner.

A complementary direction is that, instead of producing a classification in a single step, we propose to iteratively make binary sub-decisions which, when combined as a whole, aim to produce the same classification result while revealing a decision tree as thought process. In the short term, we aim to incorporate attribute information at the class-level to give the binary decisions a semantic meaning and allowing to trade-off interpretability and classification accuracy. We anticipate that the decision tree resulting from the sequence of binary decisions would reveal a hierarchical clustering of the data and can be treated as attributes learned from scratch.

In the long term, we would like to utilize deep and explainable frameworks in more mainstream tasks. For instance, deep neural perception and control networks have become key components of self- driving vehicles. User acceptance is likely to benefit from easy-to-interpret visual and textual driving rationales which allow end-users to understand what triggered a particular behavior. As a starting point, we plan to use visual (spatial) attention model to train a convolutional network end-to-end from images to steering angle commands. The attention model identifies image regions that potentially influence the network's output. We would then apply a causal filtering step to determine which input regions causally influence the vehicle's control signal. As for the human-machine communication, we would use a video-to-text language model to produce textual rationales that justify the model's decision. The explanation generator would then use a spatiotemporal attention mechanism, which is

encouraged to match the controller's attention.

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 novel computational methods for everyday eye contact detection using unsupervised gaze target discovery (Best Paper Honourable Mention Award at UIST 2017), for fully articulated eye gaze redirection in video (Best Paper Honourable Mention Award at Eurographics 2018), for forecasting user attention during everyday mobile interactions using device-integrated and wearable sensors (Best Paper Award at MobileHCI 2018), as well as a ground-breaking new approach to mobile eye tracking using multiple tiny, low-resolution cameras and learning-based gaze estimation (Distinguished Paper Award in the Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies). We have also released OpenGaze (www.opengaze.org), the first open source toolkit for camera-based gaze estimation and interaction that aims to democratize their use in HCI and therefore designed the toolkit specifically for gaze interface designers. In the area of human-computer interaction, we have contributed a detailed analysis of the past, present, and future of gaze-enabled handheld mobile devices (Best Paper Honourable Mention Award at MobileHCI 2018), a study on the identifiability of different visual representations of oneself on public displays (Best Paper Honourable Mention Award at ACM CHI 2018), an investigation into the viability of a threat model that involves multiple shoulder surfers (Best Paper Honourable Mention Award at MUM 2017), and we have introduced, prototypically implemented, and empirically evaluated the novel concept of error-aware gaze-based interfaces for robust mobile gaze interaction (Best Paper Award at ETRA 2018).

Image and Video Segmentation *Investigators: Anna Khoreva and Bernt Schiele*

Image and video segmentation is a very challenging task and one of the most crucial steps toward scene understanding. Much progress has been made in this field over the last years. To a large extent, the success can be attributed to the strong appearance models completely learned from data, in particular using deep learning methods. However, to perform best these methods require large representative datasets for training with expensive pixel-level annotations, which for many applications are prohibitive to obtain in practice. This considerably restricts the potential to transfer these models to approach different domains or various object categories. Therefore, there is a need to relax the constraint of exploiting expensive pixel-level annotations for training and to consider alternative forms of supervision, which are easier and cheaper to collect.

In the last few years we focused on developing approaches for training segmentation networks with weaker forms of supervision, such as leveraging bounding boxes (CVPR'17) or image labels (CVPR'17) as supervision, making use of unlabeled and synthetic data (ArXiv'18) for semantic and instance labelling tasks, exploiting static images (CVPR'17) and sophisticated data augmentation techniques (IJCV'19) as well as language supervision (ACCV'18) for video object segmentation. Our contributions in this area allowed us to achieve high-quality image and video segmentation results while relying on less expensive annotations, advancing the state of the art on multiple challenging benchmarks. Major achievements in the reporting period have been our spotlight presentation at CVPR'17 (acceptance rate 9%), oral presentations at GCPR'17, workshops at CVPR'17, ECCV'18 and ACCV'18 and the second place in the 2017 DAVIS Challenge on Video Object Segmentation.

In the long term, we would like to continue working on weakly supervised approaches for segmentation task. We are investigating using conditional generative adversarial networks (GANs) for data and feature synthesis, exploiting temporal coherence and dynamics in videos as supervision as well as leveraging multi-modal data for segmentation annotation generation to resolve ambiguities caused by partial observations of a specific modality. We believe that the proposed strategies can further boost the results and help to relax the data dependency constraint.

Publications, Cooperations and Awards

Chapter 36 of the Scientific Report contains a detailed report of the publications, cooperations and awards of the reporting period. From the journal publications 16 have been published or accepted at either IJCV or IEEE PAMI. From the conference publications, 55 have been published or accepted at one of the major computer vision conferences (CVPR, ICCV, ECCV). An additional 9 have been published at NIPS, ICML, ICLR or AAAI. Another 14 papers have been published at CHI, UIST, IMWUT, and ETRA. In the reporting period (2017–2019) 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, Oxford U, ETH Zurich, EPFL, KU Leuven, Lancaster University, TU Munich, U Tübingen, U Freiburg, MPI for Intelligent Systems, as well as Researchers from Deepmind, Google, Facebook, Intel, and Amazon.

Mateusz Malinowski received the DAGM MVTec and Dr. Eduard Martin Dissertation Awards for his PhD. Siyu Tang also received the DAGM MVTec Dissertation Award. Anna Rohrbach received the Otto Hahn Medal for her PhD. Several best paper awards have been awarded to members of the group including at 3DV (Mohamed Omran, Gerard Pons-Moll and Bernt Schiele), MobileHCI (Andreas Bulling), ETRA (Andreas Bulling), and IMWUT (Andreas Bulling). Bernt Schiele has been elected IEEE Fellow as well as IAPR Fellow.

2.3 D3: Internet Architecture

History, Group Organization and Development

The Internet Architecture department (D3) was established in January 2018, and group's organization structure and details of members are provided in §37.3 of the Scientific Report.

