max planck institut
informatik

Report 2013/2014

14 15 16 17 18

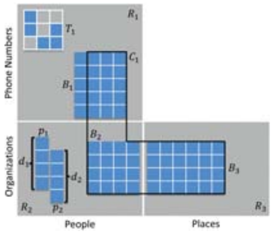

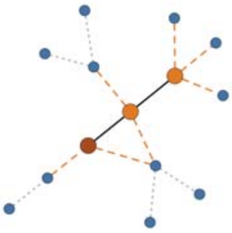

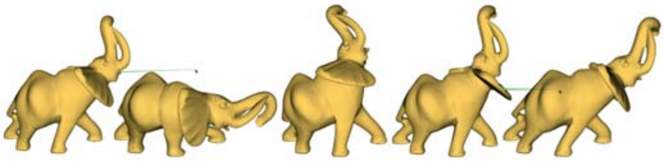

Report 2013/2014

Directors

Prof. Dr. Thomas Lengauer
Prof. Dr. Kurt Mehlhorn
Prof. Dr. Bernt Schiele
Prof. Dr. Hans-Peter Seidel
Prof. Dr. Gerhard Weikum

Scientific Advisory Board

Prof. Dr. Trevor Darrell, University of California, Berkeley, USA
Prof. Dr. Nir Friedman, Hebrew University of Jerusalem, Israel
Prof. Dr. Pascal Fua, École Polytechnique Fédérale de Lausanne, Switzerland
Prof. Dr. Thomas A. Funkhouser, Princeton University, USA
Prof. Dr. Jürgen Giesl, RWTH Aachen, Germany
Prof. Dr. Alon Halevy, Google Research, Mountain View, USA
Prof. Dr. Renée Miller, University of Toronto, Canada
Prof. Dr. Yves Moreau, Katholieke Universiteit Leuven, Belgium
Prof. Dr. Nicole Schweikardt, Humboldt University Berlin, Germany
Prof. Dr. François Sillion, INRIA Rhône-Alpes, Grenoble, France
Prof. Dr. Emo Welzl, Swiss Federal Institute of Technology, Zurich, Switzerland
Prof. Dr. Gerhard J. Woeginger, Eindhoven University of Technology, Netherlands

Curatorship Board

Hon. Adv. Professor (Tsinghua) Dr. Reinhold Achatz, Head of Corporate Technology, Innovation & Quality, Thyssen Krupp AG, Essen
Dr. Siegfried Dais, Robert Bosch Industrietreuhand KG, Stuttgart
Christiane Götz-Sobel, Head of the editorial department of Science and Technology ZDF, Munich
Prof. Dr. Joachim Hertel, CEO of Denkprozess GmbH, Saarbrücken
Prof. Dr. Henning Kagermann, President of the National Academy of Science and Technology (acatech), Munich
Annegret Kramp-Karrenbauer, Saarland Prime Minister, Saarbrücken
Prof. Dr. Volker Linneweber, President of Saarland University, Saarbrücken
Prof. Dr. Wolf-Dieter Lukas, German Federal Ministry of Education and Research, Bonn
Dr. Nelson Mattos, Vice President Engineering, EMEA, Google, Zurich, Switzerland
Prof. Dr. Wolffried Stucky, Institute of Applied Informatics and Formal Description Methods (AIFB), University of Karlsruhe
Prof. Dr. Margret Wintermantel, President of the German Academic Exchange Service (DAAD), Bonn

CONTENTS

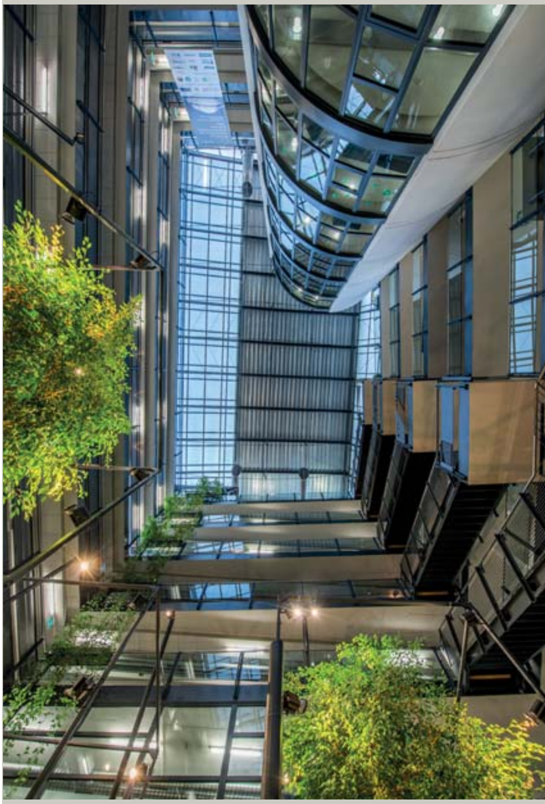

PREFACE

The Max Planck Institute for Informatics regularly publishes a report for the general public. Now again, we would like to take this opportunity to present some of the current topics, goals and methods of modern informatics. We introduce the work of our Institute and hope to arouse your interest, in the fascinating world of our science.

The Max Planck Institute for Informatics aims to be a beacon of research in informatics. We aim to have impact in the following ways: First, through our scientific work, which we disseminate mainly through publications and books, but also in the form of software and Internet services. Second, through the training of young scientists, particularly during their doctoral and postdoctoral phases. We are educating future leaders in science and industry. Third, by assuming a formative role in our field. We initiate and coordinate large research programs, and serve on important committees. Fourth, by attracting talent from within and outside the country. About half of our research staff of over 180 scientists comes from abroad. Fifth, through the transfer of our results into industry. These transfers take place through collaborative projects, spin-offs, and people. Sixth, by building a world-class competence center for informatics in cooperation with our partners: Saarland University with its departments of Informatics and Computational Linguistics as well as its Center for Bioinformatics, the German Research Center for Artificial Intelligence, and the Max Planck Institute for Software Systems. We have continued to be very successful in all of these undertakings. Our success has become more visible through the erection of many new buildings around the Platz der Informatik on the Saarbrücken campus. In the next few months, the building for the Center of Internet Security, Privacy and Accountability (CISPA) will be finished at the East corner of Stuhlsatzenhausweg. We are in the final phase of the projects funded through the German Excellence Initiative – the Cluster of Excellence “Multimodal Computing and Interaction” and the “Saarbrücken Graduate School of Computer Science”. The recent evaluation of the Institute by our Scientific Advisory Board has, once more, been very positive.

This report is structured as follows: After an overview of the Institute and its departments and research groups, we present the main areas of recent work. These topics span several departments and will also be the focus of our work in the next years. This part is followed by a section on our alumni and a brief overview of current events. The last part of the report contains a selection of recent scientific publications and a compact presentation of the Institute through key indicators.

We wish you much enjoyment reading this report.

Thomas Lengauer *Managing Director*

The Max Planck Institute for Informatics, an Overview

Information technology influences all aspects of our lives. Computer systems, hardware, software, and networks are among the most complex structures that have been constructed by man. Computational thinking is a new way of studying the universe. Basic research in informatics is needed to cope with this complexity, to lay the foundations for powerful computer systems, and to further develop computational thinking.

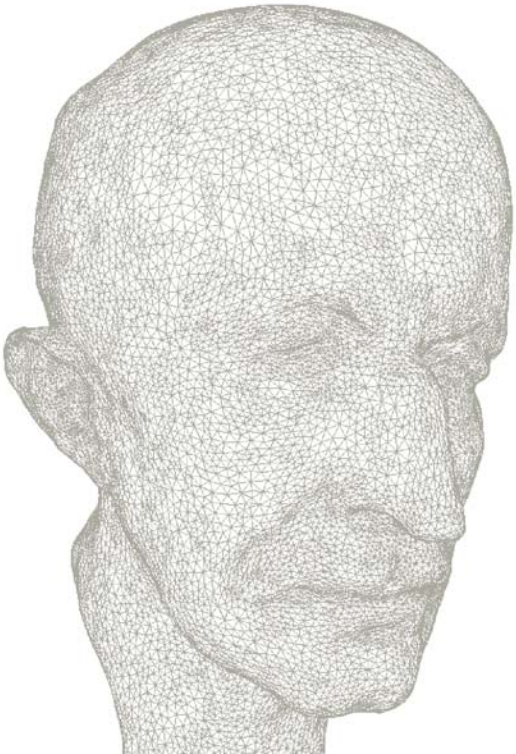

OVERVIEW

Basic research in informatics has led to dramatic changes in our everyday lives in recent years. This has become particularly clear in the last two decades: The worldwide web, search engines, compression processes for video and music, and secure electronic banking using cryptographic methods have revolutionized our lives just a few years after their discovery at universities and research institutes.

The Max Planck Society, the leading organization for basic research in Germany, reacted to these challenges by founding the Max Planck Institute for Informatics (MPI-INF) in Saarbrücken in 1990. In 2005, the Max Planck Institute for Software Systems (MPI-SWS) was established with sites in Saarbrücken and Kaiserslautern. There are departments with a strong emphasis on informatics in other Institutes of the Max Planck Society as well. The restructuring of the Max Planck Institute for Metal Research into the Institute for Intelligent Systems has further strengthened informatics within the Max Planck Society. Given the importance of the area, the establishment of further Institutes for informatics or related areas is desirable.

Goals

The Max Planck Institute for Informatics aims to be a beacon of research in informatics. We aim to have impact in the following ways:

First, through our scientific work, which we disseminate mainly through publications, but also in the form of software and internet solutions. Presently, we are concentrating on algorithms for very large, multimodal data. Multimodal includes text, speech, images, videos, graphics, and high-dimensional data.

Second, through the training of young scientists, particularly in the doctoral and postdoctoral phases. We are educating future leaders for research and business. Over 180 researchers are working at our Institute and remain with us for three years on average. In this way, we provide the society with over 60 well-trained young scientists each year.

Third, by our role in the scientific field. We initiate and coordinate large research programs and serve on important committees, e.g., the *"Wissenschaftsrat"* or DFG Review Boards. The Institute has played a significant role in forming the Excellence Cluster *"Multimodal Computing and Interaction"* and the *"Graduate School of Computer Science"*.

Fourth, by attracting talent from within and outside the country. Half of the research staff of the Institute comes from outside Germany. This strengthens the talent base in Germany and establishes bridges to foreign countries.

Fifth, by transferring our results to industry. These transfers take place through collaborative projects, spinoffs, and people. Intel founded the Intel Visual Computing Institute in 2009 together with the UdS (Saarland University), the DFKI (German Center for Artificial Intelligence), the Max Planck Institute for Software Systems, and the Max Planck Institute for Informatics.

To establish a start-up culture and support start-ups and licensing agreements, the Max Planck Society and Saarland University jointly created the IT Inkubator GmbH, located at the Starterzentrum of Saarland University. The IT Inkubator GmbH focuses on supporting scientists in founding and financing high-tech start-ups, as well as marketing the research results to industries nationally and internationally. The official kick-off event took place in March 2015 and was attended by Saarland Governor, Annegret Kramp-Karrenbauer; the President of the Max Planck Society, Prof. Dr. Martin Stratmann; and the President of Saarland University, Prof. Dr. Volker Linneweber.

Sixth, by building a world-class competence center for informatics in cooperation with our partners, Saarland university, the DFKI, and the MPI-SWS. We have been very successful in all of these endeavors in recent years.

History and Organization

The Max Planck Institute for Informatics was founded in 1990 with Kurt Mehlhorn as the founding director. He has directed the *"Algorithms and Complexity"* department since then. Also, Harald Ganzinger was involved from the very beginning and directed the *"Logic of Programming"* department until his death in 2004. A third department, *"Computer Graphics"*, followed in 1999 under the direction of Hans-Peter Seidel. Thomas Lengauer then joined in 2001 to direct the *"Computational Biology and Applied Algorithmics"* department. Gerhard Weikum has directed the *"Databases and Information Systems"* department since 2003. In the summer of 2010, the new *"Computer Vision and Multimodal Computing"* department, directed by Bernt Schiele, was formed.

In addition to the departments, the Institute is home to independent research groups: The *"Automation of Logic"* research group headed by Christoph Weidenbach as well as several junior research groups within the Excellence Cluster *"Multimodal Computing and Interaction"* and in the Max Planck and Stanford cooperation (Mykhaylo Andriluka, Roland Angst, Andreas Karrenbauer, Michael Kerber, Haricharan Lakshman, Yangyan Li and Dominik Michels).

Senior Researchers

The scientific work of the currently 29 research groups at the Institute is orthogonal to its organizational structure. These research groups are headed by so-called senior researchers. Directors and independent research group leaders automatically become senior researchers through their appointment to the Institute; moreover, postdoctoral researchers can become senior researchers by passing an acceptance process similar to the procedure for a professorial appointment to a university. The senior researchers currently with the Institute are: Ernst Althaus, Klaus Berberich, Andreas Bulling, Piotr Didyk, Mario Fritz, Martin Hoefer, Olga Kalinina, Andreas Karrenbauer, Michael Kerber, Thomas Lengauer, Christoph Lenzen, Tobias Marschall, Kurt Mehlhorn, Pauli Miettinen, Karol Myszkowski, Nico Pfeifer, Tobias Ritschel, Michael Sagraloff, Bernt Schiele, Marcel Schulz, Hans-Peter Seidel, Viorica Sofronie-Stokkermans, Jürgen Steimle, Thomas Sturm, Christian Theobalt, Jilles Vreeken, Christoph Weidenbach, Gerhard Weikum, and Andreas Wiese.

Research Topics

Algorithms are our central research subject. An algorithm is a general recipe for solving a class of problems. We enlarge the computational universe through the design of new and better algorithms. We prove the correctness of these algorithms and analyze their performance. We implement them and validate them through experiment. We make them available to the world in the form of software libraries and Internet services. We study inherent properties of computation and investigate techniques for robust software design.

We apply the algorithms to interesting application problems. Improved hardware has led to impressive gains in efficiency; improved algorithms can lead, and have led, to even larger gains. Here is an example: The state of hardware and algorithms in 1970 made it possible to calculate an optimal travelling salesman tour through 120 cities (a classic optimi-

OVERVIEW

zation problem and a recognized benchmark for computer performance). Adding one more city multiplies the runtime of the classical algorithm by 120, and adding another city multiplies it by 121; the runtime grows super-exponentially. By combining the increased hardware speed, made possible by today's technology and the classical algorithm from 1970, we can solve problems with 135 cities. It is the advance in algorithms that allows us to find an optimal route through thousands of cities today; if we were to rely solely on advances in hardware, such high performance would not even be possible in a thousand years.

The scientific problem of understanding hand-developed algorithms and their realization in computer programs has two important aspects. First, there is the question of whether the program calculates what it is intended to without "crashing", "freezing", or blocking all of the computer's resources. Second, there is also the question of whether the program is "efficient", i.e., whether the best possible algorithm has been found. The department *"Algorithms and Complexity"* concentrates on the resource requirements of algorithms. The most important resources are running time (*How long must I wait for the result of my computation?*) and space requirement (*Do I have enough storage space for my calculation?*). The group develops new algorithms with improved running times and storage requirements and also studies the basic limits of computation (*How much time and space are provably necessary for a computation?*).

The *"Automation of Logic"* research group investigates logic-based generic procedures for solving "hard" combinatorial and decision problems. Typical logic-based applications are the verification of systems with a significant discrete proportion and optimization problems.

Nowadays, computers are used to model, represent, and simulate aspects of real or virtual worlds. Since the visual sense is a key modality for humans, computer graphics has become a key technology in modern information and communication societies. The department *"Computer Graphics"* researches the entire processing chain from data acquisition to modeling (the creation of a suitable scene representation) and image synthesis (the generation of views for human consumption). The following scientific challenges emerge from this: For the input side, we want to develop modeling tools for the efficient handling and processing of large data flows, and, on the output side, we seek new algorithms for the fast computation of high-quality views; these algorithms should exploit the capabilities of modern graphics hardware.

The department *"Bioinformatics and Applied Algorithmics"* addresses the potential of computing for the life sciences. The life sciences have a high and ever-increasing demand for algorithmic support due to the recent turning of biology into a quantitative discipline. Algorithms and statistical analyses play a central role in the preparation and configuration of biological experiments and even more so in the interpretation of the biological data generated by them. As a consequence, the computer has become an indispensible tool for biology and medicine. Vast amounts of data need to be processed in modern biology, and the molecular interactions in the living organism are so complex that studying them is unthinkable without computer support. Therefore, bioinformatics methods have become essential for modern research on the diagnosis and treatment of illnesses.

The mission of the *"Databases and Information Systems"* department is the automatic acquisition of knowledge from Internet sources, text documents, and social media. The ultimate goal is a computer that knows everything Wikipedia contains in a formal knowledge representation and that can harness this for understanding natural language, deep text analytics, and logical inference. The department has been a trendsetter in the automatic construction of formal knowledge bases, which has led to major innovations like the Google Knowledge Graph and similar projects at Internet companies and other enterprises. Our knowledge base YAGO has been used in many projects around the world, including the IBM Watson System when it won the TV quiz show Jeopardy in early 2011 and leveraged YAGO for semantic type checking. Digital knowledge is a key asset for a variety of intelligent computer applications: searching for entities and complex relationships on the Internet or in enterprises, analyzing news articles and social media, summarizing entire text corpora, exploring and analyzing heterogeneous collections of Big Data, and many more.

The department for *"Computer Vision and Multimodal Computing"* investigates the processing and understanding of sensor information. Sensors range from being relatively simple, e.g., GPS and acceleration sensors, to being very powerful sensors, e.g., cameras. They are embedded in more and more devices. Although the algorithmic processing of sensor information has advanced considerably, it is still mainly limited to low-level processing. In particular, we are far from being able to fully interpret and understand sensor information. Such sensor understanding is, however, a necessary prerequisite for many areas such as man-machine interaction, the indexing of image and video databases, or for autonomous systems such as robots.

Excellence Cluster “Multimodal Computing and Interaction“

The Institute plays an important role in the Cluster of Excellence on *“Multimodal Computing and Interaction”*, which was established by the DFG (German Research Foundation) in 2007 and successfully renewed in 2012. All directors of the Institute are principal investigators of the Cluster, and Hans-Peter Seidel is the scientific coordinator.

The starting point of the Cluster’s research program was the observation that there have been dramatic changes over the last two decades in the way we live and work. Twenty years ago, most digital content was textual, whereas today, its scope has expanded to include audio, video, and graphics that are available practically everywhere. The challenge is to organize, understand, and search this multimodal information in a robust, efficient, and intelligent way and to develop reliable and secure systems that allow intuitive multimodal interaction. The Cluster addresses this challenge. The term “multimodal” describes the different kinds of information such as text, speech, images, video, graphics, and high dimensional data as well as the way in which it is perceived and communicated, particularly through vision, hearing, and human expression. The Cluster’s primary goal is to improve the ability of computer systems to efficiently and robustly collect, then process and present data of various modalities. A further goal is to analyze and interpret large volumes of distributed, noisy, and incomplete multimodal data and then organize and visualize the obtained knowledge in real-time. This is called mulitmodal processing. Everyday interpersonal communication is based on numerous different modalities; so, the Cluster’s second major goal is to facilitate a similarly natural multimodal interaction of information systems – anywhere, anytime. The systems must consider environmental context, react to language, text, and gestures, and then respond in the appropriate modalities. This research program builds on existing strengths: The Cluster comprises the departments of Computer Science, Computational Linguistics and Phonetics, and Applied Linguistics at Saarland University, our Max Planck Institute for Informatics, the German Research Center for Artificial Intelligence, and the Max Planck Institute for Software Systems. The participating institutions have agreed on a joint long-term research programas the basis of their work. The university and the state government provide special assistance to the Cluster.

A prominent goal of the Cluster is to qualify and promote young scientists. Saarbrücken has acquired a reputation over the years as an “elite school” for young scientists. For this reason, the majority of the allocated funds have been devoted to the establishment of junior research groups. This concept has proven successful in the past years, and several young scientists have since taken professorships in Germany and abroad.

OVERVIEW

Publications and Software

The scientific results of the Max Planck Institute for Informatics are distributed through presentations, publications, software, and web services. Our publications appear in the best venues of the computer science field. Most publications are freely available in the Institute's repository [http://www.mpi-inf.mpg.de/publications/].. Some of our results are available in the form of downloadable software or as a web service, for example CGAL (Computational Geometry Algorithms Library) and the clinically-used web service geno2pheno for HIV therapy support. Publications in the form of software and web services make our results available more directly and to a larger audience than classical publications.

Promotion of Young Scientists

A further goal of the Institute is to create a stimulating environment for young scientists, in which they can grow, develop their own research programs, and build their own groups. We concentrate on doctoral and postdoctoral training. Our 137 PhD students are trained in cooperation with the Graduate School of Computer Science at Saarland University and the *International Max Planck Research School for Computer Science (IMPRS-CS)*, page 102. Our postdoctoral researchers participate in international collaborations such as the *"Max Planck Center for Visual Computing and Communication"* (a cooperation with Stanford University in the area of computer graphics), page 25, the *"Indo Max Planck Center for Computer Science"* (a cooperation with leading universities in India), page 133, and in our many EU projects.

We encourage our young scientists to establish their own research programs and then move on to other institutions. Since the founding of the Institute, several researchers from the Max Planck Institute for Informatics in Saarbrücken have joined other research facilities, many of them accepting professorships. The section on the alumni of our Institute in this report represents the sustained yield of our efforts.

Structure of the Report

After a brief introduction to the departments and research groups of our Institute, we survey our work by means of representative examples.
We group the examples into subject areas, each of which spans at least two departments. This report ends with a section on alumni, a presentation of the IMPRS-CS, an overview of recent events, a presentation of the Institute in figures, infrastructure aspects, and a tabular listing of cooperations and publications.

Enjoy reading it. :::

Algorithms and Complexity

PROF. DR. KURT MEHLHORN

DEPT. 1

The group has existed since the founding of the institute. It currently has about 40 staff members and doctoral candidates. Our goals are

- **carrying out outstanding basic research in the field of algorithms**
- **implementing our fundamental work in demonstrators and generally useful software libraries**
- **promoting young scientists in a stimulating workgroup environment.**

We are successful in all three undertakings. We are effective through publications, software, and people. We publish in the best journals, present our results at the leading international conferences in the field, and our LEDA and CGAL software libraries are used worldwide. Many group alumni are in top positions domestically and abroad.

CONTACT

Algorithms and Complexity
Secretary
Ingrid Finkler-Paul | Christina Fries
Phone +49 681 9325-1000
Email infi@mpi-inf.mpg.de
chfries@mpi-inf.mpg.de

DEPT. 1

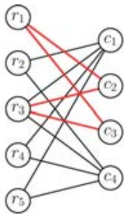

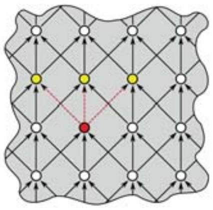

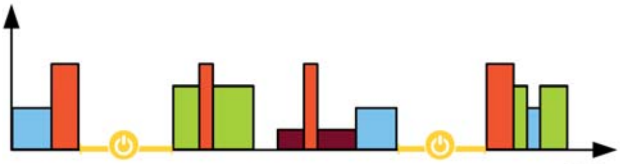

Algorithms are the heart of all software systems. We work on the design and analysis of many facets of algorithms: combinatorial, geometric, and algebraic algorithms, data structures and search processes, very different computer models (sequential, parallel, distributed, flat data or hierarchical data), exact and approximate solutions, problem-specific methods and general heuristics, deterministic and randomized solutions, upper and lower bounds, and analyses in the worst case and on the average. We develop efficient algorithms for abstract versions of application problems as well as for concrete applications, e.g., optimized control of optical displays. A part of our theoretical insight flows into the implementation of software demonstrators and software libraries; we devote a part of our practical work to collaborating with companies.

Outstanding theoretical results in the past two years include a new algorithm for the highly efficient, and at the same time, provably correct isolation of zeros, so-called certifying algorithms for combinatorial problems in graphs (i.e., algorithms that do not only compute a solution, but also a proof for the correctness of the solution) and a simple (therefore also practical) algorithm for calculating equilibrium prices in the Fischer and Walras market model. In addition to designing good algorithms, we strive to better understand fundamental algorithmic principles: we have found an explanation for how a certain slime mold can apparently determine shortest paths through mazes, we have given a mathematical explanation for the very fast dissemination of information in social networks, and we have explained why the quadratic runtime barrier for computing the Frechet distance could not be broken for more than twenty years (a faster algorithm would imply a faster algorithm for the satisfiability problem).

Notable practical results in recent years include our contributions to the Computational Geometry Algorithms Library (CGAL) as well as to the search engine Complete Search and its application in one of the most significant computer science bibliographic databases. Our software library, LEDA (begun in the 1990s) continues to enjoy strong demand in science and industry.

Our theoretical and experimental work strengthen each other. For example, the efficient implementations for dealing with curves and surfaces in CGAL are based on our theoretical work on isolating roots of polynomials and finding solutions to systems of polynomial equations. The combination of theoretical and experimental research in algorithms has become a widely accepted research direction. The DFG supported it through its priority program *“Algorithm Engineering”*. The group is involved in two international projects: the GIF project (geometric computing – with the University of Tel Aviv) and the Indo-German Max Planck Center for Computer Science (IMPECS). In addition, there is a regular international exchange through various means of support for excellent research. Our scientific staff members have received scholarships from the Humboldt Society, the European Union (Marie-Curie), and the German Academic Exchange Service (DAAD). In Germany, we participate in the trans-regional special research area AVACS (Automatic Verification and Analysis of Complex Systems).

The promotion of young scientists is an integral component of our work. We give lectures at Saarland University, which are addressed not only to students, but also to our doctoral candidates. Our annual summer school draws leading experts and distinguished international graduate students to Saarbrücken. Two of our alumni, Susanne Albers and Peter Sanders, received the Leibniz-Award of the German Research Foundation. :::

Computer Vision and Multimodal Computing

PROF. DR. BERNT SCHIELE

DEPT. 2

The department was founded in 2010 and currently includes 30 scientists. The group's main research areas are computer vision with a focus on object recognition and 3D scene description as well as multi-sensor-based context recognition in the area of ubiquitous and wearable computing.

CONTACT

Computer Vision and Multimodal Computing
Secretary
Cornelia Balzert
Phone +49 681 9325-2000
Email d2-office@mpi-inf.mpg.de

DEPT. 2

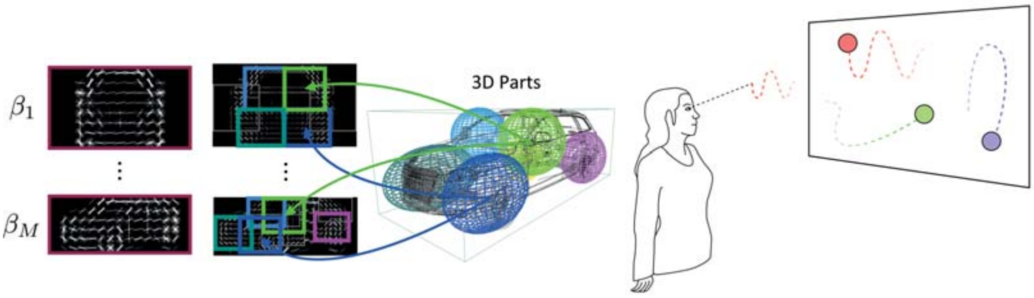

Sensors such as cameras, GPS and accelerometers are being increasingly embedded in devices and environments, and they are already helpful in various ways. Computer-controlled processing of sensor information has made enormous progress but is generally limited to simple matters. This means, in particular, that devices and computers which have access to this sensor information do not fully interpret it and thus cannot truly understand their environment. The department is therefore concerned with the understanding of sensor information, using both powerful sensors, such as cameras, and embedded sensors, such as gyroscopes and accelerometers.

In the area of computer vision, the department deals with problems such as object recognition, one of the basic problems of image understanding. In recent years, this area of computer vision has made impressive progress, and the department has played a pioneering role in presenting several innovative approaches. One of these approaches recognizes and segments the object simultaneously, leading to significantly improved results in comparison with standard approaches. Current work has presented approaches for learning 3D models. These not only enable the robust detection of objects, but also allow additional parameters that are essential for 3D scene understanding, e.g., the estimation of visual direction.

Another central theme of the department is people detection and tracking using moving cameras. This problem is not only scientifically challenging but also has a wide variety of applications, such as image and video understanding, or in robotics and the automobile industry. For example, cars equipped with such a camera may predict the movements of pedestrians and therefore react to their behavior more effectively. The department has developed approaches that robustly detect people and track them over longer periods of time, even if they might be out of sight for a long time. A recently presented approach not only describes people but also entire 3D scenes, representing a further step towards complete image and scene understanding.

In addition to computer vision, the second central research area is the processing and understanding of multimodal sensor information. The underlying observation here is that an increasing amount of computers and sensors can be found in our environment, in objects and even in our clothing. Context awareness and sensing is often seen as a means of making the computing tasks sensitive to the situation and the user’s needs. Ultimately, context awareness may support and enable seamless interaction and communication between humans and computing environments without the need for explicit interaction. In this area, the department has presented approaches to recognize long-term activities and to model personal daily routines. It was also possible to show that a person’s interruptibility can be predicted with surprising accuracy using a few sensors embedded in their clothing.

The third research area of the department is machine learning. This plays the important role as a cross-cutting theme, as the other research areas make extensive use of probabilistic modeling and inference techniques. These allow, for example, the modeling of the uncertainties that exist with any sensor processing. In addition, they allow the use of large amounts of data, and can also elegantly integrate previous knowledge. :::

Computational Biology and Applied Algorithmics

PROF. DR. THOMAS LENGAUER, PH.D.

DEPT. 3

This department dates back to October 2001 and is directed by Prof. Dr. Dr. Thomas Lengauer. The department currently comprises about 20 scientists, performing research exclusively in computational biology.

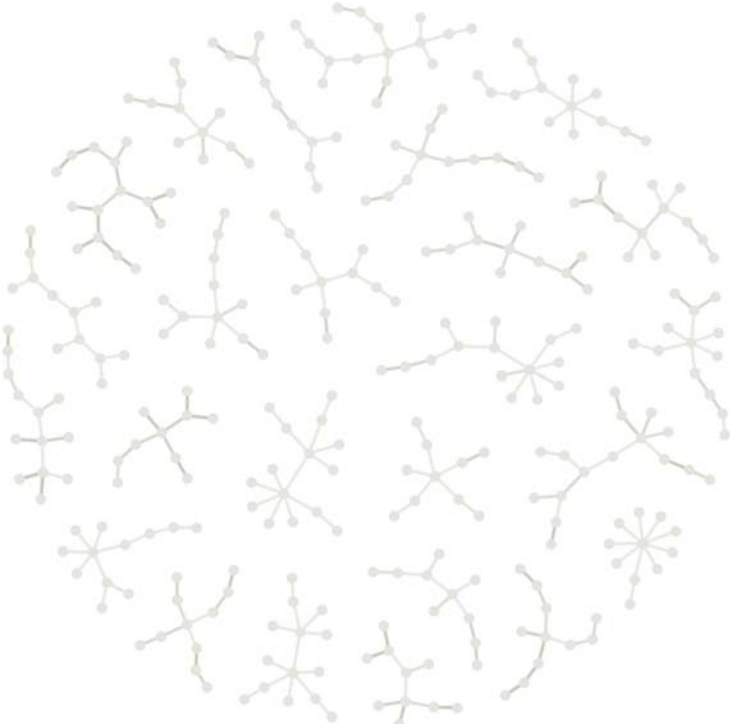

CONTACT

Computational Biology and Applied Algorithmics
Secretary
Ruth Schneppen-Christmann
Phone +49 681 9325-3000
Email ruth@mpi-inf.mpg.de

DEPT. 3

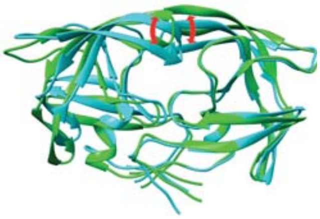

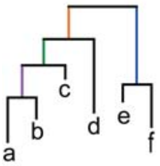

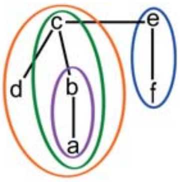

The department focuses on topics more or less closely related to the diagnosis and therapy of diseases. At the molecular level, diseases can be understood in terms of anomalies in the biochemical circuitry of an organism (*"Bioinformatics"*, page 42). The building blocks of such biochemical networks are DNA, RNA, and proteins. The circuitry is realized by transient binding of the molecules among each other and with small organic ligands. In this manner, proteins catalyze chemical reactions, regulate the expression of genes, and transduce signals in and between cells.

Our research concentrates on two groups of diseases.

On the one hand, we investigate viral infections. The molecular understanding of such diseases requires the elucidation of the function of the relevant viral proteins and their interactions with the molecules of the infected patient as well as with drug molecules. So, analysis must be carried out at the level of the three-dimensional structure of the involved molecules.

Our methods are applied to concrete diseases such as AIDS (*"Structural Insights into the Emergence of Resistance in the HIV-1 Protease"*, page 46), Hepatitis B and Hepatitis C. AIDS plays a special role; with this disease, the Max Planck Institute for Informatics goes even a step further. We analyze the resistances of HI viruses against administered combination drug treatments (*"Bioinformatical support of HIV therapy"*, page 47). Furthermore, we investigate the pattern of geographical distribution of the virus (*"Reconstructing Pathogen Transmission Networks Using Genomic Data"*, page 48, and *"Reconstructing Pathogen Transmission Networks Using Genomic Data"*, page 49).

Diseases such as cancer, neurodegenerative diseases, or immunological diseases are based on different principles. Here, the starting point is an interplay between the genomes of the patient and the environment. Therefore, an early diagnosis of a disease like cancer can be based on the analysis of genetic or epigenetic aberrations of the relevant tissue (*"Integrative Analysis of Cancer Data"*, page 50). The epigenome – the aggregate of all chemical modifications of the DNA inside the nucleus as well as the chromatin enveloping it – is the key to the complex regulation of the cell, which moves out of kilter in such diseases. Charting the entire epigenome is one of the great challenges facing molecular biology in the next years (*"Charting Epigenomes"*, page 51, and *"Analyzing the Three-Dimensional Organization of the DNA Inside the Nucleus"*, page 52).

A more general approach to supporting the diagnosis of a large variety of genetic diseases is described in *"Phenotype-Based Computational Clinical Diagnosis for Genetic Diseases"*, page 53.

A large part of the method development in the department results in software systems, which are used worldwide by many academic, clinical, and also industrial users. Examples, which are presented in this issue (*"EpiExplorer and RnBeads: Integrative Analysis of Epigenomic Data"*, page 54, and *"BiQ Analyzer HiMod for the Locus-Specific Analysis of DNA Modification Data"*, page 55), include the field of epigenetics, the analysis of protein function and protein interaction networks as well as the optimization of AIDS therapies (*"Bioinformatical Support of HIV Therapy"*, page 47).

Additional topics on which we perform basic research include the evolutionary relationships in the viral kingdom (*"Detection of Related Proteins in Unrelated Viral Species"*, page 56) and analysis methods and software for analyzing variations in individual (e.g. human) genomes (*"Structural Variation in Genomes"*, page 57).

Our department is one of the main pillars of the Bioinformatics Center Saar, an inter-faculty center at Saarland University focusing on teaching and research in the area of bioinformatics. The department is a member of the German Arevir network and the European Consortia EuResist and CHAIN, all of which contribute to bioinformatics research on viral resistance development. Moreover, the department coordinates the bioinformatics for the German Epigenome Program (DEEP), sponsored by the German Federal Ministry of Education and Research (BMBF), and is a partner in the European research programs BLUEPRINT (epigenetics) and PREDEMICS (research on viruses with considerable epidemich potential). :::

Computer Graphics

PROF. DR. HANS-PETER SEIDEL

DEPT. 4

The Computer Graphics work group was founded in 1999 and now includes 40 scientists. An important characteristic of our work is the thorough consideration of the entire processing pipeline, from data acquisition to modeling and image synthesis (Visual Computing, resp. 3D Image Analysis and Synthesis). Typical for the field area is the co-occurrence of very large data sets and the demand for fast, possibly interactive, high quality visual feedback. Furthermore, the user should be able to interact with the environment in a natural and intuitive way.

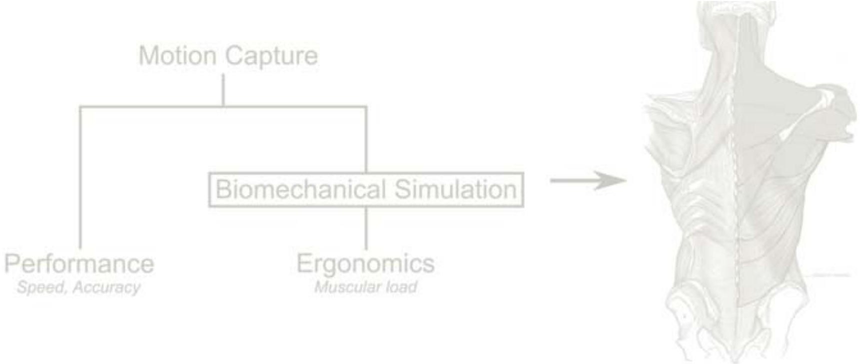

CONTACT

Computer Graphics
Secretary
Sabine Budde
Phone +49 681 9325-4000
Email budde@mpi-inf.mpg.de

DEPT. 4

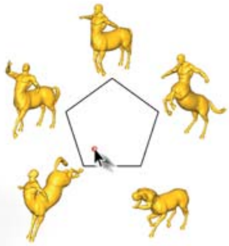

During the last three decades computer graphics has established itself as a core discipline within computer science and information technology. Ten years ago, most digital content was textual. Today it has expanded to include images, video, and a variety of graphical representations. New and emerging technologies such as multimedia, digital television, and digital photography combined with the rapid development of new sensing devices, telecommunication and telepresence, virtual reality, and 3D Internet further indicate the potential of computer graphics in the years to come. Typical for the field is the co-occurrence of very large data sets and the demand for fast, possibly interactive, high quality visual feedback. Furthermore, the user should be able to interact with the environment in a natural and intuitive way.

New approaches are needed from a scientific point of view in order to meet the above-mentioned challenges. An important characteristic in our department is, therefore, the thorough consideration of the entire pipeline, from data acquisition to modeling (creation of a suitable internal computer scene description) and image synthesis (generation of novel renderings). In our opinion, this point of view is necessary to exploit the capabilities of modern hardware, on both the input (sensors, scanners, digital photography, digital video) and the output (graphics hardware, advanced displays) sides. Moreover, we take into account existing digital footage when analyzing and synthesizing novel scene content. In particular, we seek to extract powerful priors from the abundance of digital visual data that can then assist us during the acquisition, reconstruction, editing, and image formation processes. Our long term goal is to develop methods that efficiently handle the huge amount of data during the acquisition process, extract structure and meaning from the abundance of digital data, and turn this data into graphical representations that facilitate further processing, rendering, and interaction.

The scientific activities of the Computer Graphics work group are embedded in a series of project activities on national, European, and international levels.

The Max Planck Center for Visual Computing and Communication was jointly established by the Max Planck Institute for Informatics and Stanford University in 2003 with support from the BMBF (German Federal Ministry of Education and Research). The collaboration has two intertwined goals: establish a joint research program in the area of visual computing and communication, and incorporate a strong career development component to alleviate the shortage of qualified faculty and scientists in information technology in Germany. The Max Planck Center fosters the professional development of a small number of outstanding individuals by providing these with the opportunity to work at Stanford as visiting assistant professors in the area of visual computing and communication for two years, and then return to Germany to continue their research as leaders of junior research groups at the Max Planck Institute for Informatics, and ultimately as professors at major universities.

Furthermore, the Computer Graphics group is significantly involved in the activities of the Cluster of Excellence on *“Multimodal Computing and Interaction”*. The Cluster of Excellence was established by the German Research Foundation (DFG) within the framework of the German *“Excellence Initiative”* in 2007 and was successfully renewed in 2012. Hans-Peter Seidel is the scientific coordinator of the Cluster.

Another important partner is the Intel Visual Computing Institute (Intel VCI) that was established in 2009 as a European hub for academic research in the field of visual computing. Intel VCI is a collaborative effort between on campus organizations (Saarland University, German Research Centre for Artificial Intelligence, MPI for Software Systems and MPI for Informatics) and Intel, who sponsors the research.

Altogether, since the establishment of the group fifteen years ago, the institute has produced more than 40 faculty members in computer graphics and visual computing in Germany and abroad. The exceptional quality of this academic offspring is further documented by a number of prestigious national and international awards that current and former institute members have been able to attract over the years: German Pattern Recognition Award (3x), Eurographics Young Researcher Award (7x), ERC Starting Grant (5x), Eurographics Technical Contributions Award, Humboldt Research Prize, ERC Advanced Grant, DFG Leibniz Prize (2x), Eurographics Distinguished Career Award. :::

Databases and Information Systems

PROF. DR.-ING. GERHARD WEIKUM

DEPT. 5

The department, headed by Gerhard Weikum, pursues research in three areas:

1. **Algorithmic extraction of knowledge from text and Web sources, and automatic construction of large knowledge bases.**
2. **Semantic language understanding by automatically detecting and disambiguating names and phrases and mapping them to entities and relations in data and knowledge bases.**
3. **Big Data analytics and explorative data analysis on structured data such as biological networks or product ratings as well as unstructured data such as news articles or online discussion forums.**

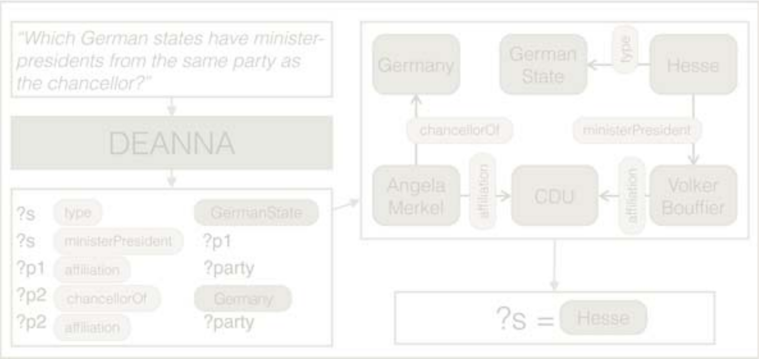

CONTACT

Databases and Information Systems
Secretary
Petra Schaaf
Phone +49 681 9325-5000
Email d5-office@mpi-inf.mpg.de

DEPT. 5

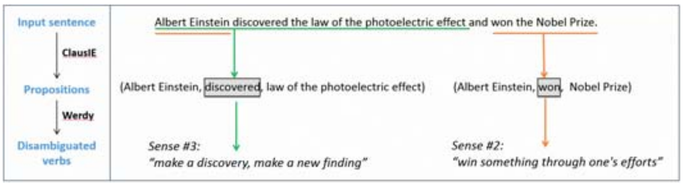

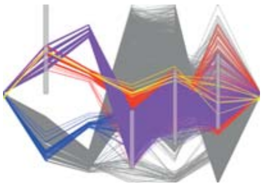

The central theme and strategic mission of the department is the automatic acquisition of knowledge from Internet sources, text documents, and social media. Envision a computer who knows everything that Wikipedia contains in a formal knowledge representation and can harness this for understanding natural language, deep text analytics, and logical inference. Such a computer would be able to give precise answers to complex questions and may even be able to pass a high-school exam in specific subjects. The ultimate goal would be to pass the Turing Test: can a computer have a conversation with a human person such that the human cannot tell whether she communicates with a machine or a person?

The department has pursued this ambitious goal since 2005 and has been a trendsetter on the automatic construction of formal knowledge bases. The trend has later led to major results like the Google Knowledge Graph and similar projects at Internet companies and other enterprises. Our knowledge base YAGO is one of the three largest publicly available knowledge graphs. YAGO comprises nearly 10 Million entities like people, places, movies, drugs, and others, as well as 200 Million facts about these entities and relationships among them. YAGO has been used in many projects around the world, including the IBM Watson System when it won the TV quiz show Jeopardy in early 2011 and leveraged YAGO for semantic type checking. We are working on improvements and extensions of YAGO and knowledge-base technology in general, especially on acquiring temporal knowledge and commonsense knowledge about everyday objects, activities, and more.

Digital knowledge is key to a variety of intelligent computer applications: searching for entities and complex relationships in the internet or in enterprises, analyzing news articles and social media, summarizing entire text corpora, exploring and analyzing heterogeneous collections of Big Data. An application where we make intensive use of machine knowledge is our AIDA project on automatic discovery and tracking of entities in text and data.

In terms of algorithmic methodology, the group pursues the combination of logic-based and statistically-based techniques. Methods of this kind are central for knowledge extraction and also key assets for explorative data analysis on both structured data and unstructured text.

Many of the tools and systems developed by the group are publicly available as open source software and are used by other research groups around the world. These include, in particular:

1. Software libraries for the automatic construction and maintenance of the YAGO knowledge base, as well as the YAGO contents itself.

2. The RDF search engine RDF-3X, which can find complex patterns in large collections of RDF triples and graph-structured data very efficiently.

3. SOFIE, PROSPERA, PATTY, ClausIE and other tools for extracting relational facts from Web contents and textual sources.

4. AIDA, a knowledge-based linguistic system for high-precision detection and disambiguation of entities in English or German texts.

5. A variety of tools for explorative analysis of multivariate data, complex networks, matrices and tensors, as well as text corpora.

The department is involved in a number of externally funded projects. These include the EU project *"LAWA: Longitudinal Analytics of Web Archive Data"*, the ERC Synergy Grant *"imPACT: Privacy, Accountability, Compliance and Trust in Tomorrow's Internet"*, and the DFG Excellence Cluster *"Multimodal Computing and Interaction"*. The group has received a Google Focused Research Award. :::

Automation of Logic

PROF. DR. CHRISTOPH WEIDENBACH

RG. 1

The independent research group "Automation of Logic", under the direction of Prof. Dr. Christoph Weidenbach, serves the complete pipeline, from basic research on new logics to industrially used automated reasoning tools.

Logics are formal languages for reasoning with mathematically precise meanings and exact rules. A simple example of a logic are linear equation systems that we know from high school, including the variable elimination method for finding their solution. Logic was developed at the end of the 19th century in order to exactly describe and calculate general mathematical arguments. Our everyday language is not suitable for exact descriptions, because of, for example, ambiguities. Since the invention of computers and information technology, logics have been further developed for computer systems so that these could also be described precisely and properties could be formulated, analyzed, and ultimately proven. By about the mid-1990s, such analytical and proof efforts required massive manual interaction and, therefore, a great deal of time and money. There have since been major advancements in the automation of the methods.

Proving or analyzing computer system properties is typically "hard", that is, even if there are automatic procedures available, the potential number of calculation steps grows at least exponentially with the size of the problem. From the mid-90s to date, however, procedures for a number of practically relevant properties have been developed to work fully automatically in reasonable time. In sum, the application of logic was extended from mathematics to properties of computer systems in the broadest sense. In addition, it has been found that many of the techniques developed are generally successful on hard problems, even if they do not originate from the formal analysis of computer systems and mathematics. The application of analysis and proof techniques has thus once again extended to general hard problems in recent years, e.g., combinatorial optimization problems.

Our research group develops automatic, logic-based procedures. These should be capable of computing solutions for a class of applications within an acceptable amount of time. We are particularly interested in the verification of controls, distributed protocols and systems, safety properties, software, solutions for decision and optimization problems, and in the computation in specific domains like in linear and non-linear arithmetic. We aim to be able to answer increasingly complex questions fully automatically and within an acceptable time period. To this end, the efficiency of the formal analysis methods used today must be raised significantly. This is the goal of our research group, "Automation of Logic". :::

CONTACT

Automation of Logic
Secretary
Jennifer Müller
Phone +49 681 9325-2900
Email jmueller@mpi-inf.mpg.de

Max Planck Center for Visual Computing and Communication

PROJECT MANAGEMENT: PROF. DR. HANS-PETER SEIDEL

In 2003, the "Max Planck Center for Visual Computing and Communication" was established with substantial sponsorship from the BMBF (German Federal Ministry of Education and Research). The Max Planck Center links two globally leading institutes in this field, namely the Max Planck Institute for Informatics in Saarbrücken and Stanford University in the USA.

The cooperation's focus is on fundamental research in the field of Visual Computing and Communication and includes, in particular, the subareas of image acquisition, image analysis, image synthesis, computer vision, visualization, human-computer interaction, and the exchange of image and video data over complex networks.

Strengthening Germany as a science and research location in the field of computer science

A major goal of the program is the development and promotion of young scientists. This is achieved by opening the way for particularly qualified young computer scientists to gain early scientific independence, with simultaneous close integration into an internationally competitive and stimulating scientific environment. In this context, particularly outstanding young postdoctoral graduates are given the opportunity to do independent research in a small work group for up to five years, under the supervision of two mentors from Germany and the United States. After a two-year stay at Stanford, where they have the status of visiting assistant professors (Phase I), the scientists return to Germany and continue their work as junior research group leaders at the Max Planck Institute for Informatics (Phase II). The second phase of the program is generally also open to outstanding postdoctoral graduates from other countries who would like to return here.

