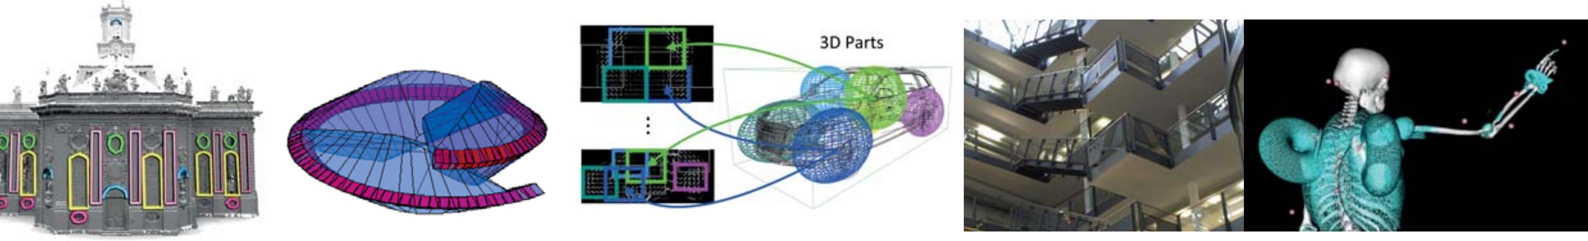

max planck institut
informatik

Report 2012

12 13 14 15 16

max planck institut
informatik

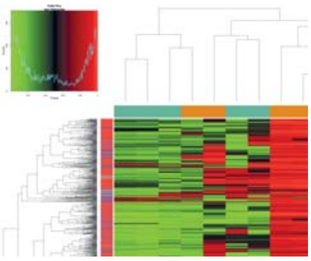

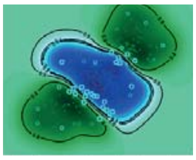

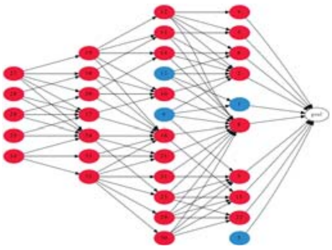

Report 2012

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. Jürgen Giesl, RWTH Aachen, Germany
Prof. Dr. Alon Halevy, Google Research, Mountain View, USA
Prof. Dr. Yannis E. Ioannidis, University of Athens, Greece
Prof. Dr. Yves Moreau, Katholieke Universiteit Leuven, Belgium
Prof. Dr. Nicole Schweikardt, Goethe University, Frankfurt am Main, Germany
Prof. Dr. François Sillion, INRIA Rhône-Alpes, Grenoble, France
Prof. Dr. Éva Tardos, Cornell University, Ithaca, USA
Prof. Dr. Demetri Terzopoulos, University of California, Los Angeles, USA
Prof. Dr. Emo Welzl, Swiss Federal Institute of Technology, Zurich, Switzerland

Curatorship Board

Hon. Adv. Professor (Tsinghua) Dr. Reinhold Achatz, Head of Corporate Technology, Innovation & Quality, Thyssen Krupp AG, Essen
Dr. Siegfried Dais, Robert Bosch Treuhand 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
Heiko Maas, Minister of Economics, Labour, Energy and Transport, Saarland, Saarbrücken
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

Each spring, the Max Planck Institute for Informatics publishes an annual report for the general public. 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 bring you, our readers, closer to the fascinating world of our science.

The Max Planck Institute for Informatics aims to be a lighthouse in 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 the doctoral and postdoctoral phases. We are educating future leaders for science and industry. Third, through our role in the 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 staff of over 190 scientists comes from abroad. Fifth, through the transfer of our results into industry. These transfers take place through cooperation projects, spin-offs, and people. Sixth, by building a world-class competence center for informatics in cooperation with our partners: Saarland University, the German Research Center for Artificial Intelligence and the Max Planck Institute for Software Systems. We have been very successful in all of these undertakings in recent years.

Our success is becoming visible in the new buildings for the Max Planck Institute for Software Systems, the Center for Bioinformatics, the Intel Visual Computing Institute, the informatics lecture halls, the informatics and mathematics library, and the Cluster of Excellence, all of which are located at the Platz der Informatik. We are very pleased about the renewals of the Cluster of Excellence “Multimodal Computing and Interaction” and the “Saarbrücken Graduate School of Computer Science” and also about the successful evaluation of our IMPRS (International Max Planck Research School).

The report is structered 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 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.

Please enjoy reading this report.

Bernt Schiele *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.

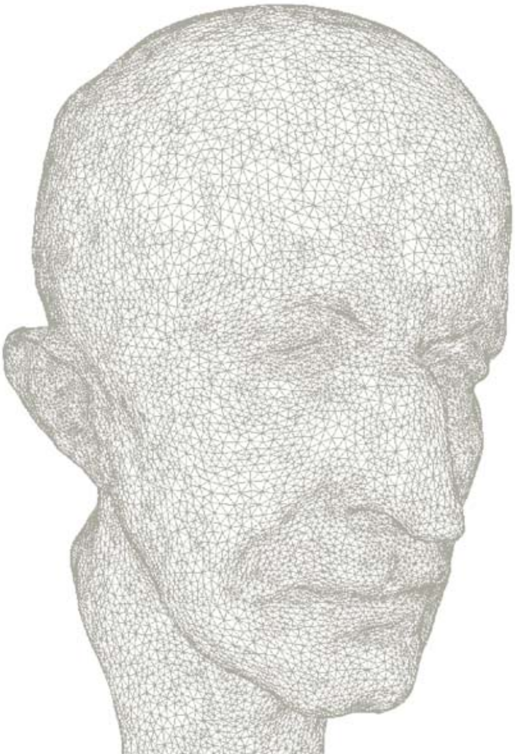

OVERVIEW

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 lighthouse in 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 solutions. At the moment, we are concentrating on algorithms for very large, multimodal data. Multimodal means text, speech, images, videos, graphs, 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 190 researchers are working at our institute and remain with us on average for three years. In this way, we provide the society with over 60 well-trained young scientists each year.

Third, by our role in the profession. We initiate and coordinate large research programs and serve on important committees, e.g., the *"Wissenschaftsrat"*. 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 cooperation projects, spin-offs, 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 MPI-SWS, and the MPI-INF. Intel is investing $12 million in a new research project with its head-quarters on the campus of the UdS. The development of future graphics and visual computing technologies are at the core of the center's work. The investment will take place over a period of five years and is the most extensive cooperation between Intel and a European university to date.

Sixth, by building a world-class competence center for informatics in cooperation with our partners, the UdS, 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 led the *"Algorithms and Complexity"* department since then. Also, Harald Ganzinger was involved from the very beginning and led 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 lead the *"Computational Biology and Applied Algorithmics"* department. Gerhard Weikum has led the *"Databases and Information Systems"* department since 2003. In the summer of 2010, the new *"Computer Vision and Multimodal Computing"* department, led by Bernt Schiele, was added.

In addition to the departments, the institute is home to independent research groups: The "Automation of Logic" research group headed by Christoph Weidenbach and the Otto Hahn research group "Ontologies" headed by Fabian Suchanek as well as several junior research groups within the Excellence Cluster "Multimodal Computing and Interaction" and in the Max Planck and Stanford cooperation (Martin Bokeloh, Adrian Butscher, Michael Kerber, Erik Reinhard, Tobias Ritschel, Michael Stark, Carsten Stoll, Thorsten Thormählen, Péter Vajda und Tino Weinkauf).

Senior Researchers

The scientific work of the 31 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 going through an acceptance process similar to the procedure for a professorial appointment to a university.

The senior researchers currently with the institute are: Mario Albrecht, Ernst Althaus, Jan Baumbach, Benjamin Doerr, Khaled Elbassioni, Tobias Friedrich, Rainer Gemulla, Ivo Ihrke, Olga Kalinina, Thomas Lengauer, Alice McHardy, Nicole Megow, Kurt Mehlhorn, Sebastian Michel, Karol Myszkowski, Antti Oulasvirta, Michael Sagraloff, Thomas Sauerwald, Ralf Schenkel, Bernt Schiele, Hans-Peter Seidel, Viorica Sofronie-Stokkermans, Rob van Stee, Fabian M. Suchanek, Martin Theobald, Christian Theobalt, Thorsten Thormählen, Michael Wand, Christoph Weidenbach, Gerhard Weikum and Tino Weinkauf.

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.

