

*Fifteenth
Biennial Status Report*

March 2019 – February 2021

Contents

1	Overview	1
1.1	Organization and Staff	1
1.2	Scientific Vision and Strategic Goals	3
1.3	Long-Term Achievements and Impact	3
1.4	Highlights 2019–2021 and New Research Directions	4
1.5	Career Mentoring	6
1.6	Collaborations and Strategic Partnerships	8
1.7	Results 2019–2021	10
2	The Research Units	13
3	D1: Algorithms and Complexity	15
4	D2: Computer Vision and Machine Learning	27
5	D3: Internet Architecture	33
6	D4: Computer Graphics	43
7	D5: Databases and Information Systems	55
8	RG1: Automation of Logic	63

1 Overview

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): Bernt Schiele (2010–2035), Computer Vision and Machine Learning (D2); Anja Feldmann (2018–2033), Internet Architecture (D3); Hans-Peter Seidel (1999–2026), Computer Graphics (D4); Gerhard Weikum (2003–2023), Databases and Information Systems (D5); and Christian Theobalt (2021–2043), heading the newly founded department on Visual Computing and Artificial Intelligence (D6). Kurt Mehlhorn (1990–2019) is director emeritus and acting scientific leader of the department on Algorithms and Complexity (D1). Thomas Lengauer retired on May 31, 2018, after more than 15 years as Director of the Bioinformatics Department. In addition to the departments, the Institute has one permanent independent research group on Automation of Logic, headed by Christoph Weidenbach, and an independent research group on Network and Cloud Systems, headed by Yiting Xia. Yiting Xia successfully applied to the tenure-track openings of the Max Planck institutes in computer science, a joint recruitment campaign. Tenure track openings are an additional instrument for recruiting talent that has been installed recently. In total, the Institute currently has 146 scientists, out of which 97 are doctoral students¹ and 49 have a doctoral degree. Due to the very recent appointment of the newest director, Christian Theobalt, he still appears as a tenured senior researcher of D4 in this report.

Senior Researchers: The Institute has five scientific ranks: director, senior researcher with tenure, senior researcher on tenure track, senior researcher, and researcher. 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. In addition to the typically non-tenured senior researcher appointment process, we have established a tenure-track senior researcher career path. Together with the MPI for Security and Privacy and the MPI for Software Systems we invite applications in all areas of computer science. In the reporting period, the Institute has been home to 21 senior researchers, in addition to its directors.²

¹Including nine students financed together with Saarland University.

²Zeynep Akata, Klaus Berberich, Jasmin Blanchette, Karl Bringmann, Balakrishnan Chandrasekaran, Andreas Karrenbauer, Marvin Künnemann, Christoph Lenzen, Dániel Marx, Karol Myszkowski (tenured), Gerard Pons-Moll, Simon Razniewski, Rishiraj Saha Roy, Thomas Sturm, Paul Swoboda, Christian Theobalt (tenured), Sophie Turret, Jilles Vreeken, Christoph Weidenbach (tenured), Yiting Xia (tenure-

Senior researcher positions are not tenure-track, in general. We expect the vast majority of 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 have received 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.

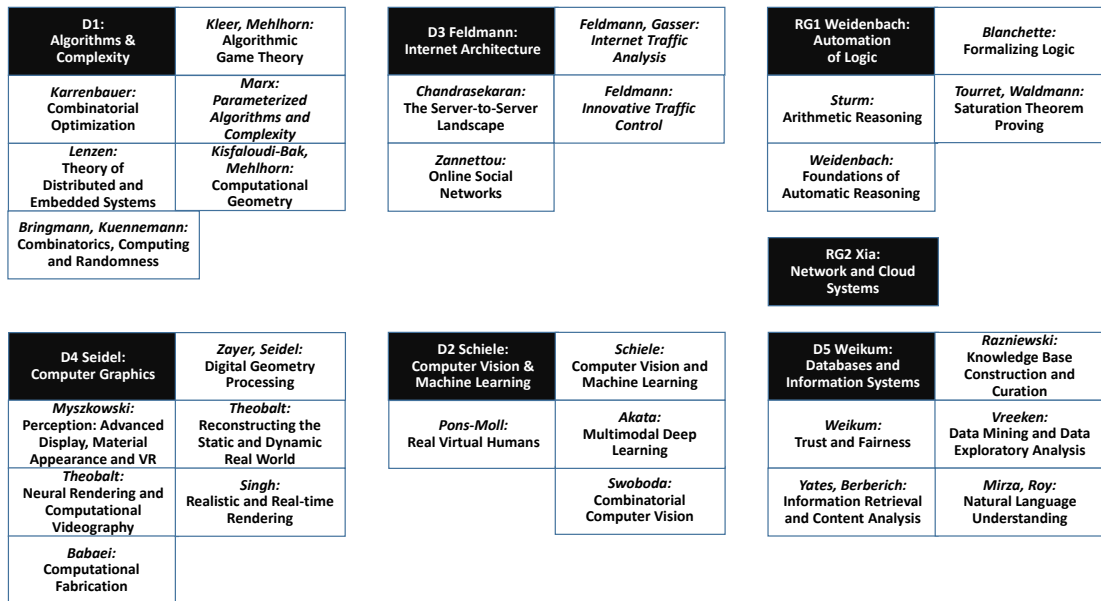


Figure 1.1: Research Areas of the Institute (as of March 2021)

Doctoral Students: As of March 2021, there are 97 doctoral students being supervised by members of the Institute, including 23 women (about 24%). 66 of the 97 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 MPI for Software Systems. Both institutes currently count 54 full-time employees including 20 IT support staff in these shared areas.

track), Andrew Yates

1.2 Scientific Vision and Strategic Goals

Algorithms and their applications are and 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 these topics.

Most of the major advances in computer science have come through the combination of new theoretical insights and application-oriented experimental validation, all driven by outstanding researchers. Our goal is, thus, to have impact through i) *publications*, ii) *software, services, and data resources* 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 performance history 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. More recently, the group has been a trendsetter on fine-grained complexity and theory of hardware.
- Bernt Schiele's group has been working on tracking multiple people for several years now. In joint work with Bjoern Andres 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 has won the Multi-Object Tracking Challenge twice.
- Anja Feldmann's research vision is to obtain insights from Internet measurements as foundation for shaping the evolution of the Internet by proposing optimizations and investigating alternative designs. As such her research group has and continues to address many challenges in *Internet measurement*, e.g., the impact of COVID-19 on Internet traffic, as well as *innovative traffic control*, e.g., by highlighting the benefits of joint optimization of ISPs/CDNs or applications/network.

- The Computer Graphics Groups stands out for its integrated view of *3D image analysis and synthesis*. Hans-Peter Seidel has 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. Christian Theobalt has conducted pioneering research on markerless performance capture and on neural rendering, and his research has led to entirely new ways of fusing and deeply integrating model-based and deep learning-based scene reconstruction. Karol Myszkowski has been and continues to be one of the pioneers and a driving force in perception-based graphics.
- Christoph Weidenbach’s group is on the next level of automated reasoning systems. While SAT solving has meanwhile found many places as a daily tool in standard processes in research and industry, we are working on lifting this success to more expressive logics.
- 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 have worked out. The sections about the five departments name some examples. Nevertheless, most of this work has resulted in novel insights (sometimes about what is, fundamentally or practically, non-viable) and often in influential publications.

1.4 Highlights 2019–2021 and New Research Directions

Kurt Mehlhorn moved to emeritus status at the end of August 2019. He is continuing as a researcher and scientific leader of D1 in the Institute. Christian Theobalt joined the Institute as a director for the department on Visual Computing and Artificial Intelligence in March 2021. The Institute is involved in the process of the Max Planck Society to establish departments in quantum computing.

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. The group shrunk in size, however, algorithm and complexity research on the campus grew with Karl Bringmann becoming full professor at Saarland University and Dániel Marx becoming tenured faculty at CISPA.
- Bernt Schiele’s group has been working on a variety of essential problems that will be also followed in the future. An important example is how to learn without or with

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. For example, we have contributed to the better understanding of privacy implications of visual data dissemination through and have also shown that compute vision model stealing is possible under weak assumptions, thereby raising concerns about the security of current machine learning techniques. Additionally, as the group is moving towards more foundational research in machine learning as well, the interpretability and explainability of machine learning models is a central theme of the research.

- Anja Feldmann’s group has been working on understanding the capabilities and vulnerabilities of today’s Internet. In recent work they have been highlighting, e.g., the resilience of the Internet (IMC 20), the evolution of the Internet topology (IMC 20), how to detect IoT devices in the wild (IMC 20), or how to steer hyper-giants’ traffic at scale. This work received the CoNEXT 19 best paper award as well as an IETF Applied Networking Research award.
- Highlight results from the Computer Graphics Group include R. Zayer’s work on subdivision-specialized linear algebra kernels (G. Enderle Award 2020), Christian Theobalt’s work on the neural rendering of humans (ACM TOG’19) and of human portrait images (StyleRig) (CVPR’20), as well as Christian’s work on human 3D performance capture (XNect) (SIG’19), and on neural scene representations (Neural Sparse Voxel Fields) (NeurIPS’20), and Karol Myszkowski’s results on foveated rendering (SIG’19), on X-Fields (SIGAsia’20), and on high-fidelity appearance reproduction (SIG’19, CGF’21). G. Singh achieved far reaching results on sampling correlations (SIG’19, CGF’20), and V. Babaei made fundamental progress on multimaterial printing (SIGAsia’20) and Laser marking (SIG’20).
- Christoph Weidenbach and his group have been working on arithmetic theory solving, higher-order formalization and mechanization, and SCL (clause learning from simple models), a new algorithmic paradigm for automated reasoning. Our SMT solver SPASS-SATT won the rational linear arithmetic category of SMT-COMP in 2019, and the superposition prover Zipperposition supporting higher-order logic aspects (CADE’19, IJCAR’20) won the higher-order category at the IJCAR system competition 2020. Martin Bromberger won the EAPLS dissertation award for his dissertation on improvements of linear arithmetic algorithms. We have made significant progress on the analysis of biological reaction networks using non-linear arithmetic reasoning (CASC’19, CASC’20). Our higher-order formalization of logical calculi and in particular CDCL is meanwhile at a level where even reasonably efficient executable code can be generated (NFM’19, LPAR’20). We have extended the SCL paradigm to first-order logic modulo theories (VMCAI’21).
- Gerhard Weikum’s group has been at the forefront of research on extracting, organizing and querying quantity properties of entities that appear in web tables and text sources. This supports answering queries such as “sprinters who ran 100m under 9.9s” or “electric cars with energy consumption below 15 kWh/100km”, use cases on which search engines

fail and knowledge bases lack coverage. Results have appeared in ICDE'19, ISWC'19, WSDM'20, WWW'21, and SIGMOD'21. In the ERC synergy grant imPACT, the focus has shifted from understanding privacy risks to enhancing the trustworthiness and controllability of recommenders and other machine-learning systems. We have developed new methods for generating user-comprehensible counterfactual explanations and for operationalizing group fairness and individual fairness in a practically viable manner. Results have appeared in WSDM'19, WSDM'20, WWW'21, ICDE'19, VLDB'20, and NeurIPS'20. Asia Biega's dissertation won the GI DBIS Dissertation Award.

New Research Directions: We continue focusing on the grand challenge of exploring algorithms that understand, search, and organize large, distributed, noisy, incomplete, and diverse information in a robust, efficient, and intelligent manner. 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 continue to explode, 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 353 doctoral students have graduated. These include 63 women, and 156 non-Germans. A number of students have won prestigious national and international awards: 15 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 granted since 2007. Maximilian Dylla won the DBIS Dissertation Award in 2014 and Asia Biega won the DBIS Dissertation Award 2021. Mateusz Malinowski (2018), Marcus Rohrbach (2015), and Siyu Tang (2018) all won the DAGM MVTec Dissertation Award. Karl Bringmann (2015) won the EATCS Distinguished Dissertation Award. Pablo Garrido and Petr Kellnhofer both won Eurographics Dissertation Awards in 2018. Martin Bromberger won the EAPLS Best Dissertation Award in 2019. Pieter Kleer was awarded the Gijs de Leve Prize for the best

PhD thesis in the area of Mathematics of Operations 2018-2020. Vladislav Golyanik won the Dissertation Award of the German Association for Pattern Recognition in 2020. Ana Serrano and Thomas Leimküler both received a Eurographics PhD Thesis Award 2020.

Young Scientist Career Advancement: A unique strength of the Institute is its successful fostering of young scientists. Our 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 81 young scientists have held such positions. 66 of them have meanwhile left the Institute. Out of these, 65 have accepted a tenured or tenure-track position offer at universities or university-like research organizations. We started tenure-track openings for senior researchers in 2019. We see this as an additional instrument for hiring talented researchers and, in particular, women. Several alumni (i. e., who graduated here or spent at least two years at the Institute and have since moved on) and current senior researchers of the Institute have won prestigious awards: Leibniz Prizes³, ERC Grants⁴, and other honors.

Support for Women: The percentage of women at our Institute is currently 24% for doctoral students and 19% 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 13 women.⁵ Out of the 10 recipients of the award who have meanwhile left the Institute, eight continued their career with a professor appointment. Our recently established senior researcher tenure-track career path constitutes an additional instrument for attracting women. Yiting Xia is the first tenure-track senior researcher at our Institute. A large number of female alumni of the Institute have become professors.⁶

³Susanne Albers 2008, Leif Kobbelt 2014, Thomas Neumann 2020, Peter Sanders 2012

⁴Susanne Albers 2016, Niko Beerenwinkel 2014, Jasmin Christian Blanchette 2016, Christoph Bock 2015, Karl Bringmann 2019, Andreas Bulling 2018, Parinya Chalermsook 2017, Piotr Didyk 2018, Jürgen Gall 2015, Bastian Goldlücke 2013, Matthias Hullin, 2018, Leif Kobbelt 2013, Christoph Lenzen 2016, 2020, Marcus Magnor 2010, Thomas Neumann 2016, Dániel Marx, 2017 Rafal Mantiuk 2016, Antti Oulasvirta 2014, Bodo Rosenhahn 2011, Thomas Sauerwald 2015, Pascal Schweitzer 2018, Jürgen Steimle 2016, Christian Theobalt 2013, 2017

⁵Anke 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 Turret 2018-2020, Jiangxin Dong 2020-2022, Erisa Terolli 2019-2021, Roohani Sharma 2020-2022

⁶Anna Adamaszek (Univ. Copenhagen, Denmark), Zeynep Akata (Univ. Tübingen), Susanne Albers (TU Munich), Iris Antes (TU Munich), Oana Balalau (INRIA), Hannah Bast (Univ. Freiburg), Carola Dörr (CNRS, France), Panagiota Fatourou (Univ. Ionina, Greece), Lilia Georgieva (Heriot-Watt Univ., UK), Qiuhong Ke (University of Melbourne), 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 Papatriantafidou (Univ. Gothenburg, Sweden), Ruzica Piskac (Yale Univ., USA), Nicoleta Preda (Univ. Versailles, France), Maya Ramanath (IIT Delhi, India), Ana Serrano (Univ. Zaragoza, Spain), Ina Schäfer (TU Braunschweig), Renate Schmidt (Univ. Manchester, UK), Viorica Sofronie-Stokkermans (Univ. Koblenz-Landau), Qianru Sun (Singapore Management University), Kavitha Teliapalli (Tata Institute, India), Yafang Wang (Shandong Univ., China), Nicola Wolpert (Univ. of Applied Sciences Stuttgart), Shanshan Zhang (Nanjing U of Science

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, 7 Leibniz Prizes (the highest scientific honor in Germany), and a total of 24 ERC Grants (at all levels).