Vision and Research Strategy

The Internet is a immensely successful human-made artifact that has fundamentally changed the society. In becoming such an immensely successful infrastructure the use of the Internet and, consequently, the Internet itself has changed and continues to change, as highlighted by some of this group's research efforts. These changes are in part driven by the user/eye-ball interests as well as by how content, including user generated content is made available. The AS-level topology of the Internet has also experienced significant changes (or "flattening") over time, i.e., it has evolved from a highly hierarchical topology to a flatter (non-hierarchical, simpler) topology. In contrast, content providers and content delivery networks are relying on sophisticated back-office infrastructures that include crawlers, caching hierarchies, as well as infrastructure to deliver advertising. The future challenges in this context are (i) continual observation of the usage as well as the underlying infrastructure, (ii) locating current performance bottlenecks within the infrastructure, (iii) understanding how novel applications interact or should interact with the infrastructure, (iv) designing network management mechanisms to minimize the need for manual configuration within the infrastructure as well as the across infrastructures and automating security mechanisms for protecting our infrastructure, and lastly (v) incentivizing efficient network usage and network upgrades.

We will be generating staggering volumes of data in a few years, if not already doing so, almost everywhere. Our analytical processing capabilities will also have further advanced, e.g., offer intelligent machine learning mechanisms, in this time frame. The increasing demands of users to have a ubiquitous access to information from anywhere at any time, in a context where data is generated in a highly distributed manner (i.e., in diverse geographical locations and networks) poses unique, extremely hard challenges to solve. A feasible solution will have to enable data streams to be processed and distributed in a coordinated manner in real time. Such an approach requires a distributed processing platform where both computations and data can move around freely as well as securely in an optimal fashion providing fast reaction times and minimal resource usage. In this context data provenance, quality criteria, and time constraints, both varying per customer, will have to be taken into account, necessitating the integration of information processing and networking into a single paradigm.

Part of the success of the early Internet is that it relied on "working code and rough consensus." The Internet infrastructure today is, unfortunately, not quite as simple as before, since the performance requirements—100 Gbps to 10 Tbps—of today's networks, particularly, in the core of the Internet require custom hardware solutions. This requirement has lead to the use of specialized hardware and software and closed-box solutions for the main components of the Internet—the routers and switches—thereby resulting in the ossification of the Internet. Software defined networking is one way in which to tackle this problem and some of the challenges in this are include (i) taking advantage of the software capabilities, (ii) supporting

wide-area data analytics, and (iii) integrating network resources, namely storage, CPU, and data, into the concept of SDN.

Our approach to tackle these challenges is to follow a data-driven systems research agenda: collect data from operational networks, analyze them using big-data analytics to identify invariants, revisit assumptions, and detect where current performance bottlenecks in the Internet are. We also employ simulation environments to validate our analysis results and to enable "what-if" studies. The insights obtained from the measurements form the foundation for shaping the future Internet via optimizations and alternative designs. Hereby, our output includes protocol enhancements, novel network management tools, and concepts for software-defined networking.

Research Areas and Achievements

Our research work spans three broad areas, (1) Internet measurement, (2) future-proofing the Internet, and (3) wide-area data analytics.

Internet Measurement

The Internet is massively heterogeneous and also continually evolving ecosystem. Naturally, no one vantage point can accurately capture the breadth of these changes. This limitation notwithstanding, there is a dire need to monitor and analyze the use of the Internet infrastructure as well as the characteristics of the infrastructure itself, especially given the constantly evolving nature of the ecosystem. Meticulously gathering measurements from diverse vantage points, and systematically analyzing these measurements to characterize the performance and operation of the Internet ecosystem, identify vulnerabilities and issues in the different components of this ecosystem, e.g., network protocols and devices, characterize the use of the infrastructure, e.g., through analyses of network traffic volume and dynamics, and understand how to evolve and upgrade the networks for the future constitutes a core research focus of our department. Analyzing the use of BGP communities in the wild and the routing vulnerabilities they introduce, leveraging BGP communities to mitigate large-scale distributed denial-of-service (DDoS) attacks, and analyzing the traffic volume asymmetry at scale are some of the ongoing research efforts in this space. In addition to providing crucial insights into the current state of the Internet, our measurement-driven systems design approaches reveal how to also upgrade and evolve the Internet infrastructure in anticipation of novel applications and changing usage patterns.

Future-proofing the Internet

The continually evolving nature of the Internet ecosystem routinely introduces new, unforeseen challenges. With the increasing adoption of sensors and Internet of Things (IoT) devices, we are generating an unprecedented volume of data. If the generated data volumes are not challenging enough, that such data now comes from diverse endpoints, widely distributed throughout the network at its edge, makes the collection and analysis of this data in real time poses a grand challenge for the networking community. Addressing the challenge necessitates the design and development of new tools and techniques, or even better computing primitives that can make scalable, accurate, real-time analyses of such data feasible in practice. The

changes in the Internet ecosystem also provide unprecedented opportunities to revisit some long standing networking problems and design practical, scalable solutions that exploit or leverage these recent changes. Redesigning protocols or applications by exploiting new opportunities facilitates the applications to better adapt to the changing network; the infrastructure in turn benefits from having well-designed protocols and applications that make more efficient use of the infrastructure.

Wide-area data analytics

We are generating an astounding volume of data in the Internet, and virtually all forecasts indicate that the amount of data we will generate, gather, store, and analyze will only continue to increase at a staggering pace. The tools and techniques to analyze this data in a systematic, scalable, and practical manner, whether for security audits or for mundane operations such as infrastructure planning, however, has not kept pace with the growth in data generation. Besides, data sources today are often widely distributed across different networks and geographies, making the already hard problem even harder. Aggregating data from widely distributed vantage points necessitate might necessitate techniques to compress or consolidate the data to minimize the volume of data transferred from different locations to a centralized location for analyses. It is also often hard to know a prior the type and nature of queries that will be run against these data sets and this limitation highlights the need for generic tools and techniques to represent data in a succinct form while still remaining amenable to be queried in a scalable manner. Alternatively, we need to also explore about how to decompose a query to run in a truly distributed manner or reorganize, reshape, and relocate data in such a way to reduce the query time without sacrificing much accuracy.