Current status

The Max Planck Center has now been in existence for more than ten years and counters the often observed „Brain Drain" to the USA by providing an attractive return perspective. The program contributes to the training and development of highly-qualified young scientists and, thus, strengthens the innovation capacity and international competitiveness of the country. Meanwhile, the Max Planck Center has gained a strong reputation as a real talent factory. Since launching the program in 2003, a total of 29 young scientists have completed the program. Of these, 21 have meanwhile received tenured faculty positions, fourteen thereof in Germany, seven in Europe, and twelve as full professorships.

The Center's success demonstrates that it is indeed possible for Germany to be successful in the global competition for the brightest talent. Some key elements of this successful program are its international focus and the flexible and highly dynamic research program, whose development strongly benefits from the input of the postdoctoral graduates. Further important elements are the early scientific independence of the young scientists with simultaneous tight integration into an internationally leading scientific environment, and last but not least the attractive return perspectives. :::

CONTACT

Prof. Dr. Hans-Peter Seidel
Secretary Sabine Budde
Phone +49 681 9325-4000
Email budde@mpi-inf.mpg.de

Excellence Cluster Multimodal Computing and Interaction

SCIENTIFIC COORDINATOR: PROF. DR. HANS-PETER SEIDEL

The Excellence Cluster "Multimodal Computing and Interaction" was established by the German Research Foundation (DFG) within the framework of the German Excellence Initiative in 2007 and successfully renewed in 2012.

The starting point of the research program in the Cluster was the observation that our living and working conditions have changed dramatically in the past two decades. Twenty years ago, most digital content was textual. Today, digital content additionally comprises the modalities of speech, audio, video, graphics, and more. Given these trends, the challenge is to organize, understand, and search multimodal information in a robust, efficient, and intelligent manner and to create reliable, privacy-friendly systems that efficiently support natural and intuitive multimodal interaction.

The Cluster of Excellence on Multimodal Computing and Interaction addresses these challenges. In this context, the term "multimodal" describes different kinds of information such as text, speech, images, video, and graphics and the way it is perceived and communicated, particularly through vision, hearing, and various forms of body language. Our first goal is to enhance the ability of computer systems to acquire, process and present different modes of data in an efficient, robust, and privacy-friendly way. We aim for systems that do the following: analyze and interpret multimodal information even when it is extensive, distributed, noisy, and incomplete; organize the obtained knowledge for powerful querying; and produce 3D visual output in real time. We refer to this type of computing as multimodal computing. Our second goal is to enable natural multimodal interaction with information systems anytime and anywhere, exploiting the wealth of modalities present in everyday human-to-human interaction. The system must be aware of each user's environment and situation, must react to speech, text, and gesture, and must respond in the appropriate output modality.

The research program builds on existing strengths. The Cluster comprises the Computer Science, Computational Linguistics and Phonetics, and Applied Linguistics departments of Saarland University, as well as the Max Planck Institute for Informatics, the German Research Center for Artificial Intelligence, and the Max Planck Institute for Software Systems. The participating institutions have agreed on a common long-term, highly interdisciplinary research agenda as the basis for this Cluster.

An integral goal of the Cluster is the promotion of young researchers. Saarbrücken has been playing a leading role in this regard for many years and has established a reputation as a talent factory. Therefore, the majority of the approved funding has been committed for the establishment of a number of junior research groups.

This concept has proven extremely successful, and a large number of junior scientists have meanwhile been appointed to professorships in Germany and abroad.

The Max Planck Institute for Informatics strongly contributes to the scientific activities of the Cluster. All directors of the Institute are principal investigators of the Cluster, and Hans-Peter Seidel serves as the scientific coordinator. :::

CONTACT

Prof. Dr. Hans-Peter Seidel
Secretary Sabine Budde
Phone +49 681 9325-4000
Email budde@mpi-inf.mpg.de

RESEARCH AREAS

Research Areas

UNDERSTANDING IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH & DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION & VISUALIZATION

OPTIMIZATION

UNDERSTANDING IMAGES & VIDEOS

Understanding images and videos is one of the fundamental problems of image processing. The scientific challenges in understanding images and videos range from the modeling of people and objects in camera systems to the reconstruction and description of 3D scenes. This area has many applications including the animation of people and the visualization of 3D scenes, the indexing of image and video material, and 3D-capturing of the environment. This research area is thus situated at the intersection between computer vision and computer graphics, resulting in numerous opportunities for cooperation within the institute.

Various research groups within the Max Planck Institute for Informatics work on aspects of image and video understanding. In the field of modeling people, one group is working on the reconstruction of animated models from multi-video data, for example. The primary goal here is the ability to model and visualize people in detail and as accurately as possible. Furthermore, we are researching methods for people detection and pose estimation, relying on monocular cameras only, which are able to detect and track many people simultaneously in complex scenes. Even if these studies pursue fundamentally different targets and use different camera configurations, they benefit from mutual similar models and algorithms.

An emerging question in recent years is how to generate natural language for images and videos. Our institute has been at the forefront of this development; for example, our researchers have proposed to formulate this problem as a language translation problem, where the input language is given by video descriptors, and the output language is natural language.

Advanced video processing is a topic at the intersection of computer vision and computer graphics. We have developed techniques well beyond the state of the art that allow us to manipulate entire scene objects, such as particular actors.

One of the fundamental problems of understanding images is the recognition of objects. With the growing omnipresence of digital image material, automatic visual object-class recognition and indexing techniques are becoming increasingly important. Even if today‘s approaches can achieve remarkable results, two significant questions remain: How can object models be learned or constructed in a scalable fashion? How can we learn models that respect the 3D nature of objects? To address these problems, research work at the institute investigates how such object models can be constructed with as little manual input as possible to enable a broad applicability. :::

CONTRIBUTIONS

Articulated Human Pose Estimation

Human body pose contains a wealth of information about a person's intention, attitude and internal state. The focus of our research is to estimate body pose in realistic conditions such as images and videos found on YouTube or captured with a mobile phone. We envision that developed approaches will become building blocks in such applications as activity recognition, marker-less motion capture and augmented reality.

We build on the recent advances in hierarchical image representations with convolutional neural networks (CNN) and explore two novel research directions: joint estimation of poses of multiple people and 3D human pose estimation from only a few mobile cameras.

Detection and pose estimation in multi-person scenes

We propose an approach that jointly solves the tasks of articulated person detection and pose estimation. Our approach infers the number of persons in a scene, identifies occluded body parts, and disambiguates the body parts of different people in close proximity of each other [see figure]. Our formulation is based on partitioning and labeling a set of body part hypotheses generated with a CNN-based body part detector. The partitioning is accomplished by solving an integer-linear program that resembles correlation clustering approaches previously proposed for image and video segmentation. One of the advantages of our formulation is that it implicitly performs non-maximum suppression, removing spurious body part detections in the background and merging multiple correct detections corresponding to the same person. We evaluate our approach on standard benchmarks showing its advantages over previously proposed strategies that operated by first detecting the people and then independently estimating their body poses.

Multi-view 3D human pose estimation

We propose a novel method for the accurate marker-less capture of articulated skeleton motion of several subjects in general scenes, indoors and outdoors, even from input filmed with as few as two cameras. Our approach unites a discriminative image-based joint detection method with a model-based generative motion tracking algorithm through a combined pose optimization energy. The discriminative part-based pose detection method, implemented using convolutional neural networks, estimates unary potentials for each joint of a kinematic skeleton model. These unary potentials are used to probabilistically extract pose constraints for tracking by using weighted sampling from a pose posterior guided by the model. In the final energy, these constraints are combined with an appearance-based model-to-image similarity term. Poses can be computed very efficiently using iterative local optimization, as CNN detection is fast, and our formulation yields a combined pose estimation energy with analytic derivatives. In combination, this enables us to track fully articulated joint angles at state-of-the-art accuracy and temporal stability with only very few cameras. :::

CONTACT

Mykhaylo Andriluka
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-2119
Email andriluka@mpi-inf.mpg.de

Leonid Pishchulin
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-1208
Email leonid@mpi-inf.mpg.de

Ten Years of Pedestrian Detection, What Have we Learned?

Object detection is one of the core problems in computer vision. Pedestrian detection is a canonical case of object detection. Due to its multiple applications in car safety, surveillance, and robotics, pedestrian detection is an active area of research, with 1000+ papers published in the last decade. Pedestrian detection also benefits from well-established benchmarks for evaluating and comparing methods.

In the last decade, numerous techniques have been developed for pedestrian detection. Using different image features, cues, and classifier types, detection performance has been steadily improving. Building larger systems that integrate more ingredients has become a reliable recipe to advance quality *(“more is more”)*.

Our research has focused on understanding which are the core ingredients for high quality pedestrian detection. By improving these core ingredients, we have been able to consistently obtain top performance on established standard benchmarks *(“less is more”)*.

Our top performing detectors are grounded on decade-old ideas (oriented gradients features, filter banks, and boosted decision trees) that we have carefully instantiated and adapted to our newest insights. Contrary to previous work, we show that top performance on pedestrian detection can be reached by using a single rigid template (no components, no deformable parts) applied as a *sliding window* (no geometry prior, scene prior, or bottom-up cues). Because of its simplicity, our approach is suitable for high-speed implementations. Our current results on pedestrian detection outperform both deep convolutional neural networks and methods that use more complex features (such as local binary patterns and covariance) by a large margin.

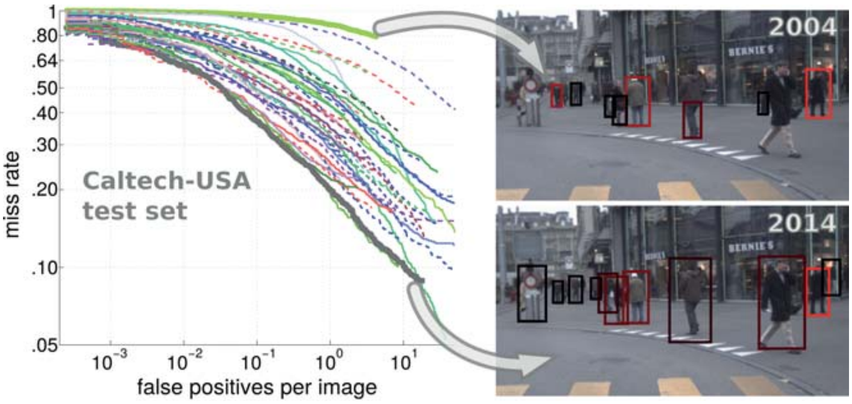

In the last ten years, pedestrian detection has shown great progress.

In the last years, our work has repeatedly advanced the top performance reached on standard benchmarks. In the last two years alone, we have improved the state-of-the-art with a 30x reduction in errors (false positive per image at 80% recall on Caltech pedestrian dataset). Our techniques developed for pedestrian detection have also shown top performance for face and traffic sign detection. For face detection, our results match the best performance reported on all four major benchmark datasets.

Although significant progress has been achieved in the last decade, we are still far from reaching human performance on this task, or from reaching the desired super-human quality for automatic operations. We expect that a shift towards a more global scene understanding will help reduce the mistakes of current methods in the future. :::

CONTACT

Rodrigo Benenson
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-2105
Email benenson@mpi-inf.mpg.de

Jan Hosang
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-2123
Email jhosang@mpi-inf.mpg.de

Markerless Reconstruction of Static and Dynamic Scenes

Reconstruction of detailed animation models from multi-video data

The development of new algorithms to reconstruct geometry and appearance models of dynamic scenes in the real world is an important goal of graphics and vision research. Captured dynamic models of humans and more general scenes are very important for computer animation and visual effects. Reconstruction of the dynamic environment with sensors is also highly important for autonomous robots and vehicles. It is also the precondition for 3D video and virtual replay production in TV broadcasting. Reconstruction of human motion, in particular, is of increasing importance in biomechanics, sports sciences, and human-computer interaction. The development of new sensors and new see-through displays also brings us closer to truly believable augmented reality. The precondition for this is also having algorithms to measure models of the real world in real-time.

The goal of our research is thus the development of new *performance capture algorithms*. These algorithms reconstruct detailed models of geometry, motion, and appearance of dynamic scenes from multi-view video data. The reconstruction is to be performed without optical interference with the scene, for instance, fiducial markers, as are commonly used in motion capture. In the past, we have pioneered the development of performance capture algorithms that are able to reconstruct detailed models of persons in general apparel or of arbitrary moving objects from multi-view video. Unfortunately, the application of these methods is constrained to indoor studios with many cameras, controlled lighting, and controlled backgrounds. In the reporting period. we therefore began to redefine the algorithmic foundations of dynamic scene reconstruction to enable, in the long run, performance capture of arbitrarily complex outdoor scenes under uncontrolled lighting and with only a few cameras.

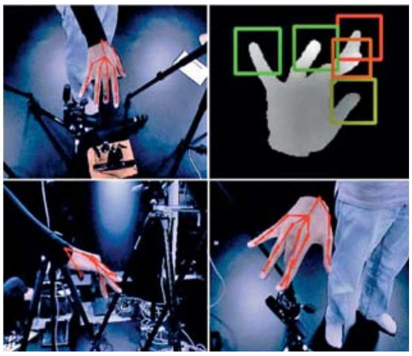

Figure 2: Real-time marker-less hand tracking

In the reporting period, we developed some of the first methods for the detailed marker-less skeletal motion capture of humans and animals in outdoor scenes with few cameras (2–3). Another milestone was the development of new inverse rendering algorithms that allow us to measure models of appearance, illumination, and detailed geometry from video footage recorded under uncontrolled conditions. This new methodology is the basis for other important results: the first approach to capture relightable 3D videos in uncontrolled scenes, the first methods to capture highly detailed dynamic face models from stereo or monocular video, and the first approach for detailed full body performance capture with a single stereo camera [see figure 1]. An important new strand of research is also the development of algorithms for marker-less hand motion capture from a few cameras, or even from a single camera view [see figure 2].

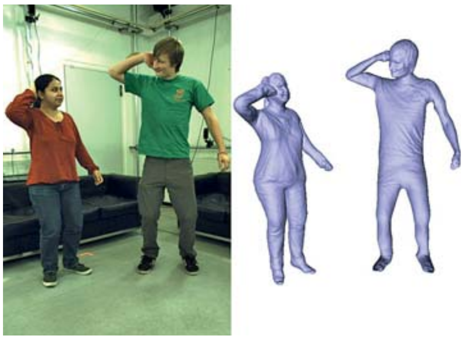

Figure 1: Performance capture with a stereo camera

Reconstruction of static and dynamic scenes with depth sensors

New types of depth cameras measure 2.5 D scene geometry in real-time. These new sensors are cheap and can be an important tool to bring reconstruction methods for static and dynamic scenes to a broader user community. Unfortunately, depth camera data is very noisy, has low resolution, and exhibits a systematic measurement bias. We have, therefore, developed methods to calibrate depth sensors, eliminate noise, and algorithmically increase the camera resolution. Furthermore, we are working on new algorithms to capture geometry models of even large scenes with such sensors.

Additionally, we research new methods to capture human skeletal motion and body shape in real-time from a single depth camera. In this context, we have also developed some of the first approaches for real-time capture of general deformable shapes with a single camera. :::

CONTACT

Christian Theobalt
DEPT. 4 Computer Graphics
Phone +49 681 9325-4028
Email theobalt@mpi-inf.mpg.de

3D Object Detection

Detecting and localizing objects in images and videos is a key component of many applications in robotics, autonomous driving, image search, and surveillance. Traditionally, the object detection task has been defined as the localization of instances of a certain object class (such as cars) in an image. That is, the input to an object class detector is an image, and the output is a two-dimensional bounding box that highlights the position of the detected object.

Recent work is making it increasingly clear that high-level applications like scene understanding and object tracking would benefit from a richer, three-dimensional object representation. Such a representation comprises, in addition to the traditional 2D bounding box, an estimate of the viewpoint under which an object is imaged as well as the relative 3D positions of object parts.

It allows us to derive scene-level constraints between multiple objects in a scene (objects cannot overlap in 3D) and between multiple views of the same object (the viewpoints of different detections must be consistent with camera movement).

3D deformable part-based models

The goal of this project has been to build a 3D object class detector that, along with a traditional 2D bounding box [see Figure 1], provides an estimate of object viewpoints as well as relative 3D part positions [see Figure 2].

Since 3D object modeling requires supervision in the form of 3D data, we leverage 3D Computer Aided Design (CAD) models as a 3D geometry proxy for the object class of interest. In addition, we use 2D real-world imagery to obtain realistic appearance models. One of the main challenges consists in bridging the gap between artificial 3D geometry and 2D real-world appearance, which we achieve through a shape-based abstraction of object appearance based on non-photorealistic rendering. The final 3D object class representation is then given by a probabilistic, part-based model, trained from both 3D and 2D data.

In our experiments, we have been able to demonstrate excellent performance in both the traditional 2D bounding box localization task and tasks aimed at scene-level reasoning, such as viewpoint estimation and ultra-wide baseline matching.

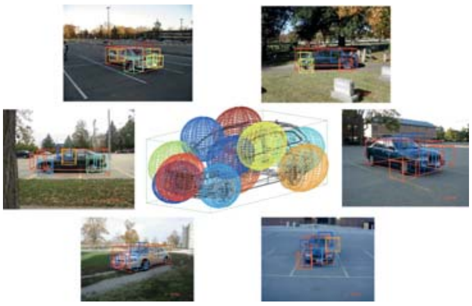

Figure 1

Occlusion patterns for object class detection

In this project, we have focused on improving object class detection in the face of partial occlusion, as it occurs when the view onto an object is blocked by other objects. This situation can be frequently observed, for example, in street scenes, where cars block the view onto other cars. Our approach is based on the intuition that the resulting occlusions often follow similar patterns that can be exploited, such as cars parked in front of other cars at the side of the road.

We have extended our 3D object class representation to include occlusion patterns in the form of object pairs. In our experiments, we have shown that our model outperforms prior work in both 2D object class detection and viewpoint estimation. :::

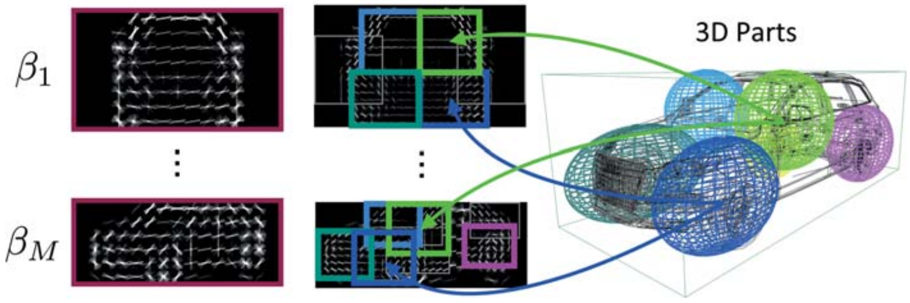

Figure 2

CONTACT

Bojan Pepik
DEPT. 2 **Computer Vision and Multimodal Computing**
Phone +49 681 9325-2208
Email bojan@mpi-inf.mpg.de

Michael Stark
DEPT. 2 **Computer Vision and Multimodal Computing**
Email stark@mpi-inf.mpg.de

Advanced Video Processing

Image processing methods are important tools during the post-processing of photographs. Many standard tools for image processing provide algorithms to fulfil typical post-processing tasks, such as standard filters for noise removal or contrast enhancement.

Post-processing of videos is a much more challenging task. Videos are not merely a temporal sequence of individual images. Modifications to videos during post-processing, therefore, need not only be spatially consistent, that is, consistent within a single image, they also need to be consistent over time, that is, over many subsequent frames. Many commercially available video editing tools, however, are based on exactly this assumption that videos can be treated like sequences of individual images. Post-Processing operations with such tools are consequently limited to the application of simple filters over a temporal window.

Many post-processing tasks in professional video and movie productions are much more complex: Entire scene elements, such as particular actors or supporting crew, may need to be removed or manipulated in a shot, or their positions in a shot may need to be changed. No existing video filtering approach would even come close to solving such a task automatically. The consequence is that such editing tasks are usually undertaken through time-consuming manual editing of individual pixels, which can easily take several weeks, even for short videos.

Further, today's personal and online video collections are extremely large. Browsing such collections is a challenge, as current approaches are often based on unreliable user-based text annotation. Ideally, one would want more semantically-based ways of exploration that exploit space-time and content relations between videos, such as the fact that videos were taken in the same environment, show a similar scene, were taken at a similar time, or deal with a similar topic. The precondition for this would be the ability to automatically extract such relations, which is a highly challenging and unresolved problem. We investigate the algorithmic foundations of both of the above problem areas.

Videos in panoramic contexts

Context-based exploration of video databases

Context-based relations between videos enable particularly interesting new exploration modes for video databases. Context relations capture, for instance, whether videos have been recorded in the same area, the same town, or even partly show the exact same location. We therefore develop new methods to extract such context relations automatically from video collections. One of our approaches computes a graph, a so-called *Videoscape*, whose nodes are portals, i.e., specific locations that videos can show, and whose edges are videos. A video is linked to a portal node if it ever shows the location represented by that portal. The automatic computation of this graph is a complex task, and we have developed new machine vision and machine learning algorithms to serve this purpose. The Videoscape can be explored interactively, for instance, one can take a tour through a city by watching videos recorded in that city. Once a frame showing a portal comes up, one can transition into another video based on a 3D reconstruction of the portal location from the video frames. We have also researched how one can overlay videos of a location with a panoramic context of that location [see figure]. This enables entirely new ways of visualizing the space-time relations between videos filmed at that location.

New video editing approaches

We developed new algorithms to compute dense correspondences between time series of images, even images exhibiting strong noise and images taken under starkly varying camera settings. These methods are the foundation for new advanced video editing operations, such as automatic dynamic background inpainting. They also enable the automatic reconstruction of high dynamic range images from a stack of images taken with different exposure times, even if there was strong motion in the scene between images. New machine learning algorithms we have developed also enable us to remove compression artifacts from images and videos, and to algorithmically increase the frame resolution in both cases. :::

CONTACT

Christian Theobalt
DEPT. 4 Computer Graphics
Phone +49 681 9325-4028
Email theobalt@mpi-inf.mpg.de

Describing Videos with Natural Language

An important aspect for automated systems is to communicate to humans what they recognize or “see”. However, current computer vision approaches typically focus on generating isolated labels for a video (e.g., “slicing”, “cucumber”, “plate”) that are not well suited for communication with humans. We thus study the problem of generating natural language descriptions of videos. A corresponding sentence for the example above could be “Someone sliced the cucumber on the plate”. Describing videos with natural language is important, e.g., for the automatic captioning of web videos, human-robot interaction, or assisting visually impaired people. In our work, we study two scenarios. First, we generate descriptions of cooking videos. Second, we study the problem of generating audio descriptions for movies to enable blind people to follow a movie without seeing it.

Generating descriptions for cooking videos

In order to study the problem of automatic video description, we propose to learn how to “translate” video snippet to a natural language sentence using techniques from statistical machine translation between two languages, e.g., French and English. For training our translation approach, we need pairs of videos and sentences depicting various cooking activities, which we have collected on a large scale for videos, e.g., open tin, stir pasta.

In contrast to related work, we do not only describe a short video snippet with a single sentence, but also a long video with multiple sentences. To ensure that we generate a consistent description, we propose to model the topic of the video shared by all short video snippets within a long video. In our kitchen scenario, the topic is a dish to be cooked, e.g. preparing pasta.

We also explore the novel task of describing videos at multiple levels of detail. All of the prior work has focused on describing videos at a fixed level of abstraction, whereas our system is able to produce detailed, short, and single sentence descriptions of cooking videos [figure 1].

Detailed: A man took a cutting board and knife from the drawer. He took out an orange from the refrigerator. Then, he took a knife from the drawer. He juiced one half of the orange. Next, he opened the refrigerator. He cut the orange with the knife. The man threw away the skin. He got a glass from the cabinet. Then, he poured the juice into the glass. Finally, he placed the orange in the sink.
Short: A man juiced the orange. Next, he cut the orange in half. Finally, he poured the juice into a glass.
One sentence: A man juiced the orange.

Figure 1: An example output of our system, which automatically generated a detailed, short and single-sentence description of a video.

Movie description

Existing video description datasets focus on short video snippets, are limited in size, or restricted to the cooking scenario. In order to overcome these limitations, we propose a new dataset of movies and associated textual descriptions. We make use of two sources of text data, namely movie scripts (screenplays) and audio descriptions available for many DVDs and Blu-rays. Audio descriptions provide linguistic descriptions of movies and allow visually impaired people to follow along with friends and family. The collected dataset additionally opens new possibilities to understand stories and plots across multiple sentences in an open domain on a large scale [figure 2]. :::

AD: Abby gets in the basket.	Mike leans over and sees how high they are.	Abby clasps her hands around his face and kisses him passionately.
Script: After a moment a frazzled Abby pops up in his place.	Mike looks down to see – they are now fifteen feet above the ground.	For the first time in her life, she stops thinking and grabs Mike and kisses the hell out of him.

Figure 2: Example textual descriptions from our dataset, aligned to three movie snippets: first, coming from audio description (AD), and second, from movie script.

CONTACT

Anna Rohrbach
DEPT. 2 **Computer Vision and Multimodal Computing**
Phone +49 681 9325-2111
Email arohrbach@mpi-inf.mpg.de

Marcus Rohrbach
DEPT. 2 **Computer Vision and Multimodal Computing**
Phone +49 681 9325-2111
Email rohrbach@mpi-inf.mpg.de

Scalability of Object Class Recognition

Due to the omnipresence of digital imagery today, automatic techniques for the visual recognition of objects are becoming more and more important. The focus here lies on the recognition of entire classes of objects, such as cars, rather than specific instances like one particular red sports car. While current state-of-the-art systems achieve remarkable recognition performance for individual classes, the simultaneous recognition of multiple classes remains a major challenge. At training time, building reliable object class models requires a sufficiently high number of representative training examples, often in the form of manually annotated images. At test time, prediction has to be efficient so that the correct label can be inferred among large sets of possible classes. Our research aims at reducing the requirements at training time in order to achieve more scalable learning procedures. Increased scalability at test time is achieved with more efficient inference schemes that share computation among classes as well as optimize for computation-time aware performance metrics.

Identifying re-usable components automatically

In order to apply the re-use of model components in practice, re-usable components have to be identified automatically. In this project, model components correspond either to attributes, i.e., distinct visual properties like color or texture, or to entire object class models. We propose to relate object classes and attributes by automatically mining linguistic knowledge bases. Examples of linguistic knowledge bases include, e.g., WordNet, Wikipedia, Flickr, and Yahoo web search [Figure 1]. The mined relations can then be used to build object class models, even for classes where there are no training examples available, relying entirely on re-used components from other models. The recognition performance of the fully automatic system is demonstrated to be en par with providing the object class-attribute relations of human subjects.

Sharing of computation between classes

Today's state-of-the-art object detectors predominantly achieve multi-class detection through a combination of multiple single-class detectors. This is undesirable because the computation time grows linearly with the number of classes. A more detailed analysis of the most common architectures reveals that such independent detection schemes perform a significant amount of redundant computation. Intuitively, car and bicycle detectors are both look for wheels, but the detection scheme does not exploit this. Therefore, we have presented shared basis representations that capture the commonalities between such parts and classes. As a consequence, we only have to compute the activations of such a basis once and can share the results among the class detectors. From this shared representation, we are able to recover the activations of the individual class detectors with high accuracy.
We have presented a detection method that leverages such decompositions, together with a hardware accelerated implementation, and have shown up to 35x speedups over individual detectors.

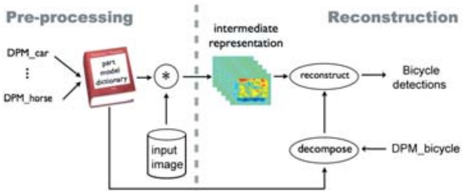

Figure 2: Efficient recognition of object classes by exploiting similarities between classes

Effective use of computation time

Real-time systems as well as large-scale recognition and retrieval systems are often limited by a computational budget. Today's approaches rarely take such constraints into consideration and do not adapt their execution strategies with regard to the expected performance gains and the computational cost of a classifier or detector. In our work, we propose a method that has access to multiple image classifiers and detectors and learns to deploy them with an optimal strategy so that an application dependent reward is optimized under computational budget constraints. Intermediate results are used to determine which classifier or detector will be executed next. In contrast to prior work, our *"reinforcement learning"*-based approach also picks actions that do not have an immediate pay-off but will generate future rewards. For instance, our image classifier does not yield any detections, but the result helps the method to make a more informed decision for the next time steps and thereby increases performance over the whole episode. We have shown that this approach yields a strategy that outperforms the best fixed execution orders under a time-sensitive performance metric. :::

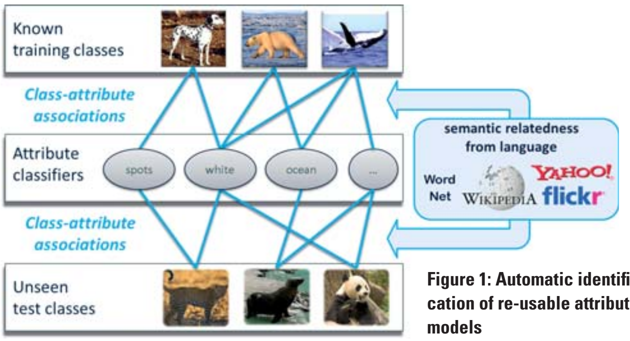

Figure 1: Automatic identification of re-usable attribute models

CONTACT

Mario Fritz
DEPT. 2 **Computer Vision and Multimodal Computing**
Phone +49 681 9325-1204
Email mfritz@mpi-inf.mpg.de

Marcus Rohrbach
DEPT. 2 **Computer Vision and Multimodal Computing**
Phone +49 681 9325-1206
Email rohrbach@mpi-inf.mpg.de
Internet http://www.d2.mpi-inf.mpg.de/nlp4vision

Topology and Data Understanding

We are living in the era of information: Huge amounts of data are generated every minute. For instance around 350 million photos are uploaded to Facebook every single day.

Data analysis is the task of extracting meaningful information content and drawing conclusions from a data collection. One example are recommendation systems for user navigation on a website like Youtube or Netflix based on the choices made. In general, data analysis becomes more challenging through the inevitable presence of noise in almost any form of real data.

In many cases, qualitative, high-level summaries are required for data analysis. For instance, a first step for analyzing a collection of images is a clustering into few categories, like images of people, buildings, landscapes, etc. The growing interest of data analysis asks for novel ways of extracting such high-level information.

Topological data analysis

A perhaps surprising connection has been established between data analysis and the *topology* of geometric shapes. Topology is a mathematical language for classifying shapes according to how they are connected. To illustrate the idea, a bagel and a Kaiser roll are topologically different, because the former contains a hole that is missing in the roll. Moreover, a pretzel is different from the former, too, again because it has more than one hole. On the other hand, a coffee mug and a bagel are, topologically speaking, the same because both have one hole, and we can transform one into the other without ever changing the connectivity of the shape.

The relation to data analysis is that many data sets can be easily interpreted as geometric data (in some high-dimensional space), and topology provides a high-level summary of that data, ignoring details like the distance between points. However, the classical notions of topology are insufficient because they are developed for the idealized situation of "clean" shapes and are therefore sensitive to noise. This problem has been overcome with the development of *persistent topology*. The rough idea is to not just count the number of topological features (like the number of holes in the example), but also to provide an "importance value" to each feature. This allows a more fine-grained analysis of the topological features, in particular a distinction between noise and relevant features of the shape.

Persistent topology has been applied to various questions in data analysis. As an example, the space of "natural images" has been shown to fit the geometric shape of a *Klein bottle*, which is a twisted version of a (hollow) bagel.

Computational challenges

In light of the increasing size of data sets, efficient ways of computing and analyzing the persistence information of data are required. Researchers at the Max Planck Institute in Saarbrücken are devoted to this goal. The computational problems connect with various classical areas in algorithmics, for instance, approximation algorithms to create a shape from data, linear algebra to compute the persistence of a shape, and combinatorial optimization to compare the topologies of different shapes. :::

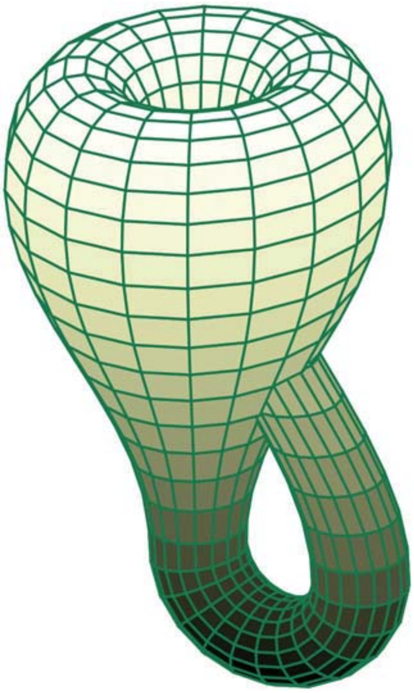

A Klein bottle embedded in three dimensions

CONTACT

Michael Kerber
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1005
Email mkerber@mpi-inf.mpg.de

Learning of Visual Representations

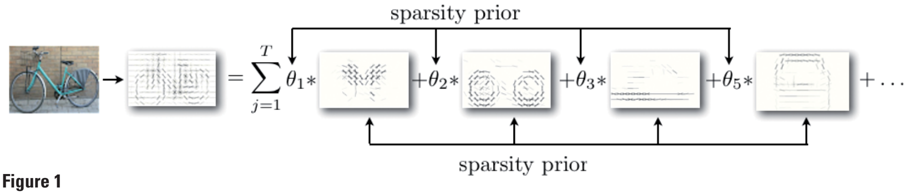

Figure 1

With the advance of sensor technology, machines can get a detailed recording of an environment. But, there is a big gap between those raw sensor readings a computer can get, e.g., from a camera to the semantic understanding a human may get by looking at the same scene. We need to bridge this gap in order to unfold the full potential of applications like autonomous robotics, content-based image search, or visual assistance for the blind.

Computer vision has come a long way by proposing representations that, step by step, try to bridge this gap. However, despite the exciting progress, we see a divergence of methods rather than convergence, and we are far away from matching the adaptivity, efficiency, and effectiveness of the human visual system. Therefore, we seek principled ways to derive visual representations in learning-based approaches in order to eventually match or even surpass human perception.

Latent additive feature models

Inspired by models from document processing, we have applied methods that automatically extract topics from text to the visual domain in order to decompose objects and image patches into parts and sub-structures. These models facilitate the learning of visual representations in a completely unsupervised fashion. This application is facilitated by an analogy between words in the text domain and local gradient structures in the visual domain. Figure 1 illustrates how the gradient distribution of the bike on the left is decomposed into sub-structures on the right. We use methods that allow us to place expectations (priors) on the unknown variables in order to regularize and solve this ill-conditioned learning problem. Our experiments show that this type of representation is very well-suited for the task of object class recognition and detection and can even deal better with effects like transparency.

Most recently, we have defined a recursive scheme that allows us to form hierarchical representations built on the same principle. Our studies have shown the importance of holistic inference for learning these representations, in contrast to the more common feedforward learning approach. We are currently extending these representations to the temporal domain.

Representations across data domains and sensors

The rapid development of the web and mobile devices provides us with rich sources of information. However, when tapping into these, we face a large heterogeneity in sensors, quality, and quantity of observations. It is upon good representations to reunite those data sources on a common ground.

We have studied approaches for domain adaptation in order to bridge the gap between images from the web, high quality and low quality sensors [see figure 2]. Our investigations show that some of the arising problems can be solved by metric learning approaches so that data from the web can be more effectively used for recognition in the real-world.

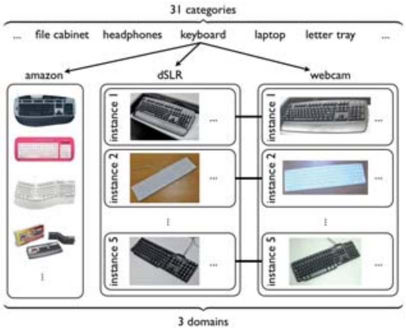

Figure 2

Better sensors are continuously added to consumer electronics. 3D sensors are probably the most prominent example right now (e.g., in mobile phones, gaming devices). Our investigation aims at understanding how 2D and 3D representations of objects can be combined best in order to improve object recognition. We have successfully employed object size constraints and have shown how they can improve classification and detection.

Representations across rendered and real data

There is an increasing interest in using rendered data for training visual classifiers in order to reduce the effort of data collection and therefore providing more training data. Despite the strong progress in rendering techniques, such synthesized images are still visually – as well as from a statistical point of view – different from real data. We have investigated this issues on the task of material recognition. We tap into publicly as well as commercially available material shaders in order to render new examples for our training procedure. We have investigated different methods for domain adaptation and data alignment in order to support material recognition from real and virtual data which result in substantial improvements in recognition performance of up to 8%. :::

CONTACT

Mario Fritz
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-1204
Email mfritz@mpi-inf.mpg.de

BIOINFORMATICS

Bioinformatics is a key discipline for progress in biological sciences, such as biotechnology, pharmacy and medicine. With computational methods Bioinformatics deepens and accelerates planning highly complex biological experiments as well as interpreting the resulting very large amounts of data.

For about 25 years, bioinformatics has significantly contributed to progress in the life sciences. Therefore, the field is a key ingredient of the revolution of biology. Bioinformatics methods and software support researchers in planning experiments, store data which come from all areas of biology, and assess this data with computer support. With the help of Bioinformatics, scientists can elucidate the molecular processes in cells – the basic units of living organisms – which form a complex system for processing matter, energy, and information. The genome harbors the construction plan for the cell and the working plan of its metabolic and regulatory processes in a complex, encoded form. In order to facilitate these cellular processes, parts of the genome must be "read off". This includes the genes, which amount to the blueprints of proteins, the molecular machinery of the cell. In turn, reading off the genes is controlled via complex molecular networks. For the synthesis of proteins, and also for their removal, there are specific molecular complexes which themselves are again subject to exquisite molecular control. The cells convert energy, they communicate with cells in their neighborhood, they take on different structures and forms, and they move. They react to changes in their environment, for example, pertaining to light, temperature and pH, and they mount defenses against invaders. Dysregulation of such processes constitutes the molecular basis of disease. Drug therapy aims to restore a tolerable molecular balance.

For more than the past two decades, classical biological research, which used to be concentrated on cellular sub-systems with limited complexity, has been complemented by high-throughput experiments. The respective screening methods afford cell-wide capture of molecular data, e.g., by measuring the frequency of all read off genes (transcription), by capturing the population of the proteins used by the cell (proteome) and of their interactions (interactome). Generating new insight into the biology of the cell as well as the basis of diseases and providing new approaches to therapy involves highly complex information processing. Bioinformatics addresses this challenge. The Max Planck Institute for Informatics performs research on many of the challenges mentioned above.

Bioinformatics has the hybrid character of a basic science which poses and follows clear application perspectives. This unique quality is highlighted by a significant number of spin-offs of bioinformatics research groups. For example, Professor Lengauer co-founded the company BioSolveIT GmbH, which develops and distributes software for drug design. Pharmaceutical companies all over the world are among the users of this software.

The Bioinformatics Center Saar, whose spokesman Professor Lengauer is, received top marks for its research in the final evaluation in 2007 of the five bioinformatics centers funded by the German Research Foundation (DFG) in the past decade. :::

CONTRIBUTIONS

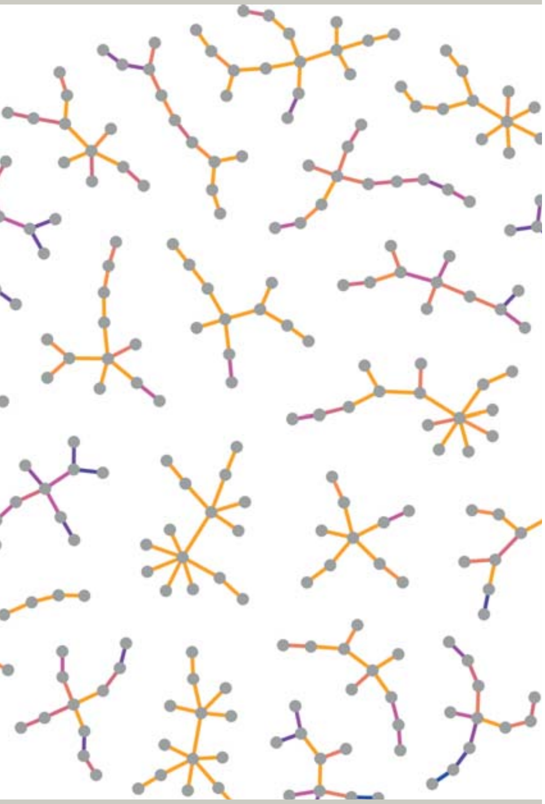

UNDERSTANDING IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH & DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION & VISUALIZATION

OPTIMIZATION

The Relation Between the Molecular Conformational Changes and the Free Energy Change

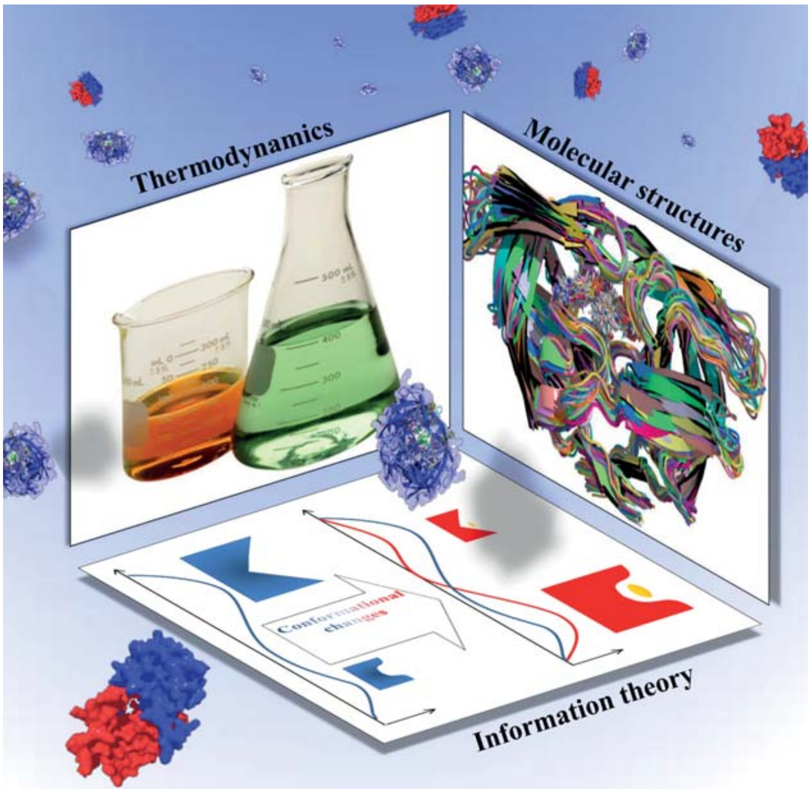

Biomolecules such as proteins are the molecular machines which work in concert to maintain living organisms. Our main interest is to develop theories to understand the relationships between the structures and the motions of the biomolecules and their roles in important biochemical interactions, such as interactions between proteins and drugs.

The atomistic three-dimensional structures of biomolecules can be obtained, e.g., by X-ray crystallography and can be considered as a snapshot of the biomolecular motion. Growing computer power affords the ability to simulate the motion of the atoms inside biomolecules and to understand the mechanistic basis of their function. Experimental thermodynamic measurements of quantities such as enthalpy, entropy, and free energy are used to quantify the biochemical processes, such as drugs binding to proteins, as well as changes of binding due to changes in the amino acids of the protein (mutations). The free energy is a result of the compensation between the molecular interactions, which are summarized by the enthalpy, and the molecular disorder, which is related to the entropy. Understanding the mechanisms by which the motions of the proteins bring about the changes in the free energy is crucial for many biochemical topics, e.g., drug design.

A striking phenomenon in biomolecular processes is that when a molecular system undergoes a change, the changes of enthalpy and entropy largely compensate each other. That is, enthalpy and entropy undergo comparatively large changes but the resulting change in free energy is comparatively small. We have provided a general theoretical explanation of this phenomenon. The theory exploits information theory to shed light on the physical basis of the molecular process. The simplicity and the generality of the concepts in information theory provide direct physical interpretations and establish a new theoretical framework to understand changes in the biochemical systems. The applications of information theory in molecular simulations are a promising approach to improving computational methods to estimate free energy changes. These methods are of high importance for predicting resistance towards certain drugs upon mutations in the amino acid sequences of proteins. :::

CONTACT
Mazen Ahmad
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 9325-3105
Email mahmad@mpi-inf.mpg.de

Attacking HIV From New Angles

According to the World Health Organization (WHO) 35 million people world-wide were living with an HIV infection in 2013. About 1.5 million died due to AIDS-related complications, including 190,000 children. The number of new HIV infections in 2013 was about 2.1 million, according to WHO. Preventive measures are an important means of combating these horrendous numbers. There have been some significant breakthroughs over the past years in HIV vaccine research. For example, Walker et al. were able to extract potent broadly neutralizing antibodies (bNAbs) that can neutralize many different HIV strains from patients with strong immune responses. Since there is no approved treatment to cure an HIV infection to date, a large part of HIV research is concerned with alleviating symptoms and extending the life-span of infected patients with the help of antiretroviral drugs.

Prediction of HIV resistance against passive antibody treatment

Treatment with a combination of bNAbs has recently proven effective against an HIV infection in humanized mice and non-human primates. A phase I clinical trial in humans for one of those bNAbs has also shown positive results. Therefore, this new kind of treatment might be available for HIV-infected patients in the near future.

The standard procedure for choosing a therapy for an HIV-infected patient in Germany is to sequence the important parts of the HIV genome and then evaluate which drugs are active (meaning that the patient does not harbor significant resistant viral populations against these drugs). If treatment with bNAbs by passive antibody transfer becomes available, there will be a need to make an informed decision analogous to this setting. Therefore, we have built prediction models that are able to classify with high accuracy whether or not a certain HIV variant is resistant against a certain bNAb. We provided additional ways to visualize the results, making them more interpretable and better understandable than standard approaches (see figure for one part of the visualizations). Even though treatment of HIV-infected patients may be a few years from clinical routine, these prediction models can already be used to assist future clinical bNAb treatment trials.

Analysis of different antibodies against HIV

An important step towards developing a vaccine against HIV is the analysis of the properties of antibodies that can neutralize the virus. To evaluate the potency of potential candidates, there exist laboratory tests that elicit how well a certain antibody can neutralize various different HIV strains. For these panels, several different HIV strains are used that are representative of the viral variation around the globe. Our analyses showed that a certain group of viral strains is less important, clinically, and that certain antibodies are significantly less effective in neutralizing these strains than other antibodies. Because current antibody panels treat all viral strains alike, they give the less important strains more attention than they should. We advocate targeting the antibodies to the important strains only, which gives room for making them more effective on these strains. Our results can be used to improve antibody tests and are consequently an important step on the way towards developing a universal vaccine against HIV.

Motif Logo for the classification of a test sample with unknown resistance capacity regarding a certain bNAb. The logo reveals residues of the test sequence that have the strongest contribution to the classification result. Here, the top 5% discriminant residues are considered. Blue upper case letters represent residues occurring in the test sequence that are strongly associated with susceptibility, whereas orange lower case letters show residues strongly associated with resistance.

Furthermore, we have been able to show that mono-treatment with certain types of bNAbs by passive antibody transfer may accelerate a change of the viral population to a type of HIV that is correlated with faster disease progression. Therefore, when deciding on treatment options with bNAbs, these effects have to be taken into account by either supplementing those antibodies with antibodies that are good against those later HIV variants or additional antiretroviral treatment. :::

CONTACT

Anna Feldmann
DEPT. 3 **Computational Biology and Applied Algorithmics**
Phone +49 681 9325-3027
Email feldmann@mpi-inf.mpg.de

Nico Pfeifer
DEPT. 3 **Computational Biology and Applied Algorithmics**
Phone +49 681 9325-3019
Email nico.pfeifer@mpi-inf.mpg.de

Structural Insights into the Emergence of Resistance in the HIV-1 Protease

The human immunodeficiency virus (HIV) causes AIDS. There is no cure or vaccine against this disease, but, in many cases, viral replication can be repressed, symptoms eased, and the disease progress slowed down by the antiretroviral drugs. One of the prominent targets for antiretroviral therapy is the HIV protease, a viral protein that is necessary during virus maturation, a process that makes viral particles fully infectious. A high mutation rate of HIV results in the emergence of viral particles with new traits that make them resistant towards a particular therapy. This in turn can lead to the necessity to switch the drug therapy, as the mutant protein becomes less likely to be blocked by the binding of the inhibiting drug. While most of major HIV protease resistance-associated mutation sites are on the binding interface between the protein and the drug molecule (in this case called the ligand), some of them are also found in more distant parts of the protein. Understanding how such mutations influence drug binding is crucial to addressing the viral resistance.