Saarland Informatics Campus: Our most important partners are the Computer Science Department of Saarland University, the Helmholtz Center for Information Security (CISPA) and the MPI for Software Systems. We have a long-standing tradition of teamwork and joint engagement in research, recruiting, and teaching. In the reporting period the senior researchers Mario Fritz, Dániel Marx, and Jilles Vreeken joined the CISPA faculty, and Karl Bringmann joined the Computer Science Department. There are numerous collaborations with faculty members from the university and with researchers from 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 Research Foundation (DFG) in the context of the German Excellence Initiative and funded for two consecutive funding periods from 2007 to 2019. Members of the Institute significantly contributed to this campus-wide effort: All directors were principal investigators, and the scientific coordinator of the Cluster was Hans-Peter Seidel. More than 230 PhD students completed their PhD work in the Cluster, and more than 200 early career researchers of the Cluster moved on to faculty positions worldwide. Researchers in the Cluster received numerous grants and awards, both on the senior and early career levels. On the senior level this includes, e.g., five ERC Advanced Grants, a DFG Leibniz prize, as well as several prestigious career awards. On the early career level this includes, e.g., 10 DFG Emmy Noether Grants, 21 ERC Starting Grants, and 6 ERC Consolidator Grants, as well as numerous PhD awards.

Max Planck Center for Visual Computing and Communication (MPC-VCC): The center was established jointly by MPG and Stanford University in 2003 and funded by the German Ministry of Research and Education (BMBF) for two consecutive funding periods from 2003 to 2020. The program offered outstanding young scientists the opportunity to spend two years at Stanford as visiting assistant professor, and then continue as leader of an independent research group at our institute. The early career researchers funded by the center received numerous awards, including 12 ERC grants, six DFG Emmy Noether grants, two Alcatel Lucent Research Awards, two Karl-Heinz-Beckurts Awards, the Dutch Prize for ICT Research, eight Eurographics Young Researcher Awards, and four DAGM Pattern Recognition Awards. More than 40 early career researchers moved on to faculty positions. The center was jointly directed by Hans-Peter Seidel (MPI-INF) and Bernd Girod and Leo

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

Guibas (Stanford). Although the program has now officially come to its end, we continue to exchange researchers between the two locations with funding from other sources on a somewhat smaller scale.

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

The Research Center on Interactive Media, Smart Systems and Emerging Technologies (RISE): This center 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, with an extension until 2022.

IMPRS for Computer Science (IMPRS-CS): The International Max Planck Research School for Computer Science (IMPRS-CS) was a joint graduate program of the MPI for Informatics and Saarland University, which was in operation from 2001 to 2019. During this time, a total of 323 graduate students obtained their doctoral degrees via IMPRS-CS, including 56 women (17%) and 159 non-Germans (49%). In addition, IMPRS-CS also supported highly talented, non-German students with Bachelor’s fellowships during their Master’s studies. Since 2001, a total of 236 students, attracted mainly from abroad, completed their Master’s degrees with IMPRS-CS support, including 76 women (32%).

IMPRS on Trustworthy Computing (IMPRS-TRUST): The International Max Planck Research School on Trustworthy Computing (IMPRS-TRUST) was established in 2020 as a joint program of the MPI for Informatics, the MPI for Software Systems, Saarland University, and TU Kaiserslautern. Currently, 97 doctoral students of MPI-INF, including 23 women (24%) and 66 non-Germans (68%), are part of the program.

CS@Max Planck Computer science centered research of the Max Planck Society has started a joint presence <https://www.cis.mpg.de/>. The institutes comprise the MPI for Informatics, the MPI for Software Systems, the MPI for Security and Privacy, the MPI for Molecular Genetics, and the MPI for Molecular Cell Biology and Genetics. The effort includes joint recruiting for tenure-track faculty positions and internships.

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 MPI for Informatics, the MPI for Software Systems, the MPI for Intelligent Systems, and top German universities.

Saarbrücken Graduate 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.

ELLIS Unit SAM: Saarbrücken Artificial Intelligence & Machine Learning: The ELLIS Unit SAM <https://www.ellis-unit-sam.de> has been founded in 2020 as part of the ELLIS Society which seeks to establish internationally visible, top-level research facilities in Europe in the area of machine learning and modern AI. The SAM Principal Investigators have agreed to jointly work on both the foundations for enhanced functionalities of Artificial Intelligence and Machine Learning (AIML) systems and the pressing needs for security, privacy, and trustworthiness that arise from the widespread use of Artificial Intelligence and Machine Learning systems. Bernt Schiele is the scientific director of the unit, Christian Theobalt is a PI, and overall, the unit brings together nine PIs from the MPI for Informatics, MPI for Software Systems, UdS, and CISPA.

1.7 Results 2019–2021

Publications, Software, Startups: In the two-year time-frame 2019–2021, the Institute published more than 500 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 AAI ICWSM 2020, AKBC 2020, CADE-27, CoNEXT 2019, CVPR 2020, EC 20, ETAPS 2020, Eurographics 2020, ILP 2019, ISAAC 19, NeurIPS 19, SoCG 2019).

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 Ministry 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 time-frame 2019–2021, 40 of our doctoral students graduated. These include 7 women. In the same time period, 18 of our researchers left the Institute to take a tenured or tenure-track faculty position ⁷. These include 4 women.

⁷Zeynep Akata (University of Tübingen, Germany - accepted, RWTH Aachen and Helmholtz Institute Jülich, Germany - declined, University of Würzburg, Germany - declined, CISPA and Saarland University, Germany - declined), Antonios Antoniadis (University of Twente, Netherlands), Oana Balalau (INRIA Paris, France), Florian Bernhard (TU Munich, Germany), Asia Biega (MPI for Security and Privacy Bochum, Germany - accepted, TU Munich, Germany - declined), Karl Bringmann (Saarland University),

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 won the Community Contribution Award at the ACM Internet Measurement Conference (IMC) 2020 and the IETF/IRTF Applied Networking Research Prize 2020; Harald Ganzinger (posthumously) and Christoph Weidenbach both received a Thoralf Skolem Award 2019; Christian Theobalt was recognized with the Eurographics Outstanding Technical Contributions Award 2020;

At the level of senior researchers, Karl Bringmann won the Heinz Maier-Leibnitz-Prize 2019 and the EATCS Presburger Award for Young Scientists 2019 and received the ERC Starting Grant TIPEA 2019; Klaus Berberich and Gerhard Weikum won the ECIR Test-of-Time Award 2020; Gerard Pons-Moll won the German Pattern Recognition Award 2019 (highest career award in Germany in the areas of Computer Vision and Machine Learning for researchers under 35) and received a Google Faculty Research Award 2019; Christoph Lenzen received an ERC Proof of Concept Grant in 2020; Zeynep Akata received an ERC Starting Grant in 2019.

At the student level, Thomas Leimkühler and Corinna Coupette (for her legal dissertation completed at the MPI for Tax Law and Public Finance) both won an Otto-Hahn-Medal 2020; Sándor Kisfaludi-Bak received the EATCS Distinguished Dissertation Award 2020 and the Best IPA Dissertation Award from the Institute for Programming Research and Algorithmics 2019; Martin Bromberger received the EAPLS Best Dissertation Award 2019 and the Dr. Eduard-Martin-Prize 2020; Pieter Kleer received the Gijs De Leve Prize for the best PhD thesis defended in the years 2018–2020 from the Dutch Network on the Mathematics of Operations Research (LNMB); Florian Streibelt won the IETF/IRTF Applied Networking Research Prize 2019; Ana Serrano and Thomas Leimkühler both won a Eurographics PhD Award in 2020; Vladislav Golyanik received the Dissertation Award of the German Association for Pattern Recognition 2020; Noshaba Cheema won the Women STEM Award 2020; Rohaani Sharma received a honorable mention for the ACM India Doctoral Dissertation Award 2021 for her PhD thesis; Asia Biega won the GI DBIS Dissertation Award 2021.

Further honors with considerable visibility include the following. Anja Feldmann was elected as a member of Acatech; Hans-Peter Seidel was appointed as a member of the DFG Senat; Bernt Schiele was appointed Fellow of European Laboratory of Learning and Intelligent Systems ELLIS.

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

Balakrishnan Chandrasekaran (University of Amsterdam, Netherlands), Renjie Chen (University of Science and Technology of China), Reuven Hodges (University of Illinois, USA), Pieter Kleer (Tilburg University, Netherlands), Stefano Leucci (University of L'Aquila, Italy), Dániel Marx (CISPA, Germany), Pranabendu Misra (National Institute of Science Education and Research, India), Gerard Pons-Moll (University of Tübingen, Germany - accepted, University of Luxembourg - declined), Helge Rhodin (UBC Vancouver, Canada), William Rosenbaum (Amherst College, USA), Ana Serrano (University of Zaragoza, Spain), Srinath Sridhar (Brown University, USA)

years, we selected three young women as recipients of this fellowship. The best postdocs are often candidates for becoming senior researchers after two years. Out of the 10 fellowship recipients who left the Institute, eight continued with a professor appointment.

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/>).

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

3 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 lead the group since its foundation; Kurt moved to emeritus status on August 31st, 2019. He is acting head of the department and continues as a scientist till today. The appointment process for a new director is on-going.

The group has shrunk over the past two years, but algorithm theory on the campus has actually become stronger. Karl Bringmann became full Professor for Computer Science and Dániel Marx became Faculty at CISPA, the Helmholtz Center for Information Security. So algorithm research is up for a bright future on campus.

As of March 1st 2019, the senior scientists and subgroup coordinators are Andreas Karrenbauer, Christoph Lenzen, Karl Bringmann (main affiliation Saarland University), Marvin Künnemann (moves to ETH in April 2021), Sándor Kisfaludi-Bak, and Pieter Klier (moves to Tilburg in April 2021)

Section 31.1 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 216; Antonios Antoniadis (Twente), Reuven Hodges (Illinois), Pieter Klier (Tilburg), Stefano Leucci (L'Aquila), Pranabendu Misra (Chennai), Tim Oosterwijk (Maastricht), and William Rosenbaum (Amhers) moved to Assistant Professorships and Karl Bringmann and Dániel Marx moved to tenured faculty positions at Saarland University and CISPA respectively. Two group members completed their PhD, see page 216. Group members received prestigious awards, see page 216, e.g., Karl Bringmann received the Heinz Maier-Leibnitz-Prize 2019 and the EATCS Presburger Award for Young Scientists 2019, Corinna Coupette received the Otto Hahn Medal of the Max Planck Society for her legal dissertation (completed at the Max Planck Institute for Tax Law and Public Finance) in 2020, Bhaskar Ray Chaudhury received the best paper with a student lead author and the exemplary paper in theory awards at the 21st ACM Conference on Economics and Computation (EC), 2020, Themis Gouleakis received the Outstanding Paper Award at the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS), 2019, Sándor Kisfaludi-Bak received the best paper award at the 30th International Symposium on Algorithms and Computation (ISAAC), 2019, and the EATCS Distinguished Dissertation Award in 2020, Pieter Kleer received the Gijs De Leve Prize for the best PhD thesis defended in the years 2018–2020 from the Dutch Network on the Mathematics of Operations Research (LNMB), Rohaani Sharma received a honorable mention for the ACM India Doctoral Dissertation Award 2021 for her PhD thesis. Some group members hold their own grants, see page 217, e.g., Karl Bringmann, Christoph Lenzen, and Dániel Marx hold ERC grants.

We have published extensively and in excellent venues, e.g., 9 papers in SODA 2020 and 2021, 4 papers in FOCS 2019 and 2020, 7 papers in ICALP 2019 and 2020, 4 papers in

STOC 2019 and 2020, 2 papers in EC 2020, and 2 papers in SoCG 2019 and 2020. We are also visible now in conferences such as NeurIPS and ICML.

Kurt continued his research on trustworthy algorithms, the Physarum dynamics, market equilibrium, and fair division.

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 parenthesis (Coordinators as of March 1st, 2021, in italics).

- Combinatorics, Computing, and Randomness (*Karl Bringmann* and *Marvin Künnemann*).
- Combinatorial Optimization (*Andreas Karrenbauer*)
- Computational Geometry and Algebra (*Sándor Kasfaludi-Bak* and *Kurt Mehlhorn*)
- Algorithmic Game Theory and Online Algorithms (*Pieter Kleer* and *Kurt Mehlhorn*)
- Theory of Distributed and Embedded Systems (*Christoph Lenzen*).
- Parameterized Algorithms and Complexity (*Dániel Marx*)

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 once a week to discuss administrative matters. Also, the subgroups meet weekly.

We teach at all levels; see page 214 for details. Frequently group members pair for lectures.

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

Main Results

We discuss some of the main results obtained in the reporting period.

A Fine-Grained Perspective on Subset Sum (see page 124) *Subset Sum* is one of the basic NP-hard problem at the intersection of theoretical computer science, optimization, and operations research. In previous reporting periods, our group developed a precise understanding of certain pseudopolynomial algorithms for Subset Sum. These results have enabled the development of improved algorithms for many different settings of Subset Sum and related problems. In the last two years, Karl Bringmann, Vasileios Nakos and Philip Wellnitz (STOC '20, SODA '21, SODA '21) greatly enhanced the understanding of Subset Sum. Here comes a representative result.

The input consists of a set of n positive integers X and a target t , and the task is to check whether some subset of X sums to t . In previous reporting periods, our group developed

a precise understanding of Subset Sum with respect to the parameters n and t : Improving upon a textbook algorithm from 1957 running in time $O(nt)$, we designed an algorithm with running time $\tilde{O}(t)$ and we established this improved running time to be near-optimal under a plausible conjecture. However, the literature on Subset Sum contains several further natural parameters such as the maximum input number $\max(X)$ and the sum of all input numbers $\Sigma(X)$. Determining the optimal time complexity of Subset Sum with respect to all these parameters is a long-standing open problem, in particular it is open whether Subset Sum can be solved in time $\tilde{O}(n + \max(X))$.

Karl and Philip (SODA '21) study the *dense* case of Subset Sum, where all of the aforementioned parameters are *polynomial* in n . In this setting, standard pseudo-polynomial algorithms solve Subset Sum in polynomial time $n^{O(1)}$. Can this be improved to near-linear time $\tilde{O}(n)$?

They provide an essentially complete dichotomy by designing improved algorithms and by proving conditional lower bounds, thereby determining essentially all settings of the parameters n , t , $\max(X)$, and $\Sigma(X)$ for which dense Subset Sum is solvable in time $\tilde{O}(n)$. For notational convenience we assume without loss of generality that $t \geq \max(X)$ (as larger numbers can be ignored) and $t \leq \Sigma(X)/2$ (using symmetry). Then their dichotomy reads as follows:

- Subset Sum is in near-linear time if $t \gg \max(X)\Sigma(X)/n^2$.
- They prove a matching conditional lower bound: If Subset Sum is in near-linear time for any setting with $t \ll \max(X)\Sigma(X)/n^2$, then plausible conjectures from fine-grained complexity theory fail.

Fine-Grained Classification of Boolean CSPs (see page 133) Under which circumstances can we find *small solutions* faster than by exhaustive search, and if so, by how much? Marvin Künnemann and Dániel Marx (CCC '20) study this question for the class of Boolean constraint satisfaction problems with size constraint exactly k , which capture a variety of problems from graph theory, combinatorial optimization, and more. Specifically, they aim to determine, for any finite constraint family, the optimal running time $f(k)n^{g(k)}$ required to find satisfying assignments that set precisely k of the n variables to 1.

Under central hardness assumptions on detecting cliques in graphs and 3-uniform hypergraphs, we give an almost tight characterization of $g(k)$ into four regimes:

1. Brute force is essentially best-possible, i. e., $g(k) = (1 \pm o(1))k$,
2. the best algorithms are as fast as current k -clique algorithms, i. e., $g(k) = (\omega/3 \pm o(1))k$,
3. the exponent has sublinear dependence on k with $g(k) \in [\Omega(\sqrt[3]{k}), O(\sqrt{k})]$, or
4. the problem is fixed-parameter tractable, i. e., $g(k) = O(1)$.