Prizes and Awards

Prof. Anja Feldmann, Ph.D.:

- Vodafone Stiftung: Vodafone Innovationspreis, 2018
- Bayerische Akademie der Wissenschaften: Schelling Prize, 2018
- TU-Berlin: Honorarprofessor, 2018
- 2019 IETF/IRTF Applied Networking Research Prize (ANRP) for paper titled “BGP Communities: Even more Worms in the Routing Can” published in the Proceedings of the ACM Internet Measurement Conference (IMC) 2018
- Best Paper award for paper titled “Taking Control of SDN-based Cloud Systems via the Data Plane” published in the Proceedings of the ACM Symposium on SDN Research (SOSR) 2018

Press and Media Coverage

Prof. Anja Feldmann, Ph.D.:

- *3SAT: Scobel on Blockchain*, http://www.3sat.de/mediathek/?obj=73309

– Interview in *Süddeutschen Zeitung*, 30.04.2018, http://www.sueddeutsche.de/digital/interview-am-morgen-jahre-www-wir-brauchen-leute-die-nicht-alles-glauben-was-ihnen-im-internet-gesagt-wird-1.3960454

Teaching

Data Networks (A. Feldmann and B. Chandrasekaran)

2.4 D4: Computer Graphics

Group Overview

The computer graphics group (D4) was established in 1999 with the appointment of Hans-Peter Seidel. Karol Myszkowski joined in 2000, Christian Theobalt joined in 2009. Over the last two decades the group graduated more than 60 PhD students, and more than 40 former group members got offers for tenured faculty position. During this period our academic offspring received a variety of prestigious grants and awards, including two ERC Consolidator Grants[6], eight ERC Starting Grants[7], four DFG Emmy Noether Fellowships[8], seven EG Young Researcher Awards[9], and three German Pattern Recognition Awards[10].

At the time of writing, the group encompasses two senior researchers with tenure (K. Myszkowski and C. Theobalt), four group leaders and senior researchers without tenure (R. Chen (GL since 2015, SR since 2018), V. Babaei (GL since 2018), G. Singh (GL since 2019), and R. Zayer (GL since 2018)), six postdocs, and 18 PhD students. Six group members completed and handed in their PhD thesis during the reporting period, and several of our young researchers and former PhD students got offers for faculty appointments[11] or postdoc positions[12]. Full details on current and recent group members are provided in Section 38.1.

Vision and Research Strategy

During the last few decades computer graphics firmly established itself as a core discipline within computer science. New and emerging technologies such as digital media, social networks, digital photography and the rapid development of new sensing devices, 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 and novel concepts from machine learning to extract powerful priors that can assist us during the acquisition, reconstruction, editing, and image formation processes.

Our vision and long term goal are completely immersive, interactive, and visually rich environments with sophisticated scene representations and the highest visual quality, fused

[6] R. Mantiuk (2016), C. Theobalt (2017)
[7] M. Magnor (2010), B. Rosenhahn (2011), B. Goldlücke (2013), C. Theobalt (2013), A. Oulasvirta (2014), J. Gall (2015), J. Steimle (2016), P. Didyk (2018)
[8] J. Kautz (2003), H. Lensch (2007), I. Ihrke (2012), J. Gall (2013)
[9] H. Lensch (2005), J. Kautz (2007), M. Goesele (2008), C. Theobalt (2009), E. Eisemann (2011), T. Ritschel (2014), B. Masia (2017)
[10] B. Rosenhahn (2009), C. Theobalt (2012), J. Gall (2014)
[11] D. Casas (Univ. Madrid, Spain, 2018), A. Chatterjee (IIT Kharagpur, India, 2017), R. Chen (University of Science and Technology of China (USTC), 2019), P. Didyk (USI Lugano, Switzerland, 2018), C. Jiang (Xi'an Jiatong University, 2019) H. Rhodin (Univ. British Columbia, Vancouver, Canada, 2019)
[12] P. Garrido (Technicolor), P. Kellnhofer (MIT), T. Leimkühler (INRIA), S. Sridhar (Stanford), M. Zollhöfer (Stanford)

seamlessly with the real world. Standard 2D screens are being replaced with high dynamic range displays, stereo and automultiscopic screens, portable and wearable displays. Imaging algorithms with embedded perceptual models ensure that the perceived quality and viewing comfort is maximized. Interaction is intuitive and light weight.

In order to make progress along the lines above, our work is both theoretical and practical with a focus on first-class research and new methods and algorithms, as well as on the integration of new algorithms into functioning software systems, and the experimental validation of systems in specific application scenarios that are of practical relevance. We also try to provide a stimulating environment for junior researchers that allows them to develop and build their own research programs and groups.

Research Areas and Achievements

Our research is currently organized into the following six research areas (coordinators in brackets):

- Digital Geometry Processing (R. Chen, R. Zayer, and H.-P. Seidel)
- Computational Videography (C. Theobalt)
- Reconstructing the Static and Dynamic Real World (C. Theobalt)
- Realistic and Real-time Rendering (G. Singh)
- Perception and Advanced Displays (K. Myszkowski)
- Computational Fabrication (V. Babaei)

However, we are not organized into disjoint subgroups, and there is little hierarchy. While each of the areas has its specific focus, some of them also have significant overlaps. Likewise, the students and researchers working in each area are dynamically formed teams rather than specifically dedicated staff. The senior researchers and group leaders together with Hans-Peter Seidel serve as an internal steering committee for the group. They also act as advisors or co-advisors of doctoral students.

Digital Geometry Processing

Digital geometry processing is concerned with the representation, analysis, manipulation, and optimization of digital shapes. During the reporting period we developed a novel sparse matrix representation for unstructured grids and demonstrated its power in different geometry processing applications, and we substantially improved the performance of the underlying linear algebra kernels. We also continued our successful work on shape deformation.