Molecular dynamics and thermodynamics

Experimentally determined three-dimensional structures of the complex of protease and bound ligand, acquired by means of X-ray crystallography, provide static information about conformations of these molecules and interactions between them. Molecular dynamics, a computer simulation technique for modeling molecular movement, can be employed to gain insight into the relevant conformations of the complex and its dynamics, starting from a crystal structure template. When used in combination with statistical mechanics methods for thermodynamic analysis, this method allows to estimate the free energy of binding between ligand and different protease mutants. The free energy describes the state of a system of molecules and is related to the probability of observing a specific conformation of this system. This knowledge of the free energy is of highest importance for understanding the behaviour of any molecular system.

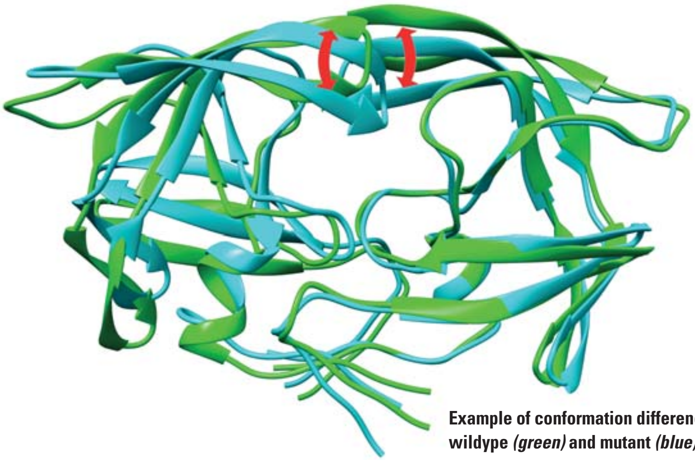

Example of conformation differences between of wildype *(green)* and mutant *(blue)* protease

Resistance mutation L76V

A particular mutation in the HIV genome, causing position 76 of the protease to change from amino acid leucine to valine (L76V), is a major resistance mutation. This means that the mutated protein is resistant towards inhibition by some protease-targeting drugs, but susceptible towards other drugs. This amino acid is located outside the drug-binding pocket. Molecular dynamics simulations of protease with different sequences, as found in some patients, together with different inhibitors were performed with both L or V variants at position 76. Thermodynamic calculations were applied to estimate the difference in free binding energy between wildtype (76L) and mutant (76V) protease when binding to an inhibitor, indicating whether the mutation renders protein-drug interaction more or less stable. Comparing wildtype and mutant proteins in complex with different drugs revealed differences in the propensities of different amino acids of the protein to make contacts with the inhibitor. This suggests possible mechanisms of how the mutation influences drug binding. Simulations of the unbound protease structures also revealed differences in the conformations assumed by the wildtype and mutant proteins [see figure]. Since it is in the unbound form in which the protease comes to interact with its natural substrates, this analysis provides insight into possible reasons behind evolutionary selection against some mutations: they may simply have consequences detrimental to the formation of the infective viral particle. Overall, this type of analysis sheds light on the effect of distant site mutations and gives possible support for drug development. :::

CONTACT

Tomas Bastys
DEPT. 3 **Computational Biology and Applied Algorithmics**
Phone +49 681 9325-3029
Email tbastys@mpi-inf.mpg.de

Bioinformatical Support of HIV Therapy

Preventing and controlling viral resistance is the central goal when administering drug therapies against viral infections. For instance, HIV exists in millions of variants. In order to suppress the replication of the virus inside the body of the patient [see figure 1], combination drug therapies are administered. These are composed of drugs from a base set of over two dozen compounds [see figure 2]. The relevant information for selecting the drugs is the viral genotype, which can be identified from blood samples of the patient with methods of genome sequencing. From the genotype one can identify the viral resistance phenotype with bioinformatics methods. There are two approaches for this purpose. The first consists of experts manually composing a set of rules for deriving the phenotype from the genotype. These rules are collected in computer-based expert systems which are utilized by the therapist. The second, more systematic approach consists of the bioinformatical derivation of the viral phenotype from the genotype based on a suitable set of clinical data about viral resistance. At the Max Planck Institute for Informatics we are following the second approach. Our research over the last decade or so has led to the geno2pheno system, which is freely available over the Internet at www.geno2pheno.org, for treating AIDS patients. Analyses offered on the geno2pheno server have found their way into the European guidelines for treating AIDS patients with certain drugs. In 2010 the research on geno2pheno was credited with the AIDS Research Award of the Heinz-Ansmann Foundation.

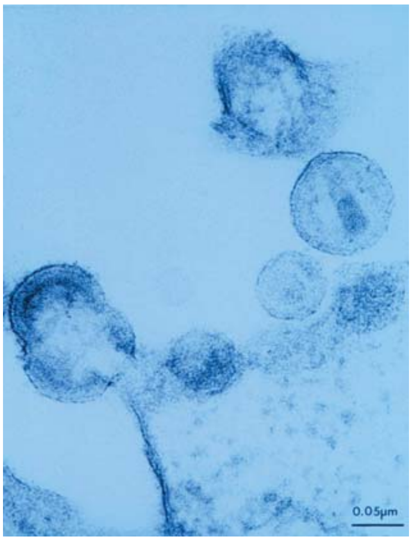

Figure 1: HIV particle budding from an infected immune cell *(lower right)* *(courtesy Prof. Schneweis, Bonn)*

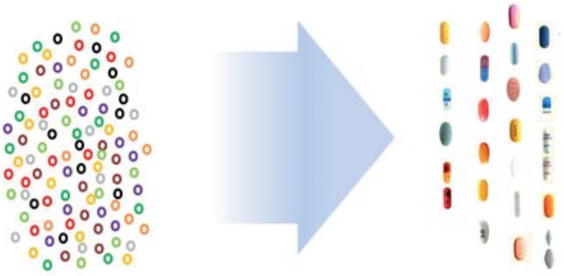

Figure 2: Millions of HIV variants are contrasted with hundreds of combination drug therapies based on over two dozen drugs. The drug selection benefits from bioinformatical support.

A class of analyses offered on the geno2pheno server that are used in clinical routine can be regarded as virtual phenotypes. A virtual phenotype is a bioinformatical, i.e. computer-based procedure which estimates the result of a laboratory experiment that is informative about the therapy of the patient and which therefore can be used as companion diagnostic for selecting the medication. Here the laboratory experiment itself is usually not accessible, be it that is too expensive, takes too much time, or is unavailable. Therefore the experiment is carried out only on a limited number of samples within a research project, in order to generate a suitable data set on which to derive the virtual phenotype. geno2pheno makes available such virtual phenotypes, which can be used as companion diagnostic for administering certain AIDS drugs.

Another option for antiviral therapy that is in development is the application of antibodies. For this still emerging approach, we have developed supporting software and performed analyses as well (*“Attacking HIV From New Angles”*, page 45).

Data collection is a substantial component in our research efforts. The resulting databases also afford analyses of networks of transmissions of the virus among AIDS patients (*“Reconstructing Pathogen Transmission Networks Using Genomic Data”*, page 48).

The molecular basis of drug resistance is a difficult mystery to which we have contributed in case studies (*“Structural Insights into the Emergence of Resistance in the HIV-1 Protease”*, page 46) as well as through basic research (*“The Relation Between the Mole-cular Conformational Changes and the Free Energy Change”*, page 44). :::

CONTACT

Thomas Lengauer
DEPT. 3 **Computational Biology and Applied Algorithmics**
Phone +49 681 9325-3000
Email lengauer@mpi-inf.mpg.de

Reconstructing Pathogen Transmission Networks Using Genomic Data

Rapidly evolving pathogens

The human immunodeficiency virus type I (HIV) has brought about one of the worst pandemics in recent human history with an estimated 2.1 million new infections in 2013. HIV is mainly transmitted through sexual contact among heterosexuals (HET) and men who have sex with men (MSM), and through the sharing of infected needles among intravenous drug users (IDU). Rapidly evolving pathogens like HIV accumulate genomic mutations as they go through a sequence of transmissions. This rapid evolution allows for the reconstruction of transmission networks using genetic data of pathogens collected from infected individuals.

Threshold-free networks

Large HIV sequence repositories, resulting from efforts to study the evolution of drug resistance in HIV, have enabled the construction of nationwide and global HIV transmission networks. Such networks have generally been constructed by thresholding evolutionary distance. A separate feature rich approach is to construct phylogenetic trees which model evolutionary history and have been widely used for inferring patterns of epidemic spread. The popularity of threshold-based networks over phylogenies is mainly due to the short computational time required for their construction.

A threshold-based approach does not seem appropriate for HIV due to variance in transmission risk by route of infection. As an alternative we construct transmission networks as minimum spanning trees which is a threshold-free approach. We used phylogenetic trees for extracting accurate estimates of evolutionary distance. In order to reconstruct phylogenetic trees in a reasonable time we employed parallel computing.

We used this approach to construct transmission networks for a collection of 27,124 HIV sequences which were aggregated from clinics and hospitals across Europe. We observed that network robustness was negatively correlated with size. This lack of robustness can be explained by the uncertainty in estimating the topology of the phylogenetic tree. This is because of a correspondence between the topology of transmission networks constructed as minimum spanning trees and the topology of phylogenetic trees. Specifically, each subtree in the phylogenetic tree corresponds to a subgraph in the transmission network as illustrated in figure 1.

Time-scaled networks

The sampling times of virus along with a phylogeny-based molecular clock model can be used to estimate the time of evolutionary events such as the most recent time point when two species diverged from a common ancestor.

While this model has been widely used to infer transmission dynamics from phylogenetic trees, it has not been used for the analysis of transmission networks.

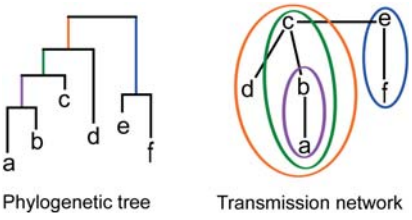

Figure 1: Topological correspondence between transmission networks and phylogenetic trees. The colored circles and branches indicate the corresponding subgraphs and subtrees.

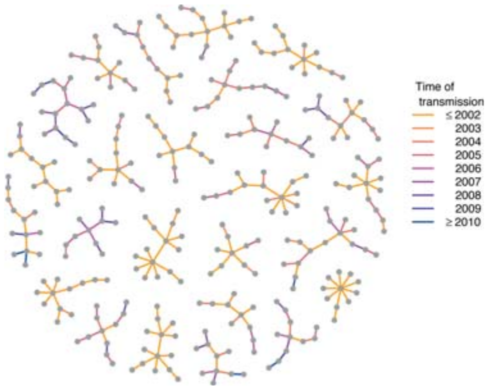

Figure 2: Transmission network with edges colored by estimated time of transmission

We applied this approach to estimate transmission time for individuals connected in the transmission network and constructed a time-scaled network as shown in figure 2. The time interval between transmission events was extracted from this network and we observed that HET intervals (34 months) were longer than MSM intervals (25 months). IDU intervals were found to be the shortest (20 months) which is probably due to the higher transmission risk involved in contact with infected blood as compared to sexual contact. :::

CONTACT

Prabhav Kalaghatgi
DEPT. 3 **Computational Biology and Applied Algorithmics**
Phone +49 681 9325-3023
Email prabhavk@mpi-inf.mpg.de

Measuring the Influence of Every Node in a Network by its Expected Force of Infection

Network models are ubiquitous in modern society. They define everything from our interpersonal relationships to the fundamental infrastructures that support our way of life. Even life itself is now frequently understood in terms of networks of biological interactions.

The importance of networks has led to a number of methods for identifying the most important nodes in a network, for example node degree or page rank. The last decade, however, has seen growing awareness that such measures are not informative for the vast majority of nodes that are not highly influential, pointing to a fundamental lack of understanding of the nature of influence in network models. Since innovation almost always originates from less influential nodes, a method for accurately quantifying the influence of any node has practical as well as intellectual utility.

We approach the question within a continuous-time epidemiological framework, building from first principles. In epidemiology the strength of an epidemic process is defined globally via its *force of infection* (*FoI*), the current rate at which individuals are becoming infected. Our task is to convert this global outbreak-specific time-varying random variable into an outbreak- and time-independent property of an individual node. Since the FoI is directly proportional to the current number of edges connecting infected and susceptible nodes, the outbreak-specific part amounts to a multiplicative constant. We resolve the time-varying aspect by only considering the early stages of transmission, because only those are influenced by individual nodes. This still leaves us with a random variable, which we summarize by taking its expected value. The resulting measure, the *expected force*, represents the expected force of infection that would be generated by an epidemic process originating from a given node.

Frankfurt has the highest *expected force* of any airport in the world.

We have tested the measure extensively on both simulated and real networks, for all basic epidemic models, in both continuous and discrete time. The *expected force* predicts epidemic outcomes with high accuracy over a broad range of network structures and spreading processes. Our most recent work evaluates its predictive power for pandemic spread over the world airline network. This study used highly realistic models which incorporate local population dynamics as well as travel patterns when simulating how a local disease outbreak would spread globally. The *expected force* of the airport servicing the outbreak's starting location showed 90 % correlation with the virulence the disease needed to become pandemic and 87 % correlation with the speed at which the outbreak achieved pandemic status.

Since the expected force depends only on local topology, it allows epidemic outcomes on the whole network to be predicted with high accuracy, even when only a small portion of the network is known. The full structure of a real network is rarely fully known; typically the network structure is inferred from indirect, incomplete, and often biased observations.

What, then, does this measure tell us about the nature of node influence? The definition of the expected force implies that influence is determined by both the degree of the node and the degree of its neighbors, and that the relative influence of these two factors is different for nodes of low versus high spreading power. Weaker nodes gain what strength they have from their neighbors, whereas more influential nodes get their strength from their large number of connections.

This confirms what we also know morally to be true. The best way to increase your personal influence is to build up the influence of the people around you. :::

CONTACT

Glenn Lawyer

DEPT. 3 **Computational Biology and Applied Algorithmics**

Phone +49 681 9325-3007

Email lawyer@mpi-inf.mpg.de

Integrative Analysis of Cancer Data

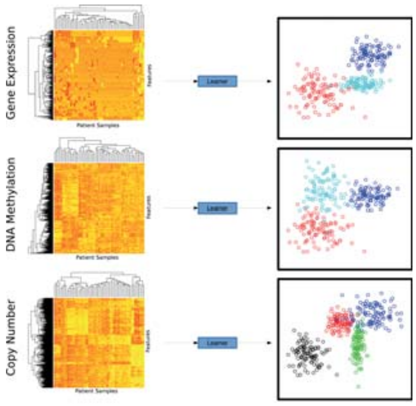

Conventional approaches analyze each data type separately.

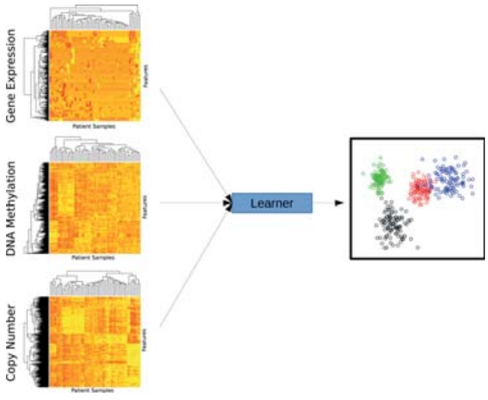

The integrative analysis of several data types can identify information that is weak in each individual data source but supported by different data sources.

Despite extensive research, cancer remains a major health threat that is difficult to treat and, in many cases, lethal. One reason for this is the heterogeneity of the disease, meaning that similar symptoms can have diverse underlying genomic or meta-genomic causes. Therefore, personalized treatments, which do not only take into consideration the tissue of the tumor but are tailored specifically to the patient, increase the chance of successful treatment. Cancer subtypes provide a helpful tool for choosing a good therapy for a patient. Here, cancer subtypes refer to groups of patients with distinct clinical outcomes, e.g., one group might benefit from a specific treatment while a different group does not. The better a subtype reflects the subtleties of a tumor, the better a therapy can be adjusted. Established subtypes are mainly based on individual data types, such as the gene expression data of patients. However, tumors often exhibit mutations on different levels of the cell. Although these changes can influence each other, there is no indivudual data type that captures all the information about a tumor. Therefore, different molecular data is measured for cancer patients, often including for example genetic and epigenetic data. In this context, one aim is the development of methods that reliably integrate these data in a biologically meaningful manner.

Data integration

One approach for the fusion of these data is regularized multiple kernel learning for dimensionality reduction (*rMKL-DR*). Using this method, weights for the different data types are learnt in an iterativ procedure. These weights reflect the information that is available from the data type such that data types with less information will have smaller weights and, therefore, less influence on the final result than data types with more information. The integrated data points, ours being cancer patients, are then projected into a low-dimensional subspace, which allows for visualisation and further analysis of the structure of the data. The distances between the data points reflect the similarities that were calculated based on the information given in all data types used. In our case, these similarities between the cancer patients can be used to identify integrated subtypes and, thus, help guide treatment decisions.

Interpretation

Conventional machine learning methods are generally incapable of explaining the underlying causes of a tumor, and they cannot adapt to unseen conditions. Therefore, we have developed a method which extracts different explanations for the molecular basis of a disease from the data; it learns which explanations are more likely, according to observed patterns in the individual samples. When a new sample is given to the method, it predicts how likely each extracted explanation is for the data partaining to this particular sample. It then gives a prediction and its interpretation accordingly. By keeping each explanation simple, in the end, the method gives a specific interpretation of its prediction, which can be used by oncologists to further investigate the specific cause of cancer in a particular patient, and also to use treatment targeted to that patient. :::

CONTACT

Adrin Jalali
DEPT. 3 Computational Biology and
Applied Algorithmics Phone
+49 681 9325-3028
Email ajalali@mpi-inf.mpg.de

Nora K. Speicher
DEPT. 3 Computational Biology and
Applied Algorithmics
Phone +49 681 9325-3027
Email nora@mpi-inf.mpg.de

Charting Epigenomes

Epigenetics is a research area which is at the center of worldwide attention today. The reason is that after the sequencing of the human genome just over 10 years ago, it became clear quite quickly that the central aspect of cell regulation cannot be inferred directly from the genomic DNA,since all cells in an organism harbor the same genome. Cell regulation comprises specifically those aspects which are different in different tissues, in which healthy cells are different from diseased cells, and which implement central biological processes such as the cellular response to stress and the process of aging. These aspects of cell regulation are implemented via the dynamic organization of the genome in the cell nucleus which, in turn, is based on chemical modifications of the genomic DNA and of the protein scaffold that envelops it. With techniques of molecular biology that have been developed in the last two decades, such modifications can be measured genome-wide. As a result, their entirety, the epigenome, can be charted.

Whereas each individual has only a single genome, each of us has many epigenomes. In principle, all of the around 200 tissue types in the human organism are characterized by different epigenomes. Furthermore, the epigenomes of healthy cells differ from those of diseased cells. The knowledge of the epigenome is a prerequisite of the understanding of the cell-regulatory processes in the corresponding cell type. For this reason, several years ago an international federation of scientists known as the International Human Epigenome Consortium (IHEC) was founded with the goal of charting at least 1000 epigenomes of human cells within the next five years.

The Max Planck Institute for Informatics contributes to IHEC in two ways. First, it is a partner in the BLUEPRINT project of the EU, which represents the European contribution to IHEC. This project, with largest ever EU funding volume for a single project in the area of biology of euro 30 million, aims at charting one hundred epigenomes of bloodline cells (haematopoietic cells), of healthy as well as diseased (mostly malignant) cells. Second, the Institute is a partner in the German Epigenome Program (Deutsches Epigenom Programm, DEEP). This project, which is funded by the German Science Ministry with about euro 20 million, comprises the German contribution to IHEC. Within DEEP, metabolic and immunologic diseases in particular are being investigated and another 70 epigenomes charted. The Max Planck Institute for Informatics coordinates the bioinformatics part of DEEP.

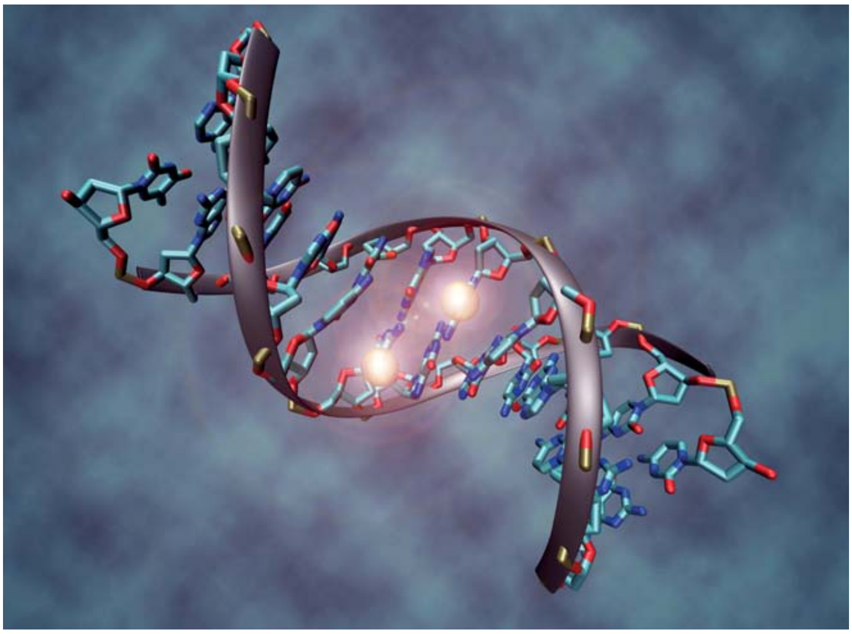

The methylation of the DNA impacts its accessibility by the transcription machinery of the cell.

In addition to charting epigenomes, we are also involved in interpreting the resulting data. This entails the identification of epigenetic patterns that are characteristic for different types and stages of disease.

For this purpose, we have been developing software systems for analyzing epigenomic data that are in worldwide use (*"BiQ Analyzer HiMod for the Locus-Specific Analysis of DNA Modification Data"* page 55, *"EpiExplorer and RnBeads: Integrative Analysis of Epigenomic Data"*, page 54).

Furthermore, the Institute is involved in validation studies that involve the application of our software in order to reap biological insight from epigenomic data. As a point in case, a study published in 2012 identified the tissues of origin of malignant tumors that could only be found in the form of metasases. Such identification is central to effective tumor therapy. :::

CONTACT

Thomas Lengauer
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 9325-3000
Email lengauer@mpi-inf.mpg.de

Analyzing the Three-Dimensional Organization of the DNA Inside the Nucleus

For any eukaryotic organism, its complete genome is packed inside the nucleus of the cell in a very dense fashion to form chromatin, which is a complex of DNA, RNA, and proteins. For example, the human genome, which comprises about 3 billion base pairs and, when stretched out, is about 2 m long, fits inside the cell nucleus with a diameter of one thousandth of a millimeter only due to this compaction brought about by forming chromatin. Albeit the compact packaging, the chromatin is arranged in three-dimensional (3D) space in a non-random manner; the arrangement, the so-called *spatial conformation of the chromosomes*, is closely correlated with the functional state of the cell (healthy or diseased), gene activities, and other factors. Thus, a better understanding of this three-dimensional structure can help gain important insights into gene regulation. Though the overall structure of the chromatin is dependent on the state of the cell, the general picture, jointly measured over many cells, of the spatial co-localization of different portions of the genome can suggest interactions between chromosomal regions that lie far away in the genome sequence, and are possibly functionally important for cells belonging to a particular tissue.

Only over the last decade have molecular biology techniques advanced to be able to identify such 'long-range' chromosomal interactions genome-wide. And these techniques have provided us with the first set of genome-wide interactions data to model the 3D structure of the chromosome. Obtaining multiple, high-resolution genome-wide interaction profiles for all the different cell-types of various organisms still remains very costly.

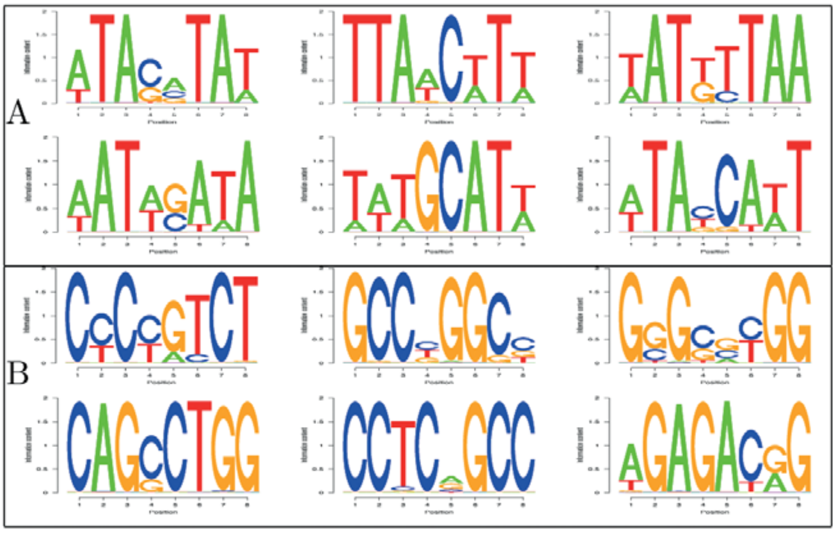

Visualization of the DNA sequence motifs that contribute most to the prediction. Panel (A) in the above figure visualizes the top-6 of a set of enriched motifs found in the regions highly interacting with a particular genomic locus of interest over a 'long-range'. Panel (B) shows the top-6 of the set of enriched motifs in the regions that are depleted for these interactions. Each sequence motif shown is visualized as a sequence logo where, at each position, the relative height of the base pairs denotes their frequency at the position, and the total height of the letters denotes the information content of the position.

In our research, we are developing statistical models that help analyze these 'long-range' interaction data. Our models let us predict the interaction between putative interacting pairs, i.e., a putative, functionally important pair of genomic loci that spatially co-localize. For example, one of these genomic loci could be a region that is a few hundred base pairs long and just flanking a gene, and another a region that lies several kilo- or megabases apart on the linear genome, but plays a role in the expression of this gene. Furthermore, we can characterize the DNA sequence motifs that are responsible for the prediction in order to analyze general underlying mechanisms. Our model can be useful in a scenario where only low-resolution data on genome-wide interactions are available, but a few selected portions of the genome have been studied at high-resolution. In this case we can build a model using the latter and help make reasonable predictions on even the low-resolution data. Based on the informative sequence signals our model picks up to make these predictions, it may be possible to further investigate the underlying folding mechanisms of the genomes. Also, our predictions may be helpful in improving existing techniques that build coarse-grained 3D models of the chromatin and suffer from the use of low-resolution genome-wide experiment data. :::

CONTACT

Sarvesh Nikumbh

DEPT. 3 **Computational Biology and Applied Algorithmics**

Phone +49 681 9325-3028

Email snikumbh@mpi-inf.mpg.de

Phenotype-Based Computational Clinical Diagnosis for Genetic Diseases

Finding the correct diagnosis for patients is challenging, especially in the field of medical genetics. A differential diagnosis is complicated by a large number of genetic disorders, each of which is characterized by many clinical symptoms that overlap with those of other diseases. Additionally, many clinical symptoms show a variable degree of expression in patients with a disease, which further complicates the diagnosis. It is essential to avoid unnecessary diagnostic procedures before a final diagnosis is made, to reduce the burden on patients and the cost of their medical care. For complex and rare genetic diseases patients sometimes what years until a correct diagnosis has been made.

Over the past years we have developed novel algorithms and statistical approaches to solve the differential clinical diagnosis problem, in cooperation with the group of Prof. Dr. Peter Robinson at the Charite hospital in Berlin. These approaches make use of the Human Phenotype Ontology (HPO) that was developed in Berlin.

The Phenomizer system for computational clinical diagnosis

The Phenomizer webserver implements our research ideas and is freely available to physicians all over the world. The user enters clinical symptoms of a patient, as defined by standardized HPO terms. Then the Phenomizer returns a ranked list of the best matching differential diagnoses from all genetic disorders in the database. Currently the Phenomizer knows more than 10 000 disease symptoms and their relationships. These are assigned to over 7500 diseases.

The Phenomizer app allows computational diagnosis of genetic diseases in clinical practice.

The Phenomizer reports for each search result the estimated probability of obtaining the same result by chance alone, obtained with efficient algorithms that we have developed. Within seconds, a doctor receives a list with the most probable results, so physicians no longer have to research in databases or books for several hours. The list supports them in detecting the correct disease more quickly or in planning additional examinations to decide between two or more genetic diseases that the Phenomizer deemed equally likely for the given query symptoms.

Going mobile for clinical practice

In order to support fast decision-making in clinical practice and due to the widespread availability of smartphones and tablets, we have developed the Phenomizer App for Android-based systems. The App allows the same functionality as the Phenomizer webserver on these devices and supports the saving and exchange of queries.

The Phenomizer server has been used over 50000 times in the last 3 years and more than one thousand people have the App installed. :::

CONTACT

Marcel H. Schulz
DEPT. 3 **Computational Biology and Applied Algorithmics**
Phone +49 681 9325-3115
Email mschulz@mpi-inf.mpg.de

EpiExplorer and RnBeads: Integrative Analysis of Epigenomic Data

Our DNA encodes how our cells behave and what they look like. However, there is a multitude of additional layers of control that determine the fate of a cell. Chemical modifications of the DNA itself (such as DNA methylation) or its molecular scaffold (such as histone modifications) regulate the ways in which the DNA is packed and processed in the nucleus of the cell and are thus responsible for the cell's biological properties (*phenotype*). Two software packages developed at the Max Planck Institute for Informatics facilitate the interpretation and integration of such *epigenetic* layers of molecular information: *EpiExplorer* allows for an integrative view and interactive exploration of genomic and epigenomic features based on technology that has been developed in the context of Internet search engines for retrieving information in a fast and intuitive fashion. *RnBeads* is a comprehensive software package that supports the analysis of experimentally obtained DNA methylation data.

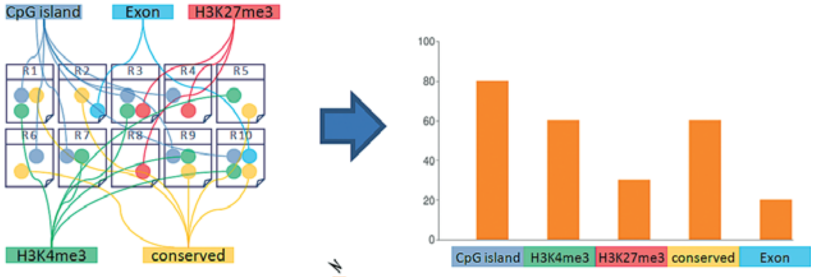

EpiExplorer facilitates fast, interactive searches in genomic and epigenomic maps and visualizes them in an intuitive fashion

EpiExplorer: A tool for the efficient search and comparison of genomic and epigenomic features

Recent technological advances are leading to an explosion in the amount of quantitative biological data available for a large number of individuals and cell types. Major bioinformatics challenges include the integration of many heterogeneous datasets comprising sequence information on the molecular architecture of the packed genome (chromatin) and on its epigenomic regulation as well as facilitating the discovery of molecular interrelations between genomic and epigenomic features. EpiExplorer is a web-based tool that enables scientists to explore a multitude of large datasets in a fast and intuitive fashion. EpiExplorer's key strength lies in supporting exploratory hypothesis generation using a broad range of genome-wide data analyses performed in real time over the internet. The tool integrates a variety of genomic and epigenomic maps and helps scientists identify landmarks along the genome sequence characterizing regions in the genome of particular interest. EpiExplorer uses a versatile text indexing scheme in order to perform comprehensive analysis and visualization over the web in a matter of seconds. Datasets can also be compared quantitatively, and users can incorporate their own datasets into the analysis.

Comprehensive analysis of DNA methylation data using RnBeads

DNA methylation comprises an important layer of epigenomic regulation of genomic programs and is probably the best studied of all epigenomic marks. It is a major contributor to embryonic development, and aberrant methylation patterns have been associated with complex diseases like cancer. Various experimental assays have been developed for genome-wide analysis of DNA methylation patterns. RnBeads is a software package that enables both wet lab scientists and computationally inclined scientists to run the entire analysis pipeline on data originating from DNA methylation experiments: Experimental biases can be detected and removed and DNA methylation fingerprints can be quantified, visualized, and compared. For instance, genomic regions that exhibit differential methylation fingerprints in cancerous tissue compared to normal cells can be identified and annotated with biological knowledge. Genomic regions of interest can be exported to EpiExplorer where comparison with other (epi)genomic maps is facilitated. While the default workflow of RnBeads requires no expert knowledge, due to its modular design it is also highly configurable for custom analysis. The tool generates comprehensive and interpretable results in HTML format, thus facilitating easy tracking and comparison of analyses as well as the exchange of results with collaboration partners. :::


CONTACT

Felipe Albrecht
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 9325-3008
Email albrecht@mpi-inf.mpg.de

Fabian Müller
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 9325-3009
Email fmueller@mpi-inf.mpg.de
Internet http://epiexplorer.mpi-inf.mpg.de

BiQ Analyzer HiMod for the Locus-Specific Analysis of DNA Modification Data

Every cell in the human body contains an identical copy of our genetic makeup, the so-called DNA (*deoxyribonucleic acid*). The DNA is seen as the blueprint for all parts of the body, yet it is obvious that due to the numerous different cells that constitute our body, e. g., skin, heart ,or liver cells, a complex machinery is required to translate the blueprint and to control human development from single egg to old age. A central component of this system is the so-called epigenome, a generic term for mechanisms that regulate the reading of genetic information and thus affect cellular development (see also *"Charting Epigenomes"*, page 51). A simple chemical modification of the DNA, called DNA methylation, is one of the best-studied epigenetic marks. DNA methylation is one of the key factors that lets the cellular machinery distinguish between active genes, i.e., readable and intended for protein production, and inactive genes.

DNA methylation as source of medical information

The importance of DNA methylation in the regulation of biological processes inside the cell is not only limited to normal human development. It is, in fact, in the focus of active biomedical research for many different diseases. Aberrant patterns of DNA methylation are known to play a role in the development and progress of severe diseases such as cancer or Alzheimer's. State-of-the-art research results in methods for the earlier diagnosis of diseases based on anomalies in DNA methylation patterns as well as more personalized treatment, e.g., an individual estimate of the chances of success of chemotherapy for certain types of cancer. To foster this progress, in particular in the field of biomedical research, modern and easy-to-use software is needed that allows analysis of DNA methylation data even for non-bioinformaticians.

Keeping pace with scientific progress

We develop software tools tailored for this specific scenario (see also *"Epi-Explorer and RnBeads: Integrative Analysis of Epigenomic Data"*, page 54). The original version of BiQ Analyzer was published in 2005 and quickly became a standard tool for visualization, quality control, and locus-specific, i.e. targeted, analysis of DNA methylation data. Within a few years, technological advances enabled scientists to produce new data faster and in higher quantities, which demanded an update of our program to accomodate data generated by these so-called high throughput technologies (BiQ Analyzer HT for *high throughput*, 2010). An important aspect of such updates is to keep established and accepted user interfaces intact, which makes it unnecessary for users to adapt new ways of operating the software. Besides technological advances, new biological insights also require software maintenance: recently discovered chemical variants of DNA methylation receive more and more attention by scientists due to their assumed biological relevance. A variant called hydroxymethylation is suspected to be a key factor in the cellular de-methylation process, i.e., the controlled removal of DNA methylation – an event that lacked an explanatory model up to now.

BiQ Analyzer HiMod [see Figure] supports the analysis of regular DNA methylation data and all its known variants. Developing a suitable update was exacerbated by the fact that standard experimental protocols do not allow to distinguish between DNA methylation and DNA hydroxymethylation during the computational analysis. Specialized laboratory protocols are necessary to reliably identify DNA hydroxymethylation later in a comparative analysis done by BiQ Analyzer HiMod. Long-term support and regular updates adapting bioinformatics software to new discoveries and technologies are essential to enable researchers not only to understand biology, but more importantly, to develop new approaches in fighting complex diseases. :::

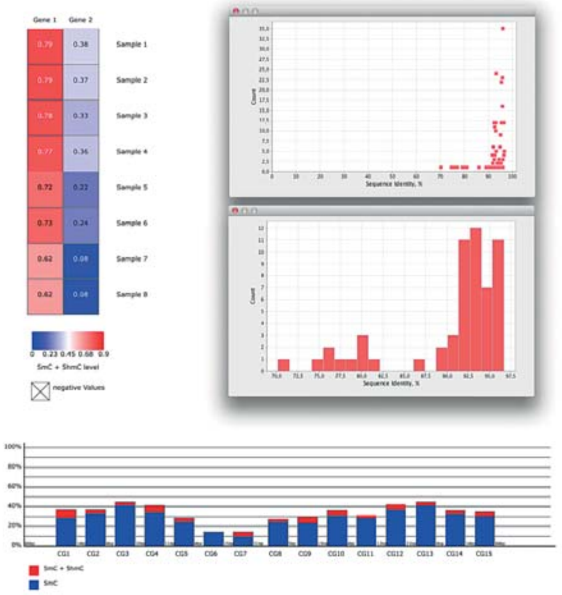

Different visualization modes of (hydroxy-)methylation data in BiQ Analyzer HiMod

CONTACT

Peter Ebert
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 9325-3008
Email pebert@mpi-inf.mpg.de
Internet http://biq-analyzer-himod.bioinf.mpi-inf.mpg.de

Detection of Related Proteins in Unrelated Viral Species

A virus is a small particle that contains a nucleic acid (DNA or RNA) harbouring the genomic information, encapsulated in a protein shell. Viruses lack ribosomes and other essential molecular machines, and thus they rely on the host cell of another organism that they infect for replication.

Unlike cellular organisms, viruses do not have a universal common ancestor, and the origin and evolution of viruses is still highly debated. Evolutionary inference in the virus kingdom is very complicated. The evolutionary relatedness of cellular organisms is usually deduced from comparison of the RNA sequence from the small ribosomal subunit, which is absent in viruses. Additionally, viruses have a very high rate of evolution, several orders of magnitude faster than that of cellular organisms. This makes sequence comparison effective only in closely related viruses, and the evolutionary relationships between different virus families are largely unknown.

In our work, we combine sequence comparison with analysis of protein structure to push further the boundaries of our understanding of viral evolution. Proteins are biological machines that implement virtually all processes in a living organism. Proteins are encoded in the genome in the form of genes, which are translated into a sequence of amino acids that fold into complex three-dimensional (3D) structures. The exact form of the protein 3D structure is dictated by the chemical properties of its amino acids. It is that form that is the basis for the biological function of the protein. Evolutionarily informative similarities between 3D protein structures can be found between much more distantly related organisms than those with detectable similarities of their genetic sequences.

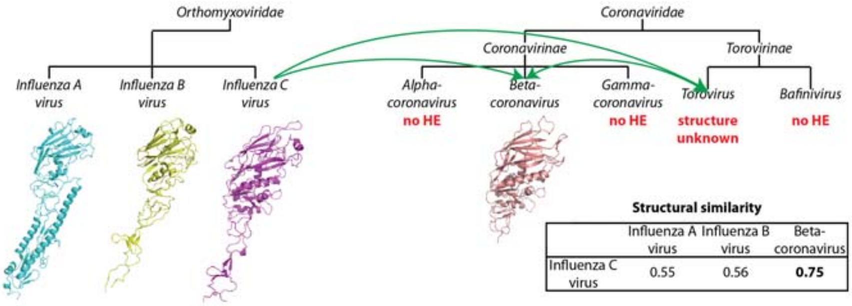

	Influenza A virus	Influenza B virus	Beta-coronavirus
Influenza C virus	0.55	0.56	**0.75**

Transfer of HE. Green arrows represent potential routes of lateral gene transfer.

We have compared all known sequences and structures of seemingly unrelated viruses in order to detect similarities that may imply evolutionary links. Particularly, we are interested in viruses with different types of genome-carrying nucleic acid (DNA or RNA, single- or double-stranded), since they exhibit the largest differences in their life cycle. We have observed two general patterns of evolution: Some viral proteins are widely distributed in very different viral species and thus are very ancient, possibly dating back to the origin of life itself. Others have been transferred between viral species via a process known as horizontal or lateral gene transfer, a process of transfer of genetic material by means other than from parent to offspring.

A very interesting example is posed by the hemagglutinin-esterases (HE) of Torovirus and Betacoronavirus, related viral lineages from the Coronaviridae family that cause gastric and respiratory syndromes in vertebrates. HEs mediate attachment of the viral particle and subsequently destroy the host cellular receptor protein involved in viral cell entry. They are similar in sequence to HE of Influenza C virus, but sequence comparison does not indicate from which of the species the protein originated. However, comparison of protein 3D structures of Influenza A, Influenza B and Influenza C (Orthomyxoviridae family) HEs with that of Betacoronavirus clearly shows that HE was introduced into Torovirus and Betacoronavirus after the split of Influenza C HE from Influenza A and B around 8000 years ago. Since other viruses from the Coronaviridae family do not have HEs, this must have been a lateral gene transfer event in Torovirus and Betacoronavirus lineages. They share hosts, so the transfer could have happened to one of them and then to the other, or independently. :::

CONTACT

Olga Kalinina
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 9325-3004
Email kalinina@mpi-inf.mpg.de

Structural Variation in Genomes

Each of our body cells harbors a copy of our *genome*, the set of all genetic material. The genome is organized in chromosomes: we inherit 23 chromosomes from our mother and father, respectively. A chromosome consists of a long DNA molecule, stabilized and spatially structured by special proteins. The DNA encodes genetic information as a sequence of its constituting bases adenine (A), cytosine (C), guanine (G), and thymine (T). In this sense, a chromosome can be viewed as a sequence of the letters A, C, G, and T. Likewise, we can represent a genome as a set of such sequences (one for each chromosome). The research field of genomics asks the question, among others, of how our genome influences our traits. Many different traits can be of interest, such as body height, the risk of contracting a certain disease, or the tolerance to a drug, for instance.

The genomes of two persons can differ in a multitude of ways, as illustrated in figure 1. Traditionally, many genetic studies focused on point mutations (SNPs). That is, one restricts oneself to differences in single letters. In recent years, however, increasing attention has also been given to larger differences, the so-called *structural variants*, such as deletions or duplications of whole DNA segments. Such variants occur frequently in humans and, hence, studying them is indispensable for capturing the full spectrum of human genetic diversity. This is enabled by new sequencing technologies in combination with novel algorithmic methods, which we develop actively.

High-throughput methods for *genome sequencing* allow for reading millions of short DNA fragments in parallel. That means, these devices *do not* return the sequence of a whole chromosome; instead, they output the sequences of short fragments, called *reads*. The analysis of the resulting data is thus comparable to solving a jigsaw puzzle: from the reads, we want to reconstruct the sequenced genome. Here, we particularly focus on detecting structural differences.

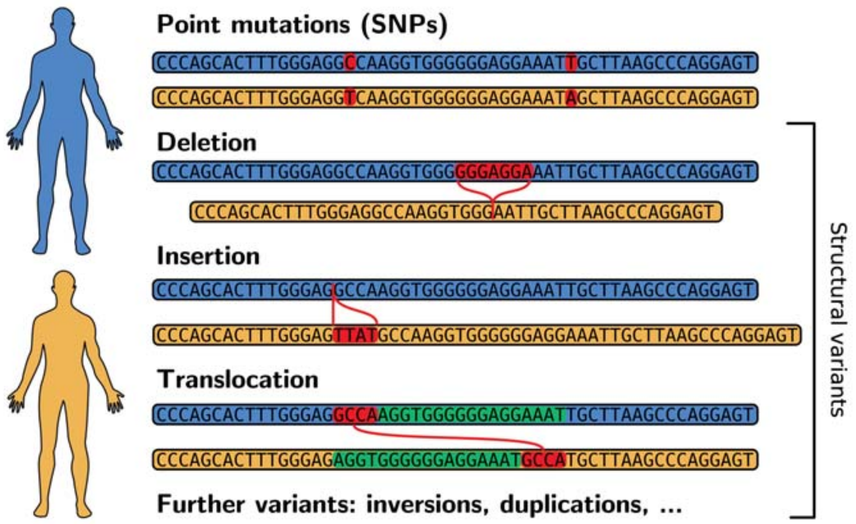

Differences between the genomic sequences of two persons

In a large study of 250 families, our computational methods have contributed to characterizing structural genetic variants in detail. In this project, the complete genomes of both parents and one child were sequenced for each family. The large number of studied families as well as the high sequencing depth render this effort a leading project worldwide. By reading the genomes of parents and children, we can determine which genetic variants found in the children were *not* inherited from the parents, but are the result of *new* mutations. The rate at which these so-called *de novo mutations* happen is of great interest as a parameter in models of human evolution.

Beyond fundamental research, studying structural genetic variants is potentially important for personalized medicine. So far, association studies that aim to establish connections between the genome of a person and clinically relevant traits were restricted to point mutations. Our recent research provides evidence that taking structural variants into account in such studies leads to new, medically relevant insights. :::

CONTACT
Tobias Marschall
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 302 70880
Email t.marschall@mpi-inf.mpg.de

GUARANTEES

Software and Hardware must be reliable. The most important criterion for reliability is correctness. Nearly as important, however, is performance: A correct answer that is not received in time is not helpful. The search for correctness and performance guarantees is a cross-departmental research topic at the institute.

Today, computers are a ubiquitous part of our lives. We use them constantly, sometimes knowingly, such as a computer on the desk, and sometimes unknowingly, such as the electronic controls in the car, airplane, or washing machine. The more our lives are dependent upon hardware and software, the more we pose the question of whether the trust we place in these products is justified. The answer to this question is especially difficult and complex because software, and meanwhile hardware, is not robust. A single change in a program, a minor error in one line of a program with several hundred thousand lines of code, or a small change in the processor's calculations can cause the overall program to crash or deliver wrong results. Or, the developers may not have considered an infrequent but possible border case, and software that has worked reliably for years can suddenly show incorrect behavior. The granting of guarantees becomes even more complex when we examine not just a single program, but a distributed system of programs in which parts work independently but also react to each other while pursuing different goals.

The article *"Dynamic Matching under Preferences"* investigates the latter situation. It examines the problem of match-making individuals pursuing different preferences, seeking answers as to which assumptions will the match-making be successful.

The articles *"Distributed Algorithms for Fault-tolerant Hardware"*, and *"Formal Verification of Peer-to-Peer Protocols"* discuss guarantees of different highly distributed systems: the former modern processor hardware, and the latter modern internet protocols.

In order to provide guarantees of programs in general, we need deductive methods that enable us to derive the desired properties from a program in the form of proofs. This requires automatic verification procedures, which are discussed in the article *"Automated Deduction"*, as are techniques to properly divide and conquer resulting proof obligations, in particular, from theories such as arithmetic. This aspect is discussed in the article *"Quantifier Elimination – Statements can also be Calculated"*. :::

CONTRIBUTIONS

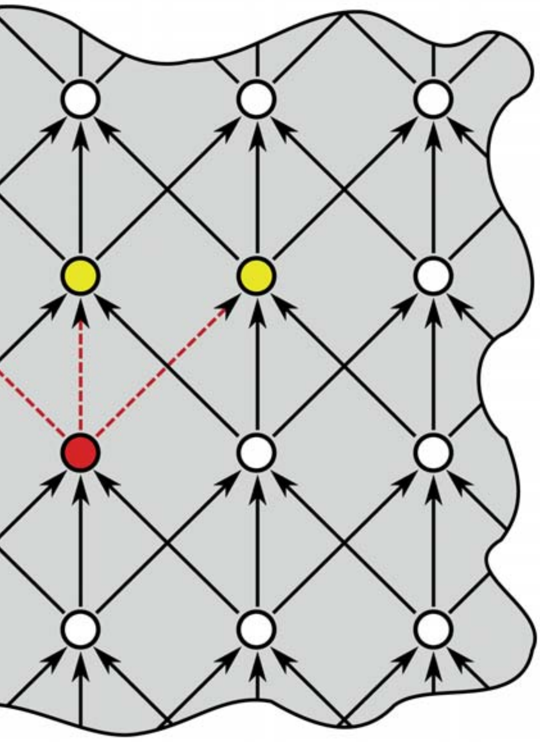

UNDERSTANDING IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH & DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION & VISUALIZATION

OPTIMIZATION

Formal Verification of Peer-to-Peer Protocols

Peer-to-peer (P2P) networks are thriving as a means of communication over the Internet. An email service, for example, may be based on the more traditional client-server model, where all information is stored in a central, dedicated location called a *server*, and *client* devices, such as phones and laptop computers, communicate with the server to get new messages or manage the inbox. The server controls the level of access users/clients have to resources and information and coordinates the communication between different clients. This centralized model has its benefits, but it is vulnerable due to the server being a single point of failure.