Optimal Approximation of Feedback Vertex Set in Tournaments (see page 146) Pran-abendu Misra and co-workers (SODA '20) consider the *Feedback Vertex Set in Tournaments (FVST) problem*, where the input is a tournament G with vertex weights $w : V(G) \rightarrow \mathbb{R}^+$, and the objective is to find a minimum weight subset of vertices $S \subseteq V(G)$ such that $G - S$ is an acyclic tournament. It is known that FVST cannot have a better than factor-2 approximation under Unique Games Conjecture and the previous best known approximation

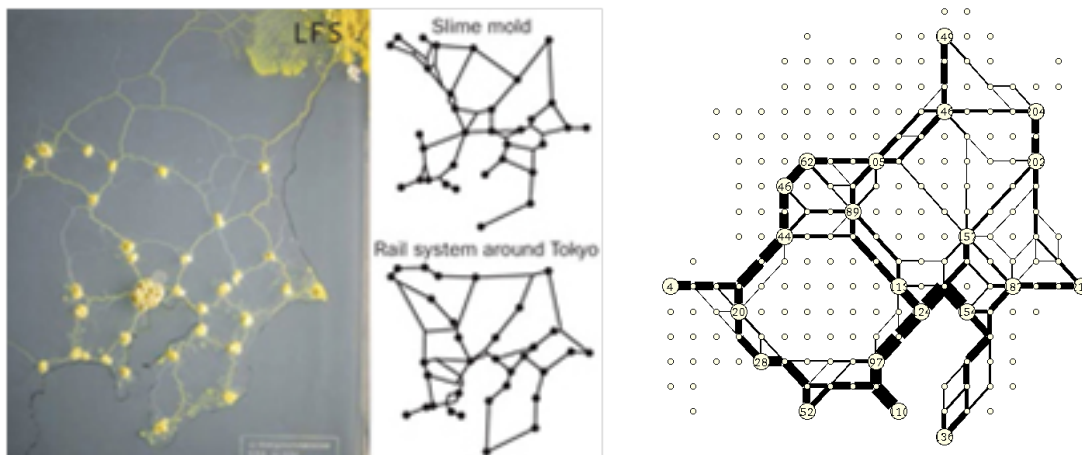


Figure 3.1: A network constructed by the slime mold and a network constructed by our dynamics. For the slime-mold experiment the Greater Tokyo region was modeled and food was provided at locations corresponding to major cities in the region. For the computer simulation on the right the Greater Tokyo regions was modeled as a grid graph (each grid node was connected to its 8 nearest neighbors) and each pair of city was chosen as a source-sink pair.

implementations in LEDA are based on algorithms whose correctness has been proven, mistakes were introduced in the implementation process. Since the mid-nineties the LEDA team (Kurt Mehlhorn, Stefan Näher, Christian Urig) adopted the principle of certifying algorithms for their work. A certifying algorithm computes in addition to its output a certificate (or witness) that proves the correctness of the particular output. The output together with the witness is checked by the checker program. Almost all algorithms in LEDA are certifying. The use of certifying algorithms greatly enhanced the reliability of the library. About 10 years ago, group members showed that the checker programs can be formally verified.

In 2019, Kurt Mehlhorn and co-workers (MFCS '19) went one step further. They proved the correctness of the blossom-shrinking algorithm for maximum cardinality matching in Isabelle. This project is a collaboration with Mohammad Abdulaziz and Tobias Nipkow, two experts in automated theorem proving.

To be fair, we should also say what they did not do. They proved the correctness of the basic algorithm and not the correctness of the most efficient implementation. Also they proved the correctness of the formulation of the algorithm in Simpl, they did not prove the correctness of a C++ implementation.

Physarum Dynamics (see page 156) Ruben Becker, Andreas Karrenbauer, Pavel Kolev, Kurt Mehlhorn, and Golnoosh Shahkarami and co-workers continued their work on the Physarum dynamics (TCS '20, TCS '20, arXiv '20). They showed convergence of the non-uniform dynamics and gave, for the first time, a theoretical model for the network design capabilities of the slime.

The input is an undirected graph. Each edge e has a length c_e and a time-varying capacity $x_e(t)$. The graph is viewed as an electrical network where the resistance of edge at time t is $r_e(t) = c_e/x_e(t)$. The authors formulate network design as the problem of constructing a network that efficiently supports a multi-commodity flow problem. There are now k source-sink pairs (s_i, t_i) . For each pair, one unit of current is sent from the source to the corresponding sink. Let $q_i = (q_{i,e})_{e \in E}$ be the resulting electrical flow. The dynamics become

$$\dot{x}_e = \frac{d}{dt}x_e(t) = a_e \left(\sqrt{\sum_{1 \leq i \leq k} q_{i,e}^2} - x_e \right),$$

i. e., the individual flows over an edge are combined through their two-norm. The authors investigate the dynamics in computer simulations and analytically. The simulations show that the dynamics are able to construct efficient and elegant networks. In the theoretical part they show that the dynamics minimize an objective combining the cost of the network and the cost of routing the demands through the network. They also give alternative characterization of the optimum solution. Figure 3.1 shows two networks, one constructed by the slime and one by a computer simulation of the dynamics.

Secretary and Online Matching Problems with Machine Learned Advice (see page 151)

In *online selection problems*, the goal is to select the “best feasible” subset out of a finite set of objects that arrive online (i. e., one by one). Arguably the most well-known problem in this area is the *classical secretary problem*: There is a set of secretaries arriving in uniformly random order, where each secretary is associated with a nonnegative value. Whenever a secretary arrives, the algorithm has to irrevocably decide whether to hire that person or not, and the goal is to select the secretary with the highest value. It is well known that the following deterministic algorithm for the problem hires a secretary of expected value at least $1/e$ times the optimal value, and that this is tight: In the *first phase* the algorithm observes the values of the first $1/e$ -fraction of the secretaries arriving, but does not hire any of them. In the *second phase*, it hires the first secretary that exceeds the best value encountered so far.

Antonios Antoniadis, Themis Gouleakis, Pieter Klier, and Pavel Kolev (NeurIPS '20) study (among other problems) the classical secretary problem in which we are given in addition a prediction (of unknown quality) on the value of the best secretary, that can be utilized to improve upon the competitive ratio of $1/e$ in case this prediction is reasonably accurate. At the same time, we do not want to lose too much in competitive ratio when this is not the case. For many online selection problems, it is possible to obtain such predictions, for example via a machine learning approach that leverages historical data. They show how to benefit from advice. If the advice is poor, the algorithm is still competitive and if the advice is good, they achieve a better approximation ratio than $1/e$.

Polyline Simplification (see page 163) Karl Bringmann and Bhaskar Chaudhury (SoCG '19) consider polyline simplification. In this problem, given a polygonal curve P as a sequence of n vertices, and an error threshold δ , the goal is to determine a subsequence P' of minimum size (say k) from P , such that the polygonal curves P and P' have a distance of at most δ . The distance between the original curve P and the simplification P' can be determined

by applying the classic distance measures between curves (Fréchet distance and Hausdorff distance) *globally*, i. e., to the whole curves P and P' , or *locally*, i. e., to each simplified subcurve and the line segment that it was replaced with separately (and then taking the maximum). This gives rise to four problem variants: Global-Hausdorff (known to be NP-hard), Local-Hausdorff (known to be in time $\mathcal{O}(n^3)$), Global-Fréchet (known to be in time $\mathcal{O}(kn^5)$), and Local-Fréchet (known to be in time $\mathcal{O}(n^3)$).

Karl and Bhaskar design an algorithm for Global Fréchet simplification, that runs in time $\mathcal{O}(n^3)$, improving the previous best running time by a factor of $\Omega(kn^2)$. This shows that all three problems (Local-Hausdorff simplification, Local-Fréchet simplification, and Global-Fréchet simplification) can be solved in cubic time.

In the same paper, Karl and Bhaskar also provide evidence that in high dimensions, cubic time is essentially optimal for all three problems (Local-Hausdorff simplification, Local-Fréchet simplification, and Global-Fréchet simplification).

Fréchet Distance (see page 164) Karl Bringmann, Marvin Künnemann, and André Nusser (SoCG '19, ESA '20) investigate fast practical algorithms for computing the Fréchet Distance of curves. A simple algorithm computes it in near-quadratic time, a strongly subquadratic algorithm cannot exist unless the Strong Exponential Time Hypothesis fails. Still, fast practical implementations of the Fréchet distance, in particular for realistic input curves, are highly desirable.

Karl, Marvin, and André present a fast implementation for deciding the Fréchet distance, in order to (1) complement its pessimistic worst-case hardness by an empirical analysis on realistic input data, and to (2) improve over the state-of-the-art implementation. They experimentally evaluate their implementation on a large benchmark consisting of several data sets, observing running time improvements of more than two orders of magnitude compared to the previous state-of-the-art.

Furthermore, consider the natural question of how to measure the similarity of curves in the plane by a quantity that is invariant under translations of the curves. Such a measure is justified whenever one aims to quantify the similarity of the curves' shapes rather than their positioning in the plane, e. g., to compare the similarity of handwritten characters or detect position-independent motion patterns in GPS trajectories.

Their implementation enables, for the first time, the use of the Fréchet distance for very long practical input curves and the Fréchet distance under translation in applications where previous algorithmic approaches would have been computationally infeasible. Furthermore, the authors hope that our combination of continuous optimization and computational geometry will inspire similar approaches for further algorithmic questions.

Approximation Algorithms for Euclidean TSP (see page 167) Sándor Kisfaludi-Bak and co-workers (SoCG '20, arXiv '20) obtain improved approximation algorithms for the Euclidean Traveling Salesman Problem. The shortest tour can be approximated within a $(1 + \varepsilon)$ factor for any fixed positive $\varepsilon > 0$ in polynomial time (Arora98 and Mitchell99). Arora's scheme in particular inspired a lot of related research in computational geometry and in approximation algorithms, and it is now often taught at the graduate level in related courses.

Sándor and co-workers revisit Arora's algorithm, and significantly improve it using an

elegant modification. The resulting running time is $2^{O(1/\varepsilon^{d-1})} n \log n$ for any fixed dimension d . Moreover, they show that the ε -dependence of their algorithm is best possible under the Gap Exponential Time Hypothesis (Gap-ETH). The algorithm is based on a new technique called *sparsity-sensitive patching*, which changes the granularity with which the tour is modified based on the local density of the tour.

Computation of (pure) Nash Equilibria (see page 176) Congestion games constitute an important class of games which capture many applications in network routing, resource allocation and scheduling. Intuitively, these games model situations in which several independent agents (or players) compete over a limited number of resources. In particular, the congestion game model of Rosenthal has received a lot of attention in the algorithmic game theory community in the last twenty years. Fundamental problems studied here are the complexity of computing (pure) Nash equilibria, as well as bounding their inefficiency compared to a socially optimal outcome.

Pieter Kleer and co-worker (Math Programming '20) provide a unifying framework for various special classes of Rosenthal's congestion game model for which polynomial time algorithm exists to compute a pure Nash equilibrium. They identify polyhedral conditions, the *integer decomposition property* and *box-total dual integrality*, that are, roughly speaking, sufficient to guarantee that a polynomial time algorithm exists for computing a pure Nash equilibrium. Furthermore, these conditions also give a quality guarantee on the computed equilibrium that is much better than known general upper bounds on the quality of pure Nash equilibria in congestion games.

Fair Division of Items (see page 174) *Fair division of items* is a fundamental area that lies at the intersection of computer science, economics and social choice theory. At a high level, given a set of agents, a set of items, and valuation functions of the agents for the items, the goal is to determine an allocation of items to agents that achieves high total welfare and makes every agent content. Bhaskar Ray Chaudhury, Kurt Mehlhorn, and Alkmini Sgouritsa and co-workers (EC '20, SODA '20, AAI '21, SODA '21) have worked on allocating *divisible* or *indivisible goods and chores*. Here, we consider only indivisible goods.

The closest analogue of envy-freeness in the context of indivisible goods is that of *envy-freeness up to any good* (EFX). An allocation is said to be EFX if no agent envies another agent following the removal of *any* single good from the other agent's bundle. Until now, it is not known whether EFX allocations exist even when agents have additive valuations. Ariel Procaccia, in an editorial note in Communications of ACM, refers to the question as "*fair division's biggest open problem*".

Bhaskar, Kurt, and Alkmini have taken steps towards solving this problem. They show that even when agents have any weakly monotone valuation function (much more general than additive valuation functions), an EFX allocation always exists if we allow a small number of goods to remain unallocated. To be precise, the number of goods not allocated is less than the number of agents, and for each agent, the value of the unallocated goods is smaller than the value of the bundle allocated to the agent.

Bhaskar, Kurt and co-worker also show that *complete* EFX allocations exist for the case of three agents with additive valuations. For four or more agents the question is open.

A Breezing Proof of the KMW Bound (see page 177) Kuhn, Moscibroda, and Wattenhofer (KMW) proved a seminal hardness result for several fundamental graph problems in the Local model: For any (randomized) algorithm, there are n -node graphs of maximum degree Δ on which $\Omega(\min\{\sqrt{\log n / \log \log n}, \log \Delta / \log \log \Delta\})$ (expected) communication rounds are required to obtain polylogarithmic approximations to a minimum vertex cover, minimum dominating set, or maximum matching. Via reduction, this hardness extends to symmetry breaking tasks like finding maximal independent sets or maximal matchings.

Despite its significance, the KMW lower bound has not inspired substantial follow-up results. This might be due to the fact that the proof is quite technical. Corinna Coupette and Christoph Lenzen (SOSA '21) provide a fully self-contained and *simple* proof of the KMW lower bound. Their key argument is algorithmic, and it relies on an invariant that can be readily verified from the generation rules of the lower bound graphs.

Computing Low Stretch Tree Metrics (see page 180) Approximating the distance metric induced by graphs by tree metrics or other simplified metrics is a powerful technique for obtaining efficient approximation algorithms. *Low-stretch spanning trees* guarantee that the *stretch*, i. e., the ratio between the distance of a pair of neighbors in the tree and the weight of their connecting edge, is small on average. Christoph Lenzen (DISC '19, ITCS '20) and co-workers designed the first Congest algorithm achieving polylogarithmic stretch in a non-trivial number of rounds.

In further work, he obtains metric tree embedding algorithms in the vein of Bartal, whose computational complexity is optimal up to polylogarithmic factors. The embeddings have the additional useful property that the tree can be mapped back to the original graph such that each edge is “used” only logarithmically many times, which is of interest for capacitated problems and simulating Congest algorithms on the tree into which the graph is embedded.

Metastability-Containing Clock Synchronization (see page 183) Distribution of a single clock signal to all parts of a hardware chip is infeasible in many cases today. Chips can be very complex and timing uncertainty accumulates over long wires, such that eventually the skew between arrival times of a single clock signal can be large.

To avoid this problem, hardware is partitioned into different computational modules, each with its own pulse giving device. However, if the clocks are uncorrelated, communication between such modules has high latency and incurs the use of large data buffers. Hence, it is advisable to synchronize the individual module clocks.

Johannes Bund, Christoph Lenzen, and Will Rosenbaum and co-workers (IEEE Trans. on Circuits and Systems '20, IEEE Symp. on Asynchronous Circuits and Systems '20) provide a solution for computational modules connected by an arbitrary network. Their solution is based on an earlier *gradient clock synchronization* algorithm of Christoph minimizing the worst-case difference between clocks of neighbors. The authors demonstrate that this difference is substantially smaller than one clock cycle even in large networks: parameters obtained from a 15 nm ASIC implementation running at 2 GHz yield worst-case bounds of 30 ps on phase offset for a 32×32 node grid network, and this bound grows only logarithmically in the network diameter.

Parameterized Counting Complexity (see page 193) A natural generalization of *decision problems* (*Is there a solution?*) are *counting problems* (*How many solutions exist?*). For many interesting (real-world) problems, such as finding specific patterns in a (host) graph or searching for answers to (conjunctive) queries in (relational) databases, even detecting the existence of a solution is already NP-hard, so fast algorithms to count all solutions are unlikely. However, as the size k of a typical pattern or a typical query is very small compared to the size n of the host graph or the database, a *parameterized* view is natural: For instance, even algorithms running in time $2^k \cdot \text{poly}(n)$ (or in general in fixed-parameter tractable time) may be fast enough for practical purposes.

Dániel Marx and Philip Wellnitz and co-workers (FOCS '20, SODA '20, ICALP '19) advanced the knowledge on counting problems on several fronts: counting answers to existential queries, finding and counting permutations via CSP, and counting graph patterns. Here is a representative result.

How often does a (small) pattern graph appear in a (large) host graph? More precisely, given a graph property Φ , on input a graph G and an integer k , the task is to compute the number $\#\text{INDSUB}(\Phi)$ of induced subgraphs of size k in G that satisfy Φ .