A key advantage of working with structured grids is the ability to directly tap into the powerful machinery of linear algebra. This is not much so for unstructured grids where intermediate bookkeeping data structures stand in the way. We developed a novel sparse matrix representation for unstructured grids and demonstrated its power in different geometry processing applications such as high performance surface tessellation and high performance mesh subdivision. Our representation reduces the memory storage requirements but also cuts down on the build of data movement from global storage to the compute units. Our framework translates geometric computations and topological modifications into concise linear algebra operations, and we capitalize on the nature of sparse matrix-vector multiplication which allows avoiding explicit transpose computation and storage (EG'17, SIGAsia'18, CoRR'18).

Linear algebra kernels and graph algorithms are the back bone of modern high performance computing, and our work above demonstrates the power of these kernels for high performance geometry processing. In a sequence of papers we have specifically addressed sparse matrix-vector product (SpMV), general sparse matrix-matrix multiplication (SpGEMM), sparse matrix assembly, and dynamic graph data structures, and substantially improved the performance of these key operations (2xHPEC'17 (Best Paper Award), 2xICS'17, SC'18, ACM PPoPP'19).

We also continued our successful work on shape deformation and mapping optimization and proposed an interactive deformation system that is driven by user specified positional constraints. Our method is based on an unconstrained minimization of isometric energies, and is guaranteed to produce smooth locally injective maps by operating within a reduced dimensional subspace of harmonic maps. We also extended the approach to triangle mesh discretizations on surfaces (SGP'17, SIGAsia'17, CGF'18).

In other works, we investigated the design and optimization of space structures and path planning with divergence-based distance functions (SIG'17, SGP'18, CAGD'18).

Computational Videography

Advanced computational videography aims to provide algorithms and tools for the realistic editing of live videos. We developed *Deep Video Portraits*, the first approach that enables photo-realistic re-animation of the entire posture and face expression of a human in a portrait video using only an input color video. The core of the approach is a generative adversarial neural network with a space-time architecture that converts computer graphics renderings of a (reconstructed) face model into photo-realistic and temporally-coherent video. We also addressed the gaze-aware reenactment of stereo video content for VR telepresence applications (*FaceVR*). The developed approach makes it possible to computationally remove the VR headset from a stereo face video of a person, while reproducing the full face expression and eye gaze "behind" the display. We also developed a novel learning-based approach for the animation and reenactment of human actor videos (SIG'18, TOG'18, TOG'19).

In a second line of research we developed new advanced approaches for inverse rendering, i.e., estimation of lighting and reflectance from monocular color and depth video. Examples are the first real-time method for BRDF estimation from monocular RGB, as well as the first real-time approach for live user-guided intrinsic video (TVCG'17, CVPR'18).

Reconstructing the Static and Dynamic World

We advanced the state-of-the-art of both static and dynamic scene reconstruction methods along several important dimensions: the generality of scenes that can be handled, the accuracy and quality of the reconstruction, the efficiency and robustness of reconstruction, and the simplicity of sensors needed for capturing. Notably, we researched entirely new ways of fusing and deeply integrating model-based and deep learning-based scene reconstruction. This paves the way for a new generation of future real world reconstruction, perception and understanding methods that can learn and continuously refine their internal algorithmic concepts and model representations on even sparsely labeled or unlabeled real world examples.

For *static scene reconstruction* we developed *BundleFusion*, the first approach for globally-consistent online reconstruction of large scale scenes with an RGB-D camera. Global consistency

in the captured geometry is achieved by real-time bundle adjustment on a truncated signed distance field that is dynamically updated and effectively stored in space and time. We also developed *CurveFusion* to scan very thin and tubular structures, and we proposed a new algorithm for point cloud denoising (EG STAR'18, TOG'17, SigAsia'18, 3DV'18).

We greatly advanced the capabilities of marker-less *human performance capture* methods, as well as methods *capturing general deformable objects* with cameras. Our *MonoPerfcap* algorithm is the first marker-less approach for temporally coherent 3D performance capture of a human with general clothing from monocular RGB video. We tackle this challenging problem by using a novel approach that employs sparse 2D and 3D human pose detections from a convolutional neural network using a batch-based pose estimation strategy and then refining the surface geometry based on fully automatically extracted silhouettes. While *MonoPerfcap* is an off-line method, *LiveCap* is the first template-based real-time human performance capture approach that reconstructs dense, space-time coherent deforming geometry of entire humans in general clothing from just a single RGB video. We also further advanced performance capture from multi-view video (3DV'18, CVPR'18, TOG'18, TOG'19, IJCV'17).

We proposed new methods for substantially improved marker-less *3D skeletal motion capture*, without surface reconstruction, from video. We developed *VNect*, the first method for accurate and real time 3D skeletal motion capture from monocular RGB video. It combines a learning based approach with generative fitting of a kinematic skeleton model. The system uses a novel fully-convolutional 3D pose estimation approach which strongly couples 3D pose inference of each body joint to its image evidence. We also proposed the first real-time approach for the egocentric estimation of 3D human body pose in a wide range of unconstrained everyday activities (PAMI'17, 3DV'17, 3DV'18, SIG'17, IEEE VR'19).

Reconstruction of high-quality 3D models of *moving human faces* is essential for the creation of digital human actors. We proposed *MoFA*, the first unsupervised approach for monocular reconstruction of the identity geometry, albedo texture, face expression, and scene illumination from a single color image. This was made possible by a new *model-based face autoencoder* which tightly integrates in a new way a 3D parametric face model, differentiable image formation and layers of a CNN. The algorithm can be trained on unlabeled community image data. In later follow-up work we substantially extended and improved the approach. A different approach to monocular reconstruction of 3D faces is presented in our *InverseFaceNet* algorithm (ICCV'17, PAMI'18, CVPR'18, SIG'18, EGSTAR'18, CommACM'19).

We have developed state-of-the-art techniques to *track hands*, in real-time without markers or gloves, and using only a single RGB-D or color camera. The deep methodological integration of model-based and learning-based reconstruction also enable some of the first methods that handle non-trivial hand-object interactions in cluttered scenes (ICCV'17, CVPR'18).