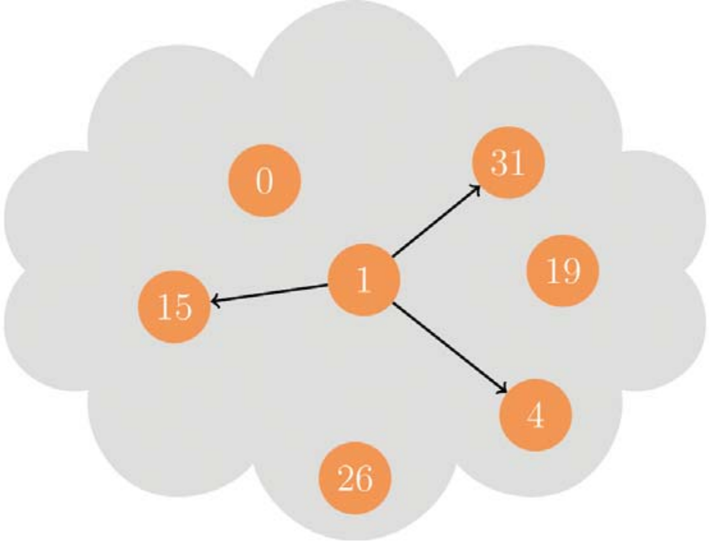

Instead, many modern Internet applications that we use every day, like Skype, Adobe Flash Player, Jaman or Tor, rely on peer-to-peer technology. Participating nodes form an overlay network on top of the Internet. In this overlay network, messages do not pass through a central location; the work of delivering a message to its destination is *distributed* among nodes. Each node maintains a *routing table* with the addresses of a few other nodes in the network. If node x does not know the address of the destination node y, it must forward the message to some node z that it does know about on the way to y. *Distributed hash tables* provide the underlying data structure for this kind of message routing. Among the most popular P2P protocols are Chord, Pastry and Kademlia.

P2P protocols are susceptible to *churn*, since new nodes may join at any time, and existing nodes may sign-off or fail without giving notice. This makes it very challenging to ensure that each node's routing table stays up to date. Inconsistency in the nodes' local information may lead to big problems. In August 2007, for example, Skype was unavailable to the majority of users for 3 days, most likely due to a failure in its underlying P2P technology. It is therefore very important to verify the correctness guarantees claimed by these protocols. Hand-written proofs are subject to mistakes, and they almost always contain some. For example, though both Pastry and Chord offer some correctness/performance guarantees for which there are hand-written proofs, research has shown that no published version of either protocol is correct.

Formal verification offers a much better, though much more tedious, alternative. *Interactive* tools allow a user to write a mathematical proof in some formal language in the form of steps. Each proof step can be sent to a *theorem prover* software for verification. The proof is accepted if the prover can verify every step. If a step is too large or difficult for the theorem prover to prove automatically at one time, it can be broken into sub-steps. The size of the final proof largely depends on how good the underlying theorem prover is, both at handling large problems, and at handling problems of a particular type or from a particular domain. For example, verifying P2P protocols requires provers that can reason about arithmetic and sets, since node addresses are typically represented by numbers from a ring. Gaining experience in the formal verification of such distributed systems, as well as devising suitable reasoning techniques for the underlying provers, is an important field of research. :::

CONTACT

Noran Azmy
RG. 1 **Automation of Logic**
Phone +49 681 9325-2908
Email azmy@mpi-inf.mpg.de

Distributed Algorithms for Fault-tolerant Hardware

Distributed Computing is concerned with systems in which *communication* – as opposed to computation – is the main obstacle to solving a given task fast. It also addresses the issue of *fault-tolerance*, i.e., the ability of a system to work correctly even if some of its components fail. As hardware development continues to follow "Moore's Law" which states that the number of basic component in off-the-shelf hardware doubles roughly every 18 months, both of these aspects become increasingly important. In other words, nowadays a common computer chip, and even more so a multi-core processor, is a distributed system in its own right.

If this claim strikes you as odd, just consider the following facts: (i) at today's clock speeds in the gigahertz range, an electrical signal cannot traverse an entire chip in a single clock cycle; (ii) by now, standard processors host billions of transistors, making it virtually impossible to guarantee that *each* of them operates correctly; and (iii) the most basic standard components of computer chips have reached a size of about 10-20 nanometers (a nanometer is a billionth of a meter!), a size at which physical effects necessitate dealing with *transient faults*, i.e., components temporarily and unpredictably operating out-of-spec.

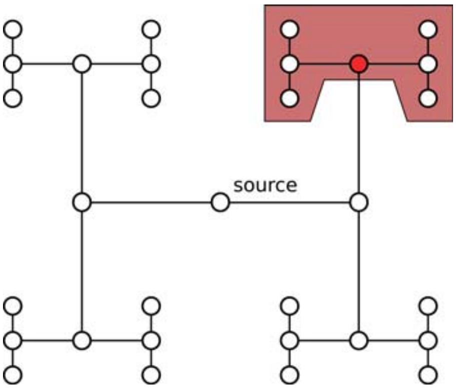

Figure 1: Clock tree. In case of a fault at the junction marked red the entire red area fails.

One of the most fundamental challenges in designing computer chips is the distribution of a well-behaved *clock signal* throughout the chip. Up to a certain scale, this can be successfully achieved by so-called clock trees [see figure 1]. Basically, a clock tree is carefully engineered to ensure that the clock signal generated by a single quartz crystal arrives at adjacent components at (almost) the same time. This permits to perform computations *synchronously*: a computational step is executed, then the results are "locked" and communicated, and then the received new values are used for the next computational step. The synchronous design paradigm offers many advantages and is therefore the de facto standard in hardware design – up to the point at which systems become too large to be efficiently clocked by a single clock tree.

We argue that using distributed algorithms, one can efficiently and reliably generate and distribute a clock signal on a significantly larger scale than possible with a single clock tree. For instance, consider Figure 2, which illustrates a simple fault-tolerant clock distribution mechanism. In contrast to a clock tree, there are different paths along which the clock signal is propagated; hence, the system as a whole can still operate correctly even if a few isolated nodes or links fail. We work towards efficient realizations of this and similar schemes to enable building larger, yet more reliable systems. :::

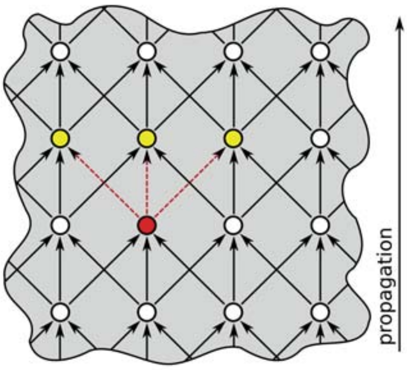

Figure 2: Part of a clock distribution grid. Each node waits for the *second* received signal before forwarding the clock signal on all outgoing links. Hence, the yellow nodes still operate correctly even if the red node or its outgoing links fail.

CONTACT

Christoph Lenzen
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1008
Email clenzen@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/departments/algorithms-complexity/research/theory-of-distributed-systems/

Automated Deduction

In order to guarantee that a piece of hardware or software is working correctly, it must be verified – i.e., its correctness must be proven formally. Proving correctness means that one has to check whether certain properties of the system follow from other properties of the system that are already known. The question of how to use computer programs for such proof tasks has been an important research topic for many decades. Ever since the fundamental theoretic results of Gödel and Turing at the beginning of the twentieth century, it has been known that not everything that is *true* in a mathematical sense is actually *provable*, and that not everything that is *provable* is *automatically provable*. Correspondingly, deduction systems differ significantly in their expressiveness and properties. For example, *decision procedures* are specialized for a certain type of data (such as real numbers) and are guaranteed to detect the correctness or incorrectness of a statement within this area. *Automatic theorem provers* for so-called *first-order logic* can handle any data types defined in a program. For such provers, it is guaranteed only that a proof will be found if one exists; if none exists, the prover may continue searching without success, possibly forever. Even more complicated problems can be handled using *interactive provers*; these provers, however, only work with user assistance and without any guarantee of completeness.

How does a theorem prover for first-order logic work? It is not difficult to write a program that correctly derives new formulas from given formulas. A logically correct derivation is, however, not necessarily a useful derivation. For example, if we convert $2 \cdot a + 3 \cdot a$ to $2 \cdot a + 3 \cdot a + 0$ and then to $2 \cdot a + 3 \cdot a + 0 + 0$, we do not make any computation errors, but we are not getting any closer to our goal either. The actual challenge is thus to find the few useful derivations among infinitely many correct derivations. The first step in this direction is easy: Obviousy it is a good idea to apply equations in such a way that the result is simplified, for example, "$x + 0 = x$" only from left to right and not the other way round.

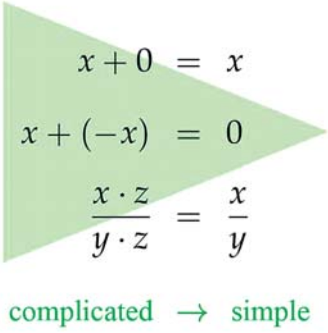

Application of equations

However, this approach is not always sufficient. A typical example is fractional arithmetic: We all know that it is occasionally necessary to expand a fraction before we can continue calculating with it. Expansion, however, causes exactly what we are actually trying to avoid: The equation "$(x \cdot z)/(y \cdot z) = x/y$" is applied from right to left – a simple expression is converted into a more complicated one. The superposition calculus developed by Bachmair and Ganzinger in 1990 offers a way out of this dilemma. On the one hand, it performs calculations in a forward direction; on the other hand, it systematically identifies and repairs the possible problematic cases in a set of formulas where backward computation could be inevitable. Superposition is thus the foundation of almost all theorem provers for first-order logic with equality, including the *SPASS* theorem prover that we have developed at the Institute.

In our research group, our current focus is on refining the general superposition method for special applications. In particular, we are developing techniques for combining the capabilities of various proof procedures. We are investigating, on the one hand, how arithmetic decision procedures can be integrated into superposition, and on the other hand, how superposition provers can improve the degree of automation of interactive provers. We are also addressing the question of how superposition can be used for business applications such as product configuration. :::

```
1[0:Inp] ||  -> F(U,skf6(U))*.
3[0:Inp] || equal(U,V) -> F(U,skf8(V))*.
2[0:Inp] || F(U,skf8(V))* -> equal(U,V).
4[0:Inp] || F(U,V)* F(U,skc1)* -> F(skf5(V),V)*.
5[0:Inp] || F(U,skf6(V))* -> F(V,skc1) F(skf7(U,W),U)*.
6[0:Inp] || F(U,skf6(V)) -> F(V,skc1) F(skf7(U,V),skf6(V))*.
7[0:Inp] || F(U,V)* F(W,V) F(U,skc1)* F(W,skf5(V))* -> .
Derived: 8[0:Res:3.1,2.0] || equal(U,V)* -> equal(U,V).
Derived: 10[0:Res:3.1,4.0] || equal(U,V)* F(U,skc1)* -> F(skf5(skf8(V)),skf8(V))*.
Derived: 9[0:Res:1.0,4.0] || F(U,skc1) -> F(skf5(skf6(U)),skf6(U))*.
Derived: 11[0:Res:1.0,5.0] ||  -> F(U,skc1) F(skf7(U,V),U)*.
Derived: 14[0:Res:11.1,4.0] || F(skf7(U,V),skc1)* -> F(U,skc1) F(skf5(U),U)*.
Derived: 13[0:Res:11.1,2.0] ||  -> F(skf8(U),skc1) equal(skf7(skf8(U),V),U)**.
Derived: 12[0:Res:11.1,5.0] ||  -> F(skf6(U),skc1) F(U,skc1) F(skf7(skf7(skf6(U),V),
Derived: 16[0:Res:9.1,4.0] || F(U,skc1) F(skf5(skf6(U)),skc1) -> F(skf5(skf6(U)),skf
Derived: 15[0:Res:9.1,5.0] || F(U,skc1) -> F(U,skc1) F(skf7(skf5(skf6(U)),V),skf5(sk
Derived: 18[0:SpR:13.1,13.1] ||  -> F(skf8(U),skc1)* F(skf8(U),skc1)* equal(U,U).
Derived: 17[0:SpR:13.1,11.1] ||  -> F(skf8(U),skc1)* F(skf8(U),skc1)* F(U,skf8(U))*.
Derived: 22[0:Res:19.1,4.0] || F(U,skc1) -> F(skf8(U),skc1) F(skf5(skf8(U)),skf8(U))
Derived: 21[0:Res:19.1,2.0] ||  -> F(skf8(U),skc1)* equal(U,U).
...
Derived: 1193[2:MRR:1192.0,1192.1,1.0,1180.0] ||  -> .
```

Output of an automatic theorem prover

CONTACT

Uwe Waldmann

RG. 1 Automation of Logic

Phone +49 681 9325-2905

Email uwe@mpi-inf.mpg.de

Dynamic Matching under Preferences

Matching, the grouping of items into disjoint pairs, is a fundamental problem in graph theory, combinatorics, and computer science. A number of important applications for matching arises when participants with preferences strive to form contacts or groups. For example, consider a set of boys and a set of girls, each with a preference over the members of the other sex. If they form couples, what matchings are desirable or likely to evolve? A natural condition is *stability*: There is no pair of boy and girl such that both prefer each other over their current partner [see figure]. Such stable matchings have many applications in uncoordinated job or housing markets, in multi-agent systems, and centralized markets like the NRMP program in the United States, where medical residents are assigned to hospitals.

Matching under preferences is an active research area in computer science and economics, documented by recent Nobel memorial awards in economics for Alvin Roth and Lloyd Shapley for their work on the subject.

Dynamics in matching markets

In uncoordinated markets, there is no authority that dictates a matching. Instead, participants deviate iteratively to preferred pairs if possible. These *matching dynamics* are known to describe a "path to stability", a sequence of matchings that converges to a stable matching.

In large markets, participants often need to learn about possible matches before they can engage in a joint activity. For example, in job markets people often discover possible jobs only via personal contacts or because they work in the same company. To capture this aspect, we studied locally stable matching and dynamics with changing visibility based on a (social) network. A new pair can form only if they *currently know about each other* and prefer to deviate.

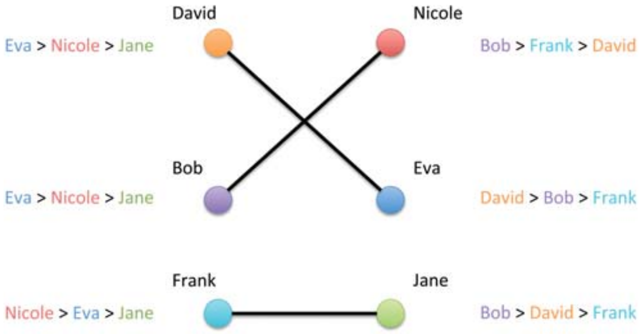

A stable matching with three boys and three girls

We showed that here a path to stability might not exist and dynamics might fail to converge. Interestingly, if each participant has a small memory and repeatedly remembers a random previous partner, convergence can be achieved, in some cases even in polynomial time. Thus, the emergence of stability crucially depends on the means to overcome locality constraints.

Matching in dynamic markets

Even in coordinated markets, dynamics can arise. For example, while a firm can make centralized decisions about jobs and employees, there might be retirements, new positions, current employees seeking promotion, etc. Here, the firm needs to *maintain* a matching under preferences in a dynamic environment. A natural goal is to avoid unnecessary re-assignments and to make only a small number of changes. It turns out that stable matchings cannot be maintained in such a way and might require changing many pairs frequently.

Another stability criterion for matchings is *popularity*. For a given $\alpha \geq 1$, a matching is α-more popular than another if there are α times more participants that prefer the former to the latter. For an α-popular matching, there is no other matching that is α-more popular. Popularity is a global criterion that says a majority of participants prefer the status quo. The *unpopularity factor* α relaxes this idea – for larger α, more unhappy participants are required to label a matching unpopular, and hence, the set of α-popular matchings expands.

We have shown how to maintain α-popular matchings using dynamics: Move to an α-more popular matching whenever possible. For large α, this strategy makes changes only rarely, but the maintained conditions are unpopular. We have provided a tight analysis on how a lower α leads to a larger rate of change. In fact, if every person is interested in at most d possible partners, our strategy maintains $O(d)$-popular matchings at $O(d)$ changes per iteration. :::

CONTACT

Martin Hoefer

DEPT. 1 Algorithms and Complexity

Phone +49 681 9325-1055

Email mhoefer@mpi-inf.mpg.de

Quantifier Elimination – Statements can also be Calculated

The understanding of *computing* on computers has historically evolved in several steps. In the early days of electronic computing equipment, in addition to dealing with large amounts of data, the numerical processing of numbers was in the foreground. The development of computer algebra since the mid 1960s generalized that processing to computation with symbolic expressions:
A simple example is the calculation of $(x-1) \cdot (x+2)+2$ with the result $x^2 + x$. The fact that computational work has indeed been performed here is evident, among other things, because the evaluation of the original expression for a concrete x requires double the amount of arithmetic operations as that of the result. Building on this, it became possible, e.g., to automatically derive real functions and even calculate indefinite integrals automatically.

In mathematics, symbolic expressions typically appear in complex statements that can quantify some of the occurring symbols, as, for example, in the following statement about real numbers: *For all x there is a y, such that* $x^2 + xy + b > 0$ *and* $x + ay^2 + b \leq 0$. What can it possibly mean to “compute” such a statement? Whether our statement is true or not, depends on the selection of a and b. We therefore determine the possible choices of a and b for which the statement is true. The result can be rephrased as a statement that no longer contains any kind of quantification: $a < 0$ *and* $b > 0$. Here again, significant computational work has been performed: based on the result, in contrast to the original formulation, one can determine with minimal effort for concrete choices of a and b whether our considered statement is true or not.

Statements about real numbers that contain arithmetic operations, comparisons, logical operations and quantification, as in our example, can always be computed in such a way that the result does not contain any quantification. In the special case where all symbols in a statement are quantified, the result contains no symbols at all; it is then either $0=0$ (“true”) or $1=0$ (“false”).

Let us consider, e.g., a sequence of real numbers as follows: The first two elements x_1 and x_2 are arbitrary. All others are created according to the rule $x_{n+2} = |x_{n+1}| - x_n$. If we start with, say, 1, 2, then we subsequently obtain 1, –1, 0, 1, 1, 0, –1, 1, 2, and we observe that the sequence repeats after the first nine items. This fact can be formulated as a “for all” statement about $x_1, \ldots, x_{11}$. From this, our method computes $0 = 0$. This way, we automatically prove that for *arbitrary starting values* the sequence will repeat after the first nine items.

If we consider statements about integers instead of real numbers, then there can be no corresponding software; and this is mathematically provable. In this sense, the methods discussed here border on what is mathematically possible. In addition to integers and real numbers, there are also many other domains in the natural sciences being researched at the Institute in terms of such methods. This even includes domains in which symbols not only stand for numbers, but also for functions. Corresponding statements can contain arithmetic and derivatives.

This theoretical research is practically implemented in the software system Redlog, which is globally considered as being a powerful and efficient computational tool by numerous scientists. The computation time for our above proof about the number sequence is only approximately 0.07 seconds with Redlog.

A recent application of Redlog this year was the analysis of an automatic speed control for cars. Thereby, speed is automatically regulated in such a way that a collision with a preceding vehicle is impossible. Redlog computed starting speeds at which the control can be safely transferred to the system.

Quantifier elimination and decision methods are also finding application in the field of the analysis of complex systems in natural sciences, for example electrical networks in physics, reaction systems in chemistry, or gene-regulatory networks in biology. :::

CONTACT

Thomas Sturm
RG. 1 Automation of Logic
Phone +49 681 9325-2920
Email sturm@mpi-inf.mpg.de

INFORMATION SEARCH & DIGITAL KNOWLEDGE

Digital information has changed our society, economy, work in science, and everyday life. Modern search engines deliver useful information for nearly every question, and the Internet has the potential to be the world's most comprehensive knowledge base. However, knowledge structures in the Internet are amorphous, and search engines rarely have precise answers to expert questions for which users would consult reference books and expert literature. A great challenge and research opportunity lies in making the leap from processing raw information to computer-aided intelligent use of digital knowledge.

In parallel to this aim for a quantum leap, we observe a complexity explosion in digital information along several dimensions: quantity, structural variety, digital history, multi-modality, semantic search and text analytics, as well as natural language access to data and knowledge.

■ In addition to billions of web pages, important information sources are: online news streams, blogs, tweets and social networks with several hundred million users, online communities for photos, music, books, health, politics, scientific topics, and last but not least encyclopedias like Wikipedia. The total volume of this data is in the order of exabytes: 10^{18} bytes – more than one million terabyte disks.

■ There is an increasing amount of more expressive data representations: tables in web pages, XML documents, RSS feeds, semantically connected RDF graphs, and much more. The richer structure and heterogeneity of the data increases the complexity of mastering this digital variety.

■ The history of digital information – for example earlier versions of our institute's web pages, which are partly conserved in the Internet Archive – is a potential gold mine for in-depth analyses along the time dimension. Sociologists and political scientists, but also media and market analysts as well as experts in intellectual property, can profit from this.

■ In addition to text-oriented and structured data, we are also experiencing an explosion of multimodal information: billions of people are becoming data producers in the web by sharing images, videos, and sound recordings with the rest of the world. A large fraction of such data arises from social media, where relationships and interactions among users are further assets for search, analyses, and recommendations.

■ Queries on the Internet and in enterprises often refer to entities and their attributes and relationships: people, companies, products, etc. Semantic search is the emerging technology to provide precise and concise answers by harnessing digital knowledge bases. In addition, this opens up great opportunities for advancing text analytics in terms of entities and relations.

■ The proliferation of smartphones entails an increasingly important user need for natural-language question answering and dialog. To this end, computational linguistics, statistical learning and digital background knowledge are essential.

Scientists at our institute address these challenges from various perspectives. The great vision of making the quantum leap towards knowledge search and discovery is pursued in work on automatic knowledge extraction from web sources such as Wikipedia. The Max Planck Institute for Informatics is a world-wide leader in this important research direction. :::

CONTRIBUTIONS

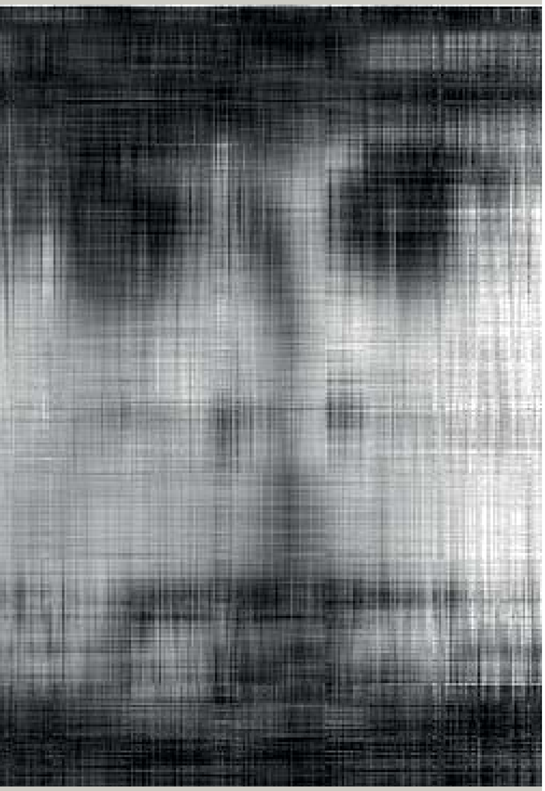

UNDERSTANDING IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH & DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION & VISUALIZATION

OPTIMIZATION

Large-Scale Mining and Analysis of Sequential Patterns

Frequent sequence mining is one of the fundamental problems in data mining with various applications. In speech recognition, for example, frequent sequences can be used to correctly identify the speech input (e.g., “get up at eight o’clock” occurs much more often than “get a potato clock”). Likewise, in customer behavior analysis, frequent sequences can be used to describe common behavior across users. As an example, an online video rental may exploit the knowledge that customers often watch Pulp Fiction shortly after having watched Kill Bill I and II to generate movie recommendations in the future.

Existing methods for frequent sequence mining typically operate on a single machine and, therefore, cannot deal with today’s vast amounts of data. We have developed novel methods for frequent sequence mining based on MapReduce, which can be deployed on clusters of many machines and scale out as the data grows. MapReduce, developed at Google in 2004, constitutes a natural environment for large-scale distributed data processing on clusters of many commodity machines and can handle hardware and software failures transparently. Our methods make use of Apache Hadoop, which is an open-source implementation of MapReduce widely used in industry and constitutes a core building block within big-data initiatives.

In more detail, we consider three variants of the frequent sequence mining problem and have devised tailored approaches for each of them. First, we consider the special case of contiguous sequences, i.e., for a sequence to be frequent, it has to occur sufficiently many times “as it is” in the data. Our second variant allows for gaps of bounded length between items constituting the sequence. This allows us to find the sequence “get up at _ o’clock” (with _ as a placeholder), although it might not frequent when plugging in any concrete number. Lastly, we allow for generalizations of items along a predefined hierarchy. For instance, when working with natural language documents, we may generalize words into their part-of-speech (e.g., article or adjective). Thus, we may find the frequent sequence “get up <IN> eight o’clock,” where <IN> is a placeholder for a preposition such as “at” or “before”.

In our experiments, we have demonstrated that our methods can compute frequent sequences on large real-world datasets. In a collection of more than 50 million documents, our methods can compute frequent sequences consisting of up to 50 words in less than an hour, using a cluster of only 11 commodity machines.

In addition to computing frequent sequences, statistics about frequent sequences can be used for several analysis and exploration tasks. For example, the evolution of the frequency of a sequence over time can be used for studying changes in the use of language or for analyzing popularity of entities (e.g., “How often was Barack Obama mentioned within each year during the last two decades?”). Further, quickly determining frequent sequences that meet certain criteria allows for targeted exploration (e.g., “Which movies are frequently watched after Pulp Fiction?”). The figure illustrates an example of targeted exploration in the field of text analytics using our software “SyntEx”. It shows frequent sequences of the form “<VERB> to <VERB> <NOUN>” sorted according to their frequency in the archive of The New York Times.

All of our methods are publicly available as open-source software. :::

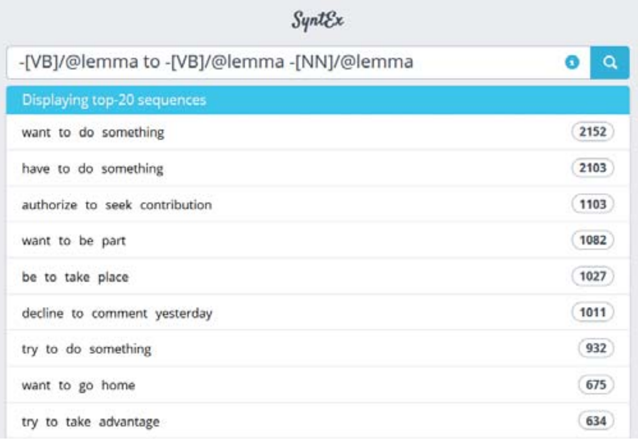

-[VB]/@lemma to -[VB]/@lemma -[NN]/@lemma

Displaying top-20 sequences

want to do something	2152
have to do something	2103
authorize to seek contribution	1103
want to be part	1082
be to take place	1027
decline to comment yesterday	1011
try to do something	932
want to go home	675
try to take advantage	634

Exploring sequences with SyntEx

CONTACT

Kaustubh Beedkar
University of Mannheim, Mannheim, Germany
Email kbeedkar@uni-mannheim.de

Klaus Berberich
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5005
Email kberberi@mpi-inf.mpg.de

Rainer Gemulla
University of Mannheim, Mannheim, Germany
Email rgemulla@uni-mannheim.de

Scalable Analysis of Very Large Datasets

The technical capabilities for data collection as well as the number of available data sources have increased tremendously in recent years, imposing new, unprecedented challenges to information management. The development of the Web 2.0 and social networks, the ubiquity of mobile devices and sensor networks, and the advances in gathering scientific data have contributed significantly to this development. The resulting flood of data is difficult, if not impossible, to manage using traditional tools for data management such as relational databases. On the one hand, the sheer size of the data requires massively-parallel processing, using hundreds or thousands of computers. On the other hand, novel methods for data analysis are necessary to extract useful information from the raw data.

At the Max Planck Institute for Informatics, we develop efficient, highly scalable methods and systems for the statistical analysis of such big datasets. Internet companies such as as Amazon, Google, or Netflix, for example, analyze data about users and their behavior to provide personalized recommendations for products, websites, news, movies, or images. From a user's point a view, such personalized recommendations enable the targeted exploration of potentially interesting items; from a provider's point of view, customer satisfaction and loyalty increase. The American movie rental company Netflix, for instance, allows its more than 20 million customers to rate movies using a 5-star system. The so-obtained ratings are used to individually recommend new, not-yet-seen movies to users. Modern recommender systems are based on an approach called *"collaborative filtering"*, i.e., the behavior of many users and user groups is analyzed in order to create recommendations for each individual user. The key challenges that recommender systems need to solve are (1) the modeling and prediction of user interests, and (2) the creation of interesting recommendations on the basis of these predictions.

Figure 1: Known incomplete matrix

Figure 2: Automatically completed matrix

A successful technique for prediction, which is also used by Netflix in production, models the available movie ratings in the form of an incomplete matrix and subsequently tries to "complete" this matrix. Every row of the matrix corresponds to a user, every column to a movie, and every entry to a rating of the respective movie by the respective user. Figure 1 visualizes such an incomplete matrix; here, black dots correspond to unknown values. We have developed parallel algorithms that can complete incomplete matrices with millions of users, millions of movies, and billions of entries within a couple of minutes. Every entry of the output matrix corresponds to a predicted rating. Figure 2 shows a completed matrix computed using our methods.

One option to create recommendations for each user is to recommend the movies with the highest predicted ratings. This is a non-trivial task, however, because the completed matrix is usually too large to be constructed and processed in reasonable time. We have developed algorithms that efficiently extract only good recommendations for each user, without computing all predicted ratings. Our methods are orders of magnitudes faster than the naive approach of fully computing the completed matrix.

In practice, we may not want to simply recommend the movies with the highest predicted rating. For example, it is important that recommendations are diverse, i.e., that each use's recommendations consist of sufficiently different movies (e.g., different genres or different directors). The availability of licenses or physical media is also crucial: if a user accepts a recommendation, then this recommendation should be delivered as fast as possible. Our techniques allow efficient, high-quality personalized recommendations that respect constraints such as those mentioned above.

Apart from personalized recommendations, our group also works on methods for the interactive exploration of large document collections, for the analysis and automated extraction of knowledge from natural language text, and for pattern mining on sequential data. :::

CONTACT

Rainer Gemulla
University of Mannheim, Mannheim, Germany
Email rgemulla@uni-mannheim.de

Christina Teflioudi
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5029
Email chteflio@mpi-inf.mpg.de

AIDA – Resolving the Name Ambiguity

Have you ever googled for your own name to find out what the Web knows about you? Chances are that you're not the only person with your name. The right Web pages are buried among others, unless you are very famous indeed! This is, of course, not the only scenario where ambiguity makes life difficult. When we read our daily news, most of the names mentioned are ambiguous. As a human being, we deal with ambiguity without thinking, the right meaning seems obvious to us. Only in the most difficult cases – take for example the sentence *"Bush was a US president."* – we notice it. Without further information we cannot know if *"Bush"* means George H. W. Bush (the 41st president) or his son George W. Bush (the 43rd president).

The knowledge about which person (organization, film, or song, etc.) is mentioned where on the Web, or indeed in any given text, is very useful for a multitude of applications. Where search engines could previously only look for a string of characters, they can now actually understand what exactly the user is looking for, and return more precise results. This allows you to specify the real meaning when you are looking for the rock group called *"Bush"* and not a US president. Or imagine you are a researcher looking for differences in the media reception of Bush Sr. and Bush Jr. You can easily get all the articles and find out how often they were mentioned in them without having to even read a single one.

Our AIDA disambiguation system resolves the ambiguity by linking names in text to a canonical entity representation in a knowledge base, for example YAGO. YAGO contains nearly 10 million unique entities, among them nearly 1 million persons, but it also contains locations, organizations, products and events – see the "YAGO – a collection of digital knowledge" article in this report for more details. The disambiguation process consists of different pieces of data and directives, each of which gives additional clues as to which entity is actually referred to in the text. Combining all of them in the right manner identifies the correct entity.

The most important insights for a correct disambiguation are the following. A name probably refers to the most prominent entity. When the name *"Paris"* is found in a text it generally refers to the French capital. There has to be really good contextual evidence to suggests otherwise. Take for example the sentence *"Paris had to steal Helen from her husband, the king of Sparta."* The words in this sentence suggest that it is a person from Greek mythology. This kind of contextual evidence is the second feature for our disambiguation mechanism. Every entity in our knowledge base is associated with a textual description in keyword form that is compared to the surrounding context of the name. The more the context and the description overlap, the better the indication for the entity. However, in some cases, words alone are not enough, especially when the contextual evidence is very limited. To deal with these cases our methodology enforces coherence among the resolved entities, preferring candidates that go well together. In the example sentence *"Paris met Helen."*, Paris and Helen from Troy are more coherent than Paris Hilton and Helen, Georgia, a small U.S. city.

The interplay of the described features can be seen very well in the sentence *"Bush did not handle the aftermath of Katrina in New Orleans very well."* New Orleans is easy to resolve, as the city is very prominently associated with the name. The first name *"Katrina"* is highly ambiguous, but the way it is placed in the context of *"aftermath"* gives a strong indication that it is not a person but indeed a natural disaster. Once both *"Katrina"* and *"New Orleans"* are resolved to the hurricane and the city, it becomes clear that "Bush" refers to Bush Jr., as the hurricane hit during his presidency.

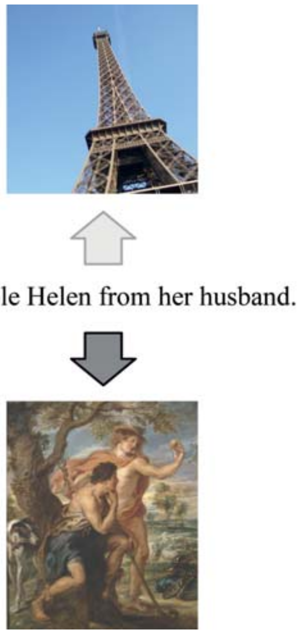

The disambiguation quality of AIDA was tested on a collection of newswire articles, and AIDA achives better results than any other existing disambiguation method. The resulting knowledge of entities in a text allows a more powerful search in these texts and additionally serves as a basis for the extraction of further knowledge about the entities, for example how they relate to each other. :::

CONTACT

Johannes Hoffart
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5004
Email jhoffart@mpi-inf.mpg.de

Gerhard Weikum
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5000
Email weikum@mpi-inf.mpg.de

Mohamed Amir Yosef
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5009
Email mampi@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/yago-naga/aida/

YAGO – A Collection of Digital Knowledge

In recent years, the Internet has developed into a significant source of information. Train schedules, news, even entire encyclopedias are available online. Using search engines, we can query this information, but current search engines have limits. Assume we would like to know which scientists are also active in politics. This question can hardly be formulated in a way that it can be answered by Google. Queries like *"politician scientist"* only return results about opinions on political events. The problem here is that the computers we use today can store a tremendous amount of data, but they are not able to relate this data to a given context, much less understand it. If it were possible to make computers understand data as knowledge, this knowledge could be helpful not only for Internet searches, but also for many other tasks, such as understanding spoken language or the automatic translation of text into multiple languages. This is the goal of the "YAGO-NAGA" project at the Max Planck Institute for Informatics.

Before a computer can process knowledge, it must be stored in a structured fashion. Such a structured knowledge collection is called an ontology. The building blocks of an ontology are entities. An entity is every type of concrete or abstract object: the physicist Albert Einstein, the year 1879, or the Nobel Prize. Entities are connected by relations, for example, Albert Einstein is connected to the year 1879 by the relation "born" [see graph]. We have developed an approach to automatically create such an ontology using the online encyclopedia Wikipedia. Wikipedia contains articles about thousands of personalities, products, and organizations. Each of these articles becomes an entity in our ontology.

There is for example an article about Albert Einstein, so the physicist can be recognized as an entity for the ontology. Each article in Wikipedia is classified into specific categories, the article about Einstein, for example, in the category *"born in 1879."* The keyword *"born"* allows the computer to store the fact that Einstein was born in 1879. Using this approach, we get a very large ontology, in which all of the entities known to Wikipedia have their place. This ontology is called YAGO (Yet Another Great Ontology, http://www.mpi-inf.mpg.de/yago-naga/yago/). At the moment, YAGO contains nearly 10 million entities and about 80 million facts.

More recent extensions of YAGO pay particular attention to the issue of multi-lingualty and the organization of entities and facts in space and time – two dimensions that are highly useful when searching in a knowledge base. As example, the great majority of the approximately 900,000 person-entities in YAGO are anchored in time by their birth and death dates, allowing us to position them in their historical context. For example, one can ask questions about important historical events during the lifetime of a specific president, emperor or pope, or ask if and when someone actually was president.

Most of the approximately 7 million locations in YAGO have geographic coordinates which place them on the Earth's surface. Thus, spatial proximity between two locations can be used as search criterion. An exemplary search using the space and time criteria could be: Which 20th century scientists were awarded a Nobel Prize and were born in the vicinity of Stuttgart. In YAGO one finds, among others, Albert Einstein, because both his lifetime (1879-1955) and his birthplace Ulm (70 kilometers from Stuttgart) are stored in the knowledge base. The multilingual nature of YAGO allow these queries to be placed in any language, for instance asking for "Wissenschaftler" who received a "Nobelpreis". :::

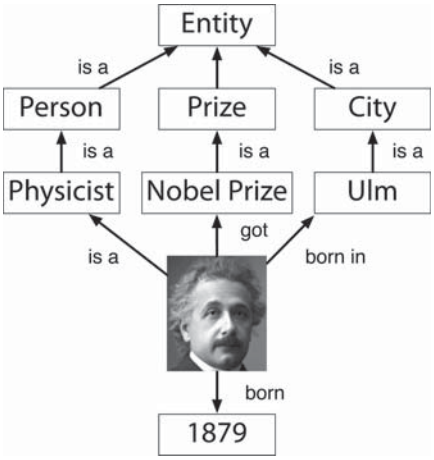

CONTACT

Johannes Hoffart
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5028
Email jhoffart@mpi-inf.mpg.de

Fabian Suchanek
DEPT. 5 Databases and Information Systems
Email fabian@suchanek.name

Gerhard Weikum
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5100
Email weikum@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/yago-naga/yago/

Fast Searches in Large RDF Knowledge Bases

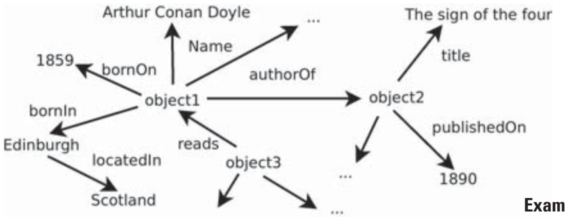

Example RDF graph

Semantic data

Our focus in this work lies in the scalable management of semantic data, that is, data that contains information about relationships between objects. Such relationships are ubiquitous. For example, a book has one or more authors, a protein takes part in specific reactions, and a user of Web 2.0 pages has connections to his friends. These relationships between things, or, more abstractly: entities, together with the entities themselves, form a network or graph structure. Using formal semantics associated with it, the data graphs form semantic networks, i.e., networks which allow for automatic reasoning.

A data format, which was specially designed for graph-structured data, is the Resource Description Framework (RDF). An RDF data collection consists of triplets, each of which corresponds to an edge and an associated node pair in the data graph. A triplet consists in the RDF writing of a subject, a predicate and an object. In the graph, this corresponds with the <predicate> labeled edge from subject to object. There are edges in the example <id1, bornOn, 1859>, which express that the subject <id1> is associated with the name "Arthur Conan Doyle", who was born in 1859.

Search in RDF graphs

Using the above notation, complex relationships can be formulated in a concise manner. The simplicity and flexibility, which is achieved through the graph structure of the data, leads to the fact that RDF searches are relatively expensive. Search queries are normally formulated in a query language such as SPARQL, for example, and describe a pattern that must be searched for in the data. If one searches, for example, for the title of books by Scottish authors, one can describe this with the following triple pattern:

?autor <geborenIn> ?stadt
?stadt <liegtIn> Schottland
?autor <autorVon> ?buch
?buch <Titel> ?titel

The parts beginning with question mark are individual variables whose values must be determined by the database engine. The portion in brackets is the value given as search conditions by the user. In order to answer this query, the database system must find all data triples that match the pattern of the query. Because the data graphs are very large, there can be many millions of candidate triples, making searches very difficult.

The two database engines we have developed for this purpose, RDF-3X *("RDF Triples Express")* and TriAD *("Triple Asynchronous Distributed")*, approach these problems on several levels. First, the data itself is stored appropriately so that individual triple patterns can be efficiently evaluated. Compression and indexing data using search trees enables quick evaluation of any given triple pattern. However, this is not enough; the users are mostly interested in larger interrelationships and thus pose queries with linked triple patterns. Therefore, we map such interrelated queries into execution strategies using algebraic operators and choose the ones that can be expected to have the lowest execution time from the different execution alternatives. The different alternatives are assessed, and then the most efficient alternative is chosen with the assistance of statistical models. For the example query, one must estimate whether there are more authors or more cities in the data set and how this affects the execution time. A careful choice of the execution strategy can often accelerate the query processing by a factor of 10 or more.

Efficiency and Scalability

Based on these ideas for indexing RDF data collections and optimizing queries, we have recently designed the TriAD engine to further scale these techniques to very large RDF collections. TriAD follows, in principle, a classical master-slave, shared-nothing architecture but employs a proprietary, asynchronous communication protocol based on the Message Passing Interface (MPI) in which all slaves can freely communicate with each other. RDF data is partitioned ("sharded") and indexed in multiple permutations at each slave. One of the novel approaches followed in TriAD is to shard triples based on their locality in the original RDF data graph. This helps in minimizing the communication costs and keeping query processing at the local slaves as efficient as in a centralized approach. However, unlike in a centralized setting, we can process a query over the different shards, which are distributed across all compute nodes in parallel. In addition to the main RDF indexes, a compact summary graph of the RDF data graph is maintained at the master node, by which we can additionally prune irrelevant partitions of the main indexes even before the actual processing at the slaves takes place. In our experiments, which we have conducted on a number of both real-world and synthetic datasets, TriAD consistently outperformed comparable approaches. For a large instance of the synthetic LUBM benchmark (1.8 B triples), we could demonstrate a performance improvement over RDF-3X installation of a factor of more than 300 (using a 10-node compute cluster). :::

CONTACT

Sairam Gurajada

DEPT. 5 Databases and Information Systems

Phone +49 681 9325-5026

Email gurajada@mpi-inf.mpg.de

Answering Natural Language Questions Logically

Modern search engines such as Google or Bing cannot answer natural language questions in general and instead return a list of documents in which the user can find answers to his question. Understanding natural language, the question as well as the sources of information, is the main problem for which there is no comprehensive solution. Many linguistic phenomena such as the ambiguity of words and phrases (for example *"bank"*: financial institute or sitting accommodation?) require much data and complex procedures.

In our research, we develop novel techniques for answering *English* questions based on the knowledge base YAGO (see *„YAGO – A Collection of Digital Knowledge"*, page 71). The two main steps are:

1) Analysis of the meaning of the question and
2) Search for an answer in YAGO

The second step has already been well researched in our group so that the focus of the current research lies on the analysis of the meaning of the question. For this purpose we cooperate with the department of computational linguistics at the Saarland University.

Analysis of the meaning of a question

The meaning of a question can be derived from the meaning of its subphrases. First, subphrases are derived from single words of the question and are used to derive larger subphrases. The procedure is repeated until the complete question has been covered and a correct meaning of it parsed. Figure 1 shows an exemplary analysis of the question *"Who saw the man with the telescope?"*. This question is ambiguous and can be interpreted in two different ways:

1) as asking for a person who saw a man through a telescope [figure 1].
2) as asking for a person who saw a man carrying a telescope [figure 2].

Such ambiguities occur very often in the English language and make natural language processing very difficult.

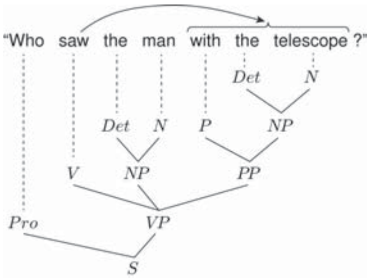

Figure 1

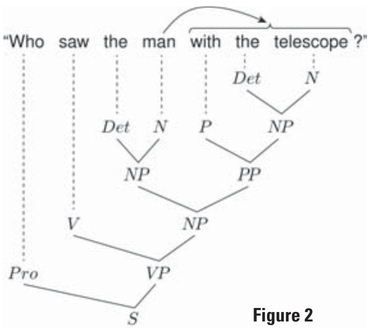

Figure 2

Semantic analysis for disambiguation

Humans can disambiguate such phrases as the above because of their life experience and knowledge about the context, i.e., additional knowledge about a given situation. A computer system does not have any life experience and needs other forms of knowledge additionally to the knowledge about the context with which the true meaning can be recognized. For example,if the computer system knows that only the referenced man from the above question owns a telescope and nobody else (knowledge about the context), then the computer system can refute the first interpretation of the question. With the following general knowledge and knowledge about the English language, the sentence *"The chair saw the table."* can be analyzed unambiguously.

General knowledge (in form of an *ontology*):

- A chair is a piece of furniture.
- A piece of furniture is not a living being.
- A chairman is a person.
- A person is a living being.

Knowledge about the English language:

- "chair" can represent a piece of furniture.
- "The chair" can represent a chairman.
- "saw" can represent a verb which requires the subject to be a living being.

With the combination of both knowledge sources, the computer system can conclude that the phrase *"The chair"* refers to a living being, a chairman, and not a piece of furniture. Finally, knowledge about the context enables the system to resolve the reference and realize which person is being referred to.

We achieve such semantic analyses by means of the *first-order logic* and corresponding methods from the domain of automated reasoning (see *„Automated Deduction"*, page 62). Our focus lies, among other things, on the early and efficient recognition of semantic contradictions. :::

CONTACT

Eugen Denerz
RG. 1 Automation of Logic
Phone +49 681 9325-2914
Email edenerz@mpi-inf.mpg.de

DEANNA – Natural Language Questions over the Web of Data

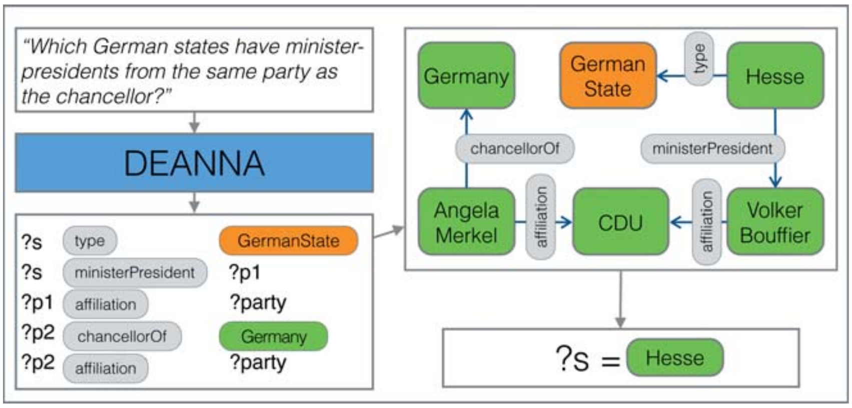

Question answering in DEANNA

When using a search engine, we are often interested in answers as opposed to links in response to our queries. Major search engines have recently started returning crisp answers in the form of entities such as people, places, and movies in response to very simple keyword queries. For example, the answer *Angela Merkel* will be returned for the query *German chancellor*. Answers to such queries come from large Knowledge Graphs that store triples, also known as facts, where a relation connects a pair of entities together (e.g., (Angela Merkel, chancellorOf, Germany)).

Existing efforts, however, leave much to be desired as they do not fully exploit the potential offered by Knowledge Graphs. Complex queries of interest to researchers and journalists can be answered by Knowledge Graphs. For example, one can find answers for the query *"What was the party of Merkel's predecessor?"*, or *"Which German states have minister-presidents from the same party as the chancellor?"*. Such questions cannot be put in keyword query form, as they completely lose their semantics, which is what makes them interesting in the first place. What we need is a system that takes questions asked by humans, maps them to the queries over the Knowledge Graph, and returns answers to the user.