OVERVIEW

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

The scientific problem of understanding the 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 was 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 via modeling (creation of a suitable scene representation) to image synthesis (generation of views for human consumption). The following scientific challenges emerge from this: For the input side, we want to develop modeling tools for efficient handling and processing of large data flows, and, on the output side, we are seeking new algorithms for 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 an increasing demand for algorithmic support due to the recent large increase in experimental data. Algorithms play a central role in the preparation and configuration of biological experiments and even more in the interpretation of the biological data generated by them. The computer is now an essential tool for biology and medicine. The understanding of biological processes on the molecular level is not possible without sophisticated information processing. Vast amounts of data need to be processed in modern biology, and the biochemical interactions in the living organism are so complex that studying them is hopeless without algorithmic support. Therefore, bioinformatics methods have become essential for modern research on the diagnosis and treatment of illnesses.

The Otto Hahn research group *"Ontologies"* is concerned with automated knowledge base construction. Their research topics include information extraction from web documents, learning the rules of data mining, the integration of web services, and ontology matching.

The *"Databases and Information Systems"* department is dedicated to the topics of search, distribution and organization of data in digital libraries, scientific data collections, and the worldwide web. Our long-term goal is the development of easy-to-use, scalable, and precise tools for intelligent searches, which actively support the user in the formulation of queries and in finding relevant information from different data formats. A special characteristic of this research is the automatic extraction of structured information from unstructured sources such as the worldwide web. Our extraction processes combine pattern recognition, linguistic methods, and statistical learning. In this way, the department has created one of the most comprehensive knowledge bases over the past few years and has made it publicly available. In addition, we are developing new methods and software tools for the search and analysis of XML documents, graphics-based RDF data, and very data-intensive internet archives. We investigate various approaches for the implementation of these methods in distributed computing systems for better scalability.

The department for *"Computer Vision and Multimodal Computing"* investigates processing and understanding sensor information. Sensors range from relatively simple, e.g., GPS and acceleration sensors, to 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. 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 exploded 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 then 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 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 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 program

as 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 proved considerably successful in the past years, and several young scientists have since taken professorships in Germany and abroad.

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/]. Part of our results is available in the form of downloadable software or as a web service. Examples are CGAL (Computational Geometry Algorithms Library), as well as 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 the creation of 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 110 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 94. 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 26, or the *“Indo Max Planck Center for Computer Science”* (a cooperation with leading universities in India), page 102, or one of our many EU projects.

We encourage our young scientists to establish their own research programs and 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.

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 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 work groups environment.**

We are successful in all three aspects. 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, our LEDA and CGAL software libraries are used worldwide, and we offer completely new search engine opportunities for efficient and intelligent searches in larger amounts of data. Many former members of the group are in top positions domestically and abroad.

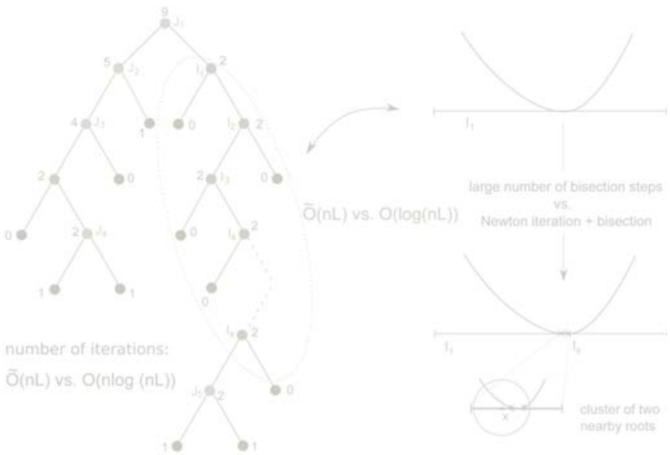

CONTACT

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

DEPT. 1

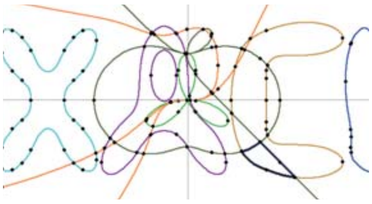

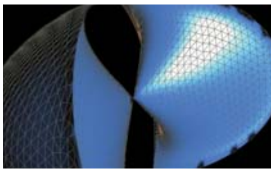

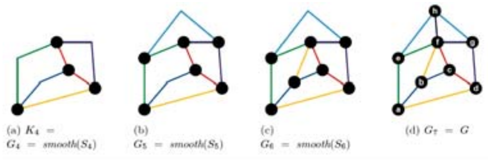

Algorithms are the heart of all software systems. We work on the design and analysis of algorithms in many facets: combinatorial, geometric and algebraic algorithms, data structures and search processes, very different computer models (sequential, parallel, distributed, flat data or hierarchical data), exact and approximate solutions, problem-specific methods and general heuristics, deterministic and randomized solutions, upper and lower bounds, analyses in the worst case and on the average. We develop efficient algorithms for abstract versions of applications problems as well as for concrete applications, e.g., resource optimization in the production of semiconductors. A portion of our theoretical insight is used to implement software demonstrators and software libraries; as a part of our practical work, we collaborate 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 validate the correctness of the solution) as well as 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.

Outstanding practical results in recent years include our contribution to the CGAL software library, which also makes the handling of non-linear objects possible, and the CompleteSearch search engine and its application to one of the most important literature databases in computer science.

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.

The combination of theoretical and experimental research in algorithms has become a widely accepted research direction. The DFG supports it via its priority program “Algorithm Engineering”.

The group is involved in 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 “Algorithm Engineering” priority program and are part of the trans-regional special research area AVACS (Automatic Verification and Analysis of Complex Systems). Siemens and Google provide funding for PhD students in the work group.

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. :::

Computer Vision and Multimodal Computing

PROF. DR. BERNT SCHIELE

DEPT. 2

The department was founded in 2010 and currently includes 16 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.

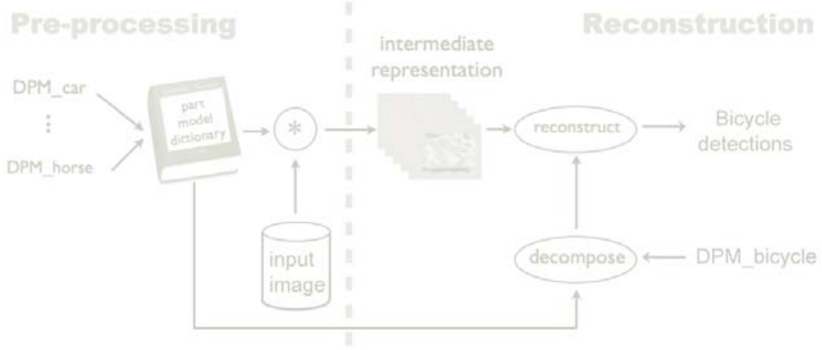

CONTACT

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

DEPT. 2

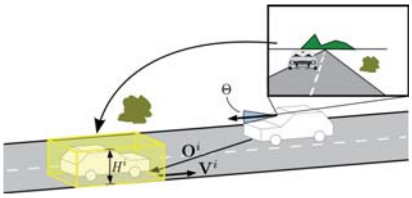

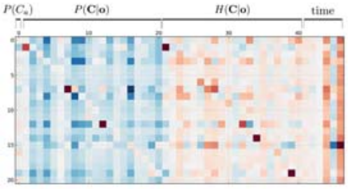

Sensors such as cameras, GPS and accelerometers are being increasingly embedded in devices and environments, and they are already helpful in various ways. Computer-controlled processing of sensor information has made enormous progress but is generally limited to simple matters. This means, in particular, 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 sup-port 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