We also investigated *foundational algorithmic questions* that are of cross-cutting relevance to key aspects of both expert designed and learned image-based reconstruction methods in particular, and visual computing or machine learning algorithms in general, including scalable multi-matching (TOG'17, CVPR'18, CoRR'18).

Realistic and Real-time Rendering

During the reporting period we have focused on the simulation of distribution effects and on the analysis and synthesis of sample correlations.

Distribution effects like depth-of-field and motion blur are an important factor to cinematic quality in synthetic images. Drawing point spread functions (PSF) for every pixel is a general algorithm for achieving high-quality distribution effects. While this is usually an order of magnitude faster than classical Monte Carlo (MC) methods, it is still not suitable for applications requiring interactive feedback. We proposed a novel approach, based on a sparse representation of the Laplacian of the PSFs, that significantly accelerates splitting of PSFs, and we suggested a novel efficient framework to produce complex distribution effects by exploiting coherency among RGB-D images rendered via pin-hole cameras (CGF'17, SIG'18).

Real-time renderers are paving their way towards full global illumination pipelines (e.g., in the game industry). However, much is left to understand regarding sample correlations that directly affect the quality of the images (with aliasing or noise) irrespective of the underlying application. We have achieved significant progress on MC based sampling techniques that are critical in approximating high dimensional light integrals representing radiant light energy. These works establish sound theoretical formulations that can represent error in closed-form and propose the first deep learning framework that can synthesize samples with good properties (low discrepancy, blue noise spectra with anisotropy and progressivity) without any special hand-crafting or mathematical intricacies (SIG'17, CGF'19, CGF'19, CoRR'18).

Other works address novel view synthesis (PAMI'17), expressive single scattering for light shaft stylization (TVCG'17), deep shading (EGSR'17), and CPU vector graphics rendering (Eurographics'19). For our previous work on stackless KD-tree traversal for GPU ray tracing we received the High Performance Graphics (HPG) Impact Award 2017.

Perception and Advanced Displays

Modeling the dynamics of perceptual mechanisms in the human visual system (HVS) offers great promise for designing novel and advanced display systems which can attain much higher visual quality and viewing comfort. We introduced an award-winning advanced near-eye display (NED) design, and we developed novel fundamental algorithms for perception-aware gaze-contingent rendering and for the automatic generation of high-quality 3D content. Finally, we developed a novel CNN-based metric for light field quality evaluation which outperforms existing solutions.

We improved the state-of-the-art in displaying high-quality visuals by introducing an advanced NED design, which addresses the lack of accommodative cues in conventional displays. Enabling accommodative cues is expecially critical in AR applications where computer-generated overlays must be properly combined with real world stimuli. Our display prototype, which has been designed in collaboration with UNC, essentially solves this problem with a novel deformable membrane half-mirror setup. The membranes are installed on air-tight chambers and the optical principle behind our design is approximating spherical concave surface whose curvature is controlled by changing the air pressure inside the chamber (changing the virtual depth level). The design provides a wide field of view, and the approach successfully resolves the well-known mergence-accommodation conflict. Our design received the Best Paper Award at IEEE VR'17, and the Emerging Technologies DCEXPO Special Prize at SIG'17.

Gaze-contingent rendering in combination with the use of eye tracking focuses the rendering budget on those areas where it is most needed. The main bottleneck is the limited sampling and refresh rates of affordable eye trackers and displays, resulting in an increased latency. We

developed a model for predicting the landing positions of saccade eye movements that solves this problem (CGF'17, SIG'17, PACMCGIT'18, TVCG'19). Moreover, we developed novel fundamental algorithms for the automatic generation of high-quality 3D content, based on Mono2Stereo and Stereo2Multiview Image conversion (SIG'17, TMM'18, TVCG'18).

As light fields are gaining in popularity, there is a need to measure and quantify light field quality in the processing pipeline. We collected an extensive dataset of reference and distorted image pairs together with user markings indicating whether distortions are visible or not. Using this dataset, we developed a novel CNN-based metric for light field quality evaluation which outperforms existing solutions (TOG'18, CVPR'19).

Computational Fabrication

Our interest is in the visual appearance of objects. Similar to 2D printers, multimaterial 3D printers are also binary devices where only a single material with a fixed concentration can be deposited in every volume element. Our "contoning" approach eliminates previous halftoning artifacts by printing layers of transparent inks around an opaque body made of white, diffuse material and mixing the colors within the volume of the object. Capturing and fabricating the object's spectral reflectance in the visible range of light is called spectral reproduction. We focused on the faithful reproduction of paintings and used a deep neural network to both predict the spectrum of a printed ink stack and optimize for the stack layout that best matches a target spectrum. Current printing materials for inkjet 3D printing suffer from unwanted translucency that results in volumetric light transport (cross-talk) between surface points and severe blurring of details. We employed a general Monte-Carlo simulation of heterogeneous scattering that predicts the color at each location on the surface (SCF'17, SIG'17, SIGAsia'17, SIGAsia'18).

In other works we also addressed the perception-aware modeling and fabrication of drawing tools and the design and analysis of directional front projection screens (SCF'17, SIG'18).

Software and Datasets

As part of our research we have developed several libraries, tools, and large corpora of reference data sets that are being made available to the research community at large. These include the GVVPerfCapEva repository of human shape and performance capture datasets, the PFSTOOLS for processing high dynamic range images and video, the LocVis dataset of locally annotated images, and the MPI light field archive. My former PhD students N. Hasler, C. Stoll and M. Richter, together with C. Theobalt, also formed the startup The Captury, a spin-off company commercializing marker-less motion and performance capture.

Some Performance Indicators

The group continues to be a leader in computer graphics, with a specific focus on the integrated view of 3D image analysis and synthesis. Our work is being referenced widely, and we are among the trend-setters in this direction.