We have developed DEANNA, a framework for question answering over Knowledge Graphs. DEANNA allows users to ask potentially complex questions and obtain answers from a Knowledge Graph. The main task DEANNA has to perform is to understand the question posed by the user with respect to the Knowledge Graph. In this setting, understanding means being able to map phrases in the question to objects in the Knowledge Graph, and then combine these to create a query. Knowledge Graph objects include entities (e.g., *Angela Merkel*), types (e.g., *political party*), relations (e.g., *chancellorOf*). The target query language is composed of triple patterns, where unknown parts of a triple are replaced with variables whose values will be determined from the Knowledge Graph. For example, the question *"Which German states have minister-presidents from the same party as the chancellor?"* translates to the triple pattern query:

?s type GermanState

?p1 presidentMinisterOf ?s

?p1 affiliation ?party

?p2 chancellorOf Germany

?p2 affiliation ?party

where the occurrence of the same variable, such as *?s*, in two triple patterns connects the two.

The main problem in question understanding in our setting is the inherent ambiguity of language. *Merkel* can refer to one of many people, a computer might think *the chancellor* refers to the 1875 Jules Verne novel although it is intended as a relation connecting a person to a country. While such ambiguities seem trivial for a human to resolve, computers need to be given a framework to resolve them, which is what DEANNA provides.

DEANNA's disambiguation framework is based on integer linear programming (ILP) with an objective function to maximize and a set of constraints. Here, the constraints prevent any nonsensical understandings of the user's question. For instance, understanding *the chancellor* as a novel prevents us from being able to connect it to any other objects mentioned in the query, and hence means that this understanding is incorrect. The objective function captures how well the different objects in an understanding of the question go well together. Intuitively, the ILP-based framework returns the most likely understanding of the question (the objective function) that makes sense (the constraints). :::

CONTACT

Mohamed Yahya
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5028
Email myahya@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/departments/databases-and-information-systems/research/yago-naga/deanna/

Open Information Extraction

The great majority of the knowledge that mankind has produced is available only in the form of text, including books, news articles, scientific papers, and web pages. Natural language is written for humans; it can be noisy, ambiguous, opinionated, or difficult to interpret without context. Our work aims to extract and represent this knowledge in a structured form amenable to automated processing by computers.

Text understanding is a difficult task and requires the combination of techniques from the machine learning, logic, linguistics, and data management communities, among others. An ideal – or perhaps idealized – system would be sufficiently powerful to capture the entire set of information represented in a given text collection, regardless of its domain (e.g., biology, history, economics). Language, as the primary representation of knowledge, already provides a systematic way of structuring information, but this structure is often oblivious to computers.

We have developed a system called ClausIE, which translates sentences into a form that is easier to process by computers. ClausIE exploits the grammatical structure of sentences, making use of the fact that statements are often expressed in terms of clauses. A clause is essentially a simple sentence that consists of a set of grammatical units, some obligatory (e.g., subject, verb), some optional (e.g., adverbials). Consider, for example, the sentence *"Albert Einstein discovered the law of the photoelectric effect and won the Nobel Prize."* This sentence contains two clauses: *"Albert Einstein won the Nobel Prize"* and *"Albert Einstein discovered the law of the photoelectric effect."* ClausIE recognizes such clauses and transforms them into so-called propositions, which reveal the structure of the clause. In our example, the propositions are *(Albert Einstein, discovered, law of the photoelectric effect)* and *(Albert Einstein, won, Nobel Prize)*. Propositions are easier to process by computers because they are simple, domain-independent, and structured.

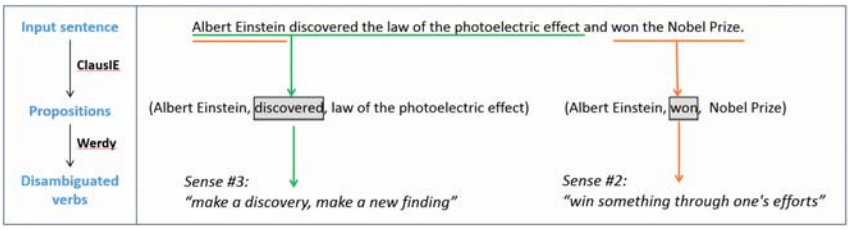

Overview of our open information extraction and disambiguation systems

A proposition does not, however, provide a complete picture of the information present. For example, in the proposition *(Albert Einstein, won, Nobel Prize)*, we would like to understand that *"Albert Einstein"* refers to the German scientist, that *"Nobel Prize"* refers to the famous award, and that *"won"* means *"having been awarded with a prize"*. The task of recognizing named entities has been well-studied in recent years. Less attention has been put to the disambiguation of the verb or verbal phrase, the main element linking the constituents of a proposition. Our Werdy system addresses this gap by automatically disambiguating verbs. The main difficulty of this task is that verbs tend to have many different meanings. For instance, the verb *"win"* has five different meanings in the electronic dictionary WordNet; other common verbs such as *"take"* can have more than 40. To determine the correct sense, Werdy exploits the observation that each verb sense occurs in only a limited set of clauses and only with a limited set of arguments. For example, the sense of *"win"* that refers to obtaining an award requires a subject and an object that can serve as an *"award"* (e.g., the Nobel Prize). By systematically leveraging this knowledge, Werdy is able to boost the precision of existing disambiguation systems.

Figure 1 illustrates our methods on our running example. ClausIE and Werdy are basic building blocks for automatic text understanding. In our ongoing work, we plan to integrate and reason about large sets of propositions in a principled, automated way. :::

CONTACT

Luciano Del Corro
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5029
Email delcorro@mpi-inf.mpg.de

Rainer Gemulla
University of Mannheim, Mannheim, Germany
Email rgemulla@uni-mannheim.de

FERRARI: Quickly Probing Graphs for the Existence of a Relationship

Entities and their connecting relationships are commonly modeled in the form of directed or undirected graphs. Prominent examples include social networks capturing personal or professional relationships among people, biological networks of interactions between biological units such as proteins, or functional dependencies in technical systems, for example call graphs in computer programs. An example for a call graph is shown in the figure. Here, functional dependencies – the relationships – are drawn as arrows between software modules (the entities). Especially over the last decade, we have witnessed substantial growth concerning both the size of individual graph-structured datasets and the number of application domains leveraging the expressive power of the graph data model.

Querying for relationships

The growing size of graphs imposes great challenges with respect to efficient processing. Graphs comprising several millions of entities and relationships can render even the most basic analyses computationally demanding; however, many applications require (near) real-time answers to certain important queries. As an example, consider a graph modeling the functional dependencies between code fragments in a large software project. In this scenario, in order to facilitate fast compilation and testing, it is an important requirement to quickly answer questions of the form: *"Is there a functional dependency of code module A on code module B?"* In graph processing, questions of this form are known under the term *reachability queries*. To answer such a question, it is required to examine the graph for the existence of a direct or indirect connection – a connection over several intermediate entities. In the example graph shown in the figure, an indirect connection (series of arrows) exists between fragments A and B. Traditional approaches for reachability queries are too slow, especially for the case of many sequentially arriving queries – the typical case in most application scenarios.

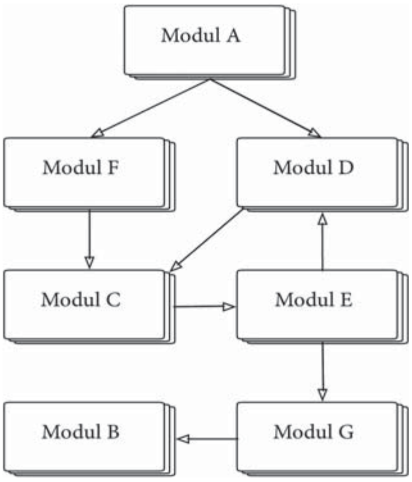

Functional depencies among software modules

The FERRARI index structure

Our research group has proposed FERRARI, an efficient *index structure* to answer questions of this form. The main idea lies in precomputing certain information with regard to the connections between entities. The construction of such an index is a one-time effort, but the resulting information is stored and can be used to speed up the answering process for individual questions later.

FERRARI precomputes and stores information about the indirect connections in the graph. As a result, questions concerning the existence of a certain connection can be answered quickly by combining standard algorithms for the reachability problem with the precomputed information in FERRARI. The algorithms can operate much faster because FERRARI has already computed parts of the answer that are common among many queries and would be costly to obtain over and over again for each newly arriving query.

Since a reachability query answers one of the most fundamental questions in graph processing – whether a connection exists between entities – it is not only used as an important application in its own right, but also as a building block in other graph analytics scenarios. In practice, FERRARI can answer reachability queries very quickly, requiring on average only a few microseconds for processing, over graphs containing many millions of entities and relationships. The time and memory requirements for computing and storing the FERRARI index structure are also lightweight. FERRARI enables the precomputation of the index in the order of a few minutes. Furthermore, any restrictions on the memory size are satisfied, either as specified by the user or mandated by the hardware of the computer. The FERRARI software has been released as an open source project. :::

CONTACT

Stephan Seufert

DEPT. 5 Databases and Information Systems

Phone +49 681 9325-5017

Email sseufert@mpi-inf.mpg.de

STICS – Search and Analysis with Strings, Things, and Cats

“Things, not Strings” has been Google’s motto when introducing the Knowledge Graph and the entity-awareness of its search engine. When you type in the keyword “Klitschko” as a query, Google still returns Web and news pages, but also explicit entities like Wladimir Klitschko and his brother Vitali (including structured attributes like date of birth, profession, and relations to other entities, from the Knowledge Graph). Moreover, while typing, the query auto-completion method suggests the two brothers in entity form with the additional hints that one is an active boxer and the other a politician.

However, the Google approach still has limitations. First, recognizing entities in a keyword query and returning entity results seems to be limited to prominent entities. Unlike the Klitschko example, a query for the Ukrainian pop singer “Iryna Bilyk” does not show any entity suggestions (neither for auto-completion nor in the search results). Second, Google seems to understand only individual entities, but cannot handle sets of entities that are described by a type name or category phrase. For example, queries like “Ukrainian celebrities” or “East European politicians” return only the usual ten blue links: Web pages that match these phrases. The search engine does not understand the user’s intention to obtain lists of people in these categories.

STICS, short for “Searching with Strings, Things, and Cats”, is a novel search engine that extends entity awareness in Web and news searches by tapping into long-tail entities and understanding and expanding phrases that refer to semantic categories. STICS supports users in searching for strings, things, and cats (short for categories) in a seamless and convenient manner. For example, when posing the query “Merkel Ukrainian opposition”, the user is automatically guided, through auto-completion, to the entity ‘Angela_Merkel’ and the category ‘Ukrainian_politicians’, which is automatically expanded into ‘Vitali_Klitschko’, ‘Arseniy_Yatsenyuk’, etc. The search results include texts like “the German chancellor met the Ukrainian opposition leader and former heavy-weight champion”, even if these texts never mention the strings “Angela Merkel” and “Vitali Klitschko”. STICS achieves this by using the named entity recognition and disambiguation system AIDA, which links ambiguous words to entities in YAGO, which in turn contains the entities’ categories. The inner workings of AIDA are detailed in *“AIDA – Resolving the Name Ambiguity”*, page 70.

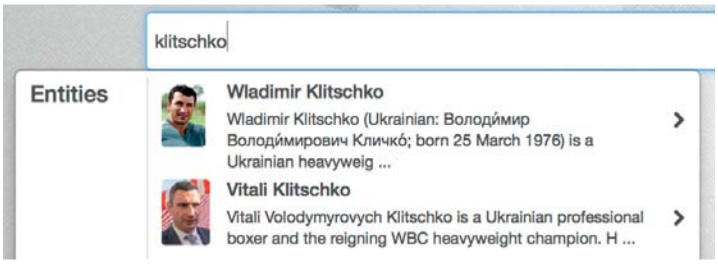

klitschko

Entities

Wladimir Klitschko
Wladimir Klitschko (Ukrainian: Володи́мир Володи́мирович Кличко́; born 25 March 1976) is a Ukrainian heavyweig ...

Vitali Klitschko
Vitali Volodymyrovych Klitschko is a Ukrainian professional boxer and the reigning WBC heavyweight champion. H ...

The same technology can also be used to improve the analysis of large archives. Consider the task of visualizing trends around the recent Ukrainian crisis, which originated from the Maidan, the square in Kiev where thousands of Ukrainians protested in early 2014. A search for “Maidan” quickly reveals that the name is highly ambiguous, as it means ‘square’ not only in Ukrainian, but also in Hindi and Arabic. Thus, simply counting the string “Maidan” will result in a large number of false positives, leading to an imprecise analysis. By specifying the canonicalized entity ‘Maidan_Nezalezhnosti’, not only do we get rid of spurious mentions of other Maidans, but we also find articles where the square is mentioned only by its English name “Independence Square”. Thus, entity-level analytics, as supported by STICS, is the only way to get accurate numbers.

Additionally, as we now have the full potential of a structured knowledge base in the background, further opportunities are opened up. In all semantic knowledge bases, entities are organized in a category hierarchy, e.g., ‘Greenpeace’ is an ‘environmental_organization’, which is in turn a subclass of a general ‘organization’. Using this category hierarchy, we can conduct analyses for entire groups of entities, for example, comparing the presence of ‘environmental_organizations’ and ‘power_companies’ in news from different parts of the world, deriving a picture of how their importance changes over time. :::

CONTACT

Johannes Hoffart
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5004
Email jhoffart@mpi-inf.mpg.de

Dragan Milchevski
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5013
Email dmilchev@mpi-inf.mpg.de

Gerhard Weikum
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5000
Email weikum@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/yago-naga/stics/

Siren – Interactive Redescription Mining

Finding alternative ways of characterizing the same entities is at the root of many scientific problems. For example, a biologist could be interested in characterizing land areas by the species inhabiting them, on the one hand, and by the climate of the area on the other. In another domain, a social scientist might be interested in characterizing voters by their socio-economical status and voting behavior. Traditionally, these questions have been answered by explicit hypothesis testing: The researcher designs a hypothesis, such as 'these species live in the following climates' or 'people in this socio-economical class vote for this party', then he collects data and tests whether the hypothesis is supported by the data.

The advent of data driven science allows us to augment the traditional approach by automatically mining the hypotheses from the data and simultaneously checking their validity. Techniques such as classification – or more generally, supervised learning – have been used to find one characterization when the other is given. The goal of *redescription mining* is to find sets of entities that can be characterized in alternative ways. In the aforementioned biological problem, the goal of redescription mining would be to find a set of species that co-exist in areas that can be succinctly characterized by their climate. For example, we can approximately describe the habitat of the moose (*Alces alces*) in Europe by the following climate: February's maximum temperature is between – 10° and 0° C, July's maximum temperature is between 12° and 25° C, and August's precipitation is between 57 and 136 mm.

Simply finding (or mining) a set of redescriptions is not enough, though. Even a small data set can contain hundreds of redescriptions, and understanding and verifying them can become a daunting task. Moreover, the users of redescription mining algorithms are usually domain experts, and their knowledge of the data and its specific characteristics should be incorporated into the mining process. Our redescription mining software *Siren* helps the domain experts by providing a single interface for mining, editing, and visualizing the redescriptions. Its interactive visualizations and brush-and-link workflow facilitate exploratory redescription mining in a powerful way.

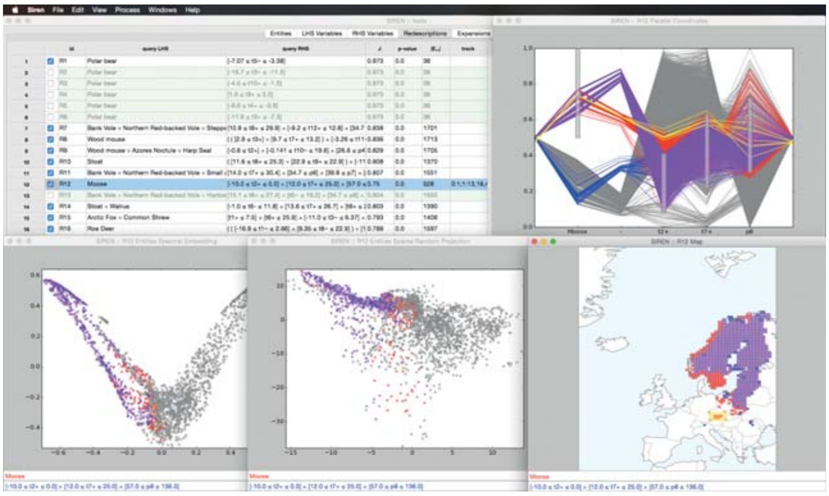

The user interface of Siren

The Figure shows the interface of Siren, where the user has selected the aforementioned redescription explaining the habitat of the moose in Europe. The map in the bottom right corner shows that there is an area in Northern Austria where the moose does live, but where the weather conditions are not as expected (red squares). The user has selected this area from the map by drawing a yellow polygon, and the data points corresponding to the selected area are highlighted in the other visualizations in yellow. The two 2D-projections left to the map do not show the selected area to be clearly different from the rest of the colored markers. The parallel coordinates view displayed above the map, however, explains the difference: the selected areas (yellow lines in the parallel coordinates plot) have a much higher maximum temperature in February than what is expected according to the redescription (second gray box from the left). The user can now decide the appropriate next step based on his/her domain knowledge: Discard the area as an anomaly, try to extend the redescription to correctly handle the area, or reject the redescription altogether.

We are currently working with biologists from Finland, Germany, and the United States to incorporate redescription mining and Siren into their data analysis workflows. We are also exploring other application areas and the potential for commercializing the technology. :::

CONTACT

Pauli Miettinen
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5012
Email pauli.miettinen@mpi-inf.mpg.de
Internet http://siren.mpi-inf.mpg.de

KnowLife: Knowledge Extraction from Medical Texts

Nowadays, the Internet is a rich and diverse repository of medical knowledge. However, since such knowledge is spread over many websites in the World Wide Web, searching for specific information is tedious. Imagine a patient newly diagnosed with diabetes. A typical response would be to search the Internet so as to learn more about the disease. The patient must piece together the big picture, combining information from different sources and identifying untrustworthy ones, which is especially important for online discussion forums. The patient's overall goal is to get a comprehensive overview of complex medical facts – disregarding redundant and false information.

With KnowLife, we present a one-stop alternative for this patient. Know Life is a knowledge base featuring over 500,000 facts automatically extracted from medical texts, where a fact is an entity-relation-entity triple such as *diabetes-hasSymptom-fatigue*. All facts together form a graph, where nodes are medical concepts and edges depict relations between these concepts. Up to now, previous works have extracted knowledge for narrow specialties such as protein-protein interactions and gene-drug relationships. In addition, they have focused on scientific publications as the de facto choice of text. In contrast, the goal of KnowLife is to build a versatile and comprehensive knowledge base: First, KnowLife spans many areas within the life sciences, covering a wide range of entities and relations involving genes, organs, diseases, symptoms, treatments, side effects, and environmental and lifestyle risk factors for diseases. Second, it taps into articles from health portal websites and online discussion forums. Third, to allow for a better interpretation of medical facts, KnowLife enriches them with context information, such as *diabetes-hasSymptom-fatigue during initial onset*.

Our fact extraction approach applies statistics-based pattern matching to recognize fact candidates and incorporates semantic assets to validate those candidates. The pattern matching module identifies word sequences in our input text that express a relation with high confidence. For example, the sequence *characteristic feature* expresses the relation (*hasSymptom*). To avoid contradictory information and prune out inconsistent fact candidates, we use rules based on the semantics of the facts and relations. The rules express constraints that the extracted facts must satisfy. One type of rule concerns entity types. For example, only a drug can have a side effect, and the side effect must be some symptom. Another type of rule concerns mutual exclusion between facts. For example, a drug that treats a symptom cannot have that symptom as a side effect at the same time. After applying the rules, the remaining facts are highly accurate at 93%.

To showcase the usefulness of KnowLife, we have implemented a web portal [http://knowlife.mpi-inf.mpg.de] that supports a number of use cases: The portal generates for every concept in our knowledge base an infobox that summarizes all the important information about the concept. Furthermore, a user can provide a webpage, and the portal will annotate on the fly any entities and facts found on that webpage and connect them to our knowledge base. This allows a user, one such as our diabetes patient, to quickly gain an overview of the various medical concepts and their relationships to each other. An expert user can read the text sources supporting the facts and form opinions based on the authority and recency of the text sources. For example, the expert will be referred to relevant clinical studies and scientific publications. Finally, the portal also provides an intuitive way to explore our knowledge base. For instance, a user can search for drug-drug interactions by providing the drug names and then explore the links between the drugs and related diseases and side effects. :::

CONTACT

Patrick Ernst
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5006
Email pernst@mpi-inf.mpg.de

Amy Siu
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5026
Email siu@mpi-inf.mpg.de

Gerhard Weikum
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5000
Email weikum@mpi-inf.mpg.de
Internet http://knowlife.mpi-inf.mpg.de

MULTIMODAL INFORMATION & VISUALIZATION

Pictures are the fastest menas to access to human consciousness. Therefore, algorithms for appropriate visualization of digital information on various display types play a significant role in informatics. Artificial and natural worlds must be presented in an ever more realistic and fast manner: in flight simulators, surgical operation planning systems, computer games or for the illustration of large amounts of data in the natural and engineering sciences. Graphical content is just one aspect of a variety of multimodal data types that are nowadays captured and generated by computer and robotic systems. Visual data processing is thus increasingly intertwined with the simultaneous processing of other data modalities. This will pave the way for entirely new approaches of man-machine interaction.

The basis of qualitatively high-value, computer-generated images are accurate scene models. At the Max Planck Institute for Informatics, we are investigating automated methods for the reconstruction of models of dynamic scenes. This is gaining more and more importance and finds many applications, for instance in computer animation, virtual reality, and in the areas of 3D TV and 3D video (see also, research reports from the area *"Understanding Images and Videos"*).

Methods to efficiently create, modify, and simulate deformable 3D models are at the heart of many visualization and interaction approaches. We research the algorithms to enable this. Biomechanical simulation and motion estimation approaches are also the precondition for developing new efficient and ergonomically optimized human-computer-interaction paradigms at our institute.

New 2D and 3D printing methods enable very quickly creation of real versions of virtual 3D models, and enable printing of new flexible sensors and displays. The algorithmic foundations of this bridge to reality are laid at the Max Planck Institute for Informatics.

In order to show lifelike models of virtual worlds on various displays (ranging from normal screens to immersive head-mounted displays), we have also developed new simulation methods for light propagation in scenes, so-called global illumination. For this, we focus on the development of real-time algorithms.

The images created through image synthesis or image capture typically have a brightness range (dynamic range) that matches the dynamic range of the real world. We research algorithms for the image processing of these so-called High Dynamic Range images (HDR) as well as methods for their portrayal on standard displays. We also research perceptually-based approaches to enhance the display quality of low and high dynamic range imagery on new types of stereoscopic displays.

Photographs and videos also represent forms of visual media. These are both aesthetically appealing and important tools for data analysis. We are therefore working on new optical systems for cameras and turn the classical camera into a calculation instrument that can extract much more information from individual images than pure light intensity, such as 3D geometry, for example.

Rendering approaches also play an important role in the visualization of complex high-dimensional data. Such data frequently arises, for example, as an output of weather simulations, flow simulations, or the analysis of large scale studies in human society. At the Max Planck Institute for Informatics, we develop new approaches to visualize this complex data, such that correlations and important features can be effectively deduced. :::

CONTRIBUTIONS

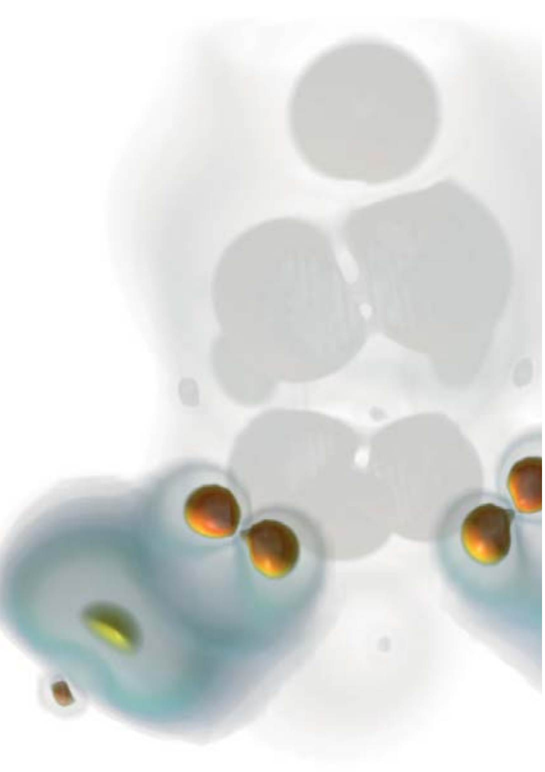

Stereo 3D and HDR Imaging: Display Quality Measurement and Enhancement

We understand human visual perception as a final and mandatory component in a visual processing pipeline. An accurate model of human visual perception is therefore crucial for improving the visual pipeline as a whole. While psychology provides a range of theoretic models of human perception, they are often not applicable for two reasons. First, too many simplifying assumptions are made about the stimuli. Second, the models are often only passive descriptions of findings. The challenge in our work is to deliver concrete models and algorithms that can predict human visual perception and be computationally inverted, for example, using optimization techniques, to improve the content.

Towards improving the stereo 3D experience

Drawing from our past experience in the image quality evaluation for regular 2D videos, we have developed techniques for comparing stereoscopic 3D content. To this end, we have conducted a number of psychophysical experiments that investigate the performance of the human visual system in perceiving depth differences between simple stimuli. Based on these studies, we have devised methods that can predict perceived differences between natural images. Such metrics find a number of applications in compression and content manipulation techniques that adjust stereoscopic content to different devices and people so that both comfort and quality are optimal. More recently, we have employed a stereo 3D eye-tracker to investigate human performance in adjusting to different depths on a stereoscopic screen. This has resulted in a model that predicts the time a human observer needs to transition between different depth levels. We have demonstrated that such a model is crucial for editing movie cuts where rapid depth changes are often unavoidable. Our technique allows us to improve subjects' task performance and provide content that is easier to view.

Overcoming physical capabilities of display devices

Apart from describing human perception, our models can be used to improve the user experience when depicting content on existing display devices. For example, a model of human retinal image integration, eye movement, and tempo-spatial information integration can be used to display a high-resolution image on a low-resolution screen such that the sensation is as close as possible to the original. The frame rate has a large impact on the perception of motion in a video.Recently, we have developed a software technique for manipulating the apparent frame rate of an image sequence without changing its actual display rate. Since we can emulate continuously varying frame rates, both in the temporal and the spatial domain, our technique enables fine control over the appearance of moving objects in the input video.

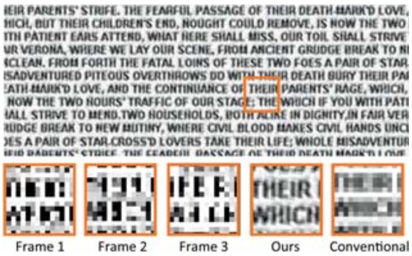

Apparent resolution enhancement for moving images: We optimize individual frames *(1-3, from bottom left)* so that they appear more detailed *(ours, frame 4)* on a human retina than conventional solutions *(bottom right).*

The possibility to combine other perceptual models of high contrast (HDR), glare effects, brightness gradients, and colors makes the building of a more and more complete computational model of human perception conceivable. This provides us with an improved way of depicting content for all users on all devices. :::

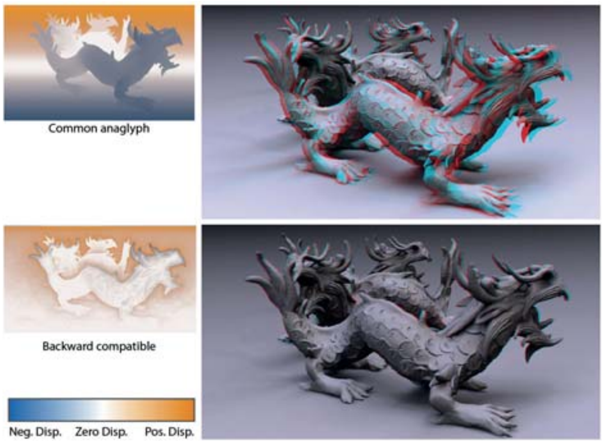

A conventional anaglyph stereo image *(top)* looks unpleasant when observed without proper anaglyph glasses, due to color artifacts. Using our backward compatible solution *(bottom)*, we can reduce the artifacts and, using anaglyph glasses, still provide a sense of depth when viewing the image.

CONTACT

Piotr Didyk
DEPT. 4 **Computer Graphics**
Phone +49 681 302 70756
Email pdidyk@mmci.uni-saarland.de

Karol Myszkowski
DEPT. 4 **Computer Graphics**
Phone +49 681 9325-4029
Email karol@mpi-inf.mpg.de

Advanced Real-time Rendering

The creation of photorealistic images using computers is a basic technology required in many visual media. Originating from special effects, nowadays, the usage of computers to render films is common. However, the required realism can only be achieved using high computational effort where, typically, the synthesis of a single frame of motion picture can take several hours. At the same time, real-time-rendering has become part of our everyday life: In computer games, geo-visualization on a cell phone or in interactive kitchen planners on desktop computers, the required images are produced instantaneously from a user's input. In order to achieve this performance, severe simplifications had to be made, which led to the development of highly specialized graphics hardware (GPUs). In our work, we aim to fill the gap between highly realistic offline-rendering and fast interactive rendering.

A typical challenge in interactive rendering are participating media, i.e. materials that do not have a hard, opaque surface, but appear transparent instead, as light scatters inside them. Classical examples are clouds and fog, but a large class of other objects such as human skin also shows translucency. Most previous rendering approaches are not able to render such effects in real time, and if, then only with certain severe limitaions, such as holding the geometry of the object fixed or assuming the scattering to be isotropic. We have proposed different approaches to overcome these limitations [figure 1]. The first extends classic photon mapping to handle slowly moving clouds. The second is based on a decomposition of the volumetric objects into its principal directions of light propagation, after which scattering becomes a diffusion process that can be efficiently computed. The third is simply based on blurring the image in a physically principled way, allowing light scattering with relatively low computational demand within a few milliseconds.


Figure 1: Real-time rendering of clouds using our method.

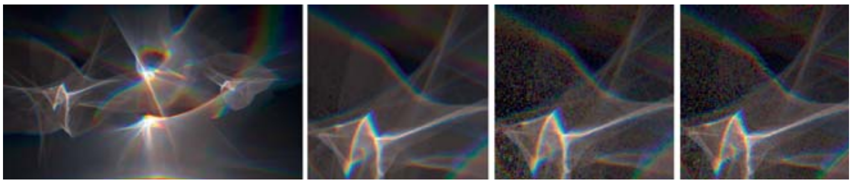
Figure 2: Real-time rendering of dispersive caustics using spectral ray differentials (first image). The three left images show a zoom into a result of our method and how the two competing approaches result in noise or banding, which we remove using very little additional computation.

A particular challenge in rendering are dispersive materials such as in a glass prism. Here, light does not follow a single ray, but is split up in space depending on the wavelength. This introduces another dimension that can be hard to sample properly, often leading to distracting visual noise. In our spectral ray differential work, we have derived formulas to describe the spatial change of a diffracted ray according to wavelength [figure 2]. Put simply, this is a local approximation to the shape of a rainbow. Using this information, we can extrapolate spectral information inside the image such that noise from diffraction is filtered out. The resulting images have no noise but a slight blurring bias, which is often visually less disturbing. We have further shown how a refined computation can guarantee even this bias to vanish completely after sufficient compute time.

Modern graphics hardware is best at performing many tasks in parallel that are simple and independent. A particular instance of such a task is an image filter. Therefore, we have continued to extend our stream of work that uses GPU image filtering to perform advanced interactive rendering (screen space shading). We have shown that scattering and diffraction of light can be understood as image filters and how such image filters can be efficiently implemented in graphics hardware. Furthermore, we have improved the general quality of screen space shading by introducing the concept of deep screen space, a point cloud representation of a scene that has an output-sensitize sampling, similar to a common framebuffer, but does not inherit its limited support for occlusion or grazing angles.

Finally, distributing rendering workload to a computationally powerful rendering server has recently received growing attention. However, the visual quality might be reduced for slow connections, for example due to a low frame rate when bandwidth drops or latency increases. In such conditions, most remote rendering software extrapolates in-between frames in order to keep an almost-constant frame rate. This is of particular importance in head-mounted displays where a low frame-rate can lead to cyber sickness. The currently-used extrapolation however tacitly assumes opaque diffuse surface. In our work, we have extended this extrapolation to handle refractive and reflective surfaces as well, leading to better visual quality in combination with a high frame rate. Applications include streaming visualization of specular objects, such as a virtual car prototype. :::

CONTACT
Tobias Ritschel
DEPT. 4 Computer Graphics
Phone +49 681 9325-4157
Email ritschel@mpi-inf.mpg.de

Feature-based Visualization

Structures repeat in both nature and engineering. An example is the symmetric arrangement of the atoms in molecules such as Benzene. A prime example for a periodic process is the combustion in a car engine where the gas concentration in a cylinder exhibits repeating patterns.

Structures repeat with variations. Sometimes these variations are clearly noticeable. For example, the weather patterns over the Atlantic Ocean are comparable every summer, but not quite the same. Similarly, in engineering, the flow fields around a smaller and a larger car are not quite the same, but they exhibit significant similarity.

One of the research challenges in visualization is to enable domain scientists to detect, quantify and compare structures in their data sets that are similar to each other.

What is a structure?

Our work focuses on comparing data sets from numerical simulations or physical measurements. For such data, the term “structure” needs to be mathematically well-defined. This is a matter of the underlying application, but often it refers to an arrangement of geometric objects such as lines or surfaces. The structures in scalar fields such as temperature or pressure are often described using the topology of the isocontours. This generic description of structures has proven useful in a large number of applications. This approach encodes the structures of the scalar field in a tree, where the size of a branch corresponds to the importance of the structure in the data. The structures in vector fields such as fluid flows or electromagnetic fields can be described using the trajectories of particles. A trajectory is the line consisting of all point locations visited by the particle over time.

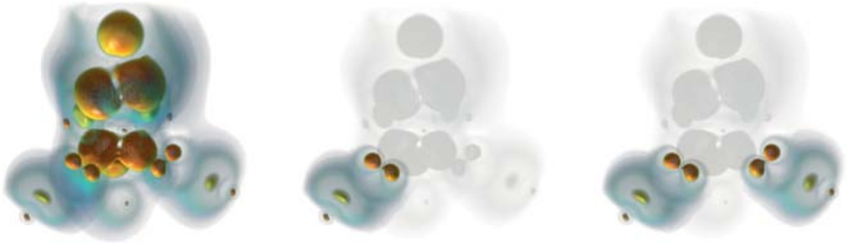

Figure 1: Similar structures in the *neghip* data set *(left)* are found by selecting a structure *(middle)* and then our algorithm finds the most similar ones in the entire volume *(right).*

How can structures be compared?

Comparing two structures requires a function that returns the difference between them. This is not necessarily simple for trees and particle trajectories. We have developed new methods in this field. Our method for comparing trees exploits the hierachy in the tree to drastically speed-up the comparison and to make it more robust against noise at the same time. As a result, we know the pairwise differences of all structures in the tree. Figure 1 shows how this works in practice.

The difficulty in comparing particle trajectories lies in the fact that particles exhibit different flow patterns at different times. As an example, one particle may first be part of a vortex and then part of a hyperbolic region, whereas another particle may exhibit only the vortex. Hence, the trajectories are only partly similar. We deal with this by segmenting trajectories into meaningful parts. This is the basis for both focusing on a local structure and detecting its similar occurrences. Figure 2 shows an example.

Future work

A notion of structural similarity gives rise to many interesting future applications. For example, one could compare different data sets with each other. This is particularly interesting for comparing data sets from different modalities such as numerical simulation and physical measurement. Another possibility would be to identify outliers in a data set, which are characterized by *not* being similar to other structures. :::

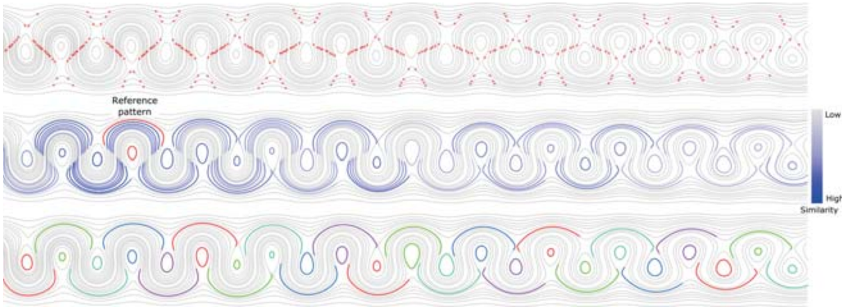

Figure 2: The trajectories at the top have been segmented into meaningful parts to facilitate comparison. The bottom image highlights all trajectories similar to the reference pattern.

CONTACT

Tino Weinkauf

DEPT. 4 Computer Graphics

Phone +49 681 9325-4020

Email weinkauf@mpi-inf.mpg.de

Internet http://feature.mpi-inf.mpg.de/

Towards a Visual Turing Test:
From Perception over Representation to Deduction

As vision techniques like segmentation and object recognition begin to mature, there has been an increasing interest in broadening the scope of research to full scene understanding. What do we mean by the “understanding” of a scene, and how do we measure the degree of “understanding”? Today, “understanding” scenes most often refers to a correct labeling of pixels, regions, or bounding boxes in terms of semantic annotations. We have developed a task that can measure progress in a more holistic understanding of visual scenes and that also extends to more complex notions of attributes and spatial concepts.

Equally strong progress has been made on natural language understanding. For example, methods have been proposed that can learn to answer questions solely from question-answer pairs training data. These methods operate on a set of facts given to the system, which is referred to as a “world”. Based on this knowledge, the answer is inferred by marginalizing over multiple interpretations of the question.

We unite these two research directions by addressing a question answering task based on real-world images. This challenge formulates a holistic task that tests a whole chain, from perception of image content and understanding natural language questions to representing the acquired information and deducing answers to the questions.

Benchmark dataset for answering questions on images

We are lacking a substantial dataset that serves as a benchmark for question answering on real-world images. Such a task relates to the “AI-dream” of building a Turing Test for vision that captures human comprehension of language and images. While we are still not ready to address completely unconstrained settings, we argue that a question-answering task on complex indoor scenes is a timely step in this direction. We introduce a novel dataset of more than 12,000 question-answer pairs produced by humans on RGBD images [see figure]. In contrast to comparable efforts in the literature, our approach requires a deep understanding of the visual scene and the asked questions while maintaining a tractable annotation effort and having an automatic evaluation protocol lends to easy evaluation and analysis.

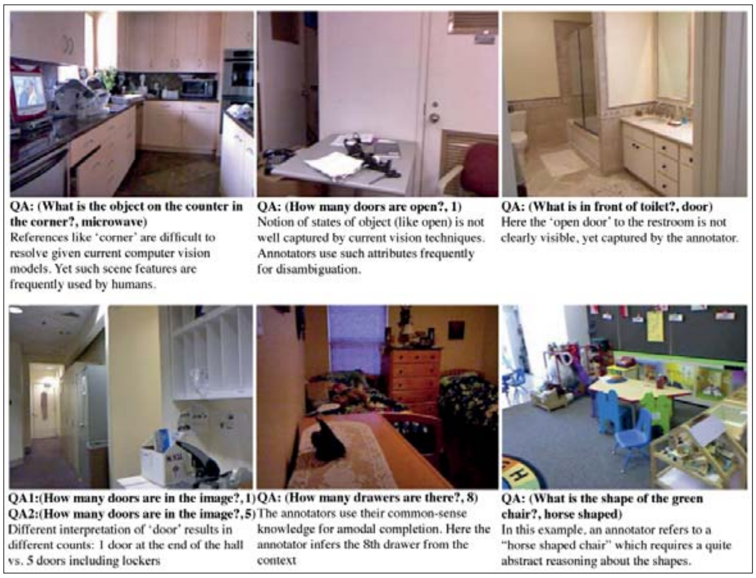

QA: (What is the object on the counter in the corner?, microwave)
References like ‘corner’ are difficult to resolve given current computer vision models. Yet such scene features are frequently used by humans.

QA: (How many doors are open?, 1)
Notion of states of object (like open) is not well captured by current vision techniques. Annotators use such attributes frequently for disambiguation.

QA: (What is in front of toilet?, door)
Here the ‘open door’ to the restroom is not clearly visible, yet captured by the annotator.

QA1:(How many doors are in the image?, 1)
QA2:(How many doors are in the image?, 5)
Different interpretation of ‘door’ results in different counts: 1 door at the end of the hall vs. 5 doors including lockers

QA: (How many drawers are there?, 8)
The annotators use their common-sense knowledge for amodal completion. Here the annotator infers the 8th drawer from the context

QA: (What is the shape of the green chair?, horse shaped)
In this example, an annotator refers to a “horse shaped chair” which requires a quite abstract reasoning about the shapes.

Combining symbolic with uncertain visual facts in order to answer questions on images

While previous approaches on question answering in the domain of natural language processing are mostly based on symbolic reasoning and traditional AI techniques, our approach has to cope with “uncertain facts” that originate from the scene analysis by computer vision methods. We propose a Bayesian formulation that marginalizes over multiple possible “worlds” that correspond to different interpretations of the scene in order to be able to process the probabilistic output of state-of-the-art scene segmentation algorithms.

In this manner, we make a step towards bridging the gap between symbolic reasoning and the uncertain output of computer vision algorithms. While human performance is at about 50 % on this challenge, our best algorithms achieve about 13 %. We have established a benchmark and an evaluation scheme in order to make systematic progress, but much work on automatic image and language understanding lies ahead of us before we will be able to match human performance. :::

CONTACT

Mario Fritz
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-1204
Email mfritz@mpi-inf.mpg.de
Internet http://www.d2.mpi-inf.mpg.de/visual-turing-challenge

Physically-based Geometry Processing and Animation

Creating digital geometric content is an important task for various applications in areas such as digital manufacturing, architecture, computer animation, and virtual reality. Modern acquisition technologies, like 3D-scanning, facilitate the accurate digitalization of detailed real-world objects, and modern fabrication technologies, like 3D printing, allow the automated realization of physical objects from digital designs. Between acquisition and fabrication lies the field of geometry processing that concerns the representation, analysis, manipulation, and optimization of digital geometric data.

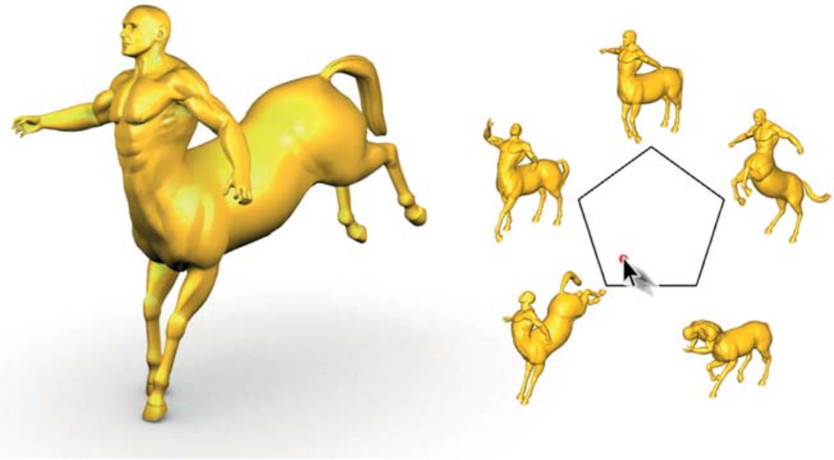

Our tool allows to interactively explore the nonlinear space of shapes spanned by a set of example shapes.

Physically-based geometry processing

Triggered by the needs in the applications, physically-based modeling is becoming more and more important for geometry processing. We are using concepts from continuum mechanics to devise methods and algorithms for geometry processing. Examples are our schemes for deformation-based shape editing and for synthesizing new shapes from examples. These processes are modeled by non-linear optimization problems that the user can control with a small number of parameters. Since the shapes are highly resolved, this leads to a fundamental problem. On the one hand, high-dimensional non-linear optimization problems have to be solved, and, on the other hand, users expect the shape processing tools to provide instant feedback. We have devised algorithms that allow for the efficient computation of approximate solutions of the complex optimization problems. The principle is to carefully construct a low-dimensional approximation of the complex problem in a preprocess and to solve only the reduced problem in the interactive phase.

In particular, we want the computational cost for solving the reduced problems to be independent of the resolution of the shapes to be processed. A highlight in this area is our scheme for nonlinear shape interpolation. It is the first algorithm that allows for computing interpolated shapes of a set of larger meshes in real-time. In our experiments, we obtain rates of 20–100 interpolated shapes per second, even for meshes with 500 k vertices per input shape.

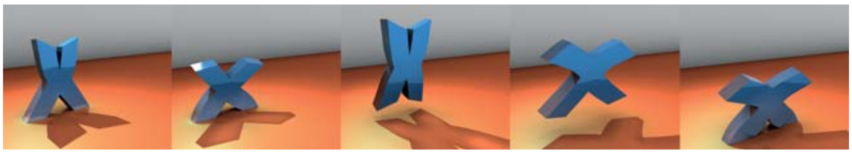

Snapshots of an animation of a soft-body character generated with our tool.

Physically-based animation

In addition to the processing of static shapes, we consider the problem of creating dynamic shapes. In particular, we consider the problem of generating motions of objects and characters that are physically plausible and follow an animator's intent. This is a key problem in computer animation. With traditional computer animation techniques, like keyframing, this is difficult since many degrees of freedom must be determined and secondary motion effects have to be modeled by hand. Physical simulation can be of great help for creating realistic and detailed animations, but one needs to determine the forces and physical parameters that produce the effect an animator wants to create. Space-time optimization, in which the keyframes serve as constraints in space-time, allows to combine the realism provided by physical simulation with the control over an animation offered by keyframing. The results are planned motions showing effects like squash-and-stretch, timing, or anticipation, which are desired by animators. We have introduced a scheme that allows for generating interesting and realistic-looking motion with rich secondary motion from minimal input. An animator only needs to generate partial (as opposed to full) keyframes and can thereby generate complex motion from rough sketches. :::

CONTACT

Klaus Hildebrandt

DEPT. 4 Computer Graphics

Phone +49 681 9325-4012

Email hildebrandt@mpi-inf.mpg.de

Eye-Based Human-Computer Interaction

Our eyes are a compelling input modality for interacting with computing systems. In contrast to other and more well-established modalities, such as touch or speech, our eyes are involved in nearly everything that we do, anywhere and at any point in time – even while we are sleeping. They also naturally indicate what we are interested in and attending to, and their movements are driven by the situation as well as by our goals, intentions, activities, and personality. Finally, eye movements are also influenced by a range of mental illnesses and medical conditions such as alzheimers, autism, or dementia. Taken together, these characteristics underline the significant information content available in our eyes and the considerable potential of using them for implicit human-computer interaction. For example, a wearable assistant could monitor the visual behavior of an elderly person over a long period of time and provide the means to detect the onset of dementia, therefore increasing the chances of an earlier and more appropriate treatment. In addition, our eyes are the fastest moving external part of the human body and can be repositioned effortlessly and with high precision. This makes the eyes similarly promising for natural, fast, and spontaneous explicit human-computer interaction, i.e., interactions in which users consciously trigger commands to a computer using their eyes.

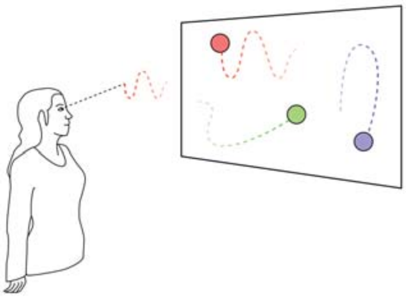

Our method using smooth pursuit matches the user's eye movement with the movement of on-screen objects.

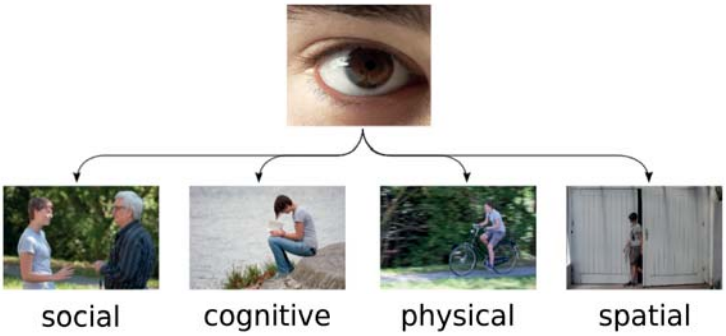

Our method infers contextual cues about different aspects of what we do, by analyzing eye movement patterns over time *(from left to right)*: social *(interacting with somebody vs. no interaction)*, cognitive *(concentrated work vs. leisure)*, physical *(physically active vs. not active)*, and spatial *(inside vs. outside a building)*

Despite considerable advances over the last decade in tracking and analyzing eye movements, previous work on eye-based human-computer interaction has mainly developed the use of the eyes in controlled laboratory settings that involved a single user, a single device, and WIMP-style (windows, icons, mouse, pointer) interactions. Previous work has also not explored the full range of information available in users' everyday visual behaviour, as described above.