DEPT. 3

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

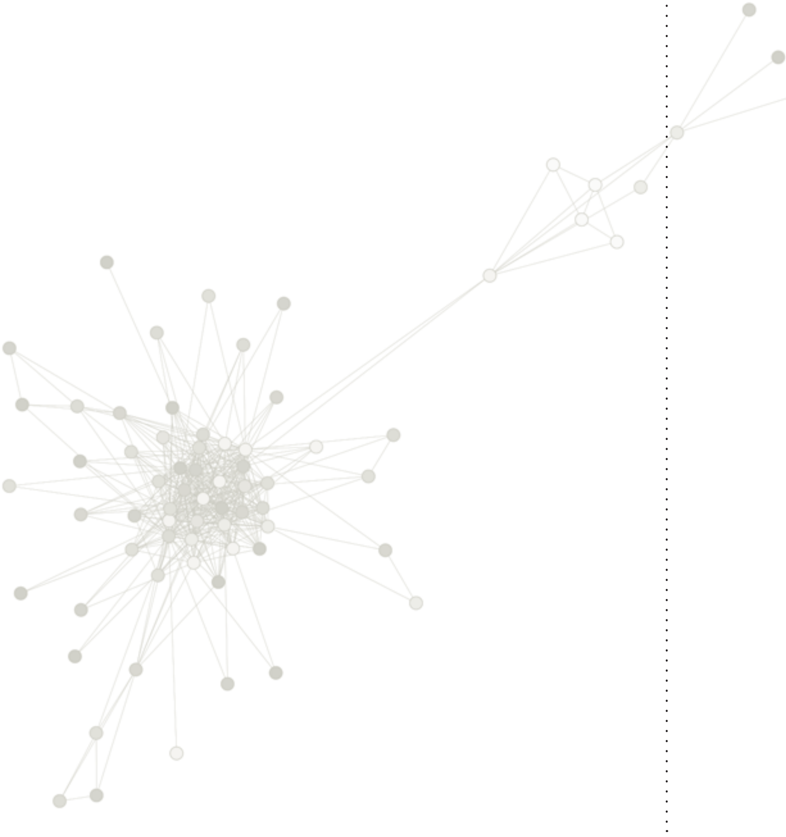

CONTACT

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

DEPT. 3

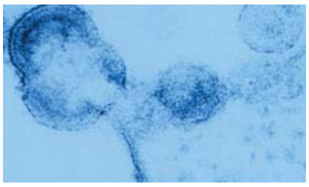

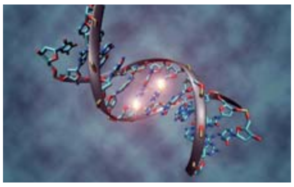

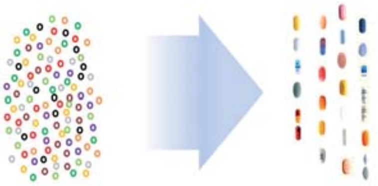

The department focuses on topics more or less closely related to the diagnosis and therapy of diseases. At the molecular level, diseases can be understood in terms of anomalies in the biochemical circuitry of an organism (*"Bioinformatics"*, page 42). The building blocks of such biochemical networks 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, together with RNA they regulate the expression of genes, and they 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. Furthermore, understanding the processes in the complex networks composed of these interactions is essential.

Our methods are applied to concrete diseases such as AIDS (*"Protein Structure and Interactions"*, page 44), Hepatitis B and Hepatitis C (*"Analysis of Viral Genomes Using Next-generation Sequencing"*, page 46). 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 45). Furthermore, we investigate the pattern of geographical distribution of the virus (*"Networks and Phylogenies: Understanding Disease Spread and Evolution"*, page 48).

Diseases like cancer, neurodegenerative diseases, or immululogical 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. 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 49).

A large part of the method development in the department results in software systems, which are used worldwide by many academic, clinical, and often industrial users. Examples, which are presented in this issue (*"EpiExplorer and RnBeads: Integrative Analysis of Epigenomic Data"*, page 90), 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 45).

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 (epigenics) and PREDEMICS (research of viruses with considerable epidemic potential). :::

Computer Graphics

PROF. DR. HANS-PETER SEIDEL

DEPT. 4

The Computer Graphics work group was founded in 1999 and includes 40 scientists today.

An important characteristic of our work is the thorough consideration of the entire pipeline, from data acquisition to modeling to image synthesis including (3D image analysis and synthesis). Typical in this area is the co-occurrence of very large data sets and the demand for a faster and more interactive presentation, wherever possible.

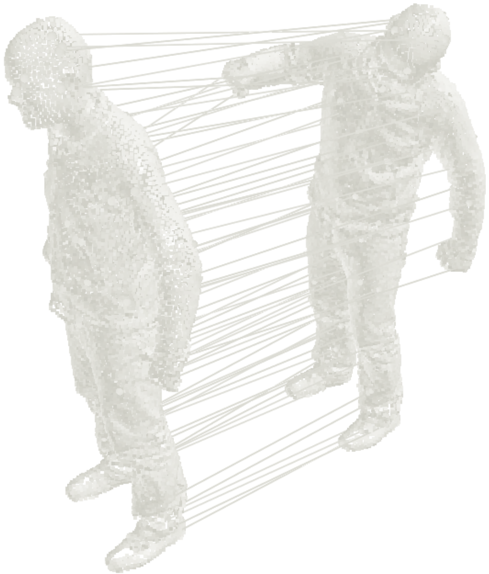

CONTACT

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

DEPT. 4

Computers are often used today to construct segments of, simulate, and present the real or a virtual world. Due to the importance of visual information for people, computer graphics have in the previous decade become a key technology for the modern information and communication society, and their future potential for application is merely hinted at by such keywords as multimedia, digital television, telcommunications, virtual reality, or 3D internet. Typical for this area is the co-occurrence of very large data sets and the demand for a faster (if possible, more interactive) visual presentation of results with high image quality on scalable displays. In addition, the user should be able to intuitively interact with his environment.

New approaches are needed from a scientific point of view in order to meet the above-mentioned challenges. An important characteristic in our work group is, therefore, the thorough consideration of the entire pipeline, from data acquisition to modeling (creation of a suitable internal computer scene description to image synthesis (generation of arbitrary views). These integrated views are necessary in order to adequately use the performance capabilities of modern hardware, the input (image generating processes) as well as the output (graphics hardware). In the meantime, the term 3D image analysis and synthesis has been coined for this integrated view. Central scientific challenges arising from this are the development of appropriate modeling tools for the efficient handling and further processing of data flows from the input side as well as the development of new algorighms for fast and qualitatively high-value presentation with close linkage to the possibilities and prospects of modern graphics on the output side.

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

Especially important is the jointly founded *“Max Planck Center for Visual Computing and Communication”*, developed with the support of the Max Planck Society and Stanford University and with significant support from the BMBF (German Federal Ministry of Education and Research). It was founded in October 2003. The aim of this bridge-building effort between the two top locations in Germany and the United States is to strengthen research efforts in this key area of modern information and communication technology. By establishing new exchange mechanisms with attractive opportunities for returning researchers, the center is making a significant contribution to training and winning back outstanding young scientists. The direction of the center is in the hands of Professor Bernd Girod, Professor Leonidas Guibas (Stanford University), and Professor Hans-Peter Seidel (Max Planck Institute for Informatics).

Furthermore, the Computer Graphics work group is significantly involved in the activities of the Cluster of Excellence on *“Multimodal Computing and Interaction”*. The excellence cluster was newly formed in 2007 within the framework of an Excellence Initiative by the German federal and state governments and has been awarded a renewal for a second funding period (2012 – 2017). The scientific coordinator of the excellence cluster is Prof. Hans-Peter Seidel.

An additional important development was the founding of the *“Intel Visual Computing Institute”* in May 2009. This new research institute is located on campus and is jointly supported by Intel, Saarland University, the DFKI (German Research Center for Artificial Intelligence), the Max Planck Institute for Informatics and the Max Planck Institute for Software Systems. Prof. Hans-Peter Seidel represents the Max Planck Institute for Informatics on its Governance Board. Professor Christian Theobalt is a member of the Steering Committee.

Over the past ten years, more than 30 young scientists from the Computer Graphics group have received professorships in Germany and abroad. The group has earned a number of prizes, including awards for young scientists, the Eurographics Distinguished Career Award for Professor Hans-Peter Seidel, and the German Pattern Recognition Prize for Professor Christian Theobalt. :::

Databases and Information Systems

PROF. DR.-ING. GERHARD WEIKUM

DEBT. 5

The Department, which is led by Gerhard Weikum, pursues research in five topic areas:

1. **Knowledge discovery on the Web with statistics- and logics-based methods for automatic fact extraction from Internet sources like Wikipedia;**
2. **Text mining for the automatic classification of documents and for the identification of interesting patterns in text corpora, especially in news and web archives over long timescales;**
3. **Ranking and inference methods for queries for which only the top k answers are important, and for coping with uncertain data (e.g. automatically extracted relations from text);**
4. **Query processing and optimization of execution plans for efficient search on structured and semistructured data (e.g. in XML or RDF format);**
5. **Methods and tools for big data analysis across large collections of structured and unstructured data (e.g., for individualized music recommendations that are automatically generated from comments and ratings in online communities, or for entity recognition in news articles and other texts).**