During the reporting period we have again made significant progress in our research along the lines above, and members of the group have actively published in the top conferences and journals in the field. As a syntactic indicator, within the two-year reporting

period from spring 2017 to spring 2019, members of the group published 18 papers at SIGGRAPH/SIGGRAPH Asia/ACM TOG, 13 papers at Eurographics/CGF/TVCG, 19 papers at ICCV/ECCV/CVPR/3DV/IJCV/PAMI, and 4 papers at ICS/SC (see Section 38.15 for details). We have actively participated in the program committees of major conferences (SIGGRAPH, SIGGRAPH Asia, Eurographics, Pacific Graphics, SGP, EGSR, SCA, HPG, Graphics Interface, CAD/Graphics, VMV, ICCV, ECCV, CVPR, 3DV, GCPR), and we hold editorial board seats with journals such as ACM TOG, ACM TAP, IEEE TVCG, IEEE CG&A, CAGD, GMOD, Int. J. Shape Modeling, and Visual Informatics. We have given numerous invited talks and tutorial presentations at major national and international events (see Section 38.11.3 for details). Our software and datasets have been successfully used in a variety of projects (see Section 38.10), and many young researchers from the group have spread out to other institutions.

Projects and 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, Princeton, UNC Chapel Hill, U. Hong Kong, U. Tokyo, Aalto U., Bar-Ilan U., U. Bath, Cambridge U., CNRS, TU Delft, U. Edinburgh, U. Erlangen, Fraunhofer IIS, IST Austria, TU Graz, Technion Haifa, UCL London, EPFL Lausanne, USI Lugano, Charles U. Prague, U. Zaragoza, ETH Zürich. In addition, we also collaborated with some leading industrial research labs, including Google, Microsoft, Occulus, Intel, Ndivia, Disney, Adobe and Technicolor. Formal cooperations exist with Stanford (Max Planck Center for Visual Computing and Communication (MPC-VCC)), Fraunhofer IIS (Perceptually-aware light field capture, processing and display), and the RISE Research Centre of Excellence in Cyprus. We also participate in the Horizon 2020 Training Networks DISTRO and RealVision. 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 Saarbrücken Graduate School of Computer Science.

Awards

Hans-Peter Seidel received the Eurographics Gold Medal (2017), and the High-Performance Graphics (HPG) Impact Award (2017). Christian Theobalt received an ERC Consolidator Grant (2017), and he was awarded the prestigious Karl Heinz Beckurts Prize (2017). Piotr Didyk received an ERC Starting Grant (2018). Karol Myszkowski was awarded the 2016 IEEE TVCG Best Associate Editor Award (announced Dec. 2017).

Pablo Garrido (2018) and Petr Kellnhofer (2018) both won a Eurographics Dissertation Award. Petr also won the Eduard Martin Dissertation Prize of Saarland University (2017). Franziska Müller received a Google PhD Fellowship (2017), and the national Women STEM Award (2017). Marc Habermann received the Günther Hotz Medal of Saarland University for the best Master students in computer science (2017). Petr Kellnhofer, Piotr Didyk, and Karol Myszkowski (IEEE VR'17 and SIGGRAPH'17 Emerging Technologies DCEXPO Special Prize), Thomas Leimkühler (EGSR'17), Franziska Müller and Christian Theobalt (ACM ISS'18), and Rhaleb Zayer (IEEE HPEC'17) all won best paper awards at major international conferences.

2.5 D5: Databases and Information Systems

D5 was established in 2003. It is headed by Gerhard Weikum and, as of May 2019, consists of 21 doctoral students, 5 senior researchers (Klaus Berberich, Simon Razniewski, Rishiraj Saha Roy, Jilles Vreeken and Andrew Yates) and 6 other post-doctoral researchers. During the reporting period 2017–2019, three senior researchers (Pauli Miettinen, Daria Stepanova, Jannik Strötgen) left the group for faculty(-like) positions at universities or research labs.

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., varying beliefs about 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 five technical areas:

- Knowledge Base Construction and Curation
- Data Mining and Exploratory Data Analysis
- Information Retrieval, Information Extraction and Text Analysis
- Natural Language Understanding
- Privacy, Trust and Fairness

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., Amazon, 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 nearly 3000 citations, and the Yago2 paper from 2013 already has more than 900 citations. The influence and value of Yago has been recognized by the research community through the AIJ Influential Paper Award 2017[13] (for the 2013 Yago2 paper in the Artificial Intelligence Journal) and the W3C Seoul Test of Time Award 2018[14] (for the original WWW 2007 paper on Yago).

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

[13] http://aij.ijcai.org/aij-awards-list-of-previous-winners
[14] https://www.iw3c2.org/updates/ToT/

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 about 700 citations. This work has spawned a startup called Ambiverse and the development of the ambiverseNLU software suite, available as open source code.

Contributions and Impact: Major Results 2017–2019

We publish our results in top-tier conferences in several communities: Web research (Web), data mining (DM), database systems (DB), information retrieval (IR) and natural language processing (NLP). In the two-year timeframe 2017–2019, the group had 17 full papers in first-rate Web venues (WWW, WSDM, ISWC, ICWSM, TWEB, IJCAI, AAAI), 12 in DM (KDD, ICDM, SDM, TKDD, NIPS), 5 in DB (VLDB, ICDE, EDBT, TODS), 6 in IR (SIGIR, CIKM, ECIR), and 9 in NLP (ACL, EMNLP, COLING, NAACL). In addition, we are successful in building prototype systems and publishing demo papers at top venues like ACL, VLDB, SIGIR, WWW, etc. Three of these publications won best paper awards at top-tier conferences (EMNLP 2017, WWW 2018, ICDM 2018).

Rule Mining from Knowledge Bases and Web Text: We successfully continued our research on learning rules for intensional world knowledge and applying them to the task of knowledge base curation (with publications in ISWC 2017, ISWC 2018, IJCAI 2018, WSDM 2019, WWW 2019). Most notably, we introduced a notion of (in-)completeness awareness to rule mining algorithms, and we incorporated latent embeddings to overcome gaps in the underlying knowledge base. The latter involved tapping into text and web sources, and we further enhanced this theme to compute evidence and generate explanations for doubtful statements which are either confirmed or refuted using rule-based inference.