From an *implicit* criterion due to Curticapean, Dell, and Marx, it was known that for each property Φ , the problem $\#\text{INDSUB}(\Phi)$ is either fixed-parameter tractable, or $\#\text{W}[1]$ -hard (where $\#\text{W}[1]$ is the counting analogue to $\#\text{W}[1]$) and cannot be solved in time $f(k) \cdot |V(G)|^{o(k)}$ for any computable function f , unless the Exponential Time Hypothesis fails. Philip and co-workers take multiple steps into the direction of making the aforementioned criterion *explicit*, that is, they find (classes of) graph properties for which the problems $\#\text{INDSUB}(\Phi)$ is hard. In particular, their results include:

- For *any* graph property Φ that is closed under the removal of vertices and edges, and that is non-trivial for bipartite graphs, the problem $\#\text{INDSUB}(\Phi)$ is $\#\text{W}[1]$ -hard and cannot be solved in time $f(k) \cdot |V(G)|^{o(k)}$ for any computable function f , unless the Exponential Time Hypothesis fails.
- For any non-trivial monotone graph property Φ , the problem $\#\text{INDSUB}(\Phi)$ is $\#\text{W}[1]$ -hard and cannot be solved in time $f(k) \cdot |V(G)|^{o(k/\log^{1/2}(k))}$ for any function f , unless the Exponential Time Hypothesis fails.

Algorithms and Data Structures: The Basic Toolbox (see page 210) The second edition of the textbook

Sequential and Parallel
Algorithms and Data Structures
The Basic Toolbox

appeared in the summer of 2019. It now covers sequential and parallel algorithms and data structures.

Ideen und Konzepte der Informatik: An Online University Course (see page 210) For the last several years, Kurt is teaching a course *Ideas and Concepts of Computer Science* for non-majors. Two years ago, he started to produce videos for the course, see <https://www.mpi-inf.mpg.de/departments/algorithms-complexity/teaching/winter20/ideen>. Recently,

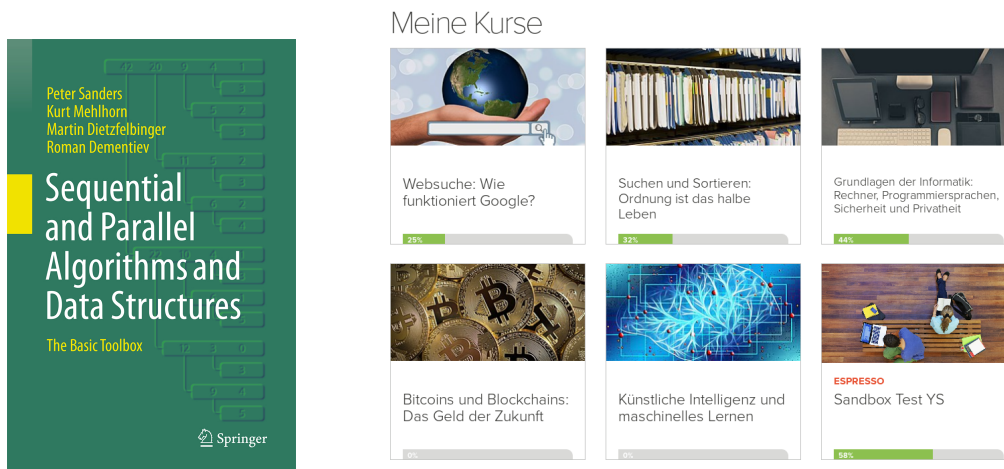


Figure 3.2: The cover of the Toolbox book and the front page of the Iversity course.

he signed a contract with Iversity to market the course and reach a wider audience. The course will become available on the Iversity platform in the summer of 2021.

4 D2: Computer Vision and Machine Learning

Group Overview

The Computer Vision and Multimodal Computing group (D2) was established in 2010 with the appointment of Bernt Schiele. In the reporting period, the group was renamed to become the Computer Vision and *Machine Learning* group for two main reasons: first, multimodal computing has become less of a research focus for our group in recent years; and second, while machine learning always was an important component of our research, its importance has increased significantly both for our research in computer vision but also in general. Therefore, renaming the group is both underlining the importance of machine learning for our research in computer vision as well as clearly stating that machine learning is an integral part of our research agenda, also resulting in an increased presence at top-tier machine learning venues.

At the time of writing, the group encompasses three senior researcher (research group leaders), two postdocs, and 21 PhD students. Among those 27% are female (one research group leader, one postdoc, and five PhD students). The senior researchers each have their own PhD-students to conduct research in their respective area. The current research group leaders are Zeynep Akata (appointed 2014, part-time), Gerard Pons-Moll (appointed 2017), and Paul Swoboda (appointed 2018). Additionally, Mario Fritz and Andreas Bulling, both senior researchers previously, have moved to faculty positions but were affiliated partly during the reporting period, as they still supervised or co-supervised students at the institute.

Nine group members completed their PhD during the reporting period. Our researchers get very good offers for faculty positions in academia¹, postdoc positions in academia², and research positions in industry³.

Vision and Research Strategy

Understanding visual information and more generally 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 sensors to the complete description of large-scale streams. At the same time we observe a tremendous increase in both

¹ Zeynep Akata (University of Tübingen, RWTH Aachen, University of Würzburg, CISPA Helmholtz Center), Qihong Ke (University of Melbourne), Gerard Pons-Moll (University of Tübingen, University of Luxembourg), Qianru Sun (Singapore Management University)

² Eldar Insafutdinov (University of Oxford), Yang He (CISPA), Jan-Hendrik Lange (University of Tübingen), Philipp Müller (Senior Researcher DFKI), Yongqin Xian (MPI for Informatics)

³ Eldar Insafutdinov (Apple), Jan-Hendrik Lange (Amazon), Tribhuvanesh Orekondy (Qualcomm), Hosnieh Sattar (Zalando), Rakshith Shetty (Amazon), Julian Steil (Bosch)

the quantity as well as the diversity of visual and sensor information due to the increasing number of sensors (such as cameras or inertial sensors) embedded in a wide variety of digital devices and environments as well as due to the increasing storage of visual and sensor data (such as surveillance data, personal and multimedia databases, or simply the Internet). While storing and indexing large amounts of visual and sensor data has made tremendous progress, understanding of this 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 visual and multi-sensory data both for online use as well as for large-scale databases.

In the reporting period, the group focused on two main areas, namely computer vision and machine learning, but also worked on multimodal sensor processing. In the area of computer vision we address some of the most fundamental problems of image and video understanding such as object class recognition, people detection and tracking, and scene understanding. In the area of machine learning, we are focusing, on the one hand, on problems at the intersection of computer vision and machine learning (such as adversarial robustness), and on the other hand on more foundational problems in machine learning (such as interpretability of deep learning). In the area of multimodal computing we currently focus on perceptual user interfaces.

Research Areas and Achievements

In the following, we report highlights from our research largely following the sub-group structure of D2. As you will notice, there are substantial cooperations among the various sub-groups with many joint topics, cooperations and publications.

Computer Vision and Machine Learning *Investigator: Bernt Schiele*

Since the establishment of D2 at the Max Planck Institute for Informatics, the sub-group headed by Bernt Schiele has been working at the intersection of computer vision and machine learning. As a result of the increased importance of machine learning, in particular for our research but also overall in computer science, D2 itself and the sub-group have been renamed in *computer vision and machine learning* during the reporting period. Therefore, in the following, we will discuss research in the area of computer vision, in the area of machine learning, and at the intersection of these two areas.

Adversarial robustness is a key property of deep neural networks that is rather poorly understood, despite the large numbers of papers in the space. We have contributed papers to disentangle adversarial robustness and generalization, to confidence calibrated adversarial training, and to bit-error robustness. Similarly, interpretable machine learning is a key focus in recent years, where we, e.g., proposed a novel neural network architecture, that lends itself far better for interpretability than previous work. Another important area of our research is learning from less supervision where we have contributed a range of meta-learning approaches, novel algorithms for domain transfer and ‘lucid data dreaming’ for video segmentation. As in previous years, topics related to people detection have been another focus, including motion segmentation, tracking, and person search. Also, we have been working on novel deblurring methods combining e.g. traditional methods like Wiener filtering with deep network learning to obtain significantly improved results. In fact, these kind of models combining classic

methods (here Wiener filtering) with deep neural networks architectures is a highly promising area.

Major achievements in the reporting period, that are related to the above topics, have been published in 2 TPAMI and 2 IJCV journal publications, as well as 11 top-tier publications in computer vision and 6 top-tier publications in machine learning. A similar number of top-tier publications, co-authored by Bernt Schiele, are discussed in the other sections as they are based on co-operations with other senior researchers, in particular with Zeynep Akata and Mario Fritz.

Real Virtual Humans *Investigator: Gerard Pons-Moll*

Digitizing human beings would redefine the way we think and communicate (with other humans and with machines), and it is necessary for many applications, for example, to transport people into virtual and augmented reality, for entertainment and special effects in movies, and for medicine and psychology. The task requires many components of artificial intelligence such as natural language processing, emotion and gesture recognition, social signal processing, 3D appearance modeling and autonomous human motion synthesis. We focus on two aspects of human digitization: 3D human appearance and semi-autonomous motion generation and perception from sensory data with models that lie at the intersection of Machine Learning, Computer Vision and Graphics.

We learn digital humans by capturing and reconstructing real humans from consumer and specialized sensors: images, videos and scans. In contrast to problems like segmentation or classification where labels are available, natural images or videos paired with full 3D human reconstructions and motion are impossible to obtain. This poses several challenges, which we address through 1) Incorporating mathematical computer graphics models in deep learning frameworks to gain interpretability and control, 2) Learning from synthetic data and transferring to real data, 3) innovating methods to capture humans jointly with their 3D environments from mutli-modal data (images, depth, 3D scans and wearable sensors), and 4) developing computational models of human motion and interaction.

Our group has pioneered the areas of modelling people in clothing, and reconstructing *detailed* people from images in CV. Some highlights include the first methods to learn detailed of 3D people including texture (CVPR19,ICCV19) by learning from synthetic 3D scans and transferring to real data, the first differentiable and easy to use model of clothing with control over the shape, style, and dynamics using mesh representations (CVPR'20) and implicit functions (CVPR'21), and the first self-supervised scan-to-model registration method (NeurIPS'20). We have also contributed a method (CVPR'20) to reconstruct 3D objects and humans from sensory input using neural implicit functions, which is the SOTA in 3D tasks (as of Sept. 2020) (most accurate in ShapeNet and winners of all challenges of an ECCV Workshop on 3D shape completion of geometry and texture). In recent work, we introduce the first method to reconstruct 3D human motion and register it with the full 3D environment, including full buildings and areas of more than $4000m^2$ (CVPR'21), which is a major advance towards learning models of human scene interaction.

Our work is published in the major vision, graphics and learning conferences. In the reporting period we published 20 x (CVPR/ICCV/ECCV) papers including 5 Orals, 1 best student paper honorable mention at CVPR'20, 2 x PAMI, 2 x TOG, 2 NeurIPS (1 Oral), 1

3DV, and workshop papers. The PI was recognized with the German Pattern Recognition Award, the highest career award in Germany in the areas of Computer Vision and Machine Learning for researchers under 35. The group receives funding from Google, Facebook, Snap through faculty awards or research sponsorships, and from the DFG through the prestigious Emmy Noether Grant.

Multimodal Learning *Investigators: Zeynep Akata and Bernt Schiele*

While deep neural networks lead to impressive successes, e.g. they can now reliably identify 1000 object classes, argue about their interactions through natural language, answer questions about their attributes through interactive dialogues, integrated interpretability is still in its early stages. In other words, we do not know why these deep learning based visual classifications systems work when they are accurate and why they do not work when they make mistakes. Enabling such transparency requires the interplay of different modalities such as images and text, whereas current deep networks are designed as a combination of different tools each optimising a different learning objective with extremely weak and uninterpretable communication channels. However, deep neural networks draw their power from their ability to process large amounts of data in an end-to-end manner through a feedback loop with forward and backward processing. Although interventions on the feedback loop have been implemented by removing neurons and back propagating gradients, a generalizable multi-purpose interpretability is still far from reach.

Deep neural networks require a large amount of labeled training data to reach reliable conclusions. For instance, the system needs to observe the driver’s behavior at the red light to be able to learn to stop at red light both in a sunny and rainy weather, both in daylight and in night, both in fog and in snow, and so on. This causes a significant overhead in labelling every possible situation. Hence, our aim is to build an explainable machine learning system that can learn the meaning of “red light” and use this knowledge to identify many other related situations, e.g. although red light may look different in darkness vs daylight, the most important aspect in such a situation is to identify that the vehicle needs to stop. In other words, we would like to transfer the explainable behaviour of a decision maker to novel situations.

In summary, we would like to develop an end-to-end trainable decision maker operating in sparse data regime with an integrated interpretability module. Our main research direction to build such a system is two folds: learning representations with weak supervision and generating multimodal explanations of classification decisions.

Major achievements in the reporting period, include one major journal publication on our pioneering work on zero-shot learning (IEEE TPAMI 2019), on hallucinating natural scenes using attributes (ACM TOG 2019), and on sketch-based image retrieval (IJCV 2020), as well as our several publications published in several top machine learning and computer vision conferences, e.g. 9 x IEEE CVPR 2019-2021, 3 x NeurIPS 2019-2020, 1x ECCV 2020 and 1x ICLR 2020, in the fields of explainable machine learning, representation learning, domain adaptation and zero- and few-shot learning.

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

As techniques from machine learning and computer vision continue to mature, many of these approaches have made their way into products and systems. With this transition, they also have an increasing societal impact and contribute to the overall attack surface of our IT infrastructure. Not surprisingly, we see an increasing risk of attacks on such systems that threaten privacy and security. On the one hand, advanced image and text analysis algorithms acquire a detailed understanding of multi-modal content and thereby also infer private information of us. We seek a better understanding of private information in these modalities as well as provide methods to inform users about privacy risk and approaches how to mitigate them. On the other hand, future intelligent systems are subject to security risk as they have shown to be susceptible to inference attacks and manipulations. Evasion attacks by adversarial perturbations have shown how the outcome of classifications can be tempered with by inducing small, imperceptible perturbations onto the image data and inference attacks can lead to leakage of sensitive information on models and the associated intellectual property.

Major achievements in the reporting period include a deep analysis of person recognition in photo-collections (TPAMI'20), a gradient sanitised approach for learning differentially private generators (NeurIPS'20), membership attacks and defenses for semantic image segmentation (ECCV'20), Knock-off nets for stealing (CVPR'19) as well as a poisoning method to defend functionality stealing in deep neural networks (ICLR'20), and automatic testing and robustification via semantic adversarial data generation (ECCV'20). This work has led to 15 top-tier publications in vision, security and machine learning, as well as 2 TPAMI publications.

Combinatorial Computer Vision *Investigator: Paul Swoboda*

Combinatorial optimization is a fundamental tool in machine learning that offers the promise to improve performance of deep learning systems further by incorporating explicit prior knowledge in the optimization. The research group Combinatorial Computer Vision studies mathematical abstractions of computer vision tasks posed as combinatorial optimization problems. One focus is on developing efficient algorithms to solve the ensuing problems and to benchmark them w.r.t. metrics defined by the application. Another field of research is on integrating combinatorial optimization problems into deep networks to jointly train the network end-to-end. Application areas include tracking, correspondence problems and geological and biomedical image analysis. Moreover, the group conducts basic research into algorithm design.

Achievements during the reporting period include: (i) The first non-trivial partial optimality criteria for multicut, a popular graph clustering formulation. This allows to fix a subset of variables to their global optimum even before starting the optimization. (ii) A new tracking formulation together with a global optimization approach has achieved leading performance for multiple object tracking on the widely used MOT17 benchmark, outperforming previous significantly. (iii) The first convex optimization algorithm for an extension of graph matching, a popular formulation for correspondence problems, to the case of multiple graphs and its embedding in a deep network for joint learning of costs for semantic matching. Our works have been published in top-tier conferences such as CVPR, ICML, ECCV and ICCV.