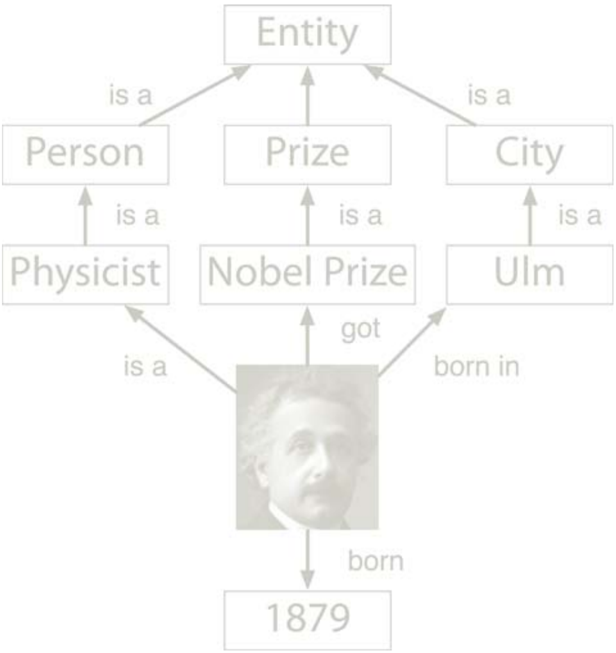

CONTACT

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

DEPT. 5

A central theme in the scientific work of the department is the automatic gathering of comprehensive knowledge bases from information on the World Wide Web as well as discovering, tracking, and analyzing patterns and individual entities (people, organisations, etc.) and their cross-relationships in dynamic web sources. In the YAGO-NAGA project, a very large knowledge collection called YAGO *(Yet Another Great Ontology)* has been created using fact extraction from Wikipedia and integrating the WordNet taxonomy. For exploration and intelligent search in YAGO, a new type of search engine called NAGA *(Not Another Google Answer)* has been developed. Both use the semantic web data model RDF (*Resource Description Framework*), and the group has developed one of the fastest RDF search engines called RDF-3X (*RDF Triple Express*). A variety of further research topics are addressed within the YAGO-NAGA framework: Inclusion of multilingual information, systematic gathering of multimodal data (e.g. photos of people), analysis of time changes in knowledge, deduction of new relationships on the basis of probable but uncertain knowledge, and more. The YAGO knowledge base is freely available and has more than 10,000 downloads. It has been used for many projects all over the world, including projects at Google, Microsoft, and in IBM's Watson project.

The vision that drives this work is the expected convergence of the *Semantic Web* with its formal ontologies and logic-oriented search and inference, the Social Web (Web 2.0) with its latent *"Wisdom of the Crowds"*, and the defacto predominating *Statistical Web* with its statistical approach to data capturing and probabilistic underpinnings of search engines. The web could be the basis of a comprehensive knowledge base, which contains the entire encyclopedic knowledge of mankind in a formal, structured, and machine-readable representation, and enables higher-level programs and intelligent services. The benefits of such a knowledge base would be enormous.

Many of the prototype systems developed in the group are publicly available and are used by other research groups worldwide. These include, in particular:

1. The XML search engine *TopX*, which has achieved top ranks in the INEX benchmarking series and provides the semantically annotated Wikipedia reference corpus for the INEX competition;

2. The RDF search engine *RDF-3X*, which can search semantic web data and other graph-structured data very efficiently for complex patterns;

3. The software prototype *RDF-Express* for entity search in combined RDF/text data.

4. The software tools that automatically create and maintain the YAGO knowledge base, as well as the YAGO contents itself;

5. The tools *SOFIE* and *PROSPERA*, which extract rational facts from arbitrary websites and text documents,

6. The software package *PATTY*, which distills semantically typed textual patterns for relational phrases as well as their paraphrases,

7. The AIDA-Software for entity recognition and disambiguation in English texts.

The department is involved in a number of externally funded projects, especially the EU projects *"Living Knowledge"* and *"Longitudinal Analytics of Web Archive Data",* the BMBF joint project *"WisNetGrid"* as well as the DFG Excellence Cluster *"Multimodal Computing and Interaction"*. The group has received a Google 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 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
Prof. Dr. Christoph Weidenbach
Secretary Jennifer Müller
Phone +49 681 9325-2900
Email jmueller@mpi-inf.mpg.de

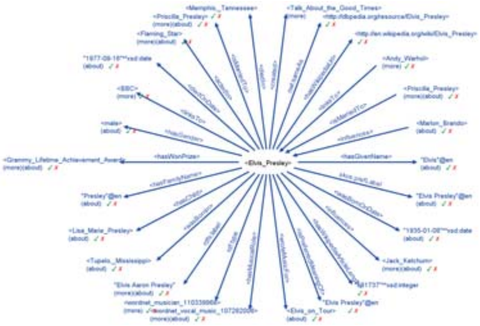

Ontologies

DR. FABIAN M. SUCHANEK

RG. 2

The independent research group "Ontologies", led by Dr. Fabian M. Suchanek, was established in January 2012. It focuses on the automatic construction of knowledge bases and their applications.

Today, computers are already able to understand many things about everyday life: Internet search engines can answer simple questions directly (for example, "When was Max Planck born?"). The Siri system in Apple's iPhone can understand spoken questions. And the Watson system from IBM recently won the American quiz show *Jeopardy!*. These things are only possible because the systems have a structured collection of knowledge at their disposal. They must know who Max Planck is, the iPhone owner's friend's telephone number, or the name of the author of *The Hitchhiker's Guide to the Galaxy*. Such a knowledge base is known as an *ontology*. Ontologies normally contain entities (such as people, cities, companies, or films) and facts about those entities (for example, which actor appears in which film, which company is headquartered in which city, and so on).

In the Ontologies research group, we investigate how such ontologies can be automatically built and maintained. The group was established at the beginning of 2012 and currently comprises the leader, a PhD candidate, and a master's student. It is financed through an Otto Hahn Grant of the Max Planck Society.

In our research, we develop algorithms that can extract information from data sources (such as Wikipedia). We also investigate how this information can be tested automatically for consistency, so that the extracted information is as accurate as possible. Our main project, which we are pursuing together with Department 5, is YAGO. YAGO is an ontology that we have generated automatically from different data sources. It currently contains about 10 million entities and 100 million facts, and is freely available at http://yago-knowledge.org .

Our group also concerns itself with different aspects above and beyond the construction of ontologies. For example, we work on the automatic discovery of logical connections in ontologies. Thus we can, for example, find out that spouses usually live in the same city, that the doctoral advisor of a scientist usually works at the university where that scientist received his or her degree, or that popes usually die in Rome.

In cooperation with the French INRIA Institute in Saclay, we are working on automatically combining different ontologies. Our work with the University of Versailles concerns the integration of web services in ontologies. Other research areas include information extraction and the proof and representation of the provenance of facts. We thus cover a whole range of research questions relating to the theme of ontologies. Our YAGO ontology, among others, is used in the aforementioned IBM Watson system. :::

CONTACT

Ontologies
Dr. Fabian M. Suchanek
Phone +49 681 9325-5014
Email suchanek@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 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 uninterrupted and quick exchange of image data and videos in 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, stimulating, and 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 Professor (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 are willing to return here.

Current status

For the past eight years, this model has rewarded Germany with attractive return perspectives, which is counter to the often observed "Brain Drain" to the USA. In this way, it contributes to the development and safeguarding of highly-qualified young scientists and thus to the sustainable strengthening of the innovation capacity and competitiveness of this location. Since its founding, the Max Planck Center has gained an international scientific reputation of being a real talent factory. Since launching the program in 2003, a total of 17 young scientists have completed the program. Of these, 14 young scientists (82%) have meanwhile received professorships, 13 thereof in Germany (76%), and eight of them being full professorships. The successful evaluation last year ensured the unabated pursuit of the program's goals.

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 owes significantly to the contributions of the postdoctoral graduates. Further important elements are the early scientific independence of the young scientists with simultaneous integration into an international, scientific environment, and also the attractive return perspectives. The great success of the Max Planck Center for Visual Computing and Communication makes the program a model for other education and research faculties. :::

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) in 2007 as part of the Excellence Initiative by the German federal and state governments. The renewal proposal last year was successful, and the second research funding phase (2012-2017) has begun.

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, digital content consisted mainly of text; today, these contents have expanded tremendously in scope to include audio, video, and graphics and are now available virtually anywhere. The challenge lies in organizing, understanding, and browsing this multimodal information in a robust, efficient and intelligent way; at the same time, reliable and safe systems that allow intuitive multimodal interactions need to be created.

The Cluster of Excellence on Multimodal Computing and Interaction has taken on these challenges. In this context, the term "multimodal" refers to the different types of information such as text, language, images, video, graphics, and high dimensional data as well as the types of perception and communication, particularly seeing, hearing, and body language. Our first goal is to improve the abilities of computer systems and to capture, process, and illustrate data from different modes efficiently and robustly.

Even large, distributed, noisy, and incomplete multimodal data are to be analyzed and interpreted; the captured knowledge is to be processed and visualized in real-time. This is called multimodal processing. In addition, daily interpersonal communication is based on a multitude of different modalities, namely spoken language, body language, gestures, and facial expressions. Our second goal is, therefore, a similarly natural and multimodal interaction with information systems anywhere and anytime. The systems must consider the context of their environment, react to language, text, and motions, and then answer in appropriate modalities.