Discovering Dependencies in High-Dimensional Data: We developed new methods for discovering approximate dependencies between attributes in multivariate datasets such as census data, economic statistics or materials properties. In contrast to traditional approaches of analyzing correlations, our methods can produce robust and interpretable outputs, this way explaining which attributes depend on which other attributes. Experimental results include studies with material scientists, and these colleagues (at the Fritz Haber Institute in Berlin) got truly excited about the findings and their interpretability. Publications on these results include papers at KDD 2017 and ICDM 2018 (winning the best paper award) as well as an article in the New Journal of Physics 2017.

Quantities in Tables, Text and Knowledge Bases: We are among the first to address the issue of extracting and semantically organizing quantities that appear in ad-hoc tables and text sources on the Web. These are numeric expressions that denote measures with units and a reference frame, such as annual revenue of companies, battery-only range of hybrid cars, carbon footprint of data centers, or conductivity of thermoelectric materials. We have developed

methods for linking noisy expressions of quantities in text and tables to a knowledge base (published in CIKM 2016), and for aligning approximate and aggregate quantities between text and tables in rich-content reports to support analysts in exploring details and spotting explanations of numeric facts (ICDE 2019 and WWW 2019). As quantities are often part of higher-arity relations (e.g., dosages of drugs for specific patient groups with specific disorders), we devised new information extraction methods that go beyond the traditional binary-relation model (WWW 2018 best paper). We also developed specific methods for mining counting quantifiers from natural language sources, in order to assess the coverage of relations in knowledge bases and to aid knowledge completion (ACL 2017 and ISWC 2018).

Question Answering over Knowledge Bases and Web Text: We revived our earlier research on answering fact-oriented questions posed in natural language (incl. telegraphic queries), with both knowledge bases and text corpora as underlying data. To tackle the vocabulary mismatch between phrases in questions and the entities, types, predicates and concepts in a knowledge base, we enhanced template-based methods for translating questions into structured queries with "never-ending learning" from user feedback (incl. pseudo-relevance) and semantic similarities among questions. Recently, we also devised a novel approach to question answering based on dynamically building quasi knowledge graphs from Web content and using advanced graph algorithms (Group Steiner Trees) to compute answers. This method is completely unsupervised, can handle a wide variety of user inputs (incl. ungrammatical utterances), and can harness both text and structured data and knowledge in a unified manner. Publications on this research theme include papers at WWW 2017, EMNLP 2017, VLDB 2018, WWW 2018, NAACL 2019 and SIGIR 2019.

Credibility and Trust of Claims: In the context of the ERC Synergy Grant imPACT, we address the issue of trust in online information, particularly, aiming to analyze and assess the credibility of textual statements in news, discussion forums and other kinds of online communities. This research led to the doctoral thesis of Subhabrata Mukherjee which won the runner-up position for the KDD dissertation award 2018. More recent work (with publications in WWW 2017, WWW 2018 and EMNLP 2018) focused on incorporating web evidence for estimating the credibility of claims made in news or by influential people (politicians, bloggers, partisan media etc.). We devised a suite of methods, including pipelined classifiers, probabilistic graphical models for joint inference on source trustworthiness and content credibility, all the way to deep-learning networks with tailored attention mechanisms.

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 Mario Boley (at Monash University, Australia), 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), Mouna Kacimi (at U Bolzano, Italy), Sebastian Michel (at U Kaiserslautern, Germany), Pauli Miettinen (at U Eastern Finland), Ndapa Nakashole (at UC San Diego, USA), Thomas Neumann (at TU Munich, Germany), Maya Ramanath (at IIT Delhi, India), Ralf Schenkel (at U Trier, Germany), Fabian Suchanek (at U Telecom ParisTech, France), Martin

Theobald (at U Luxemburg), and further alumni who have obtained faculty positions. A good fraction of our graduates joined industrial research labs, for example, Ralitsa Angelova (Google), Asia Biega (Microsoft), Maximilian Dylla (Google), Subhabrata Mukherjee (Amazon), Josiane Parreira (Siemens), Daria Stepanova (Bosch), Niket Tandon (Allen Institute for Artificial Intelligence) and Mohamed Yahya (Bloomberg). In the two-year timeframe 2017–2019, we had 10 doctoral students graduating.

Awards

Jilles Vreeken received the IEEE ICDM Tao Li Early Career Award at ICDM 2018, the data mining community's first award for recognizing the excellence and impact of young researchers at most ten years after graduation. Subhabrata Mukherjee received the ACM KDD Dissertation Runner-Up Award 2018.

Patrick Ernst (and his co-authors Amy Siu and Gerhard Weikum) won the Best Paper Award at the Web Conference (WWW) 2018. Panagiotis Mandros (and his co-authors Mario Boley and Jilles Vreeken) won the Best Paper Award at the ICDM 2018 Conference. Andrew Yates (and his co-authors Arman Cohan and Nazli Goharian) won the Best Paper Award at the EMNLP 2017 Conference. Dhruv Gupta and Thomas Tanon (and their respective co-authors) obtained honorable mentions for best demo and best student paper awards at WWW 2018 and ISWC 2017, respectively. Francesca Lisi and Daria Stepanova won the best poster award at the International Conference on Rules and Reasoning 2017.

The long-term impact of Yago has been recognized by the AIJ Influential Paper Award 2017 (for the 2013 Yago2 paper in the Artificial Intelligence Journal) and the W3C Seoul Test of Time Award 2018 (for the original WWW 2007 paper on Yago).

Teaching

The department is also committed to a fair amount of teaching at Saarland University, including the regularly offered core lecture on Information Retrieval and Data Mining. In addition, a total of 26 Bachelor's and Master's theses in the 2017–2019 timeframe were completed under the supervision of the group's senior researchers.