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 ground-breaking work on privacy-aware eye tracking (Best Paper and Best Video Awards at ACM ETRA 2019), a major journal publication on our pioneering work on deep appearance-based gaze estimation (IEEE TPAMI 2019), as well as new computational methods for emergent leadership detection and anticipation of averted gaze in multi-party interactions (ACM ICMI 2019 and ACM ETRA 2019) that has now lead into a grand challenge organised in conjunction with ACM Multimedia 2021 (www.multimediate-challenge.org).

Publications, Cooperations and Awards

Chapter 32 contains a detailed report of the publications, cooperations and awards of the reporting period. From the journal publications 13 have been published or accepted at either IEEE PAMI, IJCV, or ACM TOG. From the conference publications, 62 have been published or accepted at one of the major computer vision conferences (CVPR, ICCV, ECCV). An additional 17 have been published at major conference in machine learning (14 at NeurIPS, ICML, and ICLR; 3 at AAAI and AISTATS). Additionally, 14 publications have been published at other major conferences across areas (e.g. ETRA, WACV, MLSys)

In the reporting period (2019–2021) 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: CMU, UC Berkeley, Stanford U, U Toronto, UBC, U Tokyo, ETH Zurich, EPFL, IST Austria, INRIA, U Amsterdam, TU Munich, U Tübingen, TU Darmstadt, MPI Intelligent Systems, CISPA, as well as Google and Facebook Research.

Members of the group received a range of prizes including, an ERC Starting Grant, the German Patter Recognition Award, the young researcher award from the Werner von Siemens Ring Stiftung, winning in all categories of the SHARP ECCV workshop challenge, a Snap research gift, a number of best paper awards (best student paper award honorable mention at CVPR 2020, best paper award ACM ETRA 2020, best paper award ACM ETRA 2019, best video award ACM ETRA 2019), as well as 10 outstanding or top reviewer awards at conferences such as ICLR, ICML, NeurIPS, CVPR, and ECCV.

5 D3: Internet Architecture

History, Group Organization and Development

The Research Group Internet Architecture (D3) was established in January 2018. The group’s organizational structure and more details of the members are provided later in this report.

Despite all pandemic-related obstacles, we were able to publish at top-tier conferences, e.g., 5(+2) out of 54 papers at the ACM Internet Measurement conference 2020, win multiple awards for our research papers, receive excellent press coverage for some of our work, conduct online lectures and seminars. Additionally, we contributed to the community by organizing the Network Traffic Measurement and Analysis Conference, TMA, in April 2020 as virtual conference and taking the responsibility as Co-TPC chair of the 2020 ACM conference on Emerging Networking Experiments and Technologies (CoNEXT). TMA had a record attendance, with 275 participants from all over the world. 19 talks and key notes plus eight lightning talks for PhD students attracted many more participants than the onsite events of the previous years.

Still, due to the COVID-19 pandemic, the overall research group grew slower than intended. Hiring came to a halt in Spring 2020 with the closing of the borders and the impossibility to obtain visas for Germany. This changed in late Autumn, and several new group members could finally join the group. Unfortunately, this delayed some research projects. Still, we were able to fully staff the Research Group “Work and Cooperation in the Sharing Economy” at the Weizenbaum Institute in Berlin lead by Anja Feldmann as Principal Investigator (PI). The goal of the Weizenbaum institute is Research for the Networked Society.

Last year, MPI-INF participated in the joined W2-tenure track faculty hiring process together with MPI-SWS and MPI-SP. We were very happy that Yiting Xia accepted our offer and joined MPI-INF in Fall 2020 as independent W2-tenure track faculty member. Anja Feldmann acts as her mentor and Yiting Xia’s research group is associated with D3 for infrastructure and administrative support. Yiting Xia is in the process to establish her own research group and is looking actively for good candidates. Having worked previously for Facebook, she will focus on the field of network infrastructure for cloud computing. Her research follows a cross-layer approach and covers broad topics for optimizing the network stack, including switch hardware, network protocols, software systems, and cloud applications.

Vision and Research Strategy

The Internet is an immensely successful human-made artifact that has fundamentally changed 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 or eye-ball interests as and how content, including user-generated data, is made available. The AS-level topology of the Internet has also experienced significant changes over time: It has evolved from a highly hierarchical topology to a flatter (non-hierarchical, simpler) topology. Content providers, e.g., Youtube, and content delivery networks, e.g., Akamai, 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 of 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 across infrastructures and (v) automating security mechanisms for protecting our infrastructure, (vi) incentivizing efficient network usage and network upgrades, and, (vii) finally, understanding the interactions of the infrastructure with the society.

We are generating staggering amounts of data everyday and almost everywhere and this trend will only increase in the years to come. Our analytical processing capabilities will also have advanced, e.g., offer efficient 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 and extremely hard challenges for efficient (wide-area) data analytics. A feasible solution will 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 using minimal resources. In this context data provenance, quality criteria and time constraints, both varying on a per-customer basis, will have to be taken into account, and it necessitates 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 of today’s high-speed (100 Gbps to 10s of Tbps) networks, particularly in the core of the Internet, require custom hardware solutions. This requirement has led to the use of specialized hardware and software as well as closed-box solutions for the main components of the Internet—the routers and the switches—thereby resulting in the ossification of the Internet. Software-defined networking (SDN) is one way to tackle this problem and some of the challenges in this 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.

We follow a data-driven systems research agenda to tackle these challenges: Collect data from operational networks, analyze them using big-data analytics to identify invariants, revisit assumptions and detect and localize performance bottlenecks in the Internet. We also use simulation environments to validate our analysis and to support “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, concepts for software-defined networking, as well as studies of the impact of the technology on the society.

Research Areas and Achievements

The main research areas of the department are Internet traffic analysis (measurements), innovative traffic control to future proof the Internet, online social networks, as well as network neutrality and the sharing economy.

Internet Traffic Analysis

The Internet is a massively heterogeneous and also continuously evolving ecosystem. Naturally, not 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 this ecosystem. Therefore, our research focuses on meticulously gathering measurements from diverse vantage points. We systematically analyze these measurements to characterize the performance and the operation of the Internet ecosystem, thus, identifying vulnerabilities and issues in the different components of this ecosystem (such as network protocols and devices). Characterizing the use of the infrastructure, e.g., through analyses of network traffic volume and dynamics helps us to understand how to evolve and to upgrade the networks for the future. Analyzing the use of AS-Path Prepending and the routing vulnerabilities it introduces, understanding the evolution of the Internet topology, and detecting IoT devices in the wild are some of the ongoing research efforts in this area. In addition to providing crucial insights into the current state of the Internet, our measurement-driven systems design approaches reveal how to upgrade and evolve the Internet infrastructure in anticipation of novel applications and changing usage patterns.

Research highlight: Impact of COVID-19 on Internet Traffic

The world-wide pandemic inflicted by the Corona Virus 2019 (COVID-19), a once-in-a-generation global phenomenon, upended the lives of many and destabilized the interconnected world economy. What started as a local health emergency in Asia at the end of 2019, turned into a global event at the beginning of 2020. By March 2020, the World Health Organization declared COVID-19 as a pandemic, causing many governments around the globe to impose strict lockdowns of economic and social activities to reduce the spread of COVID-19. The lockdown regulations changed the habits of a large fraction of the global population, who now depend on residential Internet connectivity for work, education, social interaction, and entertainment. In this project we simply question whether the Internet can sustain the resulting network load. To answer this question, we review the the impact of the first year of the COVID-19 pandemic on Internet traffic: We analyze its performance using Internet traffic data from multiple locations at the core and edge of the Internet.

One year after the first lockdown measures were enforced, the aggregated traffic volume increased by around 40%, well above the typical expected annual growth, see Figure 5.1. Additionally, workday traffic patterns have rapidly changed and the relative difference to weekend patterns has almost disappeared during lockdowns. Traffic associated with applications for remote work and education, e.g., VPN and video conferencing, observed more than 200% of increase.” Overall, our study (ACM IMC’20 [1], IETF IAB Workshop ’20 [2], ACM

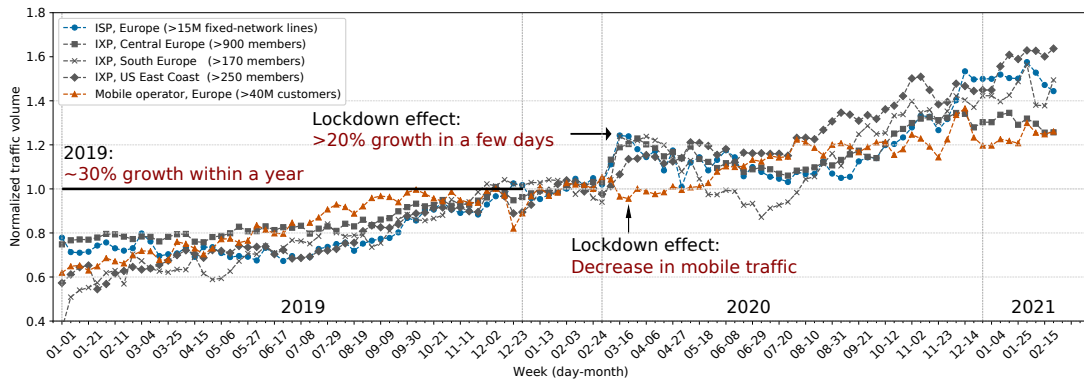


Figure 5.1: Traffic changes during the COVID-19 pandemic’s spring and fall waves at various Internet vantage points.

CACM’21 [3]) demonstrates that over-provisioning, proactive network management, and automation are key to build resilient networks, which can sustain drastic and unexpected shifts in demand such as those experienced during the COVID-19 pandemic.

Innovative traffic control/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. As if this immense amount of data was not enough, such data now comes from divers endpoints, widely distributed throughout the network (at its edge), thus, making the collection and analysis of this data in real time 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 new opportunities to revisit some long standing networking problems and design new practical, scalable solutions that exploit or leverage these recent changes. Redesigning protocols, e.g., the socket interface or TCP, or applications by exploiting new opportunities, e.g., for applications and network protocols to collaborate, 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.

High-speed networks are, on one hand, critical infrastructures, yet, they are, on the other hand, inherently fragile constructs. As such they can react badly to unforeseen events as well as targeted malicious practices. Finding ways to locate vulnerable devices as well as ensuring network stability even under duress is, thus, essential.

Research highlight: Steering Hyper-Giants’ Traffic at Scale

For an example project within this research area we revisit the topic of collaboration between content distribution networks (CDNs) and Internet Service Providers (ISPs). Namely, we

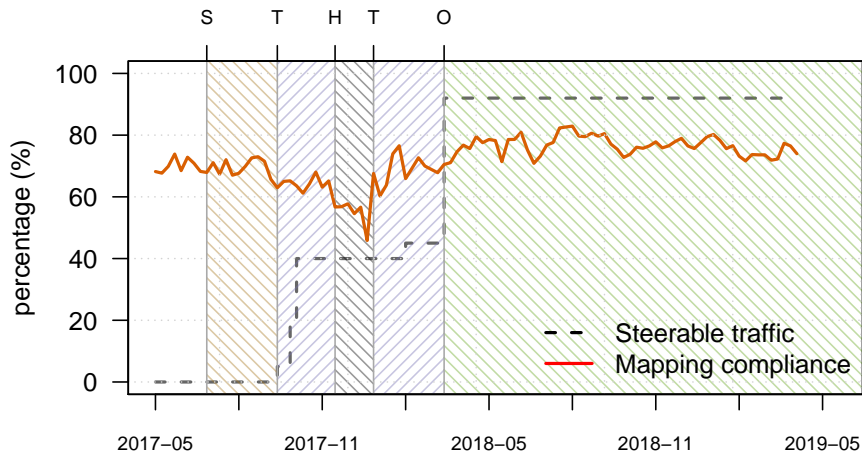


Figure 5.2: Timeline: Impact of CDN-ISP collaboration on share of optimally-mapped traffic annotated with events: Start (S/yellow), initial testing (T/blue), temporary hold (H/gray), operational (O/green).

report on our experience in building, rolling-out, and operating the first-ever large scale system to enable CDN-ISP collaboration inspired by our past research project PaDIS. As such this project shows how one can tackle the challenges of turning an idea that appeared in research papers to a fully operational system in today’s Internet. Indeed, our solution is deployed within one of the largest eyeball networks in the world with more than 50 million subscribers, and steers the traffic of one of the largest CDNs in the world. The system, enables automatic exchange of recommendations between the ISP and the CDN based on the analysis of more than 45 billion NetFlow records correlated with appropriate BGP information to provide recommendations that improve user-mapping. Our evaluation results, see, e.g., Figure 5.2, show that the compliance of the CDNs to the eyeball network’s recommendations is very high. This yields (1) close to optimal user-server mapping, and (2) reduction of this CDN’s traffic overhead on the ISP’s long-haul links of up to 15%, highlighting the mutual benefits (Best paper at ACM CoNEXT’19 [4]).

Online Social Networks

Over the past decade, Online Social Networks (OSN) have exploded in popularity, mainly because they help people in addressing their communication, information acquisition, and entertainment needs. At the same time, OSNs offer a fertile ground for the creation and amplification of important socio-technical issues like the spread of false information and hate speech. Therefore, it is imperative to analyze data from OSNs to understand, detect and mitigate these issues to minimize possible consequences both on the online and offline world (e.g., extensive dissemination of false information affecting people’s voting decisions). In this research area, we aim to analyze activity on OSNs to understand emerging socio-technical issues, develop tools/techniques to detect potentially harmful information and develop and assess the effectiveness of various mitigation strategies.

To achieve this goal, we meticulously gather large-scale datasets from multiple OSNs. Then

shared on 4chan and Gab, between 2016 and 2018, we find evidence of increasing antisemitism over time, especially in close temporal proximity with important real-world events (e.g., the 2016 US presidential election). Using word embeddings, graph analysis, and visualization techniques, we shed light on the use of language related to the Jewish community on these fringe Web communities. For instance, Figure 5.3 shows the main themes around the word “jew” on 4chan; we observe various thematic communities such as communities related to derogatory slang words (green community), nationalistic slurs (blue community), and religion (red community). Finally, using image processing techniques, we study the antisemitic Happy Merchant meme, finding that it is prevalent on these fringe Web communities and that its usage overlaps with other general-purpose and non-antisemitic memes. Taken altogether, we propose a framework for understanding and discovering hateful language/imagery and demonstrate its merits in understanding online antisemitism (Best Paper Honorable Mention Award at AAAI ICWSM’20 [5]).

Projects and Cooperations

Weizenbaum Institute:

Since 2017, Anja Feldmann is PI at the Weizenbaum Institute, Berlin. The Weizenbaum Institute for Networked Society is the German Internet Institute, a place of excellent research on the transformation and design processes of digital change. In the spirit of Joseph Weizenbaum, we research the necessary framework conditions, means, and processes for individual and social self-determination in a networked society. The project is funded by the Federal Ministry of Education and Research.

SupraCoNex:

In August 2020, we started a cooperation with Hochschule Nordhausen, Technische Universität Braunschweig, NewMedia-Net GmbH and Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V. The research project is about the development of new algorithms for sequential reinforcement learning, targeting the optimization of decision-making for resource management in wireless networks. The project is funded by the Federal Ministry of Education and Research.

Quantum Communication Development Environment (QCDE):

In August 2020, we together with (TU Munich, Humboldt University Berlin, University Bonn, Attocube systems AG and Qutools GmbH) submitted a concept for a New Major Instrumentation Initiative on the topic of “Quantum Communication Development Environment (QCDE)” which was successful. The DFG now invites proposals for the provision of Quantum Communication Development Environments (QCDE) at German universities in order to explore the scientific potential of this technology in research.

Prizes and Awards

Prof. Anja Feldmann, Ph.D.:

- German Academy of Science and Engineering (*acatech*), Elected in November 2019
- Community Contribution Award, ACM IMC 2020

- Best Paper Awards: ACM CoNEXT 2019, ACM CCR
- Elected member of the German Science Foundation (DFG) Review Board 2020
- IETF/IRTF Applied Networking Research Prize, 2020

Savvas Zannettou, Ph.D.:

- Two Best Paper Honorable Mention Awards, AAAI ICWSM 2020
- Best Reviewer Award, AAAI ICWSM 2020

Balakrishnan Chandrasekaran, Ph.D.