Together with our collaborators, we strive to use the information contained in visual behaviour in all explicit and implicit interactions that users perform with computing systems throughout the day. We envision a future in which our eyes universally enable, enhance, and support such interactions and in which eye-based interfaces fully exploit the wealth of information contained in visual behaviour. In that future, the eyes will no longer remain a niche input modality for special-interest groups, but the analysis of visual behaviour and attention will move centre stage in all interactions between humans and computers.

We work towards this vision by
1) developing sensing systems that robustly and accurately capture visual behaviour in ever-changing conditions,
2) developing computational methods for automatic analysis and modeling that are able to cope with the large variability and person-specific characteristics in human visual behaviour, and
3) using the information extracted from such behaviour to develop novel human-computer interfaces that are highly interactive, multimodal, and modeled after natural human-to-human interactions.

For example, we have developed new methods to automatically recognize high-level information about users from their visual behaviour, such as periods of time during which they were physically active, interacting with others, or doing concentrated work. In other work, we have introduced *smooth pursuit eye movements* (the movements that we perform when following moving objects, such as a car passing by in front of us) for human-computer interaction. We have demonstrated that smooth pursuit movements can provide spontaneous, natural, and robust means of interacting with ambient displays using gaze. :::

CONTACT

Andreas Bulling
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-2128
Email bulling@mpi-inf.mpg.de

Optimizing User Interfaces using Biomechanical Simulation

The design of user interfaces has become an area of prime academic and industrial interest. Two groups at the institute have pioneered data-driven solutions to a fundamental challenge: the organization of inputs and outputs in a biomechanically efficient way. Interfaces constrain the postures and movements of the human body in unique ways, and these may or may not be ergonomic and efficient.

We developed a new data-driven approach that leverages biomechanical simulation. First, optical motion capture is used to record the 3D movements of a user when interacting with a particular user interface. Second, this data steers a biomechanical simulation of the user. The result is a data set containing information about the ergonomics (e.g., muscular loads) and performance (e.g., speed and accuracy) of the captured movements. This combined output is a very rich, multidimensional description of human movement, encompassing more than 400 variables. The figure shows our computational pipeline.

Several novel approaches have been explored to leverage this complex data. We identified four types of problems:
(1) Understanding basic human factors;
(2) Optimizing design parameters;
(3) Comparative studies of different designs which aim to solve the same task.
(4) Summarization of data for design decisions.

Muscle activation clustering

We developed a novel summarization of the user performance and major muscle activations for 3D pointing tasks. The basis is a large motion data set in which we captured a large variety of pointing movements that cover the entire egocentric space of a user. A subsequent clustering of the muscle activations reveals regions with different performance and ergonomic costs. This can be used to inform the design choices of user interface designers. The work was published in a top journal of the field, ACM TOCHI.

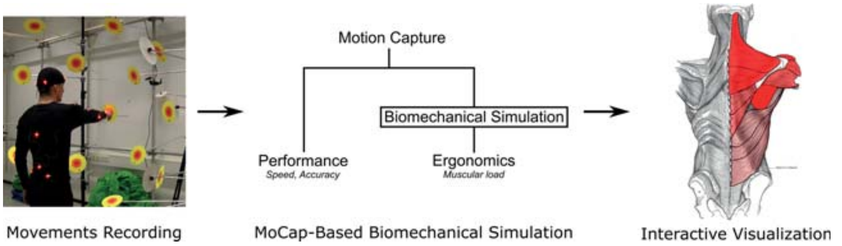

Computational pipeline from optical motion capture and biomechanical simulation to interactive visualization of the multi-dimensional data set.

We applied this approach to touchscreen interfaces. Users performed pointing tasks on five common touch surface types. The results illuminate the strengths and weaknesses of each surface as well as trade-offs among them. They show several trade-offs in terms of throughput, postures, and the expected fatigue in sustained interaction. The paper received a *Best Paper Honorable Mention* at the top HCI conference ACM CHI'15.

Interactive visualization

However, existing visualization tools do not sufficiently support designers and researchers in analyzing such data sets. We identified two shortcomings. First, appropriate visual encodings are missing particularly for the biomechanical aspects of the data. Second, the physical setup of the user interface cannot be incorporated explicitly into existing tools. We developed the versatile visualization tool *MovExp* that supports the evaluation of user interfaces using such data sets. In particular, it can be easily adapted to include the physical setup that is being evaluated, and visualize the data on top of it. Furthermore, it provides a variety of visual encodings to communicate muscular loads, movement directions, and other specifics of studies that employ motion capture and biomechanical simulation. MovExp was published in the top visualization conference IEEE VIS'14.

Optimal gestural interaction

We continued with a study of an emerging input method enabled by hand tracking: input by free motion of fingers. The method is expressive, potentially fast, and usable across many settings as it does not insist on physical contact or visual feedback. The goal was to inform the design of high-performance input methods by providing detailed analysis of the performance and anatomical characteristics of finger motion. We conducted an experiment using a commercially available sensor to report on the speed, accuracy, individuation, movement ranges, and individual differences of each finger. Findings have shown differences of up to 50% in movement times and provide indices quantifying the individuation of single fingers. They were applied to text entry by the computational optimization of multi-finger gestures in mid-air. We defined a novel objective function that considers performance, anatomical factors, and learnability. This work was published as a full paper in CHI'15. :::

CONTACT

Antti Oulasvirta
Cluster of Excellence on Multimodal Computing and Interaction
Phone +358 50 3841561
Email antti.oulasvirta@gmail.com

Tino Weinkauf
DEPT. 4 Computer Graphics
Phone +49 681 9325-4020
Email weinkauf@mpi-inf.mpg.de
Internet http://feature.mpi-inf.mpg.de/

Digital Fabrication of Flexible Displays and Touch Sensors

Graphical user interfaces offer a high degree of customizability. Graphical contents can be digitally designed and programmed, in order to tailor to the user and to the context of use. In contrast, the screens, which display these graphical contents, barely offer any possibility for customization: today's displays are typically rectangular, planar and rigid; they are relatively thick and available in only a few predefined sizes and aspect ratios.

This is getting increasingly problematic, as user interfaces of the third era of computing are increasingly embedded in the physical environment. User interfaces for the Internet of Things and for wearable computing must be customizable in size and shape to match the requirements of the physical context. Moreover, they need new properties, for instance, body-worn displays should be very thin, leightweight, and deformable.

We propose digital fabrication as a unique and novel approach to fabricating embedded user interfaces. Digital fabrication techniques, such as 3D printing and laser cutting, are already successful in commercial applications, for they enable the production of objects with highly custom shapes in industrial production quality. Our pioneering work allows us to go beyond beyond static shape alone to realizing interactive input and output components through digital fabrication. This makes it possible to produce new types of embedded or wearable user interfaces in small quantities and in a fast, easy and inexpensive way.

Our techniques allow to transform a digital vector graphics *(left)* into a functional, custom touch display for embedded user interfaces *(center)* and into a stretchable sensor for on-body input *(right)*.

Printed deformable touch-displays

With *PrintScreen*, we have contributed the first technique for digital fabrication of custom displays: Instead of buying a display off the shelf, one can now graphically design custom light-emitting and touch-sensitive displays in software, like designing conventional vector graphics. The display is then fabricated in a subsequent printing pass, similarly to printing a document.

This enables the fabrication of highly customized solutions in low volume, for instance, for prototypes and small-series products, personalized wearables, and interactive print products. The displays are around 0.1 mm thick, can be deformed, rolled, and even folded. Moreover, displays can be printed on various materials, including polymer films, paper, wood, ceramics, or leather. Only inexpensive commodity tools are required for their design and printing.

In follow-up research, we have developed *PrintSense*, a technique for printing sensor surfaces that capture additional modalities, including hovering and deformation. This enables single-pass printing of multi-modal interactive display surfaces.

Personalized touch-sensitive surfaces for on-skin user interfaces

For high usability on skin, sensor surfaces not only need to be thin, deformable, and customizable in shape, they must also be stretchable. With *iSkin*, we have contributed the first stretchable sensor surface for user interfaces on human skin. The sensor is graphically designed and then produced through digital fabrication techniques. The thin and mostly transparent sensor is made of fully biocompatible silicone. Not only the shape, size, and location of the touch-sensitive elements can be customized, but also their visual appearance, in order to tailor to the aesthetic demands of the user. This sensor enables various new types of on-body user interfaces. For instance, tattoo-like skin stickers enable the quick, direct, and discreet control of mobile computing devices. :::

CONTACT
Jürgen Steimle
DEPT. 4 Computer Graphics
Phone +49 681 302-71935
Email jsteimle@mpi-inf.mpg.de

Perceptual Fabrication

In recent years, there has been a tremendous development of new, three-dimensional (3D) printing technologies. While 3D printers are still not as popular as regular 2D printers, this may change very soon due the great possibilities that they offer. Instead of using traditional manufacturing techniques, such as cutting or injection molding, 3D printers can create complex three-dimensional objects by disposing small portions of printing material and building objects layer by layer. Besides the fine control over the created geometry, new 3D printers can use multiple materials that can be placed almost arbitrary within the printing volume. This enables not only a fine control over the final appearance of the printed objects [figure 1], but also over their mechanical properties [figure 2].

Figure 1: 3D printed model of a building with prescribed texture and subsurface scattering properties.

Figure 2: 3D printed model of a book with prescribed deformation properties.

While the fine control over the material placement provides a great flexibility in terms of what can be manufactured, it also poses some computational challenges. Given an object with specified properties (e.g., mechanical or visual), it is unclear how to combine materials available on the printer to achieve best results. Solving this problem requires optimizing for a material placement, which provides properties as close as possible to those that were specified. Due to a number of limitations of current 3D printers such as limited printing resolution, color reproduction, or material properties, there is no guarantee that such a process will succeed. This makes producing hardcopies of real objects still very challenging. In our work, we argue that this problem can be, to some extent, overcome by accounting for human perception. We believe that good understanding with regards to how the printed objects influence user experience is necessary to fully utilize this kind of technology.

An example of such a perception-driven fabrication approach is our recent work on compliance, which addresses the problem of printing deformable materials. Given a deformation property of an object that we want to print and a limited number of different materials that are available on the printer, we want to choose a material that best matches our goal. Answering this question requires building a computational model that can predict perceived differences between different materials. To this end, we first conducted a number of experiments to investigate the ability of human observers to discriminate differences between materials of varying softness. Then, based on these experiments, we built a computational model that can quantify those differences. Such models can be used later to compare perceived properties of different materials and choose the one best suited for printing the desired object. We believe that building similar models for other properties of printed objects will enable us to produce objects that are not only close replicas in a numerical sense, but also perceptual. :::

CONTACT

Piotr Didyk

DEPT. 4 **Computer Graphics**

Phone +49 681 9325-1009

Email pdidyk@mmci.uni-saarland.de

OPTIMIZATION

Optimization procedures are today of central significance for companies' effectiveness. They are employed, for example, to reduce the need for expensive resources such as work or raw materials. The challenge to science is to develop efficient procedures for solving optimization problems. These procedures should quickly find an optimal solution, or at least a solution that is close to the optimum.

Efficient optimization procedures are of core importance in various areas. For large companies, they have a decisive influence on their competitiveness. With careful planning, a large amount of resources in industrial projects can be saved, leading to lower costs. However, such planning problems are usually very complex, and different requirements have to be taken into account. This makes it difficult for computers to find optimal, or at least very good, solutions.

At the Max Planck Institute for Informatics, we are working on optimization problems arising in very different applications such as industrial optimization or medicine. On the one hand, we are developing elaborate procedures that find optimal solutions efficiently. On the other hand, if the underlying problem is too difficult to find an optimal solution quickly, we develop procedures that can at least find a solution which is close to the optimum. In addition, we are investigating how we can employ random decisions to obtain more efficient and simpler optimization procedures. In this context, we are also studying procedures that are inspired by optimization processes in nature. Such methods often allow good solutions to be found for a given problem without spending too much effort on designing a specific optimiziation procedure.

Since optimization plays a significant role in many different areas, scientists in all research areas of the institute are working on optimization problems. Optimization is nowadays a crucial technique for the design of efficient planning processes. Its importance will continue to grow in the future. :::

CONTRIBUTIONS

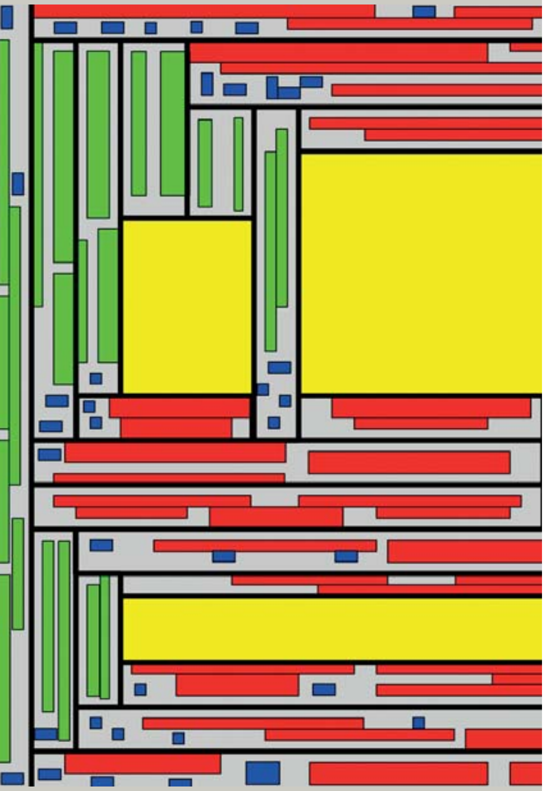

UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION
& VISUALIZATION

OPTIMIZATION

From Routing to Pricing and Learning: Why Are They Hard to Compute?

Pricing problems usually involve a setting where a store owner is interested in selling items to customers in a way that maximizes the profit. This class of problems includes many applications: a dealer who sells cars of various types, a transportation company who owns the toll roads and charges a toll to drivers, a social network owner (e.g. Facebook) who sells commercial products for their advertisers, etc. The challenges of this problem lie in striking the right prices – too low a price is not profitable; too high a price scares off customers). In an apparently unrelated setting, *network routing* problems ask for algorithms that send data (i.e. packets) over the network and guarantee certain quality criteria, e.g., route as many packets as possible in a given time. The challenges of routing lie not only in choosing the right set of requests to route but also in planning their routes along the network. These are some of the most fundamental issues in networking and optimization, so it is no surprise that these problems have shown strong connections to major development in algorithm design, complexity theory, graph theory, and optimization. As computer scientists, we are interested in efficient computational solutions for these problems, or seek an explanation for why such solutions cannot exist.

In this project, it is shown that computational aspects of routing and pricing problems share more similarities than one would imagine. Network routing problems have been popular since the 1980s. On the other hand, pricing problems only started receiving attention since 2000 – in the 90s, hardly anyone ever thought about selling their products over the internet. Computer scientists knew how to route packets much better than how to price items. We had some idea why a certain routing problem is very hard to compute, and it is better NOT to compute them than to let a computer run the task until the end of the universe. When we see a practical setting, we model the problem in a way that it is possible to efficiently compute the solution. We had much less understanding of pricing problems.

Together with my collaborators, I connect these two areas by transferring the knowledge from routing problems. Our results show that algorithm designers face a common difficulty when dealing with routing and pricing. With this discovery, now we understand that certain settings of pricing problems are computationally very hard: For instance, (even approximately) maximizing profits from the toll roads is very hard when there are many more drivers than toll segments. Theoreticians will benefit from rich technical tools developed in our work, and practitioners will learn to reformulate their problems in order to avoid these difficult settings.

We pushed this technique even further, and we ultimately succeeded in “transferring” the techniques from both pricing and routing to help understand the computational complexity of almost 20 other problems. These problems arise routinely in practical settings and we now know they all share a common source of computational difficulty. After 25 years, we now completely understand that a computer is, very likely, incapable of learning even a very simple language (i.e. it cannot learn deterministic finite automation). Computer scientists have already suspected this limitation: A ground-breaking result of Kearns and Valiant showed that any algorithm capable of such learning task would be able to break the famous cryptographic system (RSA), thus convincing a number of researchers that such an algorithm should not exist. Our result provides an even more convincing evidence: such an algorithm would not only break the cryptographic system but also solve tens of thousands of other hard computational problems (i.e., all NP-hard problems). :::

CONTACT

Parinya Chalermsook

DEPT. 1 Algorithms and Complexity

Phone +49 681 9325-1017

Email parinya@mpi-inf.mpg.de

Energy Efficient Algorithms

The efficiency of an algorithm is typically determined by the quality of the solution it returns, along with the amount of resources that it requires to return such a solution. Traditional analysis of algorithms mostly involves evaluating the performance of an algorithm in terms of the resources, *running time*, and *storage*, i.e., the number of steps it requires in order to produce the desired result, or the space it takes up in the storage device during execution. However, as *energy* is a limited and expensive resource, it is of critical importance to analyze algorithms (or the solutions returned by the algorithms) for this resource as well. To quote former Google CEO, Eric Schmidt: *"What matters most to computer designers at Google is not speed, but power, low power, because data centers can consume as much electricity as a city"*.

To address the aim of power saving, there has been a significant upsurge in the study of energy-efficient algorithms in recent years. The most obvious type of algorithmic problem that considers energy as a resource is the question of how to manage power directly at the processor level. As an example, most modern portable computers are equipped with a processor that is speed scalable and is supplied with a sleep state. Suppose we want to process a number of programs on such a computer, where each program has to be processed within a specific time frame and requires a particular amount of CPU-cycles. Then, the operating system has to schedule these programs on the processor, that is, it has to decide when to transition the processor to the sleep state, and furthermore, at which times and at what speed to process each program. These decisions can have a huge impact on the energy consumption of the processor and, in turn, on the battery life of the computer. It has been observed that it is sometimes beneficial to operate the processor faster than necessary in order to reside to the sleep state for longer periods of time, a technique commonly referred to as *race to idle*. However, operating the processor at its highest possible speed whenever there is work to be done, and transitioning it to the sleep state otherwise (as is commonly done in practice) is also suboptimal.

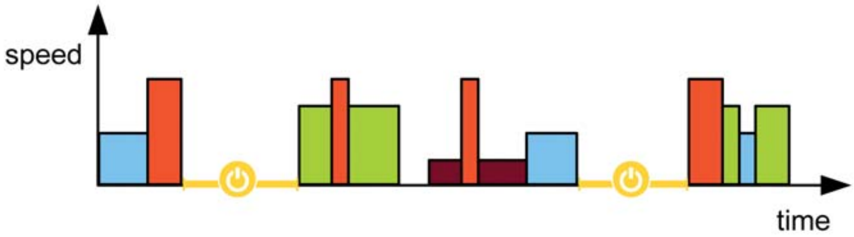

A schedule

Together with Chien-Chung Huang, we have developed an algorithm for the aforementioned setting. The schedules that our algorithm produces have an energy consumption that is guaranteed to be arbitrarily close to the optimal one. Additionally, the algorithm can compute such schedules in a running time that is polynomial in the input size, which is considered efficient. The main idea of our algorithm is to divide each program into smaller, unsplittable subprograms and identify a specific set of orderings in which these subprograms could potentially be processed. We prove that one of these orderings produces a sufficiently good solution, and that we can efficiently find the best ordering in the set by using a well-known technique called *dynamic programming*. Our result is the best possible, given universally accepted complexity-theoretic assumptions. :::

CONTACT

Antonios Antoniadis
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1025
Email aantonia@mpi-inf.mpg.de

Sebastian Ott
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1029
Email ott@mpi-inf.mpg.de

Geometric Packing: Suitcases and More

Everyone who has gone on vacation knows this problem: before the trip can start, you have to pack your suitcase at home; and, you want to take many more things with you than can fit into it (not to mention the weight limitations of airlines). So, instead of taking everything, you pick the most useful items that you can actually transport in your suitcase.

Problems of this type arise in many different situations, i.e., when loading a truck, placing advertisement banners on a website, or, with some additional side constraints, when labeling a map with names of cities and countries. Even though these problems seem very different at first glance, they can all be described as follows: Given are some geometric objects and let us assume here for simplicity that they are (two-dimensional) rectangles. Our goal is to select as many of them as possible such that we can place them in a given area (your suitcase, the loading area of a truck, etc.) without overlap. Additionally, we can assign different levels of importance to the objects in order to model that, e.g., your passport is more important than your third pair of shoes. The advantage of such an abstract description is that the solutions we find can be used for many applications – not only for packing your suitcase, but also for loading a truck.

How do we pick the most important objects in the problem described above so that they fit into the given area? There are many different solutions for problems of this type, way too many to find the best one by simply trying all the possibilities (even for the fastest computer in the world)! Therefore, we design algorithms that do not necessarily always find the best possible solution, but that are efficient and always find a solution that does not differ too much from the optimum (approximation algorithms).

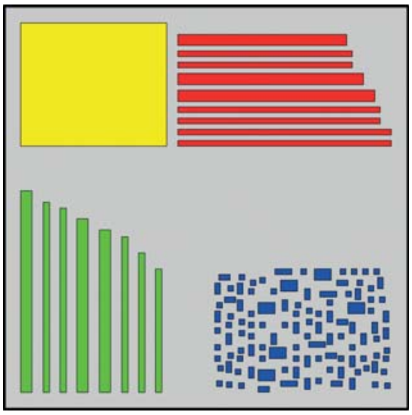

Figure 1: The four different types of recangles.

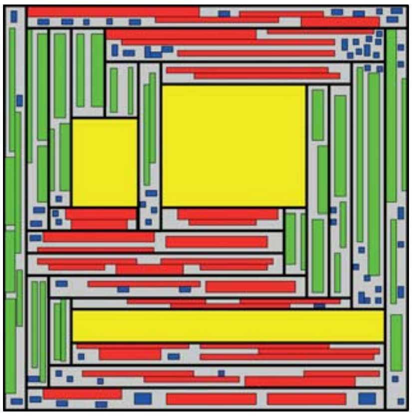

Figure 2: The partition of the given area into a few rectangular pieces.

For the problem above, we use the following strategy: Roughly speaking, we can partition the given rectangles into high and wide rectangles, high and narrow rectangles, wide and thin rectangles, and finally, rectangles that are small in both dimensions [see figure 1]. It turns out that the first three groups are the problematic ones. If there were only rectangles from one of these groups, the problem would be much easier. In particular, the rectangles of the different types can interact in the optimal solution in a very complicated way. In a nutshell, we solve this problem by partitioning the given area into a few rectangular pieces such that in each of them we allow only rectangles from one group. In the second step, we distribute the rectangles into these pieces such that in each of them we place only one type of rectangle. With mathematical methods, we can prove that with this strategy we always find a solution that differs only marginally from the best possible solution and that the resulting algorithm is efficient (in the theoretical sense). :::

CONTACT

Andreas Wiese

DEPT. 1 Algorithms and Complexity

Phone +49 681 9325-1015

Email awiese@mpi-inf.mpg.de

Rule-Based Product Configuration

In today's industry and economy, the principle of product lines plays a great role. In this context, different products, based on a preferably high number of identical components, are grouped together and jointly designed, developed, and produced. A typical example is the group-internal production of different automobile models that are based on a common platform.

A currently popular approach for describing product lines is the use of rules: They specify which components, in a specific context, or following an explicit pre-selection of properties, can be configured into an overall product. These contexts can be of a temporal or technical nature, for example. Alternatively or additionally, a product can be described based on the hierarchical structure of its assembly. An example of a component selection for a vehicle would be the selection of a color, an engine or an equipment detail such as an air conditioner. A temporal context would include a particular year of manufacturing or model year; a technical context could be the upper limit of CO2 emissions or the car's total weight. The rules themselves describe dependencies between the components and contexts: A powerful engine requires specific tire combinations based on the high speeds that can potentially be achieved, whereas trailer hitching devices exclude weak engines. The interactions are represented by so-called if-then rules, which become active through decisions by the user specifying the properties of an eventual product.

Within the framework of our research, we are interested in a formal description and verification of such a system of rules. We also want to be able to calculate mathematically precise analyses of the whole product structure. For many of the systems currently in use, the challenge does not lie in developing the concrete rules, but rather in ensuring important global properties such as consistency. The above-described rule approach is particularly characterized by its flexibility. Rules can be easily adapted to the types of products being configured and their associated components as well as to the respective areas of use. The desired level of detail can be chosen freely: Thus, a product can be regarded solely from a sales perspective, or considered down to the last little screw for engineering purposes. The approach can also be applied to the description of products composed of abstract modules. In addition to the automotive industry, applications for software products that are composed of individual modules are also possible, for example. Production lines, which consist of a sequence of individual steps, or the composition of business services from individual activities, can also be described naturally with the aid of rule systems.

Although promising results have been accomplished with rule systems on the application side, an appropriate formalization, i.e., an explicit and consistent translation of a rule-based system into a language of logic, is still the subject of basic research. However, formalization is a prerequisite for the assurance of global properties. Important properties include consistency, i.e., consistency of the system of rules, or the uniqueness of the calculated results. Further examples of these properties are the total number of possible products from one product line according to the pre-selection of components or contexts, optimality properties of products such as minimum production costs, or recognizing components that no longer contribute to an active product. For systems of rules that also describe a particular process, e.g., a production or configuration process, the confluence of the system of rules also plays a major role: Irrespective of the order of the rule executions starting from an initial state, the result is always the same.

The current research encompasses the analysis of existing rule-based systems, the development of a suitable formal language and the development of automatic methods for the calculation of properties. We have created a formalization which enables the verification of a rule-based configuration system, attaching importance to the interactivity of such a system: Is it possible that user decisions lead to inconsistencies? It essentially considers variability in the form of decisions over Boolean values (i.e., a variable in the system can have the values "true" or "false"). Building on that, procedures that allow analyzing properties of a rule-based configuration system have been implemented . Moreover, a first extension of the formal system has been completed, where additionally the variability in the form of decisions over numeric values is considered. :::

CONTACT

Ching Hoo Tang
RG. 1 Automation of Logic
Phone +49 681 9325-2919
Email chtang@mpi-inf.mpg.de

Christoph Weidenbach
RG. 1 Automation of Logic
Phone +49 681 9325-2900
Email weidenbach@mpi-inf.mpg.de

CGAL – Exact and Efficient Geometric Algorithms

The Computational Geometry Algorithms Library, CGAL, provides access to a large body of efficient and reliable geometric algorithms and data structures. Our department has been an active member of its development since its beginning in the late 1990s. CGAL reduces the gap between theoretical results and implementations that can be used in practical scenarios. It is successful in academic prototypical development and also widely employed among industrial users. The library's philosophy dictates correct results for any given input, even if intermediate round-off errors occur. This is achieved by a design that separates numerical constructions and predicates from combinatorial algorithms and data structures, which is necessary as computers actually do not compute correctly, at least not without the programmer taking extra care. Thus, naive geometrical implementations result in rounding errors that surface in real applications as crashes, non-terminating programs, or wrong results. CGAL, however, follows the exact geometric computation paradigm that ensures correct output for arbitrary instances of a problem, even handling with severe degeneracies. CGAL is developed as an open-source project by a group of research institutes and commercialized by GeometryFactory, a spin-off company that also offers consulting services for industrial users. Geometrical computations occur in various domains, as geographic information systems, recommendation systems and wireless communication planning.

When CGAL was founded, Department 1 contributed to the first geometrical kernel, and has since provided developments of algorithms for convex hulls, polygonal partitioning and Boolean operations on Nef polyhedra in two and three dimensions. The main recent contributions of Department 1 to CGAL consist of implementations for the analysis of algebraic curves and surfaces in the field of non-linear computational geometry. Such computations are needed for motion planning, for instance.1 However, they cannot be carried out reliably with standard machine precision arithmetic. The latest success in performance gain for algebraic curves is due to several ingredients: The central one consists of a cylindrical algebraic decomposition with a revised lifting step. Using results from algebraic geometry we avoid any change of coordinates and replace most of the costly symbolic operations with numerical tools. Two bottleneck tasks have been replaced with implementations that exploit the massively multi-parallelism of fast graphics hardware.

In addition, we have contributed to the combinatorial structures of CGAL in collaboration with Dan Halperin's group from Tel Aviv University in Israel: We worked on arrangements on parameterizable surfaces and provide support for polycurves, that is, curves composed of piecewise subsegments. Both build basic blocks used, for instance, in 3D printing workflows.

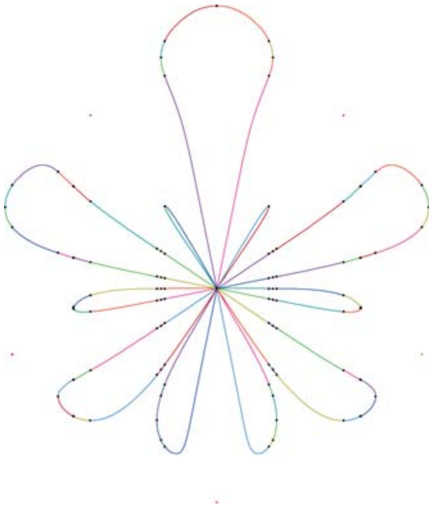
Segmenting an algebraic curve into x-monotone parts and isolated points

Moreover that, a library the size of CGAL needs constant maintenance as well as technical and human resources to keep its infrastructure alive: source code repository, adaptions to multi-core architectures, build system, external library dependencies, testing, documentation, releasing and also interactions on mailing lists and via the website. All these tasks undergo constant technological advancements; these ease the integration of new developers and also ensure industrial-strength quality. Quality assurance is provided through detailed reviews of new contributions, similar to journal submission. Finally, we use CGAL for teaching and training students in modern library design and disseminate our achievements to other researchers in the fields of mathematics and computer science at events like the International Conference for Mathematical Software. :::

CONTACT

Eric Berberich
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1012
Email eric@mpi-inf.mpg.de
Internet http://www.cgal.org

Improving Flat Panel Displays by Discrete Optimization

Our objective is to advance the state of the art in modern display technology by means of discrete optimization. Specifically, we want lower power consumption in the next generation of flat-panel displays through a reduction of their addressing time. To this end, we exploit a common feature of today's displays: the pixels are arranged in a matrix fashion with only one contact per row and per column. For the sake of simplicity, let's assume that these contacts are switches and that a pixel (i,j) shines if and only if the switches for row i and column j are closed. This implies that an arbitrary pattern cannot be displayed at the same time. Hence, several subframes are necessary to display an image entirely. Traditionally, an image is displayed row-by-row at a sufficiently high frame rate such that the eye perceives the average. This method is called Single Line Addressing since each subframe consists of a single line. This is a major cause of excessive power consumption because each pixel only shines for a very short time, and thus, it must shine very brightly to achieve a sufficient average brightness. However, it is possible to drive multiple rows at once, though only by solving a discrete optimization problem in real-time on a driving chip. Constrained by the strong competition on the display market, successful algorithms must be efficient with respect to arithmetic operations and memory consumption.

We have developed a fully combinatorial approximation algorithm for the practically relevant case in which pixels in consecutive rows are addressed simultaneously. Because the algorithm uses only addition, subtraction, and comparisons, it is well-suited for implementation in hardware. Nevertheless, our algorithm computes decompositions in real-time whose objective values are within one percent of the optimum solution in the majority of cases in practice. This invention was patented and commercialized by Dialog Semiconductor under the brand name SmartXtend and made transparent and flexible displays possible.
In particular, it is used in transparent OLED displays by Futaba (formerly TDK). The Lenovo S800 [figure 1], the Explay Crystal, and the Nexian Glaze M9090 are three mobile phones available on the international market that are driven by our technology.

Figure 1: Lenovo S800, driven by our technology

In a follow-up project funded by the German Research Foundation, we extend this research to further display technologies. Our framework will apply to devices like OLED displays, plasma screens, and e-paper.

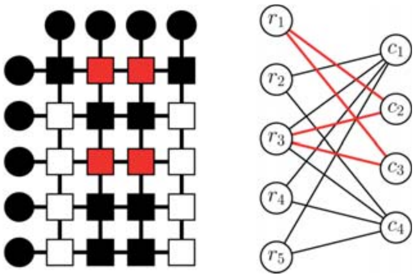

Figure 2: Representing the display matrix by a bipartite graph

We develop new sophisticated driving algorithms and thereby contribute to the theory of addressing matrix displays, which still lacks a profound understanding of the underlying computational problems. Due to the relation of binary matrices to bipartite graphs, the addressing problem is captured by the problem of decomposing the edge set of a bipartite graph into a minimum number of complete bipartite subgraphs, also called bicliques. This project will close a research gap in the area of approximation algorithms for biclique decomposition problems, and the foundational research will benefit flat panel displays and in turn other applications, for example, in graph drawing, computer security, and genetics. In fact, we have already taken a leap forward by recently proving nearly tight lower and upper bounds for approximation factors of polynomial-time algorithms for biclique covering and partition. :::

CONTACT

Andreas Karrenbauer
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325 1007
Email karrenba@mpi-inf.mpg.de

SOFTWARE

Informatics is, firstly, a discipline of basic research that deals with universal computation and problem-solving methods and investigates fundamental properties such as correctness and complexity. Secondly, it also resembles an engineering science that supports a great variety of different applications. The development of software fulfills many functions in this context. It is subject to basic research, the outcome of the implementation of new results from basic research, as well as the product of an engineering effort for the solution of a concrete application problem. Software is thus an inherent part and a connecting link in informatics research.

http://www.mpi-inf.mpg.de/software

AIDA

AIDA is a framework and online tool for entity detection and disambiguation. Given a natural-language text or a Web table, it maps mentions of ambiguous names onto canonical entities.

CGAL

CGAL is a software project that provides easy access to efficient and reliable geometric algorithms in the form of a C++ library.

ClausIE

ClausIE is an open information extractor; it identifies and extracts relations and their arguments in natural language text.

EpiExplorer

The EpiExplorer integrates multiple epigenetic and genetic annotations and makes them explorable via an interactive interface.

FunSimMat

FunSimMat is a comprehensive functional similarity database for proteins and protein families and a tool for easy and fast disease gene prioritization.

Geno2Pheno

Geno2Pheno predicts phenotypic drug resistance of HIV and of Hepatitis B and C virus from the viral genome sequence.

LEDA

LEDA is a C++ class library for efficient data types and algorithms.

NEFI

NEFI is an extensible tool for extracting graphs from images of networks.

pfstools

Tools for Reading, Writing, Manipulating and Viewing High-Dynamic Range (HDR) Images and Video Frames.

RnBeads

RnBeads is an R package for comprehensive analysis of DNA methylation data obtained with a wide variety of experimental protocols.

SPASS | SPASS+T

SPASS is an automated reasoning workbench. SPASS+T extends the theorem prover SPASS with theory support.

Waldmeister

Waldmeister is a theorem prover for unit equational logic.

YAGO

Yago is a huge semantic knowledge base knowing over 900,000 entities.

The International Max Planck Research School for Computer Science (IMPRS-CS)

The training of young scientists is fundamental to the future of science, research, and innovation in Germany. The Max Planck Society, in cooperation with German universities, has launched an initiative to promote young scientists: the International Max Planck Research Schools (IMPRS). They offer especially gifted German and foreign students the possibility to earn a doctorate within a structured program that provides excellent conditions for research. The aim is to strengthen the recruitment and training of young scientists.

IMPRS-CS

Promotion of young scientists

The IMPRS-CS is an opportunity for young scientists and scholars who are between the bachelor's or master's degree and the PhD. This includes first-class training programs, academic specialization, often with thematic linking of individual doctorates, and close collaboration between doctoral students and their academic advisors.

One focus is on international cooperation: The IMPRS-CS strives especially to attract foreign applicants to doing doctoral studies in Germany, familiarize them with the research institutions, and arouse their interest in future work at or in cooperation with German research institutions. Over 50 percent of our doctoral candidates come from abroad, with the largest contingent coming from Bulgaria, China, India and Poland.

IMPRS-CS programs

Together with Saarland University and the Saarbrücken Graduate School of Computer Science, the IMPRS-CS offers programs to establish graduate level competency and also to achieve a doctoral degree. All graduate programs are offered in close cooperation with the Max Planck Institute for Informatics, the Max Planck Institute for Software Systems and the Department of Computer Science at Saarland University. The projects are jointly supervised by the scientists of the Max Planck Institutes and their colleagues in the Department of Computer Science at Saarland University. Outstanding knowledge of English is required for all candidates.

Financial support

IMPRS-CS students receive a PhD support contract. In addition, we assist our students in finding accommodation and with administrative matters of all kinds. We offer English and German classes at several levels, joint activities, and excursions. :::

CONTACT

Jennifer Gerling
IMPRS-CS
Phone +49 681 9325-1801
Email jgerling@mpi-inf.mpg.de

Andrea Ruffing
IMPRS-CS
Phone +49 681 9325-1801
Email aruffing@mpi-inf.mpg.de
Internet http://www.imprs-cs.de

ALUMNI

SUSANNE ALBERS

CV-Milestones

- **PhD Student at MPI-INF from 1990 to 1993**
- **PhD 1993 at MPI-INF**
- **Senior Researcher at MPI-INF from 1993 to 1999**
- **Today: Professor TU Munich, Germany**

Research

My research interests have always been in algorithms, with emphasis on online and approximation algorithms. Further interests include algorithmic game theory and algorithms engineering.

My time at MPI-INF

I was always fascinated to work at MPI-INF with its extremely stimulating and lively research environment. It was inspiring to discuss and collaborate with colleagues such as Naveen Garg, Torben Hagerup, Hans-Peter Lenhof, Stefano Leonardi, Peter Sanders and Stefan Schirra (this list is by no means complete). With many of them, I am still in close touch. Moreover, it was instructive to attend the famous Noon Seminar where group members and guests reported about recent and ongoing research. The atmosphere at the institute was very open and democratic. Each group member was given the opportunity to pursue his/her own line of research and to establish a network of international contacts. It was also highly interesting to see how Kurt Mehlhorn and Harald Ganzinger built up the institute in the early days. Right from the start MPI-INF was a first-class international research institute.

In addition to the scientific work, there were social activities. Most remarkably, Wolfgang Paul, Professor at Saarland University, and Ingrid Finkler-Paul, offered a cooking course for members of the Research Training Group in their house. I still keep the photocopies of the recipes. More generally, the members of the Algorithms and Complexity Department often went out in the evening. A common meeting point was Stiefelbräu near St. Johanner Markt. I also remember great dinner evenings with Naveen Garg, Marina Papatriantafilou, Philippas Tsigas and Christos Zaroliagis. Overall, working at MPI-INF was not only scientifically challenging but also a lot of fun! :::

ANDRE ALTMANN

CV-Milestones

- **PhD Student at MPI-INF from 2005 to 2011**
- **PhD 2011 at MPI-INF**
- **Today: Senior Research Fellow at University College London, UK**

Research

My research focuses on combining high throughput sequencing and neuroimaging to better understand the disease processes in brain disorders and to identify useful biomarkers, thus achieving computer aided diagnosis and prognosis of brain disorders.

My time at MPI-INF

I joined the *"Computational Biology and Applied Algorithmics"* group of the MPI for Informatics in Saarbrücken in November 2005 – immediately after having finished my Master's degree (back then called a Diplom) in Computer Science at the RWTH Aachen. The curriculum at the RWTH demanded to attend courses in a subject of application. I selected Biology and the courses I visited sparked my interest in using computational approaches to solve problems and answer questions in Biology, specifically Molecular Biology.

Upon joining MPI-INF, I was fortunate to work on an exciting topic that went beyond pure biological questions but into supporting medical decisions using computational approaches and data mining. More precisely, we developed computational models to optimize the drug cocktail for patients infected with the human immunodeficiency virus (HIV). The project put me in contact with a truly interdisciplinary team of medical doctors, virologists, biologists and computer scientists, which gave me a good understanding on how Computer Science can be used to advance medicine. The time at the MPI-INF was very exciting and still counts towards my best research experiences. It was here where I wrote my first paper and learned the skills required for conducting high quality research. Most importantly, during the time at the MPI-INF, I met many great people with whom I am still friends. :::

ALUMNI

JAN BAUMBACH

CV-Milestones

- **PhD 2008 at Bielefeld University, Germany**
- **Senior Researcher at MPI-INF from 2010 to 2012**
- **Today: Professor at University of Southern Denmark**

Research

My current research concentrates on the combined analysis of biological networks together with OMICS data, the modeling of genetic expression pathways as well as biomarker discovery and computational methods for personalized medicine.

My time at MPI-INF

The MPI-INF not just boosted my personal career substantially. It helped me to pioneer the field of breath analysis for non-invasive biomedical diagnostics from exhaled metabolomics samples. In addition, its reputation attracted many internationally renowned bioinformaticians to talks, projects and symposia. Several of my current projects still rely on collaborations that were set during my time in Saarbrücken, some of them driving whole bioinformatics research areas. :::

NIKO BEERENWINKEL

CV-Milestones

- **PhD Student at MPI-INF from 2001 to 2004**
- **PhD 2004 at MPI-INF**
- **Today: Professor at ETH Zurich, Switzerland**

Research

My research today is at the interface of mathematics, statistics, and computer science with biology and medicine, ranging from mathematical foundations of biostatistical models to clinical applications.

The goal of our research is to support the design of medical interventions in complex and rapidly evolving systems by means of statistical modeling of high-throughput molecular profiling data, by analyzing biological networks and predicting the effect of perturbations, and by developing mathematical models of the evolutionary dynamics of pathogen and tumor cell populations.

My time at MPI-INF

After receiving my Diploma degree from the University of Bonn in 1999 with a thesis topic on algebraic number theory supervised by Florian Pop, my initial assessment of biology and its connections to math and computer science changed completely when I joined Thomas Lengauer's research group at the GMD in St. Augustin. Bioinformatics, now a well-developed discipline rooted in computer science and central to the life sciences, was still a young and less well-defined subject. It became clear to me very quickly that bioinformatics, or computational biology, is a fascinating and highly dynamic field full of unanswered (and unasked) questions as well as numerous applications, especially in biology and medicine.

When Thomas Lengauer moved to Saarbrücken in 2001 to start the "*Computational Biology and Applied Algorithmics*" Department at the MPI for Informatics, I was lucky enough to join him as his PhD student. On one of my first days at what used to be called Stuhlsatzenhausweg 85, I walked upstairs to a printer on the sixth floor, because ours on the fifth floor had not yet been installed. At the printer, I bumped into Harald Ganzinger, who recognized a new face, introduced himself, and immediately started interrogating me about my research. Although Harald took me by surprise, this early encounter left a long-lasting impression on me as it demonstrated the friendly, down-to-earth, open-minded, and stimulating atmosphere at MPI-INF that I encountered in general and enjoyed throughout my stay. At that time, I largely took the scientific excellence of the Institute for granted and it was only later when I fully appreciated it. I received my PhD in Computer Science from Saarland University in 2004, honored by the Otto Hahn Medal of the Max Planck Society, after Thomas Lengauer had introduced me to the field of bioinformatics and given me the unique opportunity to work on an exciting project concerned with the development of drug resistance in HIV. :::

ALUMNI

BRUNO BLANCHET

CV-Milestones

- **PhD 2000 at Ecole Polytechnique, Paris, France**
- **Independent Research Group Leader at MPI-INF from 2001 to 2004**
- **Today: Senior Researcher at INRIA Paris, France**

Research

My research interests deal with automatic security protocol verification both in the symbolic and in the computational model.

My time at MPI-INF

From 2001 to 2004, I was the leader of an independent research group (Nachwuchsgruppe), the Static Analysis group, at the Max-Planck Institute for Informatics. At the same time, I was also researcher at CNRS, in the abstract interpretation team of Ecole Normale Supérieure (ENS), Paris. I spent about three weeks a month at MPI-INF and one week a month at ENS. That was a very exciting period scientifically, with collaborations both in Saarbrücken and in Paris, as well as with Martín Abadi (University of California, Santa Cruz) and Cédric Fournet (Microsoft Research, Cambridge). I particularly appreciated the excellent working conditions provided by MPI-INF. :::

KASIA BOZEK

CV-Milestones

- **PhD Student at MPI-INF from 2007 to 2011**
- **PhD 2011 at MPI-INF**
- **Today: Group Leader at Okinawa Institute of Science and Technology, Japan**

Research

My field of research is in computational biology. The projects I performed can be grouped into 3 major topics: metabolomics, HIV, and circadian clock.

Throughout these projects, I applied methods ranging from comparative metabolomics/ genome/gene expression analysis, through modeling in structural biology to machine learning methods for virus classification.

My time at MPI-INF

After accomplishing my MSc in Computer Sciences at the University of Warsaw in Poland, I started my path towards computational biology. During the years 2007-2011, I was a PhD student at the MPI-INF department for *"Computational Biology and Applied Algorithmics"* with Prof Lengauer as my PhD advisor.

My stay at MPI-INF was a great time of discovering this new and exciting field of research. The stimulating and inspiring environment of the group and institute taught me scientific approach – to be creative, think critically, and explore the solutions of a given problem. It provided me a rich and unique research environment in which I could thrive and develop myself in the chosen direction.

Not only did it result in a good publication record, but also opened the doors to other opportunities. After my time at MPI-INF, I moved to Shanghai for a Post-Doc at the CAS-MPG Partner Institute for Computational Biology where I could put many of the skills and methods learnt at MPI-INF into practice. Now I am moving on to Okinawa Institute of Science and Technology in Japan to start my work as a group leader.

In all my future work, I will remember dearly the time spent at MPI-INF. This is the place where my scientific curiosity was first awakened and where I learnt the skills and tools I will use the rest of my scientific and professional life. :::

ALUMNI

ELMAR EISEMANN

CV-Milestones

- **PhD 2008 at University of Grenoble at INRIA Rhône-Alpes, France**
- **Senior Researcher at MPI-INF from 2008 to 2009**
- **Today: Professor at TU Delft, Netherlands**

Research

My research covers a broad range of topics in the field of computer graphics, including the main areas of game technology, rendering, geometry processing, and visualization.

My time at MPI-INF

After my PhD, I applied as a Senior Researcher in the Cluster of Excellence MMCI at Saarland University/Max-Planck Institute for Informatics. I initially approached Hans-Peter Seidel at Eurographics and he invited me to present my work at the MPI-INF. During the conference dinner, already the first MPI-INF students approached me and I received a glimpse of the fantastic atmosphere and team spirit.

This impression was also confirmed during my visit and the warm welcome immediately convinced me that Saarbrücken is the right next step on my career path. I still remember vividly the serious but always positive work atmosphere, combined with a literal strife for "excellence" and the various opportunities for collaboration, which clearly expanded my horizon. The many events in Saarbrücken were impressive and I got the chance to contribute to the organization of EGSR and HPG, which took place on campus in 2010.

Besides work, many other activities are still very present in my mind, ranging from an Oktober Fest, to which we went straight out of a deadline without sleep for 48h, to a visit in climbing parks and a trip to the Cologne Karneval with a small group – which, together with the fact that I impersonated a prince during the cluster launch, led to a change of my door sign into "Prinz Elmar Eisemann".

I am very grateful for the fantastic time at MPI-INF, which obviously had a strong influence on my future career steps. I enjoy many very good contacts that date back to this time and I am very thankful for their continuing support. I hope that the next 25 years (and many more) will be as successful as the last! :::

NAVEEN GARG

CV-Milestones

- **PhD 1994 at IIT Delhi, New Delhi, India**
- **Post-Doc at MPI-INF from 1994 to 1997**
- **Today: Professor at IIT Delhi, India**

Research

My research interests are: design and analysis of algorithms, designing approximation algorithms for NP-hard problems, combinatorial optimization and graph theory.

My time at MPI-INF

I first came to MPI-INF in September ’94 when I started as a Post-Doc after having completed my PhD at IIT Delhi. The institute was then housed in a small building near the “Philosophen Café”. I must admit that getting accustomed to the gray weather in Saarbrücken was a bit of a challenge but the cheer, the warmth and the camaraderie of my colleagues at MPI-INF made it possible. Over the 3 years I spent at MPI-INF, I made many friends, friends who I continue to be in touch with till this day and whose friendship I do value a lot.

One incident during my stay there that I still remember and often recount is from a day in November ’96. It started snowing around noon and since this was the very first snow that season we were all looking through our large office windows admiring the sight of things turning white. It all looked very pretty but then someone decided that the snow should not stop. So over 8 hours we had enough snow to make all buses go off-road and the roads unsuitable for biking or driving. I used to bike to the institute and my home was almost 7 km away on the other side of the town on top of a hill in a place called Bellevue (no not in the prison!). That evening I had no option but to walk and it took me over two hours to get back home. It was a beautiful and a tiring experience, one I am unlikely to ever forget.

I left Saarbrücken in ’97 to start as an Assistant Professor at IIT Delhi where I continue to this day. However my association with MPI-INF has only grown stronger over the years. My group on *“Approximation Algorithms”* at IIT Delhi became a partner group of the MPG from 2005 to 2010 and then from 2010 to 2015 I have been co-director of the Indo-German Max-Planck Centre for Computer Science (IMPECS). I believe that except for a couple of occasions, I have visited MPI-INF every summer since the time I left Saarbrücken. I also made two longer visits – one in 2012 as a “Bessel-Preisträger” and another on my sabbatical in 2006-07. While the researchers I have met and worked with have changed over the years, MPI-INF itself has been as welcoming and helpful as ever.