Cooperations

Institute and Max Planck Society: Within the institute, we collaborate with D1 on efficient algorithms for large graphs and matrices, and with D2 on background knowledge for understanding visual contents. 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 from the CISPA Helmholtz Center and Jens Dittrich from Saarland University. There are various individual collaborations with researchers at other universities and research labs across the world (including U Aberdeen, Allen Institute for AI, Amazon Research, VU Amsterdam, NU Athens, Bosch Center for AI, Northeastern U Boston, U Cyprus, IIIT Delhi, FORTH Heraklion, Georgetown U, U Melbourne, U ParisTech, Southwest Jiaotong U, Volkswagen Research, Warwick U, U Washington, 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 Machine 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. For advanced search, conversational bots, visual understanding and next-generation AI, computers need to be equipped with broader kinds of knowledge: quantitative modifiers of facts, plausibility invariants, properties of everyday objects, human activities and spatio-temporal as well as socio-cultural context.

Responsible Data Science: With the amazing advances of AI and its strong role in data science across all fields, managing this powerful technology in a responsible manner often seems to be a mere afterthought. We plan to give this neglected theme its due prominence by research on various aspects of trustworthy and responsible computing. AI systems in human-centric applications must prevent obscured, biased and unfair consequences in real-life situations. Examples where this is of utmost importance are when data-driven machine learning is used for algorithmic decision making, such as automation of loan requests, visa approvals, insurance claims, job interviews, all the way to AI components for e-government and political decisions. It is crucial that such systems be trustworthy and transparent and have explainable behavior, and that they deal with data about people in an unbiased and fair manner.

2.6 RG1: Automation of Logic

History

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 5 researchers and 4 PhD students in the group. Ching Hoo Tang, Andreas Teucke and Daniel Wand have finished their PhD theses during the reporting period and left the group. Ching Hoo Tang and Andreas Teucke joined the group's spin off L4B, while Daniel Wand accepted an offer by Bosch. Marco Voigt has submitted his PhD thesis and Martin Bromberger and Mathias Fleury will submit by the end of 2019. Thomas Sturm and Jasmin Blanchette have appointments at CNRS and VU Amsterdam, respectively, but remain associated with the group.

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. Sophie Tourret and Uwe Waldmann work on higher-order logic automation.

Vision and Research Strategy 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) drive the development of first-order, and beyond first-oder reasoning calculi, (ii) understand the mechanics of reasoning with respect to concrete models, (iii) build specific reasoning procedures for arithmetic theories, (iv) show applicability of our methods to reasoning challenges from other areas, (v) formally verify our own results, and (vi) scale the applicability of our methods to the size of real world industrial applications.

About half of our work is of a 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 results on sufficient tests for the satisfiability LIA (Linear Integer Arithmetic) constraints, Section 40.4.3 of the Scientific Report, have been implemented in SPASS-SATT, Section 40.8.5 of the Scientific Report, winning the QF_LIA category of SMT-COMP 2018; the development of subtropical methods for real constraints, Section 40.4.2 of the Scientific Report, and its implementation in the computer algebra system Redlog yielded new insights in bioinformatics. In 2012 we founded the company "Logic for Business" (L4B) providing consulting and software for the overall life-cycle management of complex products.

We have structured our research along the following topics: (i) Foundations of Automated Reasoning, (ii) (First-Order) Arithmetic, (iii) Towards Higher-Order Automated Reasoning, (iv) Formalizing Logic, (v) Logic for Machine Learning, (vi) Software. The structure does not impose a structure on the group. In fact, most of us contribute to several areas and almost all of us to the development of software.

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

New Decidable Fragments of First-Order Logic (JLC 2017, LICS 2017, CADE 2017). The idea of separated variables, already reported during the previous reporting period, has shown further potential. It yields strict, decidable extensions of the monadic fragment, the Bernays-Schönfinkel fragment, the two-variable fragment, the Ackermann fragment, the Gödel-Kalmár-Schütte fragment, the fluted fragment, the guarded fragment, the loosely guarded fragment, and the guarded negation fragment. We have also extended decidability results for two clause fragments towards recursive predicates, and non-Horn clauses, respectively.

Fast Decision Procedures for LIA Constraint Solving (FMSD 2017, IJCAR 2018). We have continued our research on fast decision procedures for linear arithmetic in the context of automated reasoning. Through the so-called *Double-Bounded-Reduction* and afterwards applied *Mixed-Echelon-Hermite* transformation, any unbounded linear arithmetic problem can be transformed in polynomial time into a fully bounded problem. Solutions of the resulting problem can be translated in polynomial time into solutions of the original problem.

Towards Higher-Order Logic Automated Theorem Proving (CADE 2017, FSCD 2017, IJCAR 2018). Lifting superposition-based first-order reasoning to higher-order logics while keeping the syntactic restrictions of the calculus poses several challenges. We have defined suitable orderings and lifted the calculus to a lambda-free higher-order logic.

Subtropical Methods for Real Constraints (FROCOS 2017, CASC 2018). We had developed *subtropical methods* for heuristically checking real constraints in the first orthant during the previous reporting period. We generalized this method from a single to finitely many ordering constraints. The resulting method performs surprisingly well both in the automated reasoning (SMT) and the bioinformatics context.

Formalizing Automated Reasoning Calculi in Logic (IJCAI 2017, FROCOS 2017, CPP 2018, JAR 2018, POPL 2019, IJCAR 2018, CPP 2019). We continue our "eat your own dog food" line of research on the formalization of (our developed) logical calculi in Isabelle and by this way testing our enhancements to Isabelle. In the current reporting period we formalized further variants of CDCL (Conflict-Driven Clause Learning) for SAT and the ordered reolution prover from Bachmair and Ganzinger's article in the Handbook of Automated Reasoning. Furthermore, we added to (Co)datatypes, (Co)recursion and Binders of the Isabelle/HOL implementation.

Projects and Cooperations

Together with Stephan Merz's group (Inria Nancy) we constitute the Inria project 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 EU FET Open action SC-Square on bringing together automated reasoning and computer algebra research, led by Thomas Sturm. In the transregional collaborative research center TRR 248, established in January 2019, we investigate automated reasoning in the context of perspicuous software systems in close cooperation with researchers from TU Dresden.