This research program builds on the existing strengths of its participating institutions and is a joint, long-term research endeavor; it is comprised of the departments of Computer Science, Computer Linguistics and Phonetics, and Applied Linguistics at Saarland University and also the Max Planck Institute for Informatics, the DFKI (German Research Center for Artificial Intelligence), and the Max Planck Institute for Software Systems. The university and the state government offer significant support to the cluster.

An explicit goal of the cluster is the qualification and promotion of young scientists. Saarbrücken has a long history of playing a leading role in this respect and has over the years acquired the reputation of an elite training hub for young scientists. Therefore, the majority of the approved funds are used for the establishment of junior research groups. In the past four years, this concept has proven to be extremely successful, and a large number of junior scientists have meanwhile been appointed to professorships in Germany and abroad.

The scientific activities of the Max Planck Institute for Informatics decisively contribute to the success of the cluster. All directors of the institute are principal investigators of the cluster, and Prof. Dr. 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

UNDERSTANDING IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH & DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION & VISUALIZATION

OPTIMIZATION

SOFTWARE

UNDERSTANDING IMAGES & VIDEOS

Understanding images and videos is one of the fundamental problems of image processing. The scientific challenges range from the modeling and tracking of people and objects in camera systems to the reconstruction and description of 3D scenes. This area has many applications; these include the animation of people and visualization of 3D scenes, the indexing of image and video material, and also 3D-capturing of the environment for automotive applications. 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 the different aspects of image and video understanding. In the field of modeling people, work is carried out 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, methods for people detection and tracking are researched that require monocular cameras only and that 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 each other, as these areas use similar models and algorithms.

A further area of understanding images and video is the reconstruction and description of 3D scenes. Again, there are different studies within the institute that benefit from each other. For example, 3D scenes are reconstructed from image series to enable a surface description of 3D objects that is highly detailed and accurate. Furthermore, the 3D environment of a moving car is described. Here, a special focus is on the complete identification and description of pedestrians and other road users. All this work shows that the integration and use of 3D scene models can significantly improve results.

One of the fundamental problems of understanding images is the recognition of objects. With the current omnipresence of digital image material, automatic visual object-class recognition techniques are becoming increasingly important. Even if today's approaches can achieve remarkable results, one of the most significant problems remains: How can object models be learned or constructed? Therefore, a central topic of various research work at the institute concerns how such object models can be constructed with as little manual input as possible to enable a broad applicability of today's approaches.

A further fundamental problem is being able to describe and understand human activities. We investigate different methods for person and object recognition and, beyond that, explore the use of linguistic descriptions of the activities. This makes it possible, for example, to recognize human activity based on a linguistic description, without having previously recorded sample data for it. :::

CONTRIBUTIONS

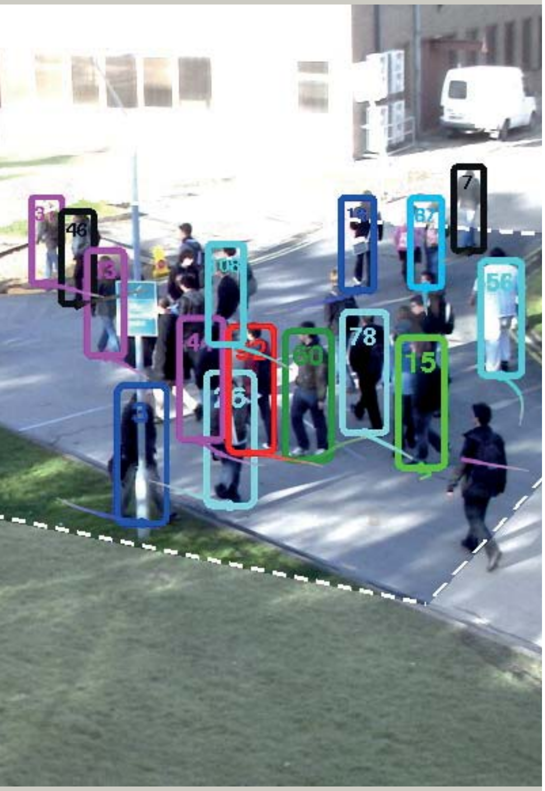

People Detection and Pose Estimation in Challenging Real-World Scenes

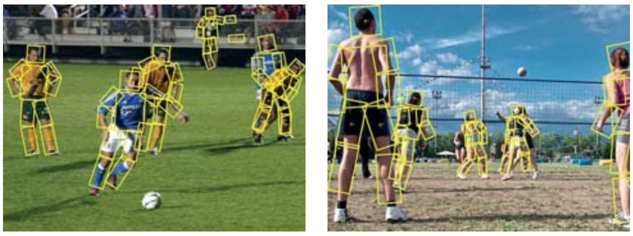

Figure 1 and 2: Example results obtained with our pose estimation model trained on combination of real and synthetically generated images

Figure 3 and 4: Example results obtained with our method for joint detection and tracking of people in crowded scenes

Detection, tracking and pose estimation of people are the key technologies for many applications such as automotive safety, human-computer interaction, robotic navigation, or indexing images and videos from the web. At the same time, they are among the most challenging problems in computer vision that remain a scientific challenge for realistic scenes.

Although state-of-the-art methods perform well for simple scenes with walking people, they often fail in scenes with people performing complex activities or in crowded scenes with multiple people that frequently partially or fully occlude each other. In this project, we address these limitations, building on the recent advances in generic object detection and computer graphics.

Training computer vision models from synthetically generated images

One of the key ingredients for the success of state-of-the-art methods is the ability to automatically learn the appearance of people from a collection of training images. We investigate how 3D human shape models from computer graphics can be leveraged to obtain synthetic training data suitable for training models in computer vision. We rely on the recent statistical model of 3D human shape and pose that is learned from a large collection of human body scans. Our approach allows us to directly control data variability while covering the major shape and pose variations of humans that are often difficult to capture when manually collecting real-world training images. Our method is able to generate synthetic images of people either by directly sampling from 3D shape model or by automatically reshaping real images of people. We validate the effectiveness of our approach on the task of articulated human detection and articulated pose estimation. The obtained results indicate that our automatically generated synthetic data helps to noticeably increase the performance on both tasks [Figure 1 and 2].

Joint detection and tracking of people in crowded scenes

We apply the idea of synthetic training data generation to the problem of people detection and tracking in crowded scenes. We combine it with the observation that in the case of partial occlusions, the joint detection of pairs of people is often easier than the detection of each person individually. We propose a joint detector that relies on the common patterns of person-person occlusion and incorporates them as features for detection. Once a joint configuration of people is detected, it can the be decoded into detections of individual persons. However, configurations of two people also exhibit an increased amount of appearance variation, which is known to make the detection problem more difficult. We compensate for that by synthetically generating examples with various degrees of partial occlusion and mutual arrangements of people. We show that our joint detector significantly increases detection performance in crowded scenes and leads to an improvement in people tracking [Figure 3 and 4]. :::

CONTACT

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

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

Siyu Tang
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-2009
Email tang@mpi-inf.mpg.de

3D Scene Analysis

Figure 1: Automatic creation of a blend *(center)* between a bicycle *(left)* and a motorcycle *(right)*.

Simple creation of new 3D objects through blending of existing models

On the Internet, more and more free or commercial databases of professionally-created 3D objects are becoming available. For a non-expert user, the complex operation of current modeling tools makes it difficult to create similarly good 3D objects. Therefore, often the only option these users have is to choose from a 3D database an existing model that most closely matches their vision. As part of our project, a tool was developed that makes it possible for a non-expert user to create a new 3D object according to his or her wishes by combining existing objects from a 3D database. To that end, the database objects must be automatically analyzed to determine which parts can be exchanged. The analysis includes firstly the segmentation of an object into its component parts. Next, contacts and symmetries among the parts are identified. The process then enables the creation of variants of the object that have similar contacts and symmetries [see figure 1].

Automatic assignment of materials for 3D objects

The material composition of a 3D object is strongly correlated with the object's geometric shape and its relation to other objects in its immediate vicinity (context). In a further project, we have sought to model these context-dependent correlations and to show that they can be learned from a database containing several hundred 3D objects with known materials. If a 3D model without materials is then passed to the software we have developed, the learned model can be used to assign plausible material parameters fully automatically (including diffuse and specular reflection, gloss, and transparency) [see figure 2].

We have also developed a user interface that provides suggestions for alternative materials. This user interface can be used to further refine the automatic suggestions [see figure 2 right]. Once a refinement has been made, the model uses this information to improve the automatic assignment of other parts of the 3D object. The approach works for 3D objects with different numbers of parts and with different topological complexity. Studies with non-expert users showed that the suggested process simplifies and accelerates the assignment of materials significantly in comparison to conventional approaches. :::

Figure 2: Given a 3D object without materials *(left)*, our approach can automatically assign appropriate materials *(second from left)*. It also suggests alternative materials to the user; these can be selected interactively to improve the automatic assignment *(right)*.

CONTACT

Thorsten Thormählen
DEPT. 4 Computer Graphics
Phone +49 681 9325-4017
Email thormae@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).

Similarly, the traditional focus on basic-level categories, such as cars, birds, and people, has been extended to finer-grained categories, such as specific car brands or plant and animal species, again providing richer constraints for scene-level reasoning.

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.

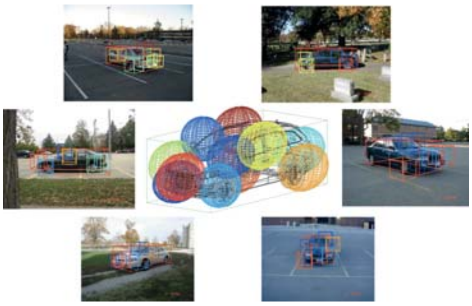

Figure 1

Fine-grained categorization for 3D scene understanding

In this project, we have followed the intuition that fine-grained object categories can encode metric information, which in turn provide valuable information for 3D scene-level reasoning. Knowing the exact brand and type of a car, for example, can be used to look up the physical size of the car in a product database, and further, to estimate its distance from the camera by relating its physical size with the size and shape of its image projection.

We have extended our 3D object class representation to encompass fine-grained category information and collected a novel dataset of car types as a testbed. In our experiments, we have shown that our model outperforms standard visual classifiers in fine-grained classification. Furthermore, we have demonstrated that fine-grained category information can aid 3D scene-level reasoning by providing metric sizes that constrain the scene geometry. :::

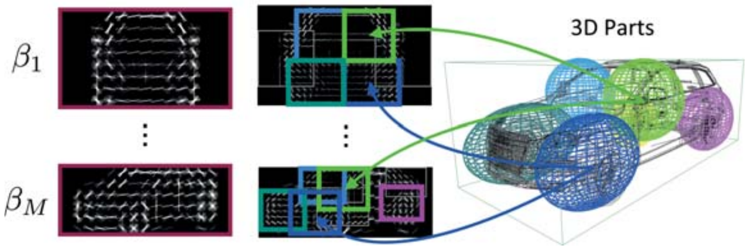

Figure 2

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
Phone +49 681 9325-2120
Email stark@mpi-inf.mpg.de

Markerless Reconstruction of Dynamic Scenes

Reconstruction of detailed animation models from multi-video data

Computer-generated people, so-called avatars, have become an important component of visual media, for example, in computer animations, films or networked virtual environments. In order to be able to present virtual humans in a convincing way, the individual characteristics of the person, such as, for example, his movement, geometry and surface texture, must be extremely realistically modeled. One way to achieve this is by defining every partial aspect of the overall appearance manually in an animation program. This is a very time-intensive and complex process. The geometry of the person must be constructed in precise detail, and each nuance of movement must be finely specified.
It is, therefore, easy to imagine that completely manually produced animations, especially with respect to the quality of movement animation, do not achieve the level of detail of a true human.

The alternative to manual modeling consists of measuring some aspects of animation from real people. Motion capture systems make it possible to reconstruct a skeletal model of a person using video image streams. Unfortunately, this method of movement capture is a very complex process, and the person being measured must often wear a special skin-tight suit with optical markers.
In addition, neither the time changing geometry nor the texture of a person is captured with such systems.

In our research, we have developed a new type of *Performance Capture Algorithms*. For the first time, it is possible to reconstruct the detailed movement, dynamic geometry, and dynamic texture of a person in complex clothing, such as, a dress or a ball gown, solely from multi-video streams. Our process requires no optical markers in the scene.

Figure 1: *left:* **one of 12 camera views;** *right:* **3D model with reconstructed high geometry detail.**

Figure 2: From video frames of only 2 cameras ***(left)*** **we can reconstruct a detailed dynamic 3D model of the face** ***(right).***

In the reporting period, we have been able to improve our performance capture methods in many ways. A new algorithm enables us, for the first time, to reconstruct detailed animation models (i.e., surface+skeleton) of several closely interacting actors in the same scene.
A further milestone is a new algorithm to estimate incident illumination in arbitrary environments from multi-view video. Based on this estimated illumination model, one can reconstruct much more detailed dynamic surface geometry than with any previous approach from the literature available to date [Figure 1].
Based on these fundamentally new theories, we developed the first approach in the literature to reconstruct detailed animation models of human faces using a single stereo camera [Figure 2]. We also developed a new approach to capture much more accurate full body poses by exploiting lighting and reflectance models that we estimated from video footage of general scenes. Another important achievement is the development of a very fast motion estimation approach that can measure complex skeleton motion from multi-view video footage captured under less controlled environment conditions (i.e., no controlled background like in studio environments).

Real-time motion capture with depth sensors

The performance capture methods for exact detailed reconstruction, which were described in the previous section, require very complex calculations and are, therefore, not real-time. In addition, these methods require several video cameras. Motion measurement from a single camera perspective is an extremely complex and highly underconstrained problem.

New types of depth cameras, such as, for example, so-called Time-of-Flight cameras measure 2.5D scene geometry in real-time. It is easier to reconstruct body movements from a single camera perspective by using both depth and video data together. 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 increase camera resolution through calculations on the original data (*super-resolution process*).

With improved depth data and through a newly-developed process from our research, all body movements of a person can be measured from one single camera perspective. The real time algorithm to measure movement combines a depth-based pose optimization procedure with a procedure for fast finding of similar poses from a large database of movement sequences. :::

CONTACT

Christian Theobalt
DEPT. 4 Computer Graphics
Phone +49 681 9325-4028
Email theobalt@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/resources/perfcap/

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 such as 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 by 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 for which no training examples are 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 by human subjects.

Sharing of computation between classes

Today's state-of-the-art object detectors predominantly achieve multi-class detection by 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 looking 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 [Figure 2]. 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.

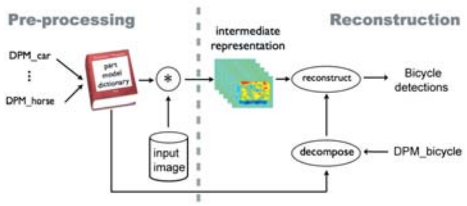

Figure 2: Efficient object class recognition by shared representations

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 is 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. :::

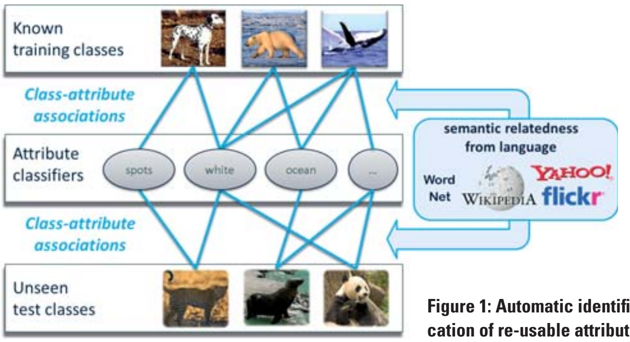

Figure 1: Automatic identification of re-usable attribute models.

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

Advanced Video Processing

Image processing methods are important tools during the postprocessing of photographies. One example is the enhancement of picture quality by applying special filters, for instance to suppress noise or to enhance contrast. Often, it is also necessary to remove certain regions from pictures and to fill-in or inpaint the hole in the image. The foundation for all such operations are mathematical methods that analyze and modify the image signal. Many standard tools for image processing provide such algorithms.

Postprocessing of videos is a much more challenging task. Videos are not merely a temporal sequence of individual images. Modifications to videos during postprocessing, 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. Postprocessing operations with such tools are, therefore, limited to the application of simple filters over a temporal window.

Many practically relevant video processing tasks are much more complex, as the following example shows. In professional video and movie productions, entire scene elements, such as certain actors or supporting crew, may need to be removed from a shot, or their positions in a shot may need to be changed during postprocessing. In consequence, entire image regions in videos that are now disclosed (or now form a hole) need to be inpainted over time. In addition, it may be necessary to adjust for the changing illumination situation in the entire video after the edit. No existing video filtering approach would even come close to solving such a complex video editing task automatically. The consequence is that such editing tasks are usually solved through time-consuming manual editing, where individual pixels in individual frames are modified by artists. Even for short scenes, this can entail weeks of necessary postprocessing time, an effort that everyday users are neither able nor willing to afford.

The goal of this research project is the development of the algorithmic theory to efficiently and effectively solve complex video editing tasks such as the one in the previous example. To this end, an algorithm needs to know more than the image signal at every frame. It is important that an algorithm develops a semantic understanding of a video: What objects are visible in a scene ? What is the illumination and reflectance of materials in the scene ? How do individual scene elements correlate and move over time ? The algorithmic challenges we face in this problem are very demanding, and so far, the algorithmic theory is still in its beginning. The further development of the algorithmic underpinnings, however, is essential to enable also inexperienced users to control advanced video editing and to produce modified videos in a mostly automatic way.

Object removal from video and dynamic background inpainting

An important subtask of many advanced video editing tasks is the inpainting of a dynamic scene background. For instance, if an object in the foreground of a scene is removed, previously occluded parts of the scene background become visible. In practice, this means that a hole in the video appears, since there is no color information available for the originally occluded regions. This hole has to be filled-in or inpainted in all frames of the video where it exists. The basic idea of this inpainting process is to look at other frames of the video (i.e., other time instants) and see if there are image regions that could be copied into the hole to fill it in. The reasoning behind this idea is that at these other frames, the region which is currently covered by the hole was actually visible. Should the hole be in an area of more or less static scene background, inpainting usually means the copying of larger coherent image regions from other frames. In contrast, if the hole is in the area of a now unoccluded dynamic object in the scene, for instance a moving person, inpainting is much more challenging. Not only does the inpainted region need to appear spatially coherent in every frame, the correct scene dynamics also needs to be synthesized, which may not be solved by a mere copying operation alone. Our methods succeed on both, video sequences filmed with static cameras and video sequences filmend with handheld cameras. In our research, we have developed the first methods in the literature to automatically solve the latter problem, namely inpainting of dynamic scene background [see Figure]. :::

Top row: **original video frame;** ***bottom row:*** **video frame in which the persons in the foreground was automatically removed and the dynamic background was automatically inpainted.**

CONTACT

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

Recognition of Kitchen Activities in Videos

The recognition of human activities in video has been receiving increasing interest in recent years. Activity recognition has a wide range of applications, including human-computer interaction, surveillance, intelligent environments, automatic extraction of knowledge from (internet) videos, and support of disabled people. However, at the same time, this research area is still confronted with a large number of challenges, such as complex human motions and interactions, very diverse activities, limited image resolution and quality as well as limited observability.

Fine-grained activities in a kitchen scenario

The current focus of many scientific studies is to distinguish very diverse activities such as running, swimming or drinking in videos with limited quality. In contrast to this, many assistance systems are required to distinguish very similar activities, e.g., it is important for a blind person to know if an opponent wants to shake hands or steal your handbag. In our work, we want to examine this problem more closely and have selected the “kitchen” scenario. In this scenario, the complexity of activities can be varied from very simple actions such as “peeling a carrot” to complex activities of several people preparing a complete meal or dish.

Figure 1: Examples of fine-grained kitchen activities

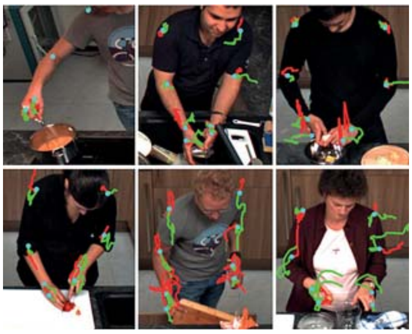

Figure 2: Trajectories of joints for activity recognition

For this research project, we built a fully functional kitchen and equipped it with several cameras [Figure 1a]. In a first dataset, we recorded twelve subjects while they prepared a diverse set of dishes over a total of eight hours of high-resolution video. On the published dataset, we evaluated different approaches for activity recognition to distinguish fine-grained activities such as cut, dice, squeeze, peel, or wash [Figure 1 b-h]. One approach estimates the trajectories of joints and learns the differences between the activities [Figure 2]. A second approach extracts trajectories of moving points in the entire video and computes image and video descriptors of moving elements. In comparison, the second approach achieves much higher performance, as the approach based on joint-trajectories does not use any visual information such as color or shape. Overall, we found that especially very similar, fine-grained activities such as “cut slices” and “cut stripes” are very difficult to distinguish from one another.

Recognition of composite activities with attributes and scripts

The previously described fine-grained activities are in most cases composed to more complex composite activities, e.g., the preparation of a dish in the kitchen scenario. Learning these composite activities is difficult because the same dish can be prepared very differently by different people, and it is nearly impossible to observe all possible combinations in the training data. We attack this problem by representing the composite activities with attributes; in our scenario, these are activities, ingredients, and kitchen tools. The representation as attributes allows recognizing new variants or completely unseen composite activities. To determine possible variants and attributes of unseen composite activities, we use textual descriptions (scripts), which can be easily collected, e.g., via crowd sourcing [Figure 3]. :::

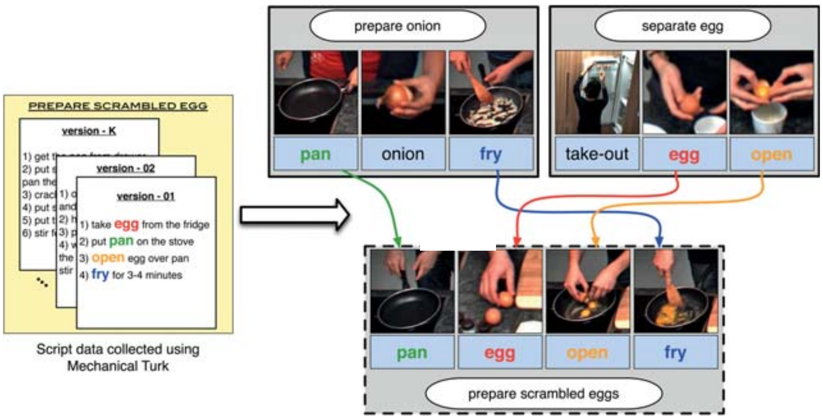

Figure 3: Composite activities represented as attributes which can be transferred using textual script information

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

3D Scene Understanding from Monocular Cameras

Scene understanding – a longstanding goal in computer vision

Inspired by the human visual system, visual scene understanding has been advocated as the “holy grail” for machine vision since its beginning. In the early days this problem was addressed in a bottom-up fashion, starting from low-level features such as edges and their orientation as the key ingredients in the derivation of a complete scene description and scene understanding. Unfortunately, the reliable extraction of saauch low-level features proved very difficult due to their limited expressive power and their inherent ambiguities. As a result and despite enormous efforts, scene understanding remained an illusive goal even for relatively constrained and simple scenes. Disappointed by these early attempts, the research community has turned to easier sub-problems and has achieved remarkable results in areas such as camera geometry, image segmentation, object detection, and tracking. As the performance for these sub-tasks starts to achieve remarkable performance levels, we believe that computer vision should reinvestigate the problem to automatically understand 3D scenes from still images and video sequences.

Robotics and automotive applications

Robotics, such as mobile service robots, as well as automotive safety applications, such as pedestrian protection systems, are clearly of major scientific and commercial interest. Therefore, we use the recognition of pedestrians and cars from a moving camera mounted on a vehicle or robot as a test case and running example for our work. We can leverage domain knowledge for both application areas. The camera, for instance, can be calibrated relatively to its surrounding area, therefore, camera height above ground and the orientation, that is rela-tively to its surrounding area, are constrained and approximately known. Moreover, it can be assumed that the ground is flat in a local neighborhood and that all objects are supported by it. Additionally, we include a Gaussian prior for pedestrian heights as well as for the heights of cars. These assumptions pose additional constraints to eliminate ambiguities and allow us to solve the scene understanding problem for such application domains more easily. See figure 1 for a sample scene.

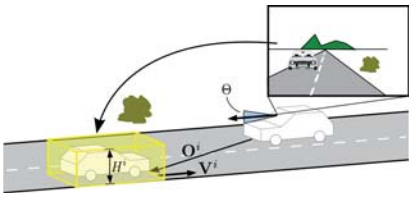

Figure 1: Our system aims to infer a 3D scene based on a combination of prior domain knowledge with videos of a monocular camera.

Integrated 3D scene model

The system that we have developed combines the above prior domain knowledge with powerful state-of-the-art *object class detectors, semantic scene labeling* and the notion of *tracklets*. Object class detectors determine 2D object locations in the image, and semantic scene labeling infers semantic classes such as roads, the sky, or objects for each pixel, whereas tracklets, due to geometric and dynamic consistency over a longer period of time, accumulate image information over several frames to achieve a more robust reasoning (see figure 2 for sample detections, segmentation, and system results).

By employing 3D reasoning, our model is able to represent complex interactions like inter-object occlusion, physical exclusion between objects, and geometric context. The results are encouraging and show that the combination of the individual components not only allows us to infer a 3D world representation but also improves the baseline given by the object class detectors. Throughout this research project we have used monocular cameras and are therefore not able to directly extract depth from the image disparity. Nonetheless, our approach is able to outperform similar systems that use a stereo camera setup. :::

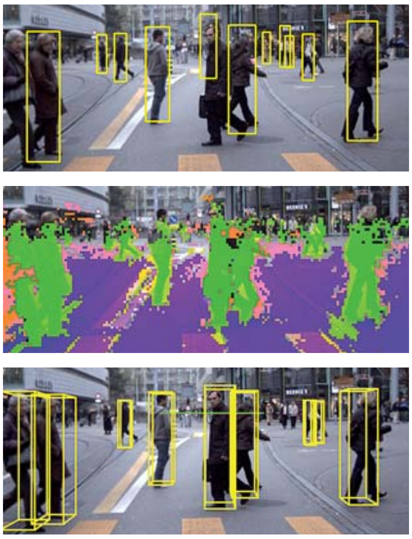

Figure 2: 2D detections, semantic scene labeling and results of our 3D scene estimation

CONTACT

Bernt Schiele
DEPT. 2 **Bildverarbeitung und multimodale Sensorverarbeitung**
Phone +49 681 9325-2000
Email schiele@mpi-inf.mpg.de

Christian Wojek
DEPT. 2 **Bildverarbeitung und multimodale Sensorverarbeitung**
Phone +49 681 9325-2000
Email cwojek@mpi-inf.mpg.de
Internet http://www.d2.mpi-inf.mpg.de/monocular-3d-scene

Learning of Visual Representations

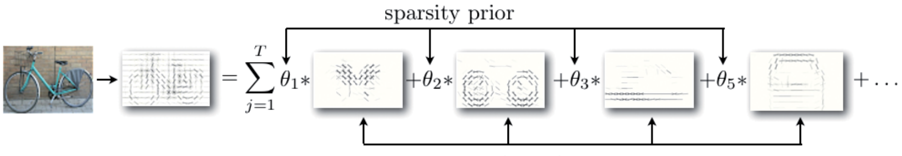

Figure 1

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.

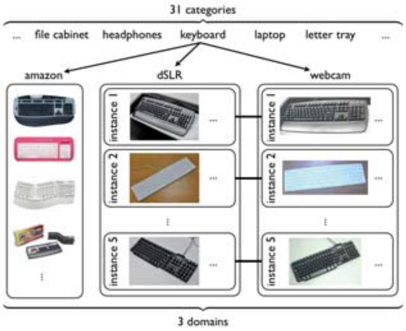

Figure 2

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 20 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 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 in-teractions (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 speaker is Professor Lengauer, received top marks for its research in the final evaluation in 2007 of the five bioinformatics centers that have been funded by the German Research Foundation (DFG) in the past decade. :::

CONTRIBUTIONS

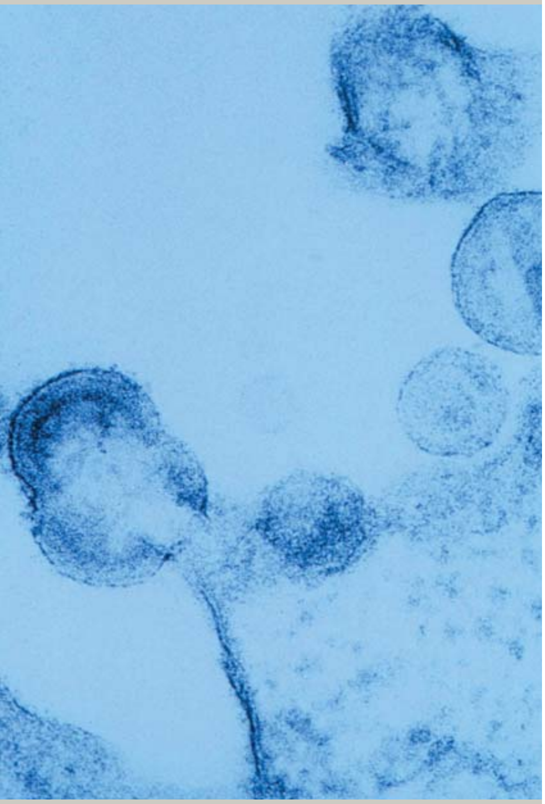

Protein Structure and Interactions

Proteins as biological machines

Proteins are biological molecules that facilitate life's processes: They translate hereditary information, encoded in the DNA, into one or more specific functional roles in the organism. They can do so because they are made up of a chain of amino acids that fold into a complicated three-dimensional structure, thus translating information from 1D to 3D. The specific way in which each protein folds depends on the physical and chemical properties of the amino acids that comprise it. Structural biology has accumulated extensive data on protein 3D structures that can be analyzed using computational methods.

If we extend the analogy described above to mechanisms, a single protein is not a machine, but a part of the conveyor belt that performs an individual action towards a common goal. That is why proteins, on many occasions, assemble into complexes to perform their task. These complexes can be a hundred times larger than a single protein. The prediction and analysis of properties of protein complexes is one of the major tasks in structural bioinformatics.

Structural bioinformatics of viral systems

Viruses invade cells of higher organisms and use them to replicate and multiply. They are small systems, typically containing 10-20 types of proteins. Using this limited toolkit, they engage in a complex network of interactions with the host cell and recruit host cell proteins to replicate. Viral mechanisms to copy their gnomes are prone to error, creating variation, which allows viruses to escape very quickly from the persistent evolutionary pressure posed by the host immune system and drug treatment. The amount of data on the variations of protein sequences in viruses, which are available from labs and clinics, is impressive. This information makes viral systems a unique target for bioinformatics studies.

Protein sequences from virus strains isolated from various patients provide valuable information for patient treatment, as demonstrated by the geno2pheno approach pioneered at our institute (see *"Bioinformatical Support for HIV Therapy"* by Thomas Lengauer). Including structural information in this analysis allows us to understand how the variations of the sequences impact the function of the protein and its interactions with partners. For example, one mutation in the protein's active site can ruin the protein's catalytic ability, whereas a different mutation nearby will merely diminish the protein's affinity to an inhibitor. A mutation that is located on an interaction interface with a host factor may prohibit the binding of the viral protein to its cellular binding partner, and the mutated virus will not replicate.

An integrated analysis of sequences and structures of viral proteins

In our work, we put sequences of viral proteins into structural context, and vice versa. We extract sequence features characteristic of viruses with specific properties (e.g., drug resistance) and analyze structures of the corresponding proteins to understand the mechanism. We have performed an analysis of HIV viruses that selectively bind to one of the two proteins (CCR5 or CXCR4) on the surface of the human immune cell, which the virus has to invade for replication. We modeled both viral and cellular proteins involved and concluded that electrostatic charge plays a pivotal role in this interaction. The charge is likely to be conferred by the posttranslational glycosylation of the viral envelope protein (attaching a sugar molecule to the protein after it has been manufactured by the host cell), which in turn has major implications for vaccine development.

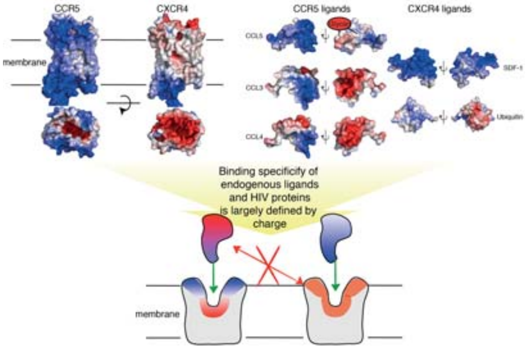

Charge-dependent interaction between CCR5 and CXCR4 and their ligands or – the HIV envelope protein

Using the wealth of data on HIV resistance, we have developed a method for the computational assessment of the binding strength between the HIV protease from different virus strains and its inhibitors (drugs that bind to the protein to block viral replication), which has the potential of substantially assisting the drug development process. :::

CONTACT

Olga Kalinina
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 9325-3004
Email kalinina@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, combination drug therapies are administered. These are composed of drugs from a base set of over two dozen compounds [see figure 1]. 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 redited with the AIDS Research Award of the Heinz-Ansmann Foundation.

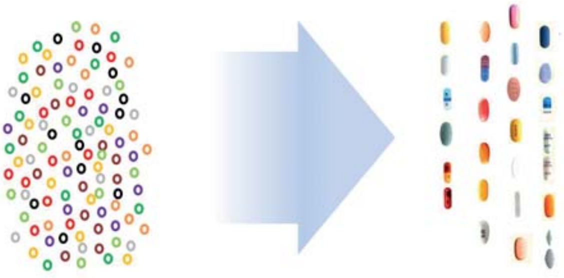

Figure 1: Millions of HIV variants are contrasted with hundreds of combination drug therapies based on over two dozen drugs. The drug selection benefits from bioinformatical support.

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 which 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.

In addition, geno2pheno offers a second generation of bioinformatical support systems for therapy selection. These systems not only estimate viral resistance against individual drugs but also assess the effectiveness