MPI-INF is celebrating its 25th birthday. Many of us who have passed through its portals owe a lot to this great institution for it has not just groomed us scientifically but also given us many of our friends and our experiences. It has been an important part of my life and as everyone around me says “my second home”. :::

ALUMNI

RAMESH HARIHARAN

CV-Milestones

- **PhD 1994 at Courant Institute of Mathematical Sciences, New York, USA**
- **Post-Doc at MPI-INF from 1994 to 1995**
- **Today: Chief Technology Officer at Strand Life Sciences, Bangalore, India**

Research

Research-wise, I work on design and analysis of algorithms with provable guarantees. In the recent past, I have been working on algorithms for connectivity problems in graphs. My PhD thesis as well as some subsequent work involved design of fast pattern matching algorithms for string, two dimensional and tree patterns.

My time at MPI-INF

20 years after I left MPI-INF, really fond memories of my time at MPI-INF remain. My most vivid memories are of vibrant discussions during talks with Kurt Mehlhorn, Naveen Garg, Sanjiv Mahajan, Torben Hagerup, Pierre Kelsen, Sunil Arya, Shiva Chaudhuri and others.

Sanjiv Mahajan's off-beat acts like singing in bars in Hindi, long runs along the Saar with Shiva Chaudhuri and Jesper Traff followed by my first marathon in Berlin, badminton with the always smiling Pierre Kelsen, learning Deutsch with Theresa, reading Hermann Hesse's Siddharta in Deutsch with Andrea Esser, and sneaking past alert guards at the French border in the pre-Schengen era to shop at a mall right across will forever be etched in my memory. :::

GEORGIANA IFRIM

CV-Milestones

- **PhD Student at MPI-INF from 2005 to 2008**
- **PhD 2008 at MPI-INF**
- **Today: Lecturer at University of Dublin, Ireland**

Research

My main research interests stem from the areas of machine learning, data mining and information extraction. My recent research focuses on designing large-scale real-time machine learning methods, for news and social media data streams. The goal is to develop novel methods to enable new and classic digital journalism applications, such as social indexing and story tracking.

My time at MPI-INF

The time at MPI-INF was truly special, the research environment was very dynamic and productive, and the cultural and social dimensions were enriched by the many nationalities that create the MPI-INF community. During the 5 years at MPI-INF, I learned German and made many friends, attended the traditional Christmas and summer trips, and worked very hard. I proudly wore the MPI-INF badge when attending international conferences, and became a part of the MPI-INF identity. MPI-INF has laid the basis for my academic career, as well as my personal life, enriching me with unique work and life experiences, a PhD and my life partner. :::

ALUMNI

NICOLE MEGOW

CV-Milestones

- **PhD 2006 at TU Berlin, Germany**
- **Post-Doc and Senior Researcher at MPI-INF from 2008 to 2012**
- **Today: Professor at TU Munich, Germany**

Research

My research deals with discrete optimization at the interface of discrete mathematics, theoretical computer science and operations research. My work focuses on developing and mathematically analyzing efficient algorithms, particularly approximation algorithms, for the kinds of combinatorial optimization problems that can occur in production, logistics or computer system control.

My time at MPI-INF

I remember the MPI-INF being a paradise for researchers. The time I spent as a Post-Doc in Kurt Mehlhorn's group was great and productive, and it strongly influenced my scientific career. The institute is a place where excellent (junior) researchers with different research areas meet in a very open and communicative environment. In addition, there are numerous short-stay guest researchers from all over the world. I met many interesting people here at the institute and built up a large research network. Kurt, those were great times! I also really enjoyed the wonderful support the administration gave us in organizing events (such as the yearly ADFOCS summer school and the MPG Career Development Seminars). Working with such a dedicated and experienced team makes everything so much easier.

All in all, I always felt that I was receiving a lot of support from the MPI-INF – and ultimately from the entire Max Planck Society. :::

SEBASTIAN MICHEL

CV-Milestones

- **PhD Student at MPI-INF from 2004 to 2007**
- **PhD 2007 at MPI-INF**
- **Senior Researcher at MPI-INF from 2009 to 2014**
- **Today: Professor at TU Kaiserslautern, Germany**

Research

My research interests are databases and (distributed) information systems, particularly Top-k query processing over data streams and multimedia information retrieval.

My time at MPI-INF

I started studying computer science at Philipps-Universität in Marburg, where I was particularly active in the "Database Systems" group led by Prof. Bernhard Seeger. Towards the end of my diploma thesis, in the autumn of 2003, I had a conversation with Bernhard and he told me that a man by the name of Gerhard Weikum would soon be taking office as a director at the Max Planck Institute in Saarbrücken, and that he would presumably have open PhD positions. Now, this might sound like a cliché, but after that conversation I did actually go back to my desk and did some online research to find out where Saarbrücken even was. In the meantime, Bernhard must have gotten in touch with Gerhard, because a few days later Gerhard sent me an invitation for an interview in Saarbrücken. At least I now knew where to find the city on a map. So in September 2003 I travelled to Saarbrücken for the interview with Gerhard, who was still working at Saarland University at the time. And – surprise! – I got the PhD position, or else you wouldn't be reading about it right now. In March 2004, I finally joined the MPI-INF.

I had a great time at the MPI-INF: nice colleagues and an extraordinary thesis advisor. Together we very successfully wrote publications very successfully, supervised students and organized seminars. In addition to working with Gerhard, my collaboration with Prof. Peter Triantafillou also had a great influence on me. Peter came here from the University of Patras in Greece and joined the MPI-INF for a research stay that lasted six months, if I remember correctly – our collaboration went on to last much longer than those six months. For me, the MPI-INF stands for outstanding research opportunities, high standards and nice people, but just as importantly it also stands for "open doors", in both the literal and figurative sense. After I received my PhD in July 2007, I became a postdoc at EPF Lausanne where I joined the working group led by Prof. Karl Aberer and had the opportunity to gain experience in the supervision of PhD candidates, teaching and project management.

In late 2008, I came across a job posting announcing open positions for junior research group leaders at Saarland University's Cluster of Excellence MMCI. I sent in my application and was selected for one of the positions. I started the new job in August 2009 and was also directly appointed Senior Researcher at the MPI-INF. While I really enjoyed the freedom I had at the MMCI, I also benefited from my involvement with the MPI-INF, especially from my renewed collaboration with Gerhard Weikum. Several aspects were new: being responsible for staff and a budget, the right to award doctorates as well as the opportunity to give my own lectures – while at the same time having a very low mandatory teaching load.

In March 2014, I was awarded a W3 professorship in databases and information systems at the University of Kaiserslautern, as the successor of Prof. Theo Härder. So now I was off to the Rhineland-Palatinate region. Upon hearing of my appointment as professor in Kaiserslautern, a colleague and friend of mine said: *"You should be used to the similar dialect by now!"* :::

ALUMNI

THOMAS NEUMANN

CV-Milestones

- **PhD 2005 at University of Mannheim**
- **Post-Doc and Senior Researcher at MPI-INF from 2005 to 2010**
- **Today: Professor at TU Munich**

Research

My research interests are: query optimization, graph-structured databases and distributed query processing.

My time at MPI-INF

At the Max Planck Institute for Informatics, work was conducted in a very productive manner. The large, mixed group and the open atmosphere fostered a high level of collaboration. A project carried out jointly with Gerhard Weikum resulted in the RDF-3X system, which was designed for efficient processing and searching in graph-structured RDF databases. This project was not only a success in itself and attracted international attention – it was also extremely useful for cooperating with other group members working on knowledge extraction projects such as YAGO, for example. The fact that the group covered such a broad spectrum of topics and that it was so easy to work with others were main hallmarks of the MPI-INF, and made for a very pleasant and productive stay. :::

ANTTI OULASVIRTA

CV-Milestones

- **PhD 2006 at University of Helsinki, Finland**
- **Senior Researcher at MPI-INF from 2011 to 2014**
- **Today: Professor at Aalto University, Finland**

Research

My group's mission is to identify and exploit optima of human-computer performance. We formulate interface design problems as optimization tasks, develop predictive modeling of interaction, and implement computational methods for interface design.

My time at MPI-INF

Moving to MPI-INF was perhaps the biggest gamble in my career. I left behind a successful research group in Helsinki that I had been co-leading for three years. I did not know much about MPI-INF beyond its excellent reputation in natural sciences, and Saarbrücken appeared to be a nice city in the middle of Europe. I would have been able to continue in Helsinki, but I felt I needed to step up and renew my research. Trained in cognitive science, neuroscience, and computer science, I felt that the work that I was leading in Helsinki was not intellectually advancing my field. I had been declined a permanent position in Helsinki – twice – and once in Zürich and was seriously thinking my options. I hence convinced my wife – an anesthesiologist – and my daughter – then 4 years – to move with me – even if neither could speak a word in German! We sold our beautiful home, packed all our stuff to a truck, and we went off.

The very first days I spent there – working on an empty desk in the MPI-INF building that was and has always been overcrowded – I realized the magnificent and outstanding character of the institute. All group leaders and professors there – many of whom would later become my long-term collaborators – were highly ambitious and brilliant researchers. The atmosphere at MPI-INF was never predatory but always collegial and supportive. It took me about one year to "calibrate myself up" to match that level. I am particularly grateful for Professor Hans-Peter Seidel, my mentor in D4. Even if he knew little about my field, his trust on me and his advice on research leadership was instrumental to my development as a group leader. At the beginning, we would meet on a weekly basis and later a bit less frequently. I remember knocking on his door on the second floor and asking *"Do you have a minute?"* He responded *"No I don't, but come in anyway."* Under his guidance, I started recruiting PhD students and a Post-Doc and carefully built a group that I still have very fond memories of. We laid the groundwork for something for which I later earned an ERC Starting Grant and am still continuing today.

During my stay at MPI-INF, I was heavily influenced by other group leaders, including in particular Dr. Tino Weinkauf, Prof. Christian Theobalt, Dr. Jürgen Steimle, Dr. Timo Kötzing, and Dr. Andreas Karrenbauer. Our family fared well, too. Although moving to a new country was stressful at the beginning, we found lots of nice escapees in the vicinity and had more time together than in Finland. My daughter eventually finished the first year class in Ost Schule and is now starting in Deutsche Schule in Helsinki – the only German-speaking school in Finland. MPI-INF's support in finding a daycare right on the campus was instrumental to my ability to organize our life in Saarbrücken. And my wife eventually reached a level in German that allows her to work in a hospital. Now, looking back to the three years in Saarbrücken, the gamble paid off. :::

ALUMNI

KNUT REINERT

CV-Milestones

- **PhD Student at MPI-INF from 1994 to 1999**
- **PhD 1999 at MPI-INF**
- **Today: Professor at FU Berlin, Germany**

Research

In my research, I aim at enabling translational research in computationally based life sciences by removing existing (communication) gaps between theoretical algorithmicists, statisticians, programmers, and users in the biomedical field. Of particular interest are devising mathematical models and algorithms for analyzing large genomic sequences, developing algorithms for analyzing data derived from mass spectrometry experiments (e.g. for detecting the differential expression of proteins between normal and diseased samples), and the implementation of methods that benefit researchers at the bench or in the clinic.

My time at MPI-INF

Being born in the Saarland it was quite natural to study computer science in Saarbrücken, back then (and still now) one of the best addresses for this subject in Germany. I started my scientific career as a Master student and relatively quickly knew that Kurt Mehlhorn's algorithms group fit my interests. Hence, I specialized in algorithms. When I started my MSc thesis with Michiel Smid it was the time when the MPI-INF was founded and Kurt Mehlhorn was named director. The Mehlhorn group moved to a temporary building where I did my MSc research with fellow MSc students, among them Hannah Bast and Michael Seel (incidentally I met there my now wife in the MPI-INF library).

During my PhD period I was looking into different topics, among those were also problems in a then relatively new field, computational biology, at the time led by Hans-Peter Lenhof. The problems of working with biological data intrigued me and I chose a topic in this field. During the research I worked with international partners, namely Gene Myers and his PhD student John Kececioglu, which should heavily influence my next career steps. :::

MARCUS ROHRBACH

CV-Milestones

- **PhD Student at MPI-INF from 2010 to 2014**
- **PhD 2014 at MPI-INF**
- **Today: Post-Doc at UC Berkeley, USA**

Research

My research interests are computer vision, computational linguistics and machine learning.

My time at MPI-INF

Even though the basic fundament had already been there at that time, getting my PhD at the MPI-INF essentially meant laying the foundation for my future in at least three different ways. First of all, it was not just me who made a fresh start in Saarbrücken after completing my Master's degree in Darmstadt – it was also a new beginning for the Computer Vision and Multimodal Computing department, or "D2", as we call it at the institute. Supervised by Director Bernt Schiele, I was, as one of the first doctoral candidates and group admins, involved in setting up the department. Secondly, my PhD became a cornerstone of my professional scientific career, just like my postdoc at UC Berkeley. And thirdly, it was at the MPI-INF that I met the woman who is now my wife, Anna Rohrbach, and she is an important part of my life's foundation.

I always enjoy coming back to visit Bernt and his group, to collaborate with them and take part in the retreats. :::

ALUMNI

BODO ROSENHAHN

CV-Milestones

- **PhD 2003 at Christian-Albrechts University of Kiel, Germany**
- **Senior Researcher at MPI-INF from 2005 to 2008**
- **Today: Professor at Leibniz-University Hannover, Germany**

Research

My research focuses on automated image interpretation. It ranges from classical computer vision to visual signal processing and bio-medical image analysis.

My time at MPI-INF

Regarding my time at MPI-INF, I would like to mention that I look back with a deep gratitude for being a part of this endeavor. I heartily thank all members of department 4 which accompanied me in Saarbrücken. It is a unique and open-minded research facility which I have never found somewhere else and it is a perfect spring board for young researchers. The first day when I arrived I had a meeting with my supervisor Hans-Peter Seidel. He told me to immediately start applying somewhere else and I thought *"wow, I haven't started here and he wants to get rid of me"*. In his second sentence, he told me that as long as my research makes progress I should not be concerned about my position. Then I realized: This is the place where I want to be. :::

PETER SANDERS

CV-Milestones

- **PhD 1996 at University of Karlsruhe**
- **Post-Doc and Senior Researcher at MPI-INF from 1997 to 2004**
- **Today: Professor at University of Karlsruhe**

Research

I work on the design, implementation and analysis of efficient algorithms; "analysis" can be both theoretical and experimental. Topics I often touch are parallel processing and communication in networks, solving problems with „irregular" structure, randomized algorithms, memory hierarchies (disks, caches) and realistic models for problems and machines.

My time at MPI-INF

My stay at MPI-INF was pivotal for my career. From day one I was integrated into a highly international team and was able to cooperate on a variety of topics with students, Post-Docs and Senior Researchers. Together with nearby Dagstuhl, one almost got the impression that travelling is not necessary since everybody important in the field will eventually come by. In particular, Saarbrücken is one of the best places to be if you want to work on graph algorithms, algorithms for memory hierarchies, or algorithm engineering. These areas became permanently important for me.

Another interesting feature at MPI-INF is the close cooperation with university. Without a mandatory teaching load, one still has the opportunity to gather experience in teaching a variety of courses. :::

ALUMNI

HOLGER THEISEL

CV-Milestones

- **PhD 1996 at University of Rostock**
- **Post-Doc at MPI-INF from 2002 to 2005**
- **Research Group Leader at MPC-VCC from 2005 to 2006**
- **Today: Professor at University of Magdeburg, Germany**

Research

My research is focusing on different aspects and subfields of visualization, namely flow visualization, information visualization, visual analytics, volume visualization and tensor visualization. In addition, I work in computer graphics, in particular in the fields of geometry processing and computer-aided geometric design.

My time at MPI-INF

Joining the MPI for Informatics in 2002 was a big chance for me. In fact, at this time I was rather uncertain about my own scientific capabilities and career chances. Working in the Computer Graphics group at MPI-INF has changed this immediately. This highly motivating and inspiring environment helped me a lot to get clearer about what I want and what I am able to do.

I remember long nights of hard work before the deadlines where everybody helped each other to finalize the research results. I also remember nice social aspects, both at the MPI-INF and at several conferences with my MPI-INF colleagues. I am still thankful to the MPI-INF (in particular to Hans-Peter Seidel) for giving me this chance; at this time, it was not clear at all which way I would go. I am proud and happy to be an alumnus of the Max Planck Institute for Informatics. :::

CHRISTIAN WOJEK

CV-Milestones

- **PhD 2010 at TU Darmstadt, Germany**
- **Post-Doc at MPI-INF from 2010 to 2011**
- **Today: Computer Vision Senior Scientist at Carl Zeiss Corporate Research, Jena, Germany**

Research

My current research interest is in the robust and reliable real-world application of computer vision methods to bio-medical (microscopy, ophthalmology) and industry domains. In particular, my focus is on multi-modal registration, multi-modal segmentation, abnormality and event detection.

My time at MPI-INF

I enjoyed the time at MPI-INF in particular due to the first class research environment. Even though I was one of the first D2 members at the time, the great support from the institute's IT department allowed me to be productive with my research from the very beginning. I also remember numerous fruitful discussions, in particular during the retreats, Friday meetings and before deadlines that made my stay at MPI-INF an inspiring and fun time. :::

NEWS

This section contains an overview of the important events of the past year: activities to promote young talent, appointments and awards, prestigious scholarships earned by institute members, cooperations, and important events surrounding the Max Planck Institute for Informatics. Significant events include scientific conferences, the visit of the Scientific Advisory Board, the Board of Curators meeting as well as the launch of the new website and the opening of the new building of the Max Planck Institute for Software Systems in Kaiserslautern.

CONTRIBUTIONS

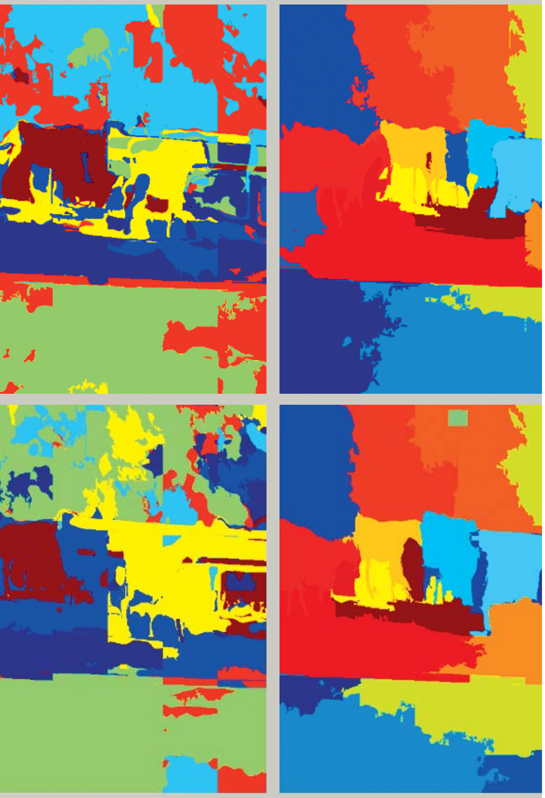

Awards, Honors, Commendations

Erasmus Medal for Kurt Mehlhorn

The Academia Europaea decorated **Kurt Mehlhorn** with its highest accolade, the Erasmus Medal for 2014. This annually appointed award attests its winners continuous internationally recognized scholarship of highest level and recognition by colleagues.

In the encomium the key role of algorithms in every computer system was emphasized and Kurt Mehlhorn was praised to be *"Mr. Algorithm"* of Europe. Among the 24 laureates he is the ninth from fundamental research, the fourth German, and the first computer scientist. :::

Doctor honoris causa of the University of Göteborg

GÖTEBORGS UNIVERSITET

The doctor honoris causa is a honorary distinction of universities or faculties that is bestowed for exceedingly academic or scientific merits. Unlike regular doctorates the doctor honoris causa is no academic degree for a certain paper but expression of appreciation.

Owing to the long-term collaboration with the University of Gothenburg, **Kurt Mehlhorn** was decorated with a doctor honoris causa. In the encomium his fundamental contributions to many different topics of computer science were recognized as well as his dedication as a professor. :::

As sign for the doctor honoris causa, Kurt Mehlhorn is wreathed with laurel by the dean.

Thomas Lengauer receives Hector Science Award

Hector Stiftung II

In 2015, **Professor Thomas Lengauer,** director at the Max Planck Institute for Informatics, was honored with the Hector Science Award. The Hector Foundation II acknowledges his achievements in the field of bioinformatics as well as his commitment to teaching. The prize is awarded annually to excellent scientists at German universities and carries a value of 150,000 EUR. The distinguished winners are also awarded the title "Hector Fellow" and become members of the Hector Fellow Academy, which promotes interdisciplinary projects and academic networks. :::

Awards, Honors, Commendations

Heinz Maier-Leibnitz Award

The German Science Foundation presents the Heinz Maier-Leibnitz Award to young scientists for outstanding results. It should support them to pursue their scientific career. The researchers must already have their own scientific profile after their Ph.D.

Nicole Megow, department of Kurt Mehlhorn, was one out of ten award winners in 2013. :::

Elisabeth Schiemann Kolleg

Within the Elisabeth Schiemann Kolleg scientists of the Max Planck Society support outstanding young female researchers on their way to tenured faculty or director positions in research institutions.

The Schiemann Kolleg promotes its female fellows' activities to establish themselves in the scientific world. Parallel to that the Kolleg provides a platform for interdisciplinary scientific exchange. The support is of aspirational not financial nature. In 2013 this support construct was established; **Nicole Megow,** from the department of Kurt Mehlhorn, was coopted to it together with four other young female scientists. :::

Lise Meitner Award

The Lise Meitner Award is given by MPI-INF to young female scientists in the field of computer science. For the year 2014 **Zeynep Akata**, INRIA Rhone-Alpes in Grenoble, was awarded. The award consists of a research stay at MPI-INF and budget for equipment. Zeynep is going to spend two years with Bernt Schiele, Department 2, focussing on the topic *"Zero-shot and Few-shots Learning Methods for Large-scale and Fine-grained Image Classification"*. [http://www.mpi-inf.mpg.de/news/employment/lise-meitner-award-fellowship/] :::

Chinese Government Award

The Chinese Government bestows with the *"Chinese Government Award for outstanding students abroad"* Chinese people who study with great success in foreign countries. The government honours their dedication and tries to strengthen their home bounds. **Chenglei Wu**, Ph.D. student with Professor Hans-Peter Seidel, was awarded in this way in 2014. :::

CREST International Research Grant

The Janpanese Science and Technology Agency, JST, awards teams of scientists within the Core Research for Evolutionary Science and Technology (CREST) program. The awards are in support of revolutionary ideas in science and technology that give rise to hope of new research fields. In 2014 **Andreas Bulling**, Dept. 2, was part of such a team. :::

Otto Hahn Medal

The Max Planck Society bestows young scientists with this distinction for outstanding results within their Ph.D.

Carola Doerr, post doctoral researcher in Kurt Mehlhorn's department until 2013, was awarded with this medal in 2013. It honoured her doctoral thesis, *"Toward a Complexity Theory for Randomized Search Heuristics: Black-Box Models"*, with which she received her Ph.D. in 2011 at Saarland University.

Ndapandula Nakashole was a Ph.D. student in the department of Professor Gerhard Weikum where she received her Ph.D. in 2012. For her thesis *"Automatic Extraction of Facts, Relations, and Entities for Web-scale Knowledge Base Population"* that was appraised by the reviewers with summa cum laude, she was awarded with the Otto Hahn Medal in 2014. :::

Eurographics Young Researcher Award

The Young Researcher Award is given each year by Eurographics (European Association for Computer Graphics) to two young researchers in the field who have already made a significant contribution and are likely to to make more. The intent of this award is to recognize people early on in their career. Eurographics nominated **Tobias Ritschel**, Dept. 4, as one of the award winners for 2014. :::

Awards, Honors, Commendations

Professor Kurt Mehlhorn, the institute's Founding Director and head of Department 1 "Algorithms and Complexity", was admitted as *Foreign Associate* by the U.S. National Academy of Engineering; he was incorporated into the class of computer science. The admission is done by decision of the members when a person has achieved outstanding results. This membership is one of the highest honours in the field of engineering science, it is often bestowed as an award for a lifework. :::

Khwarizmi International Award

Since 1987 the Iranian science organisation IROST (Iranian Research Organization for Science and Technology) bestows persons with the Khwarizmi International Award who did outstanding work, preferentially in the field of digital and mechanical development. Since 1991 also foreign scientists can be awarded. The Award is a memory to the renowned Iranian polymath Abu Jafar Mohammad Ibn Mousa Khwarizmi who's name is the etymological source of the term algorithm. His work made the Indian mathematics as well as the use of the number zero known to the western world. **Kurt Mehlhorn** was bestowed with the highest category of this award for 2013. :::

Two Awards for "The Captury"

The start up company *"The Captury"*, that was founded by three scientists from the group of Professor Christian Theobalt, got awarded twofold: In the *"Founders Competition – ICT Innovative"* they were awarded in the highest category. In this Germany-wide competition the German Government honours the foundation of innovative companies in the field of Information and Communication Technologies (ICT).
"1,2,3,GO" is the competition for the best business plan in the Large Region Saar-Lor-Lux-Trier-Wallonie; it is awarded by the Association of the Chambers of Industry and Commerce. :::

Dr. Eduard Martin Prize

Mahmoud Fouz, Ph.D. student of Benjamin Doerr, department 1, was awarded with the Dr. Eduard Martin Prize 2013 for his excellent thesis *"Randomized Rumor Spreading in Social Networks & Complete Graphs"*. This award is founded by Saarland University in memory to it's long-time president who's name it bares. It is annually bestowed by the Association of Friends of Saarland University to the best young scientists of the University. :::

LICS Test-of-Time Award

The symposium "Logic in Computerscience" of the *Association for Computing Machinery* awards annually papers that were published 20 years before and maintained influence to the field of automated logic. This award is distinguished from others as it represents the proven influence of a retrograde paper.
Uwe Waldmann was chosen for one of this awards in 2013. :::

Awards, Honors, Commendations

ERC Grant

The European Research Council awards so-called ERC Grants, the highest level of European research support, on a regular basis. It provides substantial support money as well as significant prestige for the grant holders who are chosen in a rigorous selection process.

This funding instrument is meant to support topics of *Frontier Research*, that is high-risk science at the edge of current understanding.

Scientists from the MPI-INF won two of these highly competitive awards, one *Starting Grant,* a prize for scientists who, at the beginning of their careers, want to start a new field, and one *Synergy Grant*, the greatest financial award for research teams.

European Research Council
Established by the European Commission

ERC Starting Grant

Professor Christian Theobalt, DEPT. 4, was awarded one of the prestigious prizes of the ERC. In his research project *"CapReal – Performance Capture of the Real World in Motion"*, he wants to develop new efficient ways to automatically detect and describe in 3D the motion of humans in video streams that were captured in non-artificial environments. The ERC approved an amount of 1.48 million euros for a period of five years.

ERC Synergy Grant

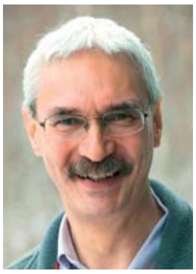

Prof. Gerhard Weikum

Prof. Michael Backes

Prof. Peter Druschel

Prof. Rupak Majumdar

Professor Gerhard Weikum won the most prestigious of the European Research Council awards together with his colleagues, **Professor Michael Backes** (Saarland University and Max Planck Fellow), and **Professor Peter Druschel** and **Professor Rupak Majumdar** (both from the Max Planck Institute for Software Systems). This team was one of only 13 winners out of approx. 450 proposals in 2013.

They will receive around. 10 million euros for their project *imPACT* which explores how to keep Internet users safe from spying and fraud, and unmasks perpetrators while maintaining commerce, freedom of speech and open access to information. *imPACT* stands for *"Privacy, Accountability, Compliance, and Trust in Tomorrow's Internet"*.

[www.impact-erc.eu/] :::

Personal Data

New Senior Researcher 2013/14

Senior Researcher are appointed by unanimous decision of the institute's directors when their performance is outstanding. Part of this is an evaluation by independent external reviewers. This appointment is one of the milestones in the academic career; it includes the right to guide Ph.D. students up to the doctorate. During the report period the following young scientists got appointed:

Klaus Berberich | DEPT. 5
Information Retrieval and Data Mining

Andreas Bulling | DEPT. 2
Machine Learning and Pattern Recognition

Piotr Didyk | DEPT. 4
Displays, Computational Fabrication, and Perception

Mario Fritz | DEPT. 2
Scalable Learning and Perception

Klaus Hildebrandt | DEPT. 4
Applied Geometry

Martin Hoefer | DEPT. 1
Disign of Algorithms and Combinatorial Optimization

Andreas Karrenbauer | DEPT. 1
Discrete Optimization

Michael Kerber | DEPT. 1
Topological and Geometric Computing

Christoph Lenzen | DEPT. 1
Theory of Distributed Computing

Tobias Marschall | DEPT. 3
Algorithms for Computational Genomics

Pauli Miettinen | DEPT. 5
Data Mining

Nico Pfeifer | DEPT. 3
Statistical Learning in Computational Biology

Tobias Ritschel | DEPT. 4
Rendering and GPUs

Marcel Schulz | DEPT. 3
High-throughput Genomics and Systems Biology

Jürgen Steimle | DEPT. 4
Embodied Interaction

Thomas Sturm | RG. 1
Arithmetic Reasoning

Sun He | DEPT. 1
Randomized Algorithms and Computational Geometry

Jilles Vreeken | DEPT. 5
Exploratory Data Analysis

Andreas Wiese | DEPT. 1
Combinatorial Optimization and Approximation Algorithms

Offers

The following young scientists of our institute received faculty offers during the report period:

Khaled Albassari | DEPT. 1
University of Abu Dhabi

Gilles Bailly | DEPT. 4
CNRS France

Syan Bhattacharya | DEPT. 1
Institute of Mathematical Sciences Chennai

Martin Cadik | DEPT. 4
Brno University of Technology

Carola Doerr | DEPT. 1
CNRS France

Philip Geevarghese | DEPT. 1
Institute of Mathematical Sciences Chennai

Rainer Gemulla | DEPT. 5
University of Mannheim

Klaus Hildebrandt | DEPT. 4
TU Delft

Matthias Hullin | DEPT. 4
University of Bonn

Arjun Jain | DEPT. 4
IIT Bombay

Kwang In Kim | DEPT. 4
University of Lancaster

Nicole Megow | DEPT. 1
TU Munich

Sebastian Michel | DEPT. 5
TU Kaiserslautern

Antti Oulasvirta | DEPT. 4
University of Aalto

Thomas Sauerwald | DEPT. 1
University of Cambridge

Jens M. Schmidt | DEPT. 1
TU Ilmenau

Marc Spaniol | DEPT. 5
University of Caen Basse-Normandie

Rob van Stee | DEPT. 1
University of Lancaster

Fabian Suchanek | Otto Hahn Group
Télécom ParisTech

Jacobo Urbani | DEPT. 5
University of Amsterdam

Magnus Wahlström | DEPT. 1
Royal Holloway University of London

Michael Wand | DEPT. 4
TU Utrecht and subsequently University of Mainz

Tino Weinkauf | DEPT. 4
KTH Stockholm

Shanshan Zhang | DEPT. 2
Nanjing University of Science and Technology

Grants

Scholarships provide a means of living wage, they are given to outstanding scientists. The purpose is focus on science. Some grants involve significant prestige and go along with a substantial step in one's career.

Some members and guests of our institute were supported by distinguished grant during the report period:

Samsung Scholarship

Seong Joon Oh got by the Samsung Scholarship the opportunity to study in the department of Professor Schiele. Since 2002 this scholarship sends most talented Korean students, of whom worldwide influence is expected in the future, to the best international research facilities. :::

Google Anita Borg Scholarship

Silke Jansen, DEPT. 4, was awarded an Anita Borg Scholarship. Since 2004, Google Inc. provides this funding for young female computer sciencist in memorial of Dr. Anita Borg. It should encourage the grantees to strive for outstanding results and leadership. Part of the scholarship is a dedicated retreat besides the financial support. :::

Humboldt Research Fellowship

The Humboldt Foundation awards outstanding post docs scholarships as long as they are at the begin of their academic career – up to four years after their Ph.D. Selected researchers from abroad are supported for 6 to 24 month so that they can continue their research in Germany. They are unrestricted in topic and also permitted to freely chose the German institution where they want to stay during that period of time. **Anna Adamaszek, Ran Duan** and **Artur Jez**, three post docs in the department of Kurt Mehlhorn as well as **Yusuke Sugano** with Bernt Schiele successfully applied for this scholarship. :::

Humboldt Research Award

Professor Wolfgang Heidrich, King Abdullah University of Science and Technology in Thuwal, Saudi Arabia was supported by the Humboldt Foundation with a Humboldt Research Award. Owing to his recent results in computer based photography and display the foundation awarded him 60,000 Euros and invited him to a research visit in Germany. As it is the rule with Humboldt Scholarships the scholar can freely chose his host. Since September 2014 the Max Planck Institute for Informatics accommodates him. :::

Alexander von Humboldt
Stiftung/Foundation

Non-academic Activities

Patents

The basic research at the Max Planck Institute for Informatics at times produces results that can be directly exploited for commercial purposes. If the results prove a high level of inventiveness, it is possible to apply for a patent on the new invention. The table below provides an overview of the activities concerning patents, licenses, and inventions for the Max Planck Institute for Informatics in the year 2014. :::

Topic	Activity	Persons involved
A Clipping-free Algorithm for Efficient HW-Implementation of Local Dimming LED-Backlight	Assignment / exploitation agreement	Karrenbauer
Videoscapes: Exploring Sparse, Unstructed Video Collections	Invention reported	Kim / Theobalt / Tompkin / Kautz
Method and System for Tracking an Object in a Sequence of Digital Video Images	Invention reported	Stoll / Theobalt / Seidel
A Perceptual Mode for Disparity	Patent management and exploitation agreement	Didyk / Ritschel / Eisemann / Myszkowski / Seidel
Method and Apparatus for Encoding High Dynamic Range Video	Patent extension	Mantiuk / Krawczyk / Myszkoski / Seidel
Apparant Display Resolution Enhancement for Moving Images	Patent application and license agreement	Didyk / Eisemann / Ritschel / Myszkowski / Seidel
Backward Compatible High Dynamic Range MPEG Video Encoding	License agreement	Efremov / Mantiuk / Krawczyk / Myszkowski / Seidel
A System for Automatic Material Suggestions for 3D Objects	Invention reported	Jain / Thormählen / Ritschel / Seidel

Cooperations

Intel Visual Computing Institute

Five years ago, the IVCI was founded as the largest Intel research institute in Europe. As there is a higher density of computer scientists in Saarland than anywhere else in the world, Intel decided that Saarbrücken was the right place for its new institute. The IVCI is headed by a Governance Board on which Intel and its four Saarland-based partners (Saarland University, the DFKI and both Max Planck Institutes – MPI-INF and MPI-SWS) are equally represented. For an initial period of five years, twelve million euros have been provided by the high-tech enterprise. The projects are selected by a specially established steering committee, with Prof. Christian Theobalt as the representative from the Max Planck Institute for Informatics. Two of his projects, *"Markerless Motion Capture"* and *"User-centric Video Processing"*, are carried out within the Intel Institute. [www.intel-vci.uni-saarland.de]. :::

Indo-German Max Planck Center for Computer Science

The *"Indo-German Max Planck Center for Computer Science"* (IMPECS) was inaugurated on February 3rd, 2010 at the Indian Institute of Technology in Delhi (IIT Delhi) by then German Federal President Horst Köhler. By encouraging close cooperation between Indian and German scientists, the center seeks to promote excellent basic research in the field of computer science. Joint research, exchanges between PhD candidates and postdocs, a Max Planck visiting professorship and a large number of workshops and schools are intended to foster the collaboration. For an initial period of five years, the IMPECS will be funded by the Max Planck Society (MPG), the German Federal Ministry of Education and Research (BMBF) and the Indian Department of Science and Technology (DST).The board of directors consists of Kurt Mehlhorn (MPI-INF) and Rupak Majumdar (MPI-SWS) from the Max Planck Society and Naveen Garg (IIT Delhi) and Manindra Agrawal (IIT Kanpur) from India.

Approximately ten Indo-German research groups are funded by IMPECS. From the Indian side, any university/institute can apply to become a partner, whereas German applicants are required to be partners of the MPI for Informatics or the MPI for Software Systems. Currently, eight binational research groups are receiving support and further groups are welcome to apply.

The IMPECS allows for an active exchange between postdocs and students in both directions. For example in summer 2011 professor Manindra Agarwal from IIT Kanpur visited the Max Planck Institute for Informatics for several month. In turn, professor Helmut Seidl of Munich, Technical University of Munich, spent two month at IISc Bangalore in fall 2011.

Furthermore, the IMPECS co-financed the *„Complexity Theory"* workshop in Kanpur and the *„Recent Trends in Algorithms and Complexity"* in Mysore. For further information, refer to the website [www.impecs.org]. :::

Bertram Somieski

CeBIT 2013/2014

Information Technology: Saarland Presents its Abilities

Every year in March the doors of Hannover's Exhibition Site open for the CeBIT. Even if the visitor records of the decade following the millenium are gone, hundreds of thousands enter to inform themselves about news in the area of digital technique. Since 1970 the CeBIT is a international competitive exhibition having high density of innovations, high media coverage, and broad public interest.

The exhibition stands of research communities, federal states, and the Bundeswehr are united in hall no. 9 that changes it's motto every year. It is tradition that the MPI-INF is part of the joint stand of Saarlands Computer Science. There, Saarland University and its associated research institutes concertedly exhibit novelties to demonstrate to the public the abilities of Saarlands IT.

The research group of Professor Christian Theobalt presented MPI-INF at the CeBIT 2013. Its exhibit *"Capturing Reality in Realtime"* is a system that can detect and follow persons in a video stream in real time. This system does not require marker points or a particular studio environment; it works stable even if two persons occlude each other partially. Federal Minister of Economics and Technology Philipp Rösler awarded *"The Captury"*, a start-up company that was founded to promote and commercially exploit this, with a grand prize of the founders competition *"IKT Innovativ"*.

Annegret Kramp-Karrenbauer, Prime Minister of the Saarland, at the exhibit of Dr. Nils Hasler.

A PIN appears in the Google Glass headset display…

… and even the police are stunned.

In 2014 another exhibits was highly noticed that got jointly developed by MPI-INF and CISPA: By means of the smart glasses *"Google Glass"* the wearer is shown a PIN that can not be seen by anyone else. By this theft of PINs is prevented. The other parts of the of this exhibition that were provided by MPI-INF also draw attention. A computer program developed in the group of Tobias Ritschel arranges pictures according to artificial aspects into a composition. The department of Gerhard Weikum presented a new search engine that scannes the news and links its contents by semantic algorithms.

When top politicians of the Saarland visit the CeBIT, a stop at the stand in hall no. 9 is mandatory; Saarland's computer Science is considered to have Germany-wide importance. :::

Bertram Somieski

"The Captury" is awarded one of the grand prizes inthe startup competition *IKT Innovativ.*

IOI Training 2014/2015

Strict and Demanding Program Prepares Students for Olympiad in Informatics

The Max Planck Institute for Informatics, together with the German Federal Ministry of Education and Research,the German Informatics Society (GI), and the Fraunhofer IuK Group, constitute the responsible body of the German National Computer Science Competition [BWINF, http://www.bwinf.de/]. We are especially committed to educating successful students of the national competition so as to qualify them for international competitions. It is for this reason that we, together with Dr. Pohl, manager of the BWINF, also organized the final selection and training for the 2015 International Olympiad in Informatics (IOI).

Simon Bürger, Tobias Lenz, and Gregor Matl, all former successful Olympiad participants, trained the five selected students from May 26 to 29, 2015. As is the case every year, the program was strict and demanding, lasting from nine o'clock in the morning until late in the evening. We are very grateful to the organizers from Gästehaus Braunshausen, which is funded by both the Saarland Gymnastics and Soccer Associations, for providing us with accommodation and infrastructure at short notice.

A typical challenge of the Olympiad in Informatics consists of, first, developing an algorithm to a problem and, then, faultlessly implementing it. The participants are even required to have background knowledge that extends to topics only taught in graduate level courses.

A challenge of this year's training was the following: Player A thinks of a number between 1 and n, and player B has to guess the number. He is only allowed to ask questions of the type "Is the number less than k?". Player A answers with either *"Yes"* or *"No"*, but is allowed to lie once in the game. A *"Yes"* from player A costs x euros, a *"No"* y euros. As a function of x, y, and n, what is the cost optimal algorithm for player B to guess the number?

The training incorporates not only finding solutions to such problems, but also holding subsequent discussions of the participants' solutions and, thus, the acquiring new knowledge and the learning of new methods.

The workshop not only aims to help the participants be successful at the Olympiad, but seeks also to excite them about computer science research. For this purpose, academic presentations and discussions complemented the IOI-specific training. This time, the topics addressed were *"What and how can quantum computers perform?"* by Prof. Markus Bläser, and *"Open Problems in Automated Reasoning"* by Prof. Christoph Weidenbach. :::

Christoph Weidenbach

ADFOCS 2013 and 2014

Summer School: Young Talents Deepen Their Knowledge in Computer Science

The ADFOCS, Max Planck Advanced Course on the Foundations of Computer Science, has meanwhile become a classic among summer schools. Every summer master's and PhD students meet in Saarbrücken to deepen their knowledge in computer science theory. During five days in August, the young scientists are intensely schooled by university teachers of international reputation:

2013: **Lukasz Kowalik**, Institute of Informatics, University of Warsaw, **Dániel Marx,** Computer and Automation Research Institute, Hungarian Academy of Sciences (MTA SZTAKI), **Saket Saurabh**, Department of Informatics, University of Bergen

2014: **Sanjeev Khanna**, University of Pennsylvania, USA, **Aleksander Madry**, EPFL, Switzerland, **Yurii Nesterov,** Catholic University of Louvain, Belgium

Prof. Kurt Mehlhorn initiated this event of advanced education in 2000 to bring together young talents from all over the world with established top scientists. The latter are asked to include their current research topics in their lectures. This was meanwhile the 15th time this summer school has been held, an attest for the continuous interest of young scientists in the course. :::

Bertram Somieski

Computer Science Research Days

Challenging and Supporting Talented High School Students in Computer Science

Every year, in collaboration with the Computer Science Department at Saarland University, the German Research Center for Artificial Intelligence (DFKI), the Center for Bioinformatics and the companies AbsInt GmbH and Logic4Business GmbH, the Max Planck Institute for Informatics organizes the Computer Science Research Days for gifted pupils.

In 2013, 2014, and 2015, the most successful participants of the German Federal Competition in Computer Science (*"Bundeswettbewerb Informatik"*) as well as Rhineland-Palatinate's best recent high school graduates who had computer science as one of their main subjects were invited to the event. The students became acquainted with, and actively participated in, current computer science research topics.

The participants spent the first evening getting to know each other while bowling. They then attended lectures on specific themes, such as *"Malware Detection"*, and also lectures from the current curriculum of the computer science department. In different workshops, the participants were given the opportunity to build and test 3D scanners, or for example learn how to avoid bugs with disastrous consequences through static program analysis. The group work covered all areas of computer science, from the fundamentals, to programming, bioinformatics, graphics, and computer vision; a wide range of current computer science research was covered.

The motto of the Computer Science Research Days is: *"Challenging and supporting"*. With this event, the Max Planck Institute for Informatics aims to raise young people's enthusiasm for the subject of computer science and seeks to discover and support the development of new talent. :::

Jennifer Müller

Opening of the New Building of the Max Planck Institute for Software Systems in Kaiserslautern

The Max Planck Institute for Software Systems was founded in 2005 as the second institute of the Max Planck Society that is entirely devoted to computer science. The institute was set up from the beginning with two sites, Kaiserslautern and Saarbrücken. The Nürtingen based architect's office Weinbrenner-Single-Arabzadeh won with a double project, although each building had a separate competition. After the inauguration of the the first building in October 2012 in Saarbrücken, the Kaisers-lautern building was opened ceremonially on 2 July 2013.

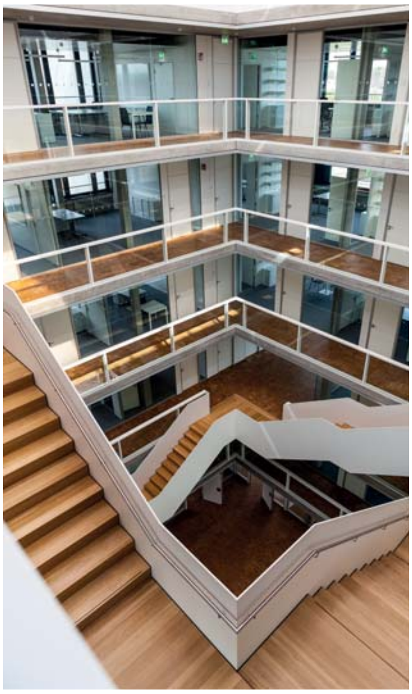

As required, the building's architectonic scheme focusses on the support of communication: Meandering internal flights of stairs connect six floors. Each floor consists of peripherical offices that encircle a square based atrium that is flooded with light. An innovative ceiling construction provides additional light. In addition to the approximately 100 office workplaces and 6 conference rooms, the atrium can serve as a large event hall. This, combined with up-to-date communication equipment, provides the institute with all that is needed for a modern building. Located at the entrance of the university campus, the building is an architectonic landmark. :::

Bertram Somieski

Meeting of the Board of Curators

Support, Inspiration, Discussions and Suggestions

Most Max Planck Insitutes have a Board of Curators. It consists of representatives of politics, science, and societal actors. This board is a link between science and society. During the yearly meetings the institute's matters are brought to the curators as well as strategic research topics explained. These information about the institute's occurrences are confidence-building measures to ensure autonomously organized research in Germany. The names of the current curators are listed on page 4 of this report.

The last meeting of the Board of Curators was held on 05th February 2013. Besides Professor Wolffried Stucky as chairman participated the curators Annegret Kramp-Karrenbauer, Saarland's Prime Minister, Dr. Siegfried Dais, partner of Robert Bosch Industrietreuhand KG, and Christiane Götz-Sobel, Chief Editor Science and Technique of Second German Television. Professor Bernt Schiele, managing director of the MPI-INF, referred the institute's current status in science and society and presented a strategic outlook to further activities. He thanked the Board of Curators for support and inspiration, for criticism and suggestions.

Professor Stucky presenting Nicole Megow the certificate as a Senior Researcher.

The institute's activities to establish a self-sustaining science center were referred about by Kurt Mehlhorn; Hans-Peter Seidel informed on Saarland's part of the German Excellence Initiative.

Four young scientists that were advanced to Senior Scientists during the previous year presented their scientific topics in short talks. Olga Kalinina, Antti Oulasvirta, Thomas Sauerwald, and Nicole Megow demonstrated to the curators the scientific width and depth of the institute's research. Subsequently, they received their certificate of appointment by the chairman. The scientific part of the meeting was completed by a poster session where young scientists of the MPI-INF presented their results.

The Board thanked Simone Bischoff, Team Leader of the institute supervisors at the Administrative Headquarters and Curator Board Meeting attendee for many years, for her service and welcomed her successor, Dr. Christoph Freudenhammer. Bernt Schiele, Managing Director, also expressed words of gratitude for a long-standing, trustful and productive collaboration. :::

Bertram Somieski

Scientific Review 2015

Scientific Lectures, Poster Sessions and Meetings with the Senior Researchers

The Scientific Advisory Board and the Senior Researchers

The Max Planck Society is mainly funded by the German federal and state governments. However, it is exclusively obliged to strive for scientific excellence in science, which is checked and evaluated on a two-yearly basis by the Scientific Advisory Board for each institute. The Scientific Advisory Board is composed of leading scientists from the fields of research in which the respective institute specializes. New members are regularly appointed to the Scientific Advisory Board: The maximum term of office of each member is six years.

Poster session: presentation of the scientists' work

The members of the Scientific Advisory Board, appointed by the President of the Max Planck Society, are:

Prof. Dr. Trevor Darell
of Berkeley University,
Prof. Dr. Nir Friedman
of the Hebrew University of Jerusalem,
Prof. Dr. Pascal Fua
of the Swiss Federal Institute of Technology in Lausanne,
Prof. Dr. Jürgen Giesl
of RWTH Aachen,
Prof. Dr. Alon Halevy
of Google Research in Mountain View, California,
Prof. Dr. Yves Moreau
of Leuven University,
Prof. Dr. Nicole Schweikardt
of Humboldt University Berlin,
Prof. Dr. François Sillion
of INRIA Grenoble Rhône-Alpes and
Prof. Dr. Emo Welzl
of ETH Zurich.