- Community Contribution Award, ACM IMC 2020
- Best DataSet Award, PAM 2020

Dr. Volker Stocker:

- Südwestmetal Förderpreis for the best Ph.D. thesis.

Dr. Apoorv Shukla:

- Best Paper Award ACM Sigcomm NetAI

Florian Streibelt:

- IETF/IRTF Applied Networking Research Prize, 2019

Patents

Prof. Anja Feldmann, Ph.D., Georgios Smaragdakis, Said Jawad Saidi:

- International Patent Application No. PCT/EP2018/064779 “Distributed and Timely Network Flow Summarization at Scale”

Press and Media Coverage

At the start of the global pandemic, the group’s attention immediately focused on the question how the Internet would deal with the drastically changing demands, with many people now suddenly working from home. The result of this research was widely covered in the media.

- Die Zeit: “Internet verkraftet Wachstum des Datenverkehrs durch Corona”, <https://www.zeit.de/news/2020-12/10/internet-verkraftet-wachstum-des-datenverkehrs-durch-corona>
- T-online: “Robustes Netz – Internet verkraftet Wachstum des Datenverkehrs durch Corona”, https://www.t-online.de/digital/id_89098572/robustes-netz-internet-verkraftet-wachstum-des-datenverkehrs-durch-corona.html
- Süddeutsche Zeitung: “Internet verkraftet Wachstum des Datenverkehrs durch Corona”, <https://www.sueddeutsche.de/service/internet-internet-verkraftet-wachstum-des-datenverkehrs-durch-corona-dpa.urn-newsml-dpa-com-20090101-201210-99-640708>

- European Union Agency for Cybersecurity: “Telecom Security During a Pandemic”, <https://www.enisa.europa.eu/publications/telecom-security-during-a-pandemic>
- The Berlin Institute for the Foundations of Learning and Data (Bifold), <https://bifold.berlin/new-study-by-bifold-researchers-how-did-covid-19-impact-internet-traffic/>
- New study by BIFOLD Researchers: How did COVID-19 impact internet traffic? – Bifold <https://bifold.berlin> for the CHIP magazine: issue April 2021, pages 46-49
- RBB24.de: “Der Internet-Verkehr hat sich drastisch geändert”; https://www.rbb24.de/studiocottbus/panorama/coronavirus/beitraege_neu/2021/01/btu-cottbus-senftenberg-corona-internet-auslastung-untersuchung.html
- Intelligente Welt.de: “Der Lockdown-Effekt: Wie Corona die Internet-Nutzung verändert” <https://intelligente-welt.de/?s=covid-10>
- Heise.online: “Wie Corona das Internet verändert hat” <https://www.heise.de/news/Wie-Corona-das-Internet-veraendert-hat-4985463.html>
- In-depth discussion of the paper at RIPE blog: https://labs.ripe.net/Members/oliver_gasser/the-lockdown-effect-implications-of-the-covid-19-pandemic-on-internet-traffic
- Handelsblatt: “Internet verkraftet Wachstum des Datenverkehrs durch Corona”

In addition, group members participated in the following activities:

- IdeenExpo Talk by Ranga Yogeshwar: “Yes we can – Wie gestalten wir den digitalen Wandel”, Jun 2019. <https://unternehmerhildesheim.de/aktuelles/ranga-yogeshwar-macht-mut-zu-industrie-40-0>
- Herrenhäuser Forum: “Blockchain – Unser Bankensystem der Zukunft”, Sep 2019. <https://unternehmerhildesheim.de/aktuelles/ranga-yogeshwar-macht-mut-zu-industrie-40-0>

Teaching

Winter Semester 2019

- Data Networks Lecture (Anja Feldmann, Balakrishnan Chandrasekaran)

Summer Semester 2020

- Data Networks Online-Lecture (Anja Feldmann, Balakrishnan Chandrasekaran, Oliver Gasser, Lars Prehn)

Winter Semester 2020/2021

- Hot Topics in Data Networks Seminar (Anja Feldmann, Oliver Gasser, Yiting Xia)

Summer Semester 2021

- Data Networks Online-Lecture (Anja Feldmann, Oliver Gasser, Devashish Gosain, Savvas Zannettou)

Master and Bachelor Theses

- Dil Paul José: Applying Machine Learning Techniques for Quality Assessment in Live Streaming Master thesis, 2021 (Supervisor: Balakrishnan Chandrasekaran, Mirko Palmer, Anja Feldmann)
- Evgeniya Khasina: The CoCo-Beholder: Enabling Comprehensive Evaluation of Congestion Control Algorithms, Master thesis, 2019 (Supervisor: Anja Feldmann, Keon Jang)
- Malte Appel: Exploiting Video Characteristics to Improve Video Streaming Master thesis, 2020 (Supervisor: Balakrishnan Chandrasekaran, Mirko Palmer, Anja Feldmann)

References

- [1] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis. The lockdown effect: Implications of the COVID-19 pandemic on internet traffic. In *IMC '20, 20th ACM Internet Measurement Conference*, Virtual Event, USA, 2020, pp. 1–18. ACM.
- [2] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis. A view of internet traffic shifts at ISP and IXPs during the COVID-19 pandemic. In *COVID-19 Network Impacts Workshop*, Virtual Workshop, 2020. IAB.
- [3] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis. A year in lockdown: How the waves of covid-19 impact internet traffic. In *Communication of the ACM*, 2021. ACM. Accepted for publication.
- [4] E. Pujol, I. Poese, J. Zerwas, G. Smaragdakis, and A. Feldmann. Steering hyper-giants' traffic at scale. In *CoNEXT '19, 15th International Conference on Emerging Networking Experiments And Technologies*, Orlando, FL, USA, 2019, pp. 82–95. ACM.
- [5] S. Zannettou, J. Finkelstein, B. Bradlyn, and J. Blackburn. A quantitative approach to understanding online antisemitism. In *Proceedings of the Fourteenth International Conference on Web and Social Media (ICWSM 2020)*, Atlanta, GA, USA (Virtual Event), 2020, pp. 786–797. AAAI.

6 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.

The most significant group development during the reporting period has clearly been the appointment of Christian Theobalt to the position of scientific director. Christian will chair the newly established Department on Visual Computing and Artificial Intelligence, and the institute will grow to six departments in the process for the foreseeable future. We are grateful to the Max Planck Society for providing the resources to make this appointment possible.

Over the last two decades the Computer Graphics 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 eight ERC Starting Grants¹, two ERC Consolidator Grants², four DFG Emmy Noether Fellowships³, seven EG Young Researcher Awards⁴, three German Pattern Recognition Awards⁵, and six EG PhD Thesis Awards⁶.

At the time of writing, the group encompasses two senior researchers with tenure (K. Myszkowski and C. Theobalt), three group leaders and senior researchers without tenure (V. Babaei, G. Singh, and R. Zayer), six postdocs, and 20 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 appointments⁷, or postdoc and senior postdoc positions⁸. Full details on current and recent group members are provided in the main part of the report.

¹ M. Magnor (2010), B. Rosenhahn (2011), B. Goldlücke (2013), C. Theobalt (2013), A. Oulasvirta (2014), J. Gall (2015), J. Steimle (2016), P. Didyk (2018)

² R. Mantiuk (2016), C. Theobalt (2017)

³ J. Kautz (2003), H. Lensch (2007), I. Ihrke (2012), J. Gall (2013)

⁴ H. Lensch (2005), J. Kautz (2007), M. Goesele (2008), C. Theobalt (2009), E. Eisemann (2011), T. Ritschel (2014), B. Masia (2017)

⁵ B. Rosenhahn (2009), C. Theobalt (2012), J. Gall (2014)

⁶ T. Ritschel (2011), T. Aydin (2012), P. Garrido (2018), P. Kellnhofer (2018), T. Leimkühler (2020), A. Serrano (2020)

⁷ F. Bernard (Univ. Bonn, Germany, 2021), R. Chen (Univ. Science and Technology (USTC) of China, 2019), H. Rhodin (Univ. British Columbia, Vancouver, Canada, 2019), A. Serrano (Defense Univ. Center, University of Zaragoza, Spain, 2021), S. Sridhar, Brown University, USA, 2019

⁸ Y. Gryaditskaya (Univ. Surrey, UK), T. Leimkühler (INRIA), A. Meka (Stanford), O. Tursun (Univ. Lugano, Switzerland), M. Zollhöfer (Facebook Reality Labs)

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 television, digital photography and the rapid development of new sensing devices, telecommunication and telepresence, virtual and augmented reality further indicate its potential and pose new challenges in the years to come.

To address these challenges, and in particular to seamlessly blend real and synthetic footage, we have adopted a new and more integrated scientific view of computer graphics as *3D Image Analysis and Synthesis* that takes into account the whole image processing pipeline from scene acquisition to scene reconstruction to scene editing to scene rendering. We also take into account human perception on all levels of the pipeline, and we exploit the abundance of digital visual data 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 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. Zayer)
- Neural Rendering and Computational Videography (C. Theobalt)
- Reconstructing the Static and Dynamic Real World (C. Theobalt)
- Realistic and Real-time Rendering (G. Singh)
- Perception: Advanced Displays, Material Appearance, and VR (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 researchs 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. We develop theoretical models, numerical schemes, and algorithms for a variety of problems in digital geometry processing.

Subdivision surfaces have become an invaluable asset in production environments. We demonstrated how the complete Catmull-Clark subdivision scheme can be abstracted in the language of linear algebra. Integrating domain knowledge about the mesh matrix data structure, we replace costly general linear algebra operations like matrix-matrix multiplication by specialized kernels. By further considering innate properties of Catmull-Clark subdivision, like the quad-only structure after refinement, we achieve an additional order of magnitude in performance and significantly reduce memory footprints. The versatility of the sparse matrix linear algebra abstraction underlying our work is further demonstrated by extension to other schemes such as $\sqrt{3}$ and Loop subdivision (EG'20, G. Enderle Award).

The rich and evocative patterns of natural tessellations endow them with an unmistakable artistic appeal and structural properties which are echoed across design, production, and manufacturing. Unfortunately, interactive control of such patterns-as modeled by Voronoi diagrams, is limited to the simple two dimensional case and does not extend well to freeform surfaces. Building on our previous work on layered fields for natural tessellations on surfaces, we developed a modular approach for creating and editing Voronoi-like cellular structures on surfaces. We focused on interactive modeling aspects and attempted to prioritize modeling over optimization. In this sense, our approach fits into the classical CAD workflow, allowing the user to freely design the partitions and interact with them in a meaningful manner. The approach also opens up a novel way for 3D printing decomposing the input surface into a cellular skeletal structure with a set of overlay shells (SIGAsia'19, EG'21).

We developed a method for volumetric shape interpolation with unique shape preserving features. The input to our algorithm are two or more 3-manifolds, immersed in 3-space and discretized as tetrahedral meshes with shared connectivity. The output is a continuum of shapes that naturally blends the input shapes, while striving to preserve the geometric character of the input (SIGAsia'19).

Neural Rendering and Computational Videography

We investigate the foundations of methods for computational videography, i.e., new approaches able to extract important higher level information about the scenes and the light transport in scenes seen in casually captured video. These methods are the foundation for advanced new ways to edit or enhance videos. In this context, we also made pioneering contributions to the area of neural rendering, an emerging new field investigating how new neural network-based image formation models and scene representations for photo-real image and video rendering can be created. In the latter area, in particular new ways to combine explicit models and representations with new neural representations has enabled us to make important steps ahead.

Inverse rendering deals with the estimation of different physical appearance properties (such as geometry and material properties) of a scene from image or video observations of the scene. In our work on deep reflectance fields, we developed the first method to estimate the full 4D reflectance field from two images of a person’s face captured under spherical gradient illumination conditions. The requirement of only two images allows for dealing with dynamic performances, enabling photo-realistic relighting of non-static faces at a higher quality compared to the state of the art. We also demonstrated the real-time global illumination decomposition of videos. (SIG’19, EG’20 STAR, ACM TOG’21).

Photo-realistic rendering of humans with the control over important scene aspects, such as body pose, appearance, viewpoints and lighting, is an important problem in computer graphics. In our research we combine traditional graphics approaches with new neural network based image synthesis strategies in novel ways. This enables photo-realistic rendering of dynamic performances with explicit controls of body pose, appearance, viewpoints and lighting, while only requiring monocular video or imagery as input. During the reporting period we developed a new method for generating video-realistic animations of real humans under user control only using monocular videos for training and testing, notably without the need of a production-quality 3D human model. This is achieved by training a neural network that implements a new strategy to translate approximate synthetic images of a human character into realistic imagery. To further improve the synthesis quality of neural human video rendering, we also proposed a novel method that explicitly disentangles the learning of time-coherent fine-scale details from the embedding of the human in screen space. We also developed the first capture and rendering framework that learns representations of appearance, viewpoint, and lighting of humans in arbitrary clothing from multi-view imagery captured in a controllable light stage setup, and we proposed a new algorithm for monocular neural re-rendering of a dressed human under a novel user-defined pose and viewpoint (ACM TOG’19, IEEE TVCG’20, SIGAsia’20, ECCV’20).

Realistic synthesis and control of face imagery is a highly relevant problem in computer graphics. We combine the strength of deep adversarial neural networks with model-based 3D reconstruction for novel neural rendering methods of human faces. Our techniques work on content captured by commodity RGB cameras and produce renderings of unprecedented photo-realism. We developed methods for processing still images and videos. For still images, we presented *StyleRig* the first approach for editing the head pose, facial expressions and scene illumination of portrait images generated with a pretrained generative face image model. Without our new method, the latter generative model completely lacks any semantic control of the generated output. Our approach maps the control space of a 3D morphable face model to the latent space of StyleGAN, which is the pretrained generative model we employ. In a follow-up work we extended the editing to images of real people. We also developed several techniques for video-based processing. In particular, we presented a novel approach for video-based dubbing (maintaining the talking style of the target actor), and we developed the first approach for high-quality text-based editing of video portraits. We also developed *Neural Voice Puppetry*, the first approach for high-quality audio-driven video reenactment. Finally, we developed *Egocentric Videoconferencing*, an approach that allows novel videoconferencing applications through the means of a single RGB camera mounted to the side of an eye-glass frame (SIG’19, ACM TOG’19, CVPR’20, ECCV’20, 2 x SIGAsia’20).

Reconstructing the Static and Dynamic Real World

We investigate advanced methods for scene reconstruction that capture high-quality models of shape, appearance and lighting of difficult smaller scale and larger scale real world scenes, ideally with only a single camera, or from casually captured community data. We also research 4D capture, where the goal is to reconstruct highly detailed 4D (3D+time) models of geometry, motion, appearance and illumination of real world scenes. Of particular interest is the reconstruction of humans. But the reconstruction of all deformable shapes outdoors, in the real world, at high quality, ideally with only a handful or even just one camera, is the long term goal. Characteristic for many of our recent works is that we developed new ways to deeply integrate model-based and deep learning-based methods and representations for reconstruction. In that context we also made pioneering contributions to the emerging field of neural scene representations. We also maintain one of the largest repositories of reference data sets for different sub-tasks in general static and dynamic scene reconstruction, as well as marker-less motion and performance capture.

An important challenge in visual computing is the reconstruction of static 3D environments, and the visual realism of virtual models critically depends on the quality of the shape models and their surface appearance. We have studied new methods of capturing high quality models of static real world scenes that only need a single color or depth camera. This includes a new learning-based method that produces consistent and highly-accurate depth estimations of indoor scenes from a monocular color video input by exploring the regularity of surface normal vectors, a method for robust and precise reconstruction of thin objects, such as wires, from an input video sequence (by introducing an effective scheme for detecting and handling self-occlusions of complex thin structures), and the first self-supervised approach for outdoor scene relighting from monocular images (2 x ECCV'20, SIG'20).

3D models of humans are nowadays used to augment the real world with our digital double or to create immersive effects in a fully virtual world. We managed to reconstruct detailed 3D human avatars including clothing and hair from a few RGB images. Our work combines bottom-up inference and top-down refinement leading to a significantly reduced reconstruction time from previous two hours to under ten seconds and we only require eight images compared to 120 images used by previous work. We also developed a simple yet effective method to infer detailed full human body shape from a single RGB image, and we presented the first approach to estimate 3D clothing and body shape under clothing as separate meshes using a only few RGB images. This method enables real-world applications like 3D virtual try-on and texture transfer as now body shape and clothing can be separated (CVPR'19, 2x ICCV'19, ECCV'20, NeurIPS'20)

We significantly advanced the state of the art in markerless 3D human pose estimation, and in human performance capture. Obtaining labeled training data for learning-based monocular 3D human pose estimation is quite cumbersome. To alleviate this issue, we proposed a novel 3D pose estimation algorithm that allows the neural network to make use of 2D-only in-the-wild labeled data which are much easier to obtain than the 3D labels. We also presented XNect, the first real-time CNN-based approach to perform multi-person 3D motion capture from a single camera. Human performance capture methods perform space-time coherent tracking of a deforming surface of a human from video data. We presented the first real-time human performance capture approach that reconstructs dense, space-time

coherent deforming geometry of full human bodies in general everyday clothing from just a single RGB video (LiveCap). We developed a novel deep learning approach for monocular dense human performance capture (DeepCap). The method is trained in a weakly-supervised manner based on multi-view supervision completely removing the need for training data with 3D ground-truth annotations (CVPR'20 best student paper honorable mention). We also presented the first approach for 3D capturing of high-speed human motions using a single event camera (ACM TOG'19, CVPR'19, 2x CVPR'20, SIG'20, SIGAsia'20, IEEE TVCG'19).

Markerless hand capture is essential for a variety of applications in graphics and HCI. Hand reconstruction from a single camera is particularly challenging, since this problem is highly ill-posed. We developed a new state of the art method for hand reconstruction from monocular depth camera imagery. It uses a new model-based training loss that allows self-supervision of a learning-based hand pose estimation method from depth maps. We also developed the first method to reconstruct the geometry of two interacting hands in real time with a single depth camera. (SIG'19, 2x CVPR'20, ECCV'20, SIGAsia'20).

Facial performance capture, i.e., inferring the 3D geometry and appearance information of faces from image observations has many applications in graphics. We presented an extensive survey on 3D morphable face models and contributed new methods for lightweight facial performance capture from monocular image information at greatly improved quality, efficiency, and robustness (CVPR'19, ACM TOG'20).

Capturing 3D geometry of deformable objects from a single monocular RGB camera is difficult, especially when no strong prior assumptions about the scene can be made. Non-rigid structure from motion (NRSfM) addresses monocular non-rigid 3D reconstruction by factorizing point tracks across multiple monocular views of a scene into camera poses and individual non-rigid states. The first neural dense NRSfM approach proposed by us parameterizes 3D shapes with an auto-decoder and learns a deformation model from dense input point tracks in an unsupervised manner. Due to the representation power of neural networks with fully-connected layers, we outperform previous competing NRSfM methods with engineered priors. Furthermore, the new type of constraints imposed on the latent shape variables in the Fourier domain allows automatic segmentation according to the shape similarity and detection of the sequence periods. We also introduced a learning-based method for non-rigid point set alignment (DispVoxNets) (2x CVPR'20, ECCV'20).

Classically, 3D objects are represented using meshes, point clouds, or voxel grids as basic, discrete data structures. Recently, neural representations, i.e. representations encoded in a trained neural network, have been proposed that, in certain cases, allow for a continuous representation, better handling of topology, and often have a learned focus of capacity on difficult parts. We developed Neural Sparse Voxel Fields, a neural representation for static 3D scenes where we condition a neural radiance field on learned local features that are attached to a sparse voxel grid. This enables us to disregard empty space from the cast rays, thereby speeding up the render time, a major bottleneck of prior work, by a factor of 10 and improving the results' quality. We also investigated deep mesh autoencoders for non-rigidly deforming objects (DEMEA). Our work learns a very compact deformation representation of only 8 or 32 latent dimensions by auto-encoding meshes using graph convolutions. (NeurIPS'20, 2x ECCV'20).

We made several contributions to foundational questions of cross-cutting relevance in 3D reconstruction. We investigated new ways of training and designing neural networks.

We proposed new solutions to challenging alignment and correspondence problems, and we particularly investigated rigid gravitational approaches. Other proposed methods address matching of multiple graphs. We are also one of the first teams to investigate how to formulate important foundational 3D representation problems such that they map to emerging new computing technology, in particular quantum computers (adiabatic quantum annealers) (2x CVPR'19, 2x ICCV'19, 2x CVPR'20).

Realistic and Real-time Rendering

We develop state-of-the-art rendering algorithms specialized for both offline (movie production) and real-time (games, virtual reality) applications. Since the basic skeleton of rendering is sampling, we develop algorithms to design sample correlations that helps not only in rendering but also in visualization and neural network training. We also investigated efficient visible set computations for streaming rendering, and we extended volumetric rendering algorithms for 3D printing.

Using Fourier tools we thoroughly analyzed the impact of sample correlations on variance during Monte Carlo integration, and we developed a novel theoretical framework for perceptual error optimization for Monte Carlo rendering that leverages kernel-based models from the halftoning literature. To make our research on sample correlations easily accessible to other fields in graphics and vision, we developed a deep learning pipeline to design novel sample correlations by simply feeding the network with a target power spectrum. No training data are required. The network can also be optimized for numerical integration problems for a given class of functions. We also demonstrated the importance of sample correlations in data plot visualization, and we studied the impact of sample set discrepancy on the training efficiency of implicit field regression networks (farthest point sampling vs. random sampling to select training points) (EG'19, SIGAsia'19, CGF'20, ECCV'20).

In the context of streaming rendering we revisited efficient potential visibility computations on graphics hardware and introduced the camera offset space (COS). The COS describes under which camera movement a sample location is covered by a triangle. Constructing the COS for all pixels of a rendered view leads to a complete potentially visible set (PVS) for complex scenes. Our GPU PVS implementation works without preprocessing, on-the-fly, on unconnected triangles (SIGAsia'19).

We also investigated modern deep learning approaches to robustly denoise Monte Carlo images (using a neural bilateral grid, to learn to denoise noisy inputs in the bilateral space), and for neural light field 3D printing. The developed method prints a light field as an attenuation-based volumetric display. At the core of the method is an encoding of the input light field as a continuous-space implicit representation in a neural network, and we leverage a fully differentiable formulation of volume rendering for end-to-end optimization (CGF'20, SIGAsia'20).

Perception: Advanced Displays, Material Appearance, and VR

Existing display devices introduce a number of physical constraints, which make it difficult to realistically reproduce the appearance of real-world scenes. Enabling such effects as glass-free stereoscopic 3D vision, eye accommodation, and motion parallax imposes not only technical

challenges, but also adds significant complexity to rendered content that should essentially reproduce complete lightfields. During the reporting period, our research focused (i) on eye accommodation and foveated rendering, (ii) on a novel implicit neural network representation (X-fields) that efficiently encodes temporally varying lightfields together with other sampled image representations, and enables their real-time rendering, (iii) on improving the visual experience in interactive VR and game setups, and (iv) on reducing the perceptual gaps between displayed and fabricated objects.

Foveated rendering aims to reduce the rendering load by lowering the image quality in the peripheral vision (outside the zone gazed by the fovea). We demonstrated that the computational savings of foveated rendering can be significantly improved if the local luminance contrast of the image is analyzed. To this end, we studied the resolution requirements as a function of luminance patterns and developed a novel low-cost predictor of the relevant parameters that can efficiently drive foveated rendering. Building on our successful collaboration with H. Fuchs, D. Dunn (UNC), and P. Diddy (USI) on near-eye display architectures that are capable of triggering eye accommodation cues, we also demonstrated that the capabilities of current and future display designs combined with efficient perception-inspired content optimizations can be used to improve human task performance beyond the human capabilities in the natural world. (SIG'19, IEEE TVCG'19, ISMAR'20).

Current and future sensors capture images of a scene from different points (video), from different angles (light fields), under varying illumination (reflectance fields), or subject to many other possible changes. In theory, this information will allow exploring time, view or light changes in Virtual Reality (VR). Taking an abstract view on all those dimensions and simply denoting any set of images conditioned on parameters as an *X-Field* (where *X* can stand for any combination of time, view, light), we developed a novel universal, compact, and interpolable representation for X-Fields, based on neural networks. The network represents the input to the rendering as an implicit map, that for any view, time, or light coordinate and for any pixel can quantify how it will move if view, time or light coordinates change (Jacobian of pixel position with respect to view, time, illumination, etc.). Our X-Field representation can be trained for a scene within minutes, leading to a compact set of trainable parameters and hence real-time navigation in view, time and illumination. In other works we investigated a no-reference metric to detect image-based rendering artifacts, and we demonstrated an application to sparse light field capturing, where a robotic arm with a camera is driven by our metric (CGF'19, SIGAsia'20).

Virtual Reality (VR) systems increase immersion by reproducing users' movements in the real world. However, several works have shown that this real-to-virtual mapping does not need to be precise in order to convey a realistic experience. In our work we measured detection thresholds for lateral translation gains of virtual camera motion in response to the corresponding head motion under natural viewing, and in the absence of locomotion, so that virtual camera movement can be either compressed or expanded while these manipulations remain undetected. The virtual expansion allows for the completion of virtual tasks within a reduced physical space and simulator sickness may be alleviated in simple scenarios when our compression method is applied. Immersion and motion sickness are also strongly dependent on the performance of VR applications. We proposed a CNN-based visibility metric, which detects artifacts caused by size reduction of textures. The metric has been applied in the game development pipeline and allows to prepare the optimal resolution of the textures

(CGF'19, SIGAsia'20).

Modern 3D printers allow for detailed control of pigment deposition in a volume. While this opens up exciting new opportunities for manufacturing synthetic objects with well-defined appearance characteristics, exact reproduction of a target appearance is hampered by the strong subsurface scattering within common print materials that causes non-trivial volumetric cross-talk between printer voxels. Consequently, the produced prints generally appear significantly different from the target models and lack high-frequency texture details. In our research we gradually developed a full software pipeline for high-fidelity appearance reproduction in 3D prints. In earlier we had shown how an iterative optimization scheme, in conjunction with Monte-Carlo simulation of light transport within a 3D print, can be used to find a spatial distribution of print materials that closely approximates a given target appearance; however, it was critically limited to color reproduction on planar surfaces. During the reporting period we successfully extended the method to arbitrary 3D shapes, enabling high-fidelity color texture reproduction on 3D prints by effectively compensating for internal light scattering within arbitrarily shaped objects. A remaining drawback of the repeated Monte-Carlo invocation within the refinement loop was the required computational effort. We succeeded in training a deep neural network to predict the scattering within a highly heterogeneous medium. With this update, our method performs around two orders of magnitude faster than Monte Carlo rendering while yielding optimization results of similar quality level. This for the first time enables full heterogeneous material optimization for 3D-print preparation within time frames in the order of the actual printing time (SIG'19, Optics Express'21, EG'21).

Computational Fabrication

The recent, wide availability of advanced manufacturing hardware has triggered huge interest in academia and industry. Most of the time, hardware offers capabilities which are never utilized, because the necessary software and algorithms are underdeveloped. We have a particular interest in the visual appearance of objects and strive for better algorithms for appearance 3D printing and for novel digital manufacturing platforms, such as laser marking.

We already reported on high-fidelity color reproduction in the previous section. Following color, gloss is arguably the most prominent visual attribute. During the reporting period we specifically addressed existing limitations for gloss control and proposed a new printing hardware based on piezo-actuated needle valves capable of jetting highly visous varnishes. Based on the new hardware setup we then developed a complete pipeline for controlling the gloss of a given 2.5 D object, from printer calibration to manufacturing the models with spatially-varying gloss. Another promising line of research is multi-material printing: Real-life objects are made of a variety of materials, most of which cannot be used by digital fabrication. To this end, we proposed to adapt the limited channels of digital printers to the appearance of the input. Particularly, given a specific painting, we asked what is the best ink set, from within a large ink library, to optimally reproduce this painting, and developed a novel mixed-integer linear programming solution for spectral reproduction whose continuous relaxation is convex and which scales gracefully (2 x SIGAsia'20, Electronic Imaging'20).

We are also interested in methods for achieving interesting optical effects using novel digital manufacturing platforms. During the reporting period we specifically investigated

the fabrication of Moiré patterns on curved surfaces (by superposing a partly absorbing layer and a layer formed by an array of cylindrical lenses), and computational laser marking. In contrast to more established technologies, laser marking is currently still in its infancy, and the relationship between the device's design space (laser parameters) and performance space (e.g., marked colors) is essentially unknown. We proposed a measurement-based, data-driven performance space exploration of the color laser marking process by formulating this exploration as a search for the Pareto optimal solutions to a multi-objective optimization problem (and solving it using an evolutionary algorithm augmented by a Monte Carlo approach). The explored set of diverse colors can then be utilized, e.g., to mark high-quality, full-color images on stainless steel (SIG'20, Optics Express'20).

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. Our 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 on high performance geometry processing and shape deformation, neural rendering and computational videography, markerless performance capture from both multiview and monocular video, the analysis and design of sampling patterns for rendering, perception-aware rendering and advanced displays, appearance fabrication is being referenced widely, and we are among the trend-setters in these areas.

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 2019 to spring 2021, members of the group published 25 papers at SIGGRAPH/SIGGRAPH Asia/ACM TOG, 16 papers at Eurographics/CGF/TVCG, and 33 papers at ICCV/ECCV/CVPR/3DV/IJCV/PAMI (see the main part of the report for details). We have actively participated in the program committees of major conferences (SIGGRAPH, SIGGRAPH Asia, Eurographics, Pacific Graphics, EGSR, 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 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 main part). Our software and datasets have been successfully used in a variety of projects (see main part), 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, UNC Chapel Hill, UBC, U. Hong Kong, U. Tokyo, Bar-Ilan U., U. Bath, TU Braunschweig, Cambridge U., CNRS, TU Delft, U. Edinburgh, U. Erlangen, Fraunhofer IIS, IST Austria, TU Graz, Technion Haifa, UCL London, EPFL Lausanne, USI Lugano, TU Munich, Charles U. Prague, U. Zaragoza, ETH Zürich. In addition, we also collaborated with some leading industrial research labs, including Google, Microsoft, Facebook Reality Labs, Intel, Nvidia, 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, the Saarbrücken Graduate School of Computer Science, the IMPRS-CS, the IMPRS-TRUST, and the Max Planck Graduate Center for Computer and Information Science.

Awards and Selected Academic Activities

Hans-Peter Seidel became a member of the DFG Senate (German Research Foundation (DFG)), and a member of the ACM SIGGRAPH Academy Committee. Karol Myszkowski was selected as Program Chair of ACM SIGGRAPH Asia 2020. Christian Theobalt received the EG Outstanding Technical Contributions Award 2020, and he was selected as Panel Chair for the ERC Consolidator Grants in Computer Science (PE6) 2021. Gurprit Singh is Co-Chair of the EG Symposium on Rendering 2021.

Thomas Leimkühler (2020) and Ana Serrano (2020) both won a Eurographics Dissertation Award. Thomas also won the Otto-Hahn-Medal of the Max Planck Society (2020). Vladislav Golyanik won the DAGM Dissertation Award 2020 of the German Association for Pattern Recognition (DAGM). Rhaleb Zayer won the EG Günther Enderle Award 2020.

7 D5: Databases and Information Systems

History and Group Organization

D5 was established in 2003. It is headed by Gerhard Weikum and, as of May 2021, consists of 19 doctoral students, 5 senior researchers (Klaus Berberich, Simon Razniewski, Rishiraj Saha Roy, Jilles Vreeken and Andrew Yates) and 4 other post-doctoral researchers.

Scientific Vision and Research Areas

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

Our research has been driven by the overarching vision of automatically constructing, growing and curating large-scale and high-quality knowledge bases from Internet sources. We have spearheaded this research avenue and refer to it as *knowledge harvesting*. To boost search, data analysis and language understanding, machines need to be equipped with comprehensive knowledge about the world's entities, their semantic properties and their relationships among each other. In addition to such encyclopedic facts, machines should also have commonsense knowledge about properties of everyday objects and human activities, and should even capture socio-cultural contexts of propositions.

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 and Content Analysis
- Natural Language Understanding
- Trust and Fairness

Contributions and Impact: Long-Term 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 (<https://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 (or knowledge graphs, as industry calls them) 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) as well as domain-specific knowledge bases (e.g., Amazon, Alibaba, Bloomberg, Mayo Clinic). 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 4000 citations, and the Yago2 paper from 2013 has more than 1200 citations. The influence and value of Yago has been recognized by the research community through the AIJ Influential Paper Award 2017 (<http://aij.ijcai.org/aij-awards-list-of-previous-winners>) for the 2013 Yago2 paper in the Artificial Intelligence Journal, and the W3C Seoul Test of Time Award 2018 (<https://www.iw3c2.org/ToT>) for the original WWW 2007 paper on Yago. We published a retrospective on knowledge harvesting in the LNCS 10000 anniversary volume, and a 350-page survey article is to appear in June 2021 (preliminary version: <https://arxiv.org/abs/2009.11564>).

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

Further Highlights: The Yago-Naga theme also spun off side projects which have been very influential. One of these is the RDF-3X database engine for efficient storage and querying

of subject-predicate-object triples, primarily developed by Thomas Neumann who is now a professor at TU Munich and has received the VLDB Early Career Award 2014 for this work (he also won the DFG Leibniz Award in 2019). The four main papers on this research together have nearly 2000 citations, and the RDF-3X system has been 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 1000 citations; it is among the most-cited papers on named entity disambiguation. This work has spawned a startup called Ambiverse and the development of the ambiverseNLU software suite, available as open source code.

In terms of methodology, the Yago-Naga project involved bridging the worlds of structured data (DB methods) and unstructured text (IR methods). One of our foundational works, on language models for temporal expressions, won the ECIR Test of Time Award in 2020.

Contributions and Impact: Major Results 2019–2021

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 2019–2021, the group had 23 papers in first-rate Web venues (WWW, WSDM, ISWC, ICWSM, IJCAI, AAAI), 9 in DM (KDD, ICDM, ECML/PKDD, NeurIPS), 5 in DB (VLDB, ICDE, SIGMOD, EDBT), 15 in IR (SIGIR, CIKM, ECIR, ICTIR), and 14 in NLP (ACL, EMNLP, NAACL). In addition, we are successful in building prototype systems and publishing demo papers at top venues like ACL, VLDB, SIGIR, WWW, etc.

Commonsense Knowledge: We advanced the classical AI theme of automatically extracting commonsense knowledge from Internet contents and organizing it into a knowledge base of everyday object properties, human activities and other concepts that humans are familiar with. Specifically, we devised new ways of tapping into question-answering forums and search-engine auto-completions, to infer crisp statements from Why questions. Our latest work, the ASCENT project, extends the state of the art by considering multi-word phrases as concepts with hierarchical refinement, and by adding an expressive set of facets to the usual form of subject-predicate-object statements. The benefits of our knowledge base are demonstrated by enhancing contextual language models on the task of masked-word prediction. This work was published in CIKM, AKBC and WWW. Simon Razniewski also gave tutorials at WSDM 2021 and AAAI 2021.

Discovering Dependencies and Subgroup Patterns 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 produce concise and interpretable outputs, based on information-theoretic principles and branch-and-bound optimization. This way, data scientists obtain insight and explanations about which attributes

depend on which other attributes and value conditions. Applications, using algorithms for subgroup discovery, include studies with material scientists (at the Fritz Haber Institute in Berlin) who got excited about the findings and their interpretability. Publications on these results include papers at KDD and ICDM (winning the best paper award), as well as an article in Nature Communications (jointly with material scientists).

Neural Information Retrieval: Deep neural learning, with Transformer networks and BERT-like language models, has become the state of the art for matching keyword queries and ranking answers. This works well for simple queries where small proximity windows are sufficient to identify the best documents, but it faces limitations on input size (short windows only) and it incurs high computational costs. We developed the PARADE framework to handle complex queries where answer-relevance cues are spread across wider passages in longer documents. PARADE improves prior works on both effectiveness (answer-ranking quality) and efficiency (run-time and memory footprint). Aggregating signals from multiple passages allows this approach to replace computationally-expensive Transformer decoders with light-weight networks. This line of work has led to publications in SIGIR, WSDM, CIKM, EMNLP and ECIR. Andrew Yates also co-authored a survey and accompanying tutorial at WSDM 2021 on the role of contextual language models for IR.

Entities with Quantities: We have been 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. Quantities are numeric expressions that denote financial, physical and other 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 developed methods for extracting quantity-centric facts from text corpora and ad-hoc web tables, and for answering queries with quantity-filter conditions such as “sprinters who ran 100 meters under 9.9 seconds”. Internet search engines support looking up quantities for a given entity, but often fail on evaluating quantity filters (lacking the proper understanding of queries and web contents). Knowledge bases, such as Wikidata, are rich in entities but very sparse in capturing quantity properties. Our work has been published in ICDE, ISWC, WSDM, SIGMOD and WWW. Demonstrator systems with large-scale data collections are publicly accessible on our website (<https://qsearch.mpi-inf.mpg.de/>).

Question Answering over Knowledge Bases and Web Text: We advanced 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. We devised a novel approach that can uniformly process structured KBs, text collections and their combination, by constructing graph representations and using graph algorithms to compute answers. This method is completely unsupervised, and can handle a wide variety of user inputs (incl. ungrammatical utterances) from simple to complex questions. We further extended the approach to cope with conversational QA where the user asks a series of questions, often in an incomplete style with implicit context. Our latest work on the latter topic employs reinforcement learning to leverage implicit feedback from users, when questions are reformulated after unsatisfactory system responses. Publications on this project include

papers at SIGIR, NAACL and CIKM. Rishiraj Saha Roy also gave tutorials on the topic at SIGIR 2020 and ICTIR 2020. Several prototype systems are publicly accessible on our website (<https://qa.mpi-inf.mpg.de/>).

Trust and Fairness: In the context of the ERC Synergy Grant imPACT, we address the broad issue of trust in online information, aiming to detect misinformation and disinformation. One topic that we have investigated in-depth is the transparency and explainability of recommendations to online users. We developed methods for generating user-comprehensible counterfactual explanations, which refer solely to the user’s own action history (clicks, likes, purchases, ratings etc.). This eliminates risks of de-anonymizing other users to which traditional explanations may refer (e.g., “you like X which was also liked by user U . . .”). A second topic of great societal importance is fairness in machine learning and algorithmic decision-making. We were among the first to address fairness issues in search rankings; our SIGIR 2018 paper is well cited. Recent work specifically focuses on models and methods for operationalizing principles of individual fairness and group fairness. These methods lower the complexity and effort for human intervention, while expanding the applicability and utility of fairness-aware learning. Results from this line of research have been published in SIGIR, EMNLP, WSDM, WWW, ICDE, VLDB and NeurIPS.

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 Oana Balalau (at INRIA, France), Joanna Biega (at MPI for Security and Privacy, Germany), Mario Boley (at Monash University, Australia), Gerard de Melo (at Hasso Plattner Institute, Germany), 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), Maximilian Dylla (Google), Patrick Ernst (Amazon), Subhabrata Mukherjee (Amazon, now Microsoft Research), Josiane Parreira (Siemens), Daria Stepanova (Bosch), Niket Tandon (Allen Institute for Artificial Intelligence) and Mohamed Yahya (Bloomberg). In the two-year timeframe 2019–2021, we had 9 doctoral students graduating.

Awards

Klaus Berberich (and his co-authors Omar Alonso, Srikanta Bedathur and Gerhard Weikum) won the ECIR Test of Time Award 2020 (for their influential paper “A Language Modeling Approach for Temporal Information Needs” at ECIR 2010). Joanna Biega received the GI DBIS Dissertation Award 2021 (for her doctoral thesis “Enhancing Privacy and Fairness in

Search Systems”). 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). Jilles Vreeken received the IEEE ICDM Tao Li Early Career Award 2018, for recognizing the excellence and impact of young researchers at most ten years after graduation.

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 14 Bachelor’s and Master’s theses in the 2019–2021 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, SIGIR and VLDB. 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 and in Nature Communications.

External Partners and Competitive Grants: D5 participated in the DFG-funded Collaborative Research Center (Sonderforschungsbereich) on Methods and Tools for Understanding and Controlling Privacy (from 2016-2019), where we collaborated 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, NU Athens, Bosch Center for AI, Northeastern U Boston, IIT Delhi, Georgetown U, Google Research, Hasso Plattner Institute, L3S Hannover, U Melbourne, U ParisTech, Rutgers U, Southwest Jiaotong U, Volkswagen Research, U Waterloo, and others).

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, extended until 2022. 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.

8 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 logics and their automation up to applications in research and industry.

There are currently 6 researchers and 2 PhD students in the group. Martin Bromberger, Mathias Fleury, and Marco Voigt have finished their PhD theses during the reporting period. Marco has joined the company “Logic for Business” L4B that evolved from the group. Mathias has joined the group of Armin Biere as a PostDoc and Martin is staying as a PostDoc. Sophie Tourret has taken a position at INRIA Nancy and remains associated with the group. Florian Frohn has taken an appointment at AbSint, a former startup of Saarland University concerned with static analysis techniques. Thomas Sturm and Jasmin Blanchette have appointments at CNRS and VU Amsterdam, respectively, and remain associated with the group.

Thomas Sturm, working in computer algebra, Jasmin Blanchette, working in interactive theorem proving and its mechanization, Sophie Tourret working on first-order logic reasoning are together with Christoph Weidenbach, the senior researchers of the group. Uwe Waldmann works with Jasmin Blanchette and Sophie Tourret on higher-order logic automation.

Vision and Research Strategy

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-order 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 of LIA (Linear Integer Arithmetic) constraints have been implemented in SPASS-SATT, Section 36.7.5, winning the QF_LIA category of SMT-COMP 2018, QF_LRA category of SMT-COMP 2019 and scoring second in the QF_LIA category of SMT-COMP 2019, only defeated by a portfolio solver including SPASS-SATT. 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) Arithmetic, (iii) Saturation Theorem Proving, (iv) Formalizing Logic, (v) Software and Applications. 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:

Model-Based Reasoning for Expressive Logics (CADE 2019, FROCOs 2019, ISOLA 2020, JAR 2020, VMCAI 2021) Two reasons for the success of CDCL are learning of only non-redundant clauses and not wasting computation time. Both rely on explicit partial model building: a trail of propositional literals either constitutes a model or it enables the learning of a non-redundant clause. Generalizing those principles to logics relying on infinite domain models is one of our major goals. Our general approach is to restrict CDCL-style reasoning to only finitely many domain elements at any point in time. This then either leads to learning a non-redundant clause or what we call a *stuck* state. A stuck state may either indicate a potentially infinite model or the necessity to extend the number of considered domain elements. For first-order logic without equality and for the combination of the Bernays–Schönfinkel fragment with a compact theory we could show that our approach is sound, complete, and generates only non-redundant clauses. For certain decidable theories, e.g., the extension of the first-order Bernays–Schönfinkel fragment with bounded difference constraints we could even show that it is sufficient to consider only an a priori bounded number of elements for effectively deciding satisfiability. The Bernays–Schönfinkel fragment extended with linear arithmetic (called SuperLog) is the specification, verification, and execution language for supervisor modeling in the collaborative research center “Perspicuous Systems”.

Qualitative Analysis of Reaction Networks and Epidemic Models (CASC 2019, CASC 2020, Math. Comput. Sci.) Chemical reaction network theory provides constructive mathematical modeling approaches for a wide range of biological networks, which include reaction networks as well as epidemic models. Examples are the well-studied MAP kinase and SEIR models, respectively. Variants of SEIR are used, e.g., for various subtypes of influenza A and recently also for Covid-19. The dynamics of the networks are described by ordinary differential equations (ODE) in time, assuming mass action kinetics. Traditional approaches have been based on numerical simulations of the ODE for fixed measured or estimated parameters. Our principal approach in the course of the SYMBIONT project¹, in contrast, combines symbolic methods and automated reasoning with model reduction methods aiming at a *qualitative* analysis of biological networks.

Going beyond numerical simulations aiming at the identification of particular steady states, we have investigated *structural properties of steady state regions*. For ODE with polynomial vector fields and fixed parameters those regions are complex or real varieties of the corresponding polynomial ideals. Existing approaches focus on the syntactic shape of the ideal generators in Gröbner bases, focusing on \mathbb{C} . We consider instead geometric properties of the actual varieties alternatively over \mathbb{C} or \mathbb{R} . We have systematically investigated 129

¹<https://www.symbiont-project.org/>

models from the BioModels.net repository and found a significant number of the varieties to be tori. Technically, we apply quantifier elimination methods to first-order characterizations of toricity. More recent work on *multi-timescale reduction* can be viewed as a generalization of the sharp distinction between steady and non-steady states. The given ODE of a biological model is reduced to a chain of several simpler ODE, linked together via formally precise hyperbolic attractivity conditions. Each obtained ODE is subject to a timescale factor so that the reaction speed decreases by orders of magnitude when stepping from one ODE to the next in the chain. Technically, all computationally hard parts of our quite comprehensive algorithmic framework are formulated as satisfiability modulo theories problems over various theories, viz. QF_LRA, QF_NRA, QF_LIA. Off-the-shelf SMT solvers performed well with our computations on models from BioModels.net.

Higher-Order Mechanization (CADE 2019, IJCAR 2020, CPP 2021) While interactive proof assistants are based on expressive higher-order logics, the most widely employed calculi for automated theorem proving have originally been developed for first-order logic. If one wants to increase the degree of automation of interactive proof assistants by integrating automated theorem provers, this mismatch necessitates an obfuscating sequence of preprocessing steps to reach first-order logic. We have continued our investigation of “graceful” generalizations of high-performance first-order proof systems to higher-order logic: systems that work natively in higher-order logic but do not sacrifice the efficiency of state-of-the-art first-order proof systems on the first-order part of the input. We have developed superposition calculi for higher-order logic with λ -abstractions and with or without interpreted Booleans. The calculi have been implemented in the Zipperposition prover, enabling Zipperposition to win the higher-order category of the CASC prover competition at IJCAR 2020, by a margin of 21 percentage points ahead of the next competitor. The completeness of our calculi has been proven following a framework that we also developed during that time and whose range goes beyond higher-order superposition to cover most known saturation-based calculi.

Formalizing CDCL (NFM 2019, LPAR 2020) CDCL is a very successful paradigm to prove satisfiability in propositional logic. We have formalized the calculus in the proof assistant Isabelle/HOL, making it possible to prove for example that CDCL does not relearn clauses. In order to go from the abstract specification to executable code, we first added a critical data structure, the two-watched-literals scheme. Then we define several intermediate programs, with increasingly efficient data structures and code, and formally prove that the improved programs behave the same as the original ones. On the last refinement level, we use a high-performance memory representation, an *arena*, that makes it possible to represent clauses and their headers consecutively in memory – a technique widely used by hand coded state-of-the-art CDCL solver implementations.

We also used our base CDCL to express and prove correctness of extensions that are based on CDCL with branch-and-bound with applications to model covering and the computation of weight optimal models.

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. In the interdisciplinary project SYMBIONT we apply reasoning in non-linear arithmetic to problems from biology. We constitute the core of the ERC starting grant Matryoshka on the mechanization of higher-order logic. In the transregional collaborative research center TRR 248, "Perspicuous Systems", established in January 2019, we investigate automated reasoning in the context of perspicuous software systems in close cooperation with researchers from TU Dresden.

Prizes and Awards

Our software SPASS-SATT and Zipperposition won awards at the respective system competitions. In addition to several best paper awards, Martin Bromberger won the EAPLS dissertation award and Christoph Weidenbach the Thoralf Skolem award of CADE.

Teaching

We regularly teach and supervise bachelor and master students at Saarland university. We are responsible for the core lecture "Automated Reasoning" of the master curriculum. In addition, we provided the lectures "Concrete Semantics with Isabelle/HOL" and "Competitive Programming" in the reporting period.