2015 they were completed by:

Prof. Dr. Renée J. Miller
Toronto University,
Prof. Dr. Gerhard J. Woeginger
of Eindhoven University, and
Prof. Dr. Thomas A. Funkhouser
of Princeton University.

The review started with presentations of the directors and selected senior researchers. It was followed by one-to-one conversations between the members of the Scientific Advisory Board and the senior researchers of the institute. During two poster sessions, the work of all scientists was presented to the Scientific Advisory Board. The picture shows the Scientific Advisory Board together with the senior researchers at the end of the first day of the review. :::

Christoph Weidenbach

Visit of the Ambassador of India and the Indian General Consul

Fostering Contact with the Institute and with Indian Compatriots Living in Germany

The Indian Ambassador, Sujata Singh, and the Indian General Consul in Frankfurt/M, Raveesh Kumar, visited the Saarland on 17 April 2013 and 04 February 2014. The purpose of their visit was to meet with their fellow countrymen and maintain business and academic contacts. In open discussions, they asked for feedback about life circumstances, social experiences, and working conditions of Indian students in the Saarland. :::

Bertram Somieski

Industrial Fair

Computer Science Explained in a Comprehensible Way

Fundamental research has the reputation of being hard to understand or difficult to apply – communicating the progress of certain areas of computer science is hence a great challenge. Current findings and developments should be, at least in part, transparently conveyed to the public at an industrial fair on 25 September 2013, where local and regional companies are the target audience. These are offered cooperations with science institutions on Saarland University's Campus and given inspiration for their own developments. The fair was also open to the public.

In addition to the institute's focus on fundamental research, this event should present applicable results of this research. Hence, our institute presented itself together with the computer science departments of the University and the DFKI in an big industrial exhibition on 25 November 2013. By 56 exhibits, displayed in the foyers of both Max Planck Institutes, the building of the Cluster of Excellence, and the DFKI, the thematic diversity of the entire Saarland computer science campus was presented. The fair was opened by a talk entitled *"Multimodal Computing and Interaction"* given by Professor Hans-Peter Seidel. :::

Bertram Somieski

Science exhibition in Saarbrücken

“Digital – An Index Ahead” – Research You Can Touch at Hauberrisser Hall

As part of the science year *“Digital Society”* an exhibition titled *“Digital – An index ahead”* was organized by the city of Saarbrücken. The Hauberrisser Hall in the town hall provided an appealing location for the computer science related institutions of the universities an affiliated institutions to display their topics.The focus was on application-oriented research results that were adopted into easy understandable exhibits, giving easy insights to non-professionals.

For a whole month, from 25th September to 25th October, random guests, visitors of the town hall, or registered tour groups could rend sorting algorithms by themselves, try their skills by optimizing package volumes, or compete with a computer to find the shortest round-trip route. Six exhibits presented the MPI-INF, depicting a colorful spectrum of research topics. Every Tuesday an *“open internet meeting for senior citicens”* was offered, including a guided tour to the exhibition.

A panel discussion part was scheduled to be part of the exhibition. Owing to the then-current situation, the chosen topic was *“Pandemic hazard by infected air travelers?”* The results of Glenn Lawyer, postdoc in Thomas Lengauer’s department, made clear this threat really exists, however the ebola epidemic at that time was no danger to Central Europe. :::

Bertram Somieski

Eine Initiative des Bundesministeriums
für Bildung und Forschung

Wissenschaftsjahr 2014

DIE DIGITALE
GESELLSCHAFT

Launch of the Institute’s New Website

After intense preparation and design, the 22th May 2014 saw the launch of the institute’s new website. The concept was realized in cooperation with a local advertising company, creating a modern display of our institute. The result is a easy-to-navigate and clear structured website that reflects the institute’s structure, offerings, and scientific results. :::

Bertram Somieski

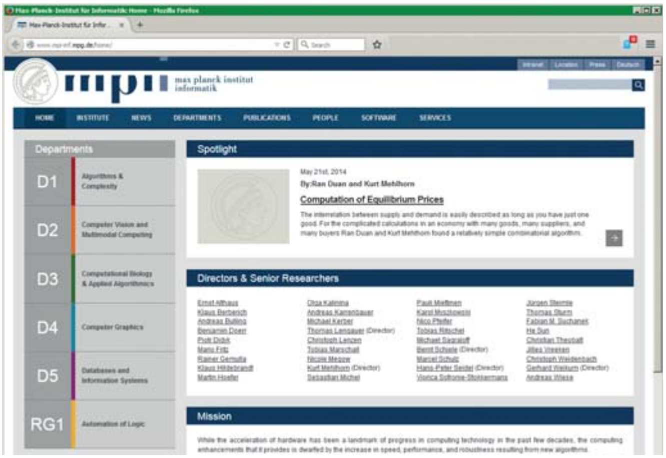

Spotlight

May 21st, 2014

By:Ran Duan and Kurt Mehlhorn

Computation of Equilibrium Prices

The interrelation between supply and demand is easily described as long as you have just one good. For the complicated calculations in an economy with many goods, many suppliers, and many buyers Ran Duan and Kurt Mehlhorn found a relatively simple combinatorial algorithm.

Directors & Senior Researchers

Ernst Althaus
Klaus Berberich
Andreas Bulling
Benjamin Doerr
Piotr Didyk
Mario Fritz
Rainer Gemulla
Klaus Hildebrandt
Martin Hoefer
Olga Kalinina
Andreas Karrenbauer
Michael Kerber
Thomas Lengauer (Director)
Christoph Lenzen
Tobias Marschall
Nicole Megow
Kurt Mehlhorn (Director)
Sebastian Michel
Pauli Miettinen
Karol Myszkowski
Nico Pfeifer
Tobias Ritschel
Michael Sagraloff
Bernt Schiele (Director)
Marcel Schulz
Hans-Peter Seidel (Director)
Viorica Sofronie-Stokkermans
Jürgen Steimle
Thomas Sturm
Fabian M. Suchanek
He Sun
Christian Theobalt
Jilles Vreeken
Christoph Weidenbach
Gerhard Weikum (Director)
Andreas Wiese

Mission

While the acceleration of hardware has been a landmark of progress in computing technology in the past few decades, the computing enhancements that it provides is dwarfed by the increase in speed, performance, and robustness resulting from new algorithms

Summer Party IMPRS-CS

Celebrating and Enjoying a Culturally Diverse Environment

The International Max Planck School for Computer Science supports and guides foreign students during their studies of computer science at Saarland University. Each summer the IMPRS-CS invites the members of both Max Planck Institutes to the summer party that is held on the Platz der Informatik. Complementary to barbecue and beverages originating in the Saarland, the international party guests cook dishes of their home countries, provide baked goods, and offer exotic salads.

Guests were entertained by a colourful programme that had activities for toddlers, live music and some contributions of scientists; ball games or dancing covered physical interests. An especially challenging and sudorific affair was the one-hour Zumba performance – four young woman demonstrated a choreography on-stage. :::

Bertram Somieski

The Saarland Business Run

The „Planck Quanten“ had a Good Run

The Saarland Business Run enjoys great popularity since its founding in 2005. It is the sports event in the Saarland attracting the highest number of participants by assembling a field of 15.000 out of bare 1 million Saarlandians. Companies and housing communities, medical offices and schools send teams of four who's attitudes vary between scores, fancy costumes, or plain fun.

Since many years running enthusiasts from both Max Planck Institutes participate in this event displaying strong spirit. In defiance of the huge crowd on the starting line and four live bands, amusing and diverting runners on the course, they resist depending on the weather heat, drenches, or cold winds during the circuit through Dillingen's city and the adjacent steel mill. In summer 2014 team “Planck Quanten #2” scored place 93 out of 1276 in the category mixed teams! Lili Jiang topped that with #60 among all appr. 5000 women. :::

Bertram Somieski

The Institute in Figures

Budget without third-party funds 2011 to 2014

Staff Current as of 1/1/2015

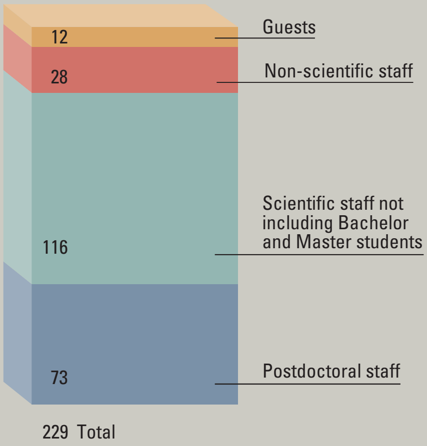

Ratio of scientific to non-scientific staff members Current as of 1/1/2015

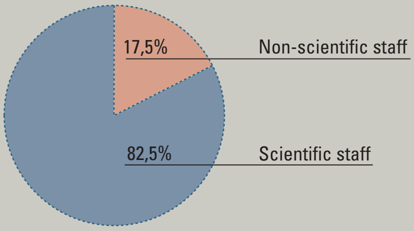

Third-party funding from 2011 to 2014 Numbers and distribution

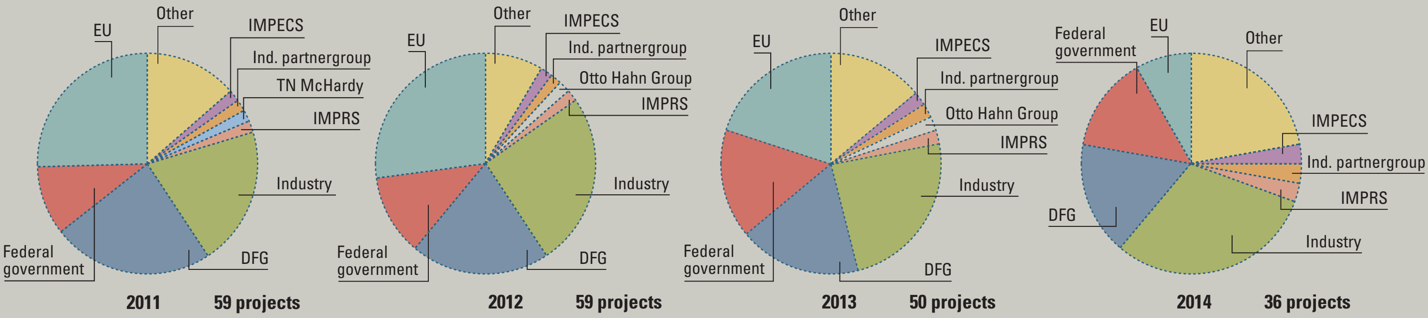

Third-party funding from 2011 to 2014

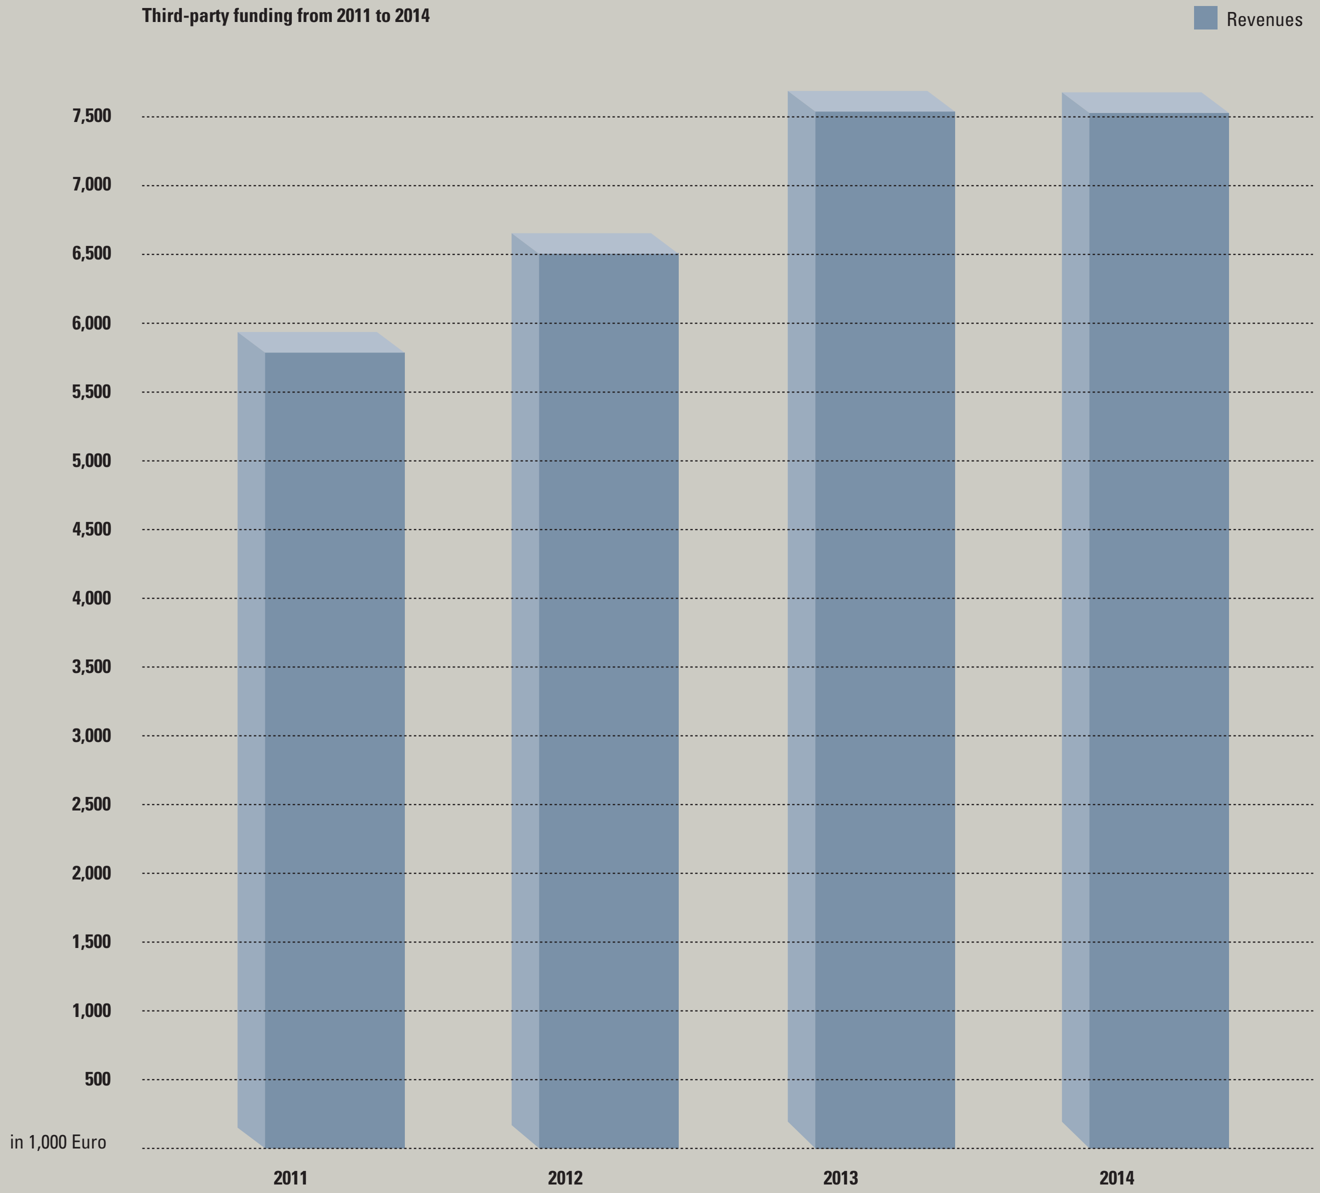

CONTACT

Volker Maria Geiß
Joint Administration
Phone +49 681 9325-7000
Email geiss@mpi-inf.mpg.de

Joint Administration of the Max Planck Institutes for Informatics and Software Systems

MPI for Informatics Saarbrücken

JOINT ADMINISTRATION

MPI for Software Systems Saarbrücken

MPI for Software Systems Kaiserslautern

When the Max Planck Institute for Software Systems was founded, the two Max Planck Institutes decided to have a joint administration. This administration comprises all services for both institutes: personnel and financial services, general services, public relations, building services, and the library. The library is shared by all computer science education and research institutions on the university campus.

In 2011, the information services and technology group was integrated into the joint administration as well.

Excellent research requires excellent service. This guiding principle is the joint administration's highest maxim in all administrative and scientific service activities that it provides for both the Max Planck Institute for Informatics and the Max Planck Institute for Software Systems Kaiserslautern/Saarbrücken.

The joint administration offers the scientists of both institutes a wide range of support on matters surrounding their research activities. Special support goes to scientists beginning and ending their tenure at the institute.

Fostering an open and friendly interaction of scientists and guests from around the world is another high standard the administration strives to uphold. This encompasses support in bureaucratic and social matters for foreign scientists and visitors living in Germany.

This all-round support is only possible because all administrative staff members are committed to practising the idea of service far beyond their required duties. For this, they deserve the deep appreciation of all institute members. :::

CONTACT

Volker Maria Geiß
Joint Administration
Phone +49 681 9325-7000
Email geiss@mpi-inf.mpg.de

Information Services and Technology

The aim of the Institute to produce first-class research results is fostered by unhindered global cooperation and communication in a motivating environment with a flexible, high-quality, reliable, and easy-to-use equipment.

These properties apply as well to our IT infrastructure: we operate a versatile system that can adapt to rapidly evolving requirements while providing consistency and reliability at the same time. The system does not neglect security, despite the openness it requires to support international cooperation.

Multiplicity of tools

Research on the front line in computer science often implies using innovative systems from various sources, including prototype and cross-platform setups in hard- and software. As a consequence, we provide heterogeneous systems with hardware from various manufacturers and all major operating systems: MacOS for clients, Solaris for file servers and Linux and Windows for client and server systems. To improve the homogeneity of these systems and their ease of use, we try to make user data and the most important software packages available independently of platform and OS.

Automation as a guarantee for reliability

Regular updates and upgrades of software and hardware for all supported platforms pose substantial demands on the reliability of the installation and administration. The mere variety of hardware and software components precludes image based installation approaches. Instead, we use an extensive, automated and partly self-developed package-based installation, configuration, and adminstration system that is designed to take the requirements of the researchers and their projects into account.

Once implemented, results can be repeated as required and used very quickly across the whole infrastructure. This implementation work requires effort and, in some cases, slows down the realization. The advantages for operating security, total expenses and time until larger changes are ready, however, clearly outweigh this disadvantage.

The system is flexibly structured, so it can be customized to suit changing requirements (such as hardware requirements) quickly through specific extensions of the installation system. However, it is only used for repetitive tasks. Protoypical installations are done manually.

Security and protection

Fixed mechanisms that protect against sabotage and espionage are not possible in open systems. They greatly limit usability. The security guidelines can therefore be nothing more than a compromise that flexibly follows the requirements.

Some direct hazards can be deflected through the structure of the network, the firewall, the encryption for external access, or virus scanners. Indirect hazards, such as the connection to virus-infected computers in the Intranet or faults in externally acting software systems, must be opposed by keeping our local software installations and virus scanners up to date.

Operational reliability

Power supply and cooling are designed in such a way that server operation can be maintained, even during a power outage. A generator with a power output of roughly one megawatt guarantees uninterrupted operation of the core infrastructure even during long outages.

Monitoring systems based on open source software provide information via email and text messaging on critical states of the server systems, the network, and on malfunctions of complex processes in terms of operational and safety aspects.

Cooperation and communication

Our network is divided into various areas, according to organizational and security-relevant points. Externally accessible, but anonymous services such as DNS (Internet address book), WWW, FTP (data transfer), and SMTP (email) are combined at the institute's firewall in several demilitarized zones (DMZs), which are differentiated according to their meaning and their hazard potential.

Guest scientists and students can connect their own devices by wire or via WiFi (BYOD). The network infrastructure automatically treats them as external machines.

International cooperation demands external access to internal resources in the infrastructure (Intranet). Here, we offer secure login and access to email, important databases, and other services. Cooperation in software development is supported by protected access to several revision control systems (software repositories: Subversion, Git, and Github).

This logical structure is spread over a multi-10 GB backbone, consisting of various floor and data center switches. The backbone is extended with Gigabit and 10 GBit leased lines to each of our locations and several partners on the campuses in Saarbrücken and Kaiserslautern. Internet connectivity is realized with a 10 GB connection to the German National Research and Education Network (DFN) that is used jointly with the university in Saarbrücken.

Workstations and notebooks are powered via their floor switch with Gigabit Ethernet. Central servers, server farms and compute clusters are using 1GB or 10GB connections according to their bandwidth requirements. Besides workstations and notebooks, most of these connections are set up to be fail-safe via multiple redundant lines.

Compute service

We operate several larger systems that work with up to 64 closely coupled processors and have up to one and a half terabyte of main memory. These machines are used for scientific research and applications which require high parallelism and uniform access to a large main memory.

Our largest cluster to date, consisting of 128 systems with 8 cores and 48 GB of main memory each, was brought into operation in mid-2010. It is operated under the Grid Engine (GE). The automatic distribution of processes

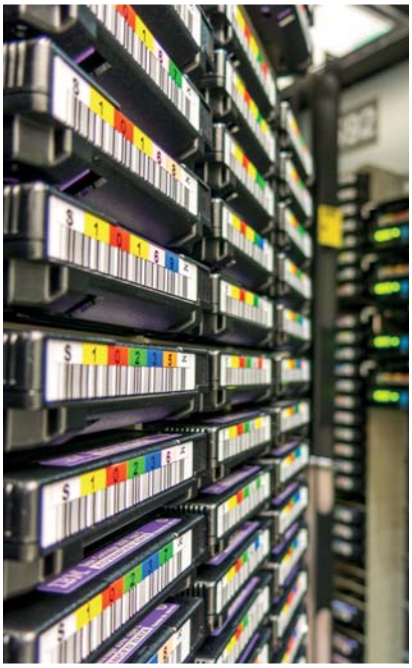

on the clusters' individual computers allows us to reach high utilization on the whole system. We can flexibly integrate the above-mentioned larger systems into this system. In order to allow their optimal distribution, jobs are categorized according to their specific requirements. The prioritization of specific processes helps with the appropriate and fair distribution of resources across all the waiting jobs.

In 2014, this setup was extended by a new cluster , which consists of 12 systems , each with 20 cores and 256 GB of main memory. Additionally, each of these systems is equipped with two server GPU cards (Graphic Processor Unit) equipped with 12 GB of memory each. Specially designed software can thus utilize the high computing power and high throughput provided by this technology.

File service and data security

The roughly 560 TB of institute data is made available via NFS and CIFS (SMB). Currently, 37 file servers distribute the volumes of about 70 RAID systems with more than 4000 disks. All RAID systems are connected to a redundant Storage Area Network (SAN). In order to prevent data loss due to a failed RAID system, the RAID volumes are operated in a paired mirroring by the ZFS file system. ZFS utilizes checksums to protect data against creeping changes. It decides during the read operation whether a data block is correct or not and in case it is not correct, it decides to use the corresponding block of the second part of the mirror to correct it. This implies as well that time consuming file system checks prior to the access of the data are unnecessary. File systems are inspected during normal service operation.

The two mirrored halves are accommodated in at least two different fire compartments so that a fire would have little chance of causing data loss, even at the deepest level. Three file servers share a view of a disk's status and can substitute for each other within a very short period of time using virtual network addresses and SAN technology. This also allows for server updates without any perceivable interruptions to service.

Our data security is based on two different systems: a conventional tape backup, which stores the data directly on a tape robot, and an online disk backup system that minimizes space requirements by using data comparison and always keeps the data online (open source system BackupPC). Because no data in this system is actually held with unnecessary redundancy, about 450 TB of gross data can be reduced to about 30 TB in the different backup runs.

The tape robot currently has the ability to access 700 tapes with a total capacity of 700 TB (uncompressed). The robot is located in a specially-prepared room to maximize the protection of its high-value data. We recently extended the system by a second robot with a capacity of almost 5 PB (uncompressed) on about 1200 tapes. The new robot operates at a remote site to protect the data against disaster situations. This is also achieved by the storage of tapes in a special fire-proof safe.

Special systems

Diverse special systems are required for special research tasks, especially in the area of computer graphics. Available systems include a video editing system, several 3D scanners, multivideo recording systems, and 3D projection systems. External communication and collaboration is supported through the operation of several videoconference installations.

Flexible support of scientific projects

The described services, servers, and computing clusters are deployed for many scientific projects in very different scenarios, even for collaborations such as the 'The German epigenome programme "DEEP"' [http://www.deutsches-epigenom-programm.de]. In order to meet the various requirements, the IST (Information Services and Technology) offers specifically tailored degrees of support that extend from pure server hosting in individual cases to applications support. The division of tasks is in most cases dictated by the intersection of project requirements and the IST portfolio. For problems, for example, with self-administered systems in hosting scenarios, the IST only provides consultation.

Responsibilities

IT procurement, installation, administration, operation, application development and support, and continuation of the described systems and techniques are the duty of the IST, a subdivision of the Joint Administration of the Institutes for Software Systems and Informatics. Due to the Institutes' cooperation with several departments and institutes of the university, the IST is also responsible for the campus library ('Campusbibliothek für Informatik + Mathematik') and the Cluster of Excellence 'Multimodal Computing and Interaction' (MMCI).

The IST is divided into one core and several institute-specific support groups. Put simply, the core group is responsible for services that are identical for both institutes or even operated together. The support groups, accordingly, cover the specific needs of the individual institutes and the services for their respective researchers.

Staff structure

In addition to management and procurement (two positions), six scientific staff members and two technicians work on "core" issues. There is also one employee for each institute who takes care of administrative tasks to support the joint library and MMCI.

For their service desk, the institute's support group (two scientific staff members and one technician) is assisted by a team of students. The service desk is reachable either by email, through a web interface, or in person during business hours. In addition to processing questions concerning the use of the infrastructure, this group also maintains information systems such as a documentation wiki. :::

CONTACT

Jörg Herrmann
Joint Administration
IT Core
Phone +49 681 9325-5800
Email jh@mpi-klsb.mpg.de
Internet http://www.mpi-klsb.mpg.de/gvw/ist.html

Selected Cooperations

UNDERSTANDING IMAGES & VIDEOS

::: German Research Center for Artificial Intelligence, *Kaiserslautern, Germany*

::: ETH Zurich, *Zurich, Switzerland*

::: IST Austria, *Klosterneuburg, Austria*

::: Katholieke Universiteit Leuven, *Leuven, Belgium*

::: Max Planck Institute for Intelligent Systems, *Tübingen, Germany*

::: Microsoft Research, *Cambridge, UK*

::: RWTH Aachen, *Aachen, Germany*

::: Stanford University, *Stanford, USA*

::: Technicolor Paris, *Paris, France*

::: TU Darmstadt, *Darmstadt, Germany*

::: Tsinghua University, *Beijing, China*

::: University of British Columbia, *Vancouver, Canada*

::: University College London, *London, UK*

::: University of California, *Berkeley, USA*

::: Saarland University, *Saarbrücken, Germany*

::: University Tel Aviv, *Tel Aviv, Israel*

BIOINFORMATICS

::: BioSolveIT GmbH, *Sankt Augustin, Germany*

::: CeMM Research Center for Molecular Medicine, *Vienna, Austria*

::: Heinrich Heine University, *Düsseldorf, Germany*

::: Informa s.r.l, *Rome, Italy*

::: Karolinska Institute, *Stockholm, Sweden*

::: Research Center for Natural Sciences, Hungarian Academy of Sciences, *Budapest, Hungary*

::: Siemens Healthcare Diagnostics, *Berkeley, USA*

::: University of Basel, *Basel, Switzerland*

::: University of Heidelberg, *Heidelberg, Germany*

::: University of Cologne, *Cologne, Germany*

::: Saarland University, *Saarbrücken, Germany*

::: Viiv Healthcare, *Brentford, UK*

::: Wolfgang-Goethe-Universität, *Frankfurt, Germany*

GUARANTEES

::: INRIA Nancy, *Nancy, France*

::: Massachusetts Institute of Technology, *Cambridge, USA*

::: Tata Institute of Fundamental Research, *Mumbai, India*

::: TU Munich, *Munich, Germany*

::: TU Vienna, *Vienna, Austria*

::: University of Freiburg, *Freiburg, Germany*

::: University of Oldenburg, *Oldenburg, Germany*

::: University Tel Aviv, *Tel Aviv, Israel*

INFORMATION SEARCH & DIGITAL KNOWLEDGE

::: Carnegie-Mellon University, *Pittsburgh, USA*

::: Google, *Mountain View, USA*

::: Google, *Zurich, Switzerland*

::: IBM Almaden Research Center, *San Jose, USA*

::: IIT Bombay, *Bombay, India*

::: IIT Delhi, *Delhi, India*

::: Max Planck Institute for Software Systems, *Kaiserslautern and Saarbrücken, Germany*

::: Microsoft Research, *Redmond, USA*

::: Paris Telecom University, *Paris, France*

::: Siemens AG, *Munich, Germany*

::: Tsinghua University, *Peking, China*

::: University of Amsterdam, *Amsterdam, Niederlande*

::: University of California, *Santa Cruz, USA*

::: University of Glasgow, *Glasgow, UK*

::: Saarland University, *Saarbrücken, Germany*

MULTIMODAL INFORMATION & VISUALIZATION

::: Albert-Ludwigs-University Freiburg, *Freiburg, Germany*

::: Carnegie Mellon University, *Pittsburgh, USA*

::: Centre National de la Recherche Scientifique, *Lyon, France*

::: Czech Technical University, *Prag, Czech Republic*

::: ETH Zurich, *Zurich, Switzerland*

::: Harvard University, *Cambridge, USA*

::: Keio University, *Minato, Japan*

::: KIT, *Karlsruhe, Germany*

::: Lancaster University, *Lancaster, UK*

::: LMU Munich, *Munich, Germany*

::: Massachusetts Institute of Technology, *Cambridge, USA*

::: Mircosoft Research, *Seattle, USA*

::: Stanford University, *Palo Alto, USA*

::: Technical University Delft, *Delft, Netherlands*

::: TU Dresden, *Dresden, Germany*

::: TU Munich, *Munich, Germany*

::: University of Erlangen-Nuremberg, *Erlangen, Germany*

::: University of Helsinki, *Helsinki, Finland*

::: University of Magdeburg, *Magdeburg, Germany*

::: University of Cambridge, *Cambridge, UK*

::: University College London, *London, UK*

::: Zuse Institute Berlin, *Berlin, Germany*

OPTIMIZATION

::: Aalto University, *Aalto, Finland*

::: Carnegie Mellon University, *Pittsburgh, USA*

::: IIT Delhi, *Delhi, Indien*

::: KTH Royal Institute of Technology, *Stockholm, Sweden*

::: Logic4Business GmbH, *Saarbrücken, Germany*

::: Siemens AG, *Vienna, Austria*

::: University of Santiago de Chile, *Santiago, Chile*

::: University of Tel Aviv, *Tel Aviv, Israel*

::: University of Valparaíso, *Valparaíso, Chile*

Selected Publications

[1] *Spectral Graph Reduction for Efficient Image and Streaming Video Segmentation*, 2014, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Oral.

[2] A. Adamaszek and A. Wiese. Approximation schemes for maximum weight independent set of rectangles. In *Proceedings of the 54th Annual Symposium on Foundations of Computer Science (FOCS)*, 2013, pp. 400–409. IEEE.

[3] M. Ahmad, V. Helms, T. Lengauer, and O. V. Kalinina. How Molecular Conformational Changes Affect Changes in Free Energy. *Journal of Chemical Theory and Computation,* 11(7):150629134059007, 2015.

[4] A. Antoniadis, N. Barcelo, M. Nugent, K. Pruhs, and M. Scquizzato. Energy-efficient circuit design. In *Innovations in Theoretical Computer Science, ITCS'14, Princeton, NJ, USA, January 12-14, 2014,* 2014, pp. 303–312.

[5] A. Antoniadis, C. Huang, and S. Ott. A fully polynomial-time approximation scheme for speed scaling with sleep state. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015,* 2015, pp. 1102–1113.

[6] Y. Assenov, F. Müller, P. Lutsik, J. Walter, T. Lengauer, and C. Bock. Comprehensive analysis of DNA methylation data with RnBeads. *Nature Methods,* 11(11):1138–1140, 2014.

[7] M. Bachynskyi, G. Palmas, A. Oulasvirta, and T. Weinkauf. Informing the design of novel input methods with muscle coactivation clustering. *ACM Transactions on Computer-Human Interaction (TOCHI),* 21(6):30, 2015.

[8] G. Bailly, A. Oulasvirta, D. P. Brumby, and A. Howes. Model of visual search and selection time in linear menus. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* 2014, pp. 3865–3874. ACM.

[9] J. Batra, N. Garg, A. Kumar, T. Mömke, and A. Wiese. New approximation schemes for unsplittable flow on a path. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 47–58.

[10] U. Bauer, M. Kerber, and J. Reininghaus. Distributed computation of persistent homology. In *2014 Proceedings of the Sixteenth Workshop on Algorithm Engineering and Experiments, ALENEX 2014, Portland, Oregon, USA, January 5, 2014,* 2014, pp. 31–38.

[11] P. Baumgartner and U. Waldmann. Hierarchic superposition with weak abstraction. In M. P. Bonacina, ed., *Automated Deduction – CADE-24*, Lake Placid, NY, USA, 2013, LNAI 7898, pp. 39–57. Springer.

[12] D. Becker, P. Lutsik, P. Ebert, C. Bock, T. Lengauer, and J. Walter. BiQ Analyzer HiMod: an interactive software tool for high-throughput locus-specific analysis of 5-methylcytosine and its oxidized derivatives. *Nucleic acids research*, 42(Web Server issue):W501–7, 2014.

[13] E. Berberich, P. Emeliyanenko, A. Kobel, and M. Sagraloff. Exact symbolic-numeric computation of planar algebraic curves. *Theoretical Computer Science*, 491: 1–32, 2013.

[14] E. Berberich, D. Halperin, M. Kerber, and R. Pogalnikova. Deconstructing approximate offsets. *Discrete & Computational Geometry,* 48(4):964–989, 2012.

[15] K. Berberich and S. Bedathur. Computing n-gram statistics in MapReduce. In *Advances in Database Technology (EDBT 2013)*, Genova, Italy, 2013, pp. 101–112. ACM.

[16] S. Bhattacharya, M. Hoefer, C.-C. Huang, T. Kavitha, and L. Wagner. Maintaining nearpopular matchings. In *Proc. 42nd Intl. Colloq. Automata, Languages, Programming (ICALP),* 2015, vol. 2, pp. 504–515.

[17] K. Bringmann. Why walking the dog takes time: Frechet distance has no strongly subquadratic algorithms unless SETH fails. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, 2014, pp. 661–670.

[18] A. Bulling, U. Blanke, and B. Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys,* 46:1–33, 2014.

[19] A. Bulling, U. Blanke, and B. Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys,* 46(3):33:1–33:33, 2014.

[20] P. Chalermsook, B. Laekhanukit, and D. Nanongkai. Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA,* 2013, pp. 370–379.

[21] P. Chalermsook, B. Laekhanukit, and D. Nanongkai. Pre-reduction graph products: Hardnesses of properly learning dfas and approximating EDP on dags. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014,* 2014, pp. 444–453.

[22] C. Chen and M. Kerber. An output-sensitive algorithm for persistent homology. *Computational Geometry: Theory and Applications,* 46(4):435–447, 2013.

[23] D. Chen, D. I. Levin, P. Didyk, P. Sitthi-Amorn, and W. Matusik. Spec2fab: a reducertuner model for translating specifications to 3d prints. *ACM Transactions on Graphics (TOG),* 32(4):135, 2013.

[24] D. Dhungana, C. H. Tang, C. Weidenbach, and P. Wischnewski. Automated verification of interactive rule-based configuration systems. In E. Denney, T. Bultan, and A. Zeller, eds., *28th IEEE/ACM International Conference on Automated Software Engineering (ASE 2013),* Palo Alto, CA, USA, 2013, pp. 551–561. IEEE.

[25] R. Duan and K. Mehlhorn. A combinatorial polynomial algorithm for the linear Arrow-Debreu market. *Information and Computation,* 243:112–132, 2015. a preliminary version appeared in ICALP 2013, LNCS 7965, pages 425-436.

[26] H. Errami, M. Eiswirth, D. Grigoriev, W. M. Seiler, T. Sturm, and A. Weber. Detection of hopf bifurcations in chemical reaction networks using convex coordinates. *J. Comput. Physics,* 291:279–302, 2015.

[27] A. L. Fietzke and C. Weidenbach. Superposition as a decision procedure for timed automata. *Mathematics in Computer Science,* 6(4):409–425, 2013.

[28] A. Geiger, M. Lauer, C. Wojek, C. Stiller, and R. Urtasun. 3d traffic scene understanding from movable platforms. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 36(5):1012–1025, 2013.

[29] M. Ghaffari, A. Karrenbauer, F. Kuhn, C. Lenzen, and B. Patt-Shamir. Near-optimal distributed maximum flow: Extended abstract. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, Donostia-San Sebastián, Spain, July 21-23, 2015*, 2015, pp. 81–90.

[30] D. Günther, A. Jacobson, J. Reininghaus, H.-P. Seidel, O. Sorkine-Hornung, and T. Weinkauf. Fast and memory-efficient topological denoising of 2D and 3D scalar fields. *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE VIS)*, 20(12):2585–2594, 2014.

[31] M. Hoefer and L. Wagner. Locally stable marriage with strict preferences. In *Proc. 40th Intl. Colloq. Automata, Languages, Programming (ICALP),* 2013, vol. 2, pp. 620–631.

[32] J. Hoffart, F. M. Suchanek, K. Berberich, and G. Weikum. YAGO2: A spatially and temporally enhanced knowledge base from Wikipedia. *Artificial Intelligence*, 194:28–61, 2013.

[33] J. Hosang, R. Benenson, P. Dollár, and B. Schiele. What makes for effective detection proposals? *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI),* In Press.

[34] J. Kappes, B. Andres, F. Hamprecht, C. Schnörr, S. Nowozin, D. Batra, S. Kim, B. Kausler, T. Kröger, J. Lellmann, N. Komodakis, B. Savchynskyy, and C. Rother. A comparative study of modern inference techniques for structured discrete energy minimization problems. *International Journal of Computer Vision,* pp. 1–30, 2015.

[35] S. Karayev, M. Fritz, and T. Darrell. Anytime recognition of objects and scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. oral.

[36] T. Lengauer, N. Pfeifer, and R. Kaiser. Personalized HIV therapy to control drug resistance. *Drug Discovery Today: Technologies*, 11:57–64, 2014.

[37] Y. Li, H. L. Nguyen, and D. P. Woodruff. Turnstile streaming algorithms might as well be linear sketches. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31-June 03, 2014,* 2014, pp. 174–183.

[38] M. Malinowski and M. Fritz. A multi-world approach to question answering about real-world scenes based on uncertain input. In *Neural Information Processing Systems (NIPS),* 2014.

[39] S. Metzler and P. Miettinen. Clustering boolean tensors. *Data Mining and Knowledge Discovery,* pp. 1–31, 2015.

[40] S. Mukherjee, G. Weikum, and C. Danescu-Niculescu-Mizil. People on drugs: Credibility of user statements in health communities. In S. A. Macskassy, C. Perlich, J. Lescovec, W. Wang, and R. Ghani, eds., *KDD'14, 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* New York, NY, USA, 2014, pp. 65–74. ACM.

[41] C. H. Nguyen, T. Ritschel, and H.-P. Seidel. Data-driven color manifolds. *ACM Trans. Graph.*, 34(2), 2015.

[42] S. Olberding, M. Wessely, and J. Steimle. Printscreen: Fabricating highly customizable thin-film touch-displays. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology [Best Paper Award – top 1%],* New York, NY, USA, 2014, UIST ,14, pp. 281–290. ACM.

[43] B. Pepik, M. Stark, P. Gehler, and B. Schiele. Multi-view and 3d deformable part models. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI),* 2015.

[44] M. Pilipczuk, P. Sankowski, and E. J. van Leeuwen. Network sparsification for steiner problems on planar and bounded-genus graphs. In 55*th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014,* 2014, pp. 276–285.

[45] M. Regneri, M. Rohrbach, D. Wetzel, S. Thater, B. Schiele, and M. Pinkal. Grounding action descriptions in videos. *Transactions of the Association for Computational Linguistics (TACL)*, 1(Mar): 25–36, 2013.

[46] K. Rematas, T. Ritschel, M. Fritz, and T. Tuytelaars. Image-based synthesis and re-synthesis of viewpoints guided by 3d models. In *Proc. CVPR,* 2014.

[47] M. Rempfler, M. Schneider, G. D. Ielacqua, X. Xiao, S. R. Stock, J. Klohs, G. Székely, B. Andres, and B. H. Menze. Reconstructing cerebrovascular networks under local physiological constraints by integer programming. *Medical Image Analysis,* 25(1):86 – 94, 2015.

[48] H. Saikia, H.-P. Seidel, and T. Weinkauf. Extended branch decomposition graphs: Structural comparison of scalar data. *Computer Graphics Forum (Proc. EuroVis),* 33(3): 41–50, 2014.

[49] S.-E. Schelhorn, M. Fischer, L. Tolosi, J. Altmüller, P. Nürnberg, H. Pfister, T. Lengauer, and F. Berthold. Sensitive detection of viral transcripts in human tumor transcriptomes. *PLoS computational biology,* 9(10):e1003228, 2013.

[50] C. Schulz, C. von Tycowicz, H.-P. Seidel and K. Hildebrandt. Animating deformable objects using sparse spacetime constraints. *ACM Transactions on Graphics (SIGGRAPH)*, 33(4):109:1–109:10, 2014.

[51] N. K. Speicher and N. Pfeifer. Integrating different data types by regularized unsupervised multiple kernel learning with application to cancer subtype discovery. *Bioinformatics (Oxford, England),* 31(12):i 268–i275, 2015.

[52] S. Sridhar, F. Mueller, A. Oulasvirtta, and C. Theobalt. Fast and robust hand tracking using detection-guided optimization. In *IEEE Intl. Conf. on Computer Vision and Pattern Recognition (CVPR),* 2015.

[53] M. Stikic, D. Larlus, S. Ebert, and B. Schiele. Weakly supervised recognition of daily life activities with wearable sensors. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI),* 33(12):2521–2537, 2011.

[54] T. Sturm. Subtropical real root finding. In *Proceedings of the 2015 ACM on International Symposium on Symbolic and Algebraic Computation,* New York, NY, USA, 2015, ISSAC ,15, pp. 347–354. ACM.

[55] S. Tang, M. Andriluka, and B. Schiele. Detection and tracking of occluded people. *International Journal of Computer Vision (IJCV)*, 2014.

[56] K. Templin, P. Didyk, K. Myszkowski, M. M. Hefeeda, H.-P. Seidel, and W. Matusik. Modeling and optimizing eye vergence response to stereoscopic cuts. *ACM Transactions on Graphics (Proc. SIGGRAPH),* 33(4), 2014.

[57] K. Templin, P. Didyk, K. Myszkowski, and H.-P. Seidel. Perceptually-motivated stereoscopic film grain. *Computer Graphics Forum (Proceedings Pacific Graphics 2014, Seoul, South Korea),* 33(7):349–358, 2014.

[58] C. von Tycowicz, C. Schulz, H.-P. Seidel, and K. Hildebrandt. An efficient construction of reduced deformable objects. *ACM Transactions on Graphics (SIGGRAPH Asia),* 32(6):213:1–213:10, 2013.

[59] M. Weigel, T. Lu, G. Bailly, A. Oulasvirta, C. Majidi, and J. Steimle. iskin: Flexible, stretchable and visually customizable on-body touch sensors for mobile computing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems [Best Paper Award – top 1%],* New York, NY, USA, 2015, CHI ,15, pp. 2991–3000. ACM.

[60] C. Wu, M. Zollhöfer, M. Niessner, M. Stamminger, S. Izadi, and C. Theobalt. Real-time shadingbased refinement for consumer depth cameras. *ACM Transactions on Graphics (Proc. SIG-GRAPH Asia),* 2014.

[61] M. A. Yosef, J. Hoffart, M. Spaniol, and G. Weikum. AIDA: An online tool for accurate disambiguation of named entities in text and tables. In H. V. Jagadish, J. Blakeley, J. M. Hellerstein, N. Koudas, W. Lehner, S. Sarawagi, and U. Röhm, eds., *Proceedings of the 37th International Conference on Very Large Data Bases (VLDB 2011),* Seattle, USA, 2011, Proceedings of the VLDB Endowment, pp. 1450–1453. VLDB Endowment.

[62] Y. Zhang, A. Bulling, and H. Gellersen. Sideways: A gaze interface for spontaneous interaction with situated displays. In *Proc. of the 31st SIGCHI International Conference on Human Factors in Computing Systems (CHI 2013),* 2013, pp. 851–860.

Directions

The Max Planck Institute for Informatics (Building E1 4) is located on the campus of Saarland University, approximately 5 km north-east of downtown Saarbrücken, in a wooded area near Dudweiler.

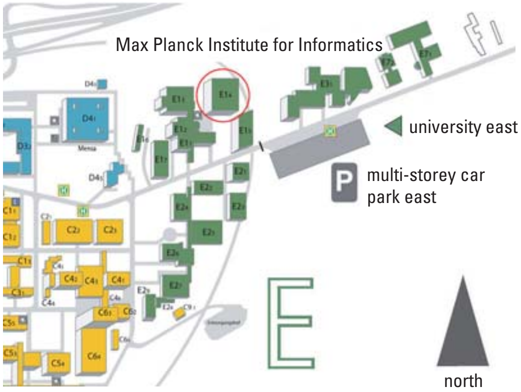

Saarbrücken has its own airport (Saarbrücken-Ensheim) and is well-connected to the Frankfurt and Luxemburg airports by shuttle bus, train, or car. There are good hourly train connections from Saarbrücken to cities all over Germany, and trains to Metz, Nancy and Paris leave at regular intervals. The express highways – or Autobahns – provide fast routes to Mannheim, Frankfurt, Luxemburg, Trier, Köln, and Straßburg as well as Metz, Nancy, and Paris.

The campus is accessible...

from the airport at Saarbrücken-Ensheim
by taxi in approximately 20 minutes

from Saarbrücken central train station
by taxi in approximately 15 minutes
by bus in approximately 20 minutes
direction *“Dudweiler-Dudoplatz”* or *“Universität Campus”*
exit *“Universität Mensa”*, or
alternatively exit *“Universität Busterminal”*

from Frankfurt or Mannheim on Autobahn A6
exit *“St.Ingbert-West”*
follow the white signs for *“Universität”*, leading to the campus

from Paris on Autobahn A4
exit *“St.Ingbert-West”*
follow the white signs for *“Universität”*, leading to the campus

IMPRINT

Publisher
Max Planck Institute for Informatics
Campus E1 4
D-66123 Saarbrücken

Editors & Coordination
Klaus Berberich
Eugen Denerz
Volker Maria Geiß
Jan Hosang
Thomas Kesselheim
Thomas Lengauer
Kurt Mehlhorn
Jennifer Müller
Nico Pfeifer
Bernt Schiele
Nathalie Schultes
Hans-Peter Seidel
Bertram Somieski
Christian Theobalt
Uwe Waldmann
Christoph Weidenbach
Gerhard Weikum

Proofreading
Krista Ames
Kamila Kolesniczenko
Nathalie Schultes

Contact
Max-Planck-Institut für Informatik
Phone +49 681 9325-0
Fax +49 681 9325-5719
Email info@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de

Report Period
1st of January, 2013 to 31st of December, 2014

Design
Behr Design | Saarbrücken

Printing
Druckerei Wollenschneider | Saarbrücken

max planck institut
informatik

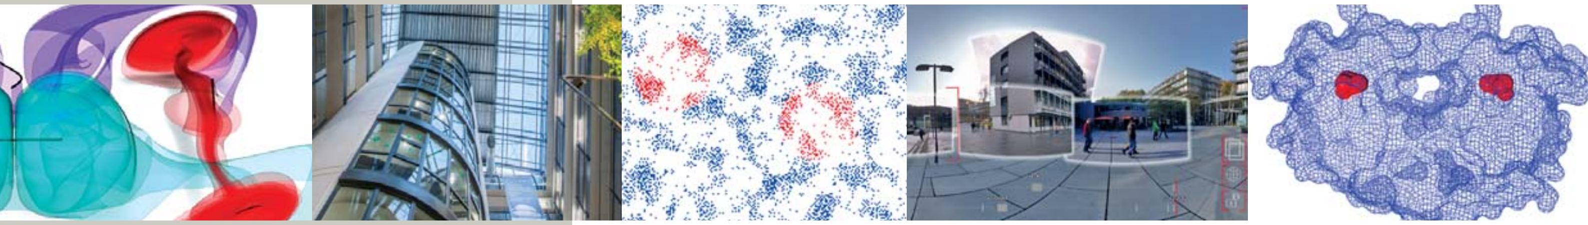

Max Planck Institute for Informatics
Campus E1 4
D-66123 Saarbrücken

Phone +49 681 9325-0
Fax +49 681 9325-5719
Email info@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de

max planck institut
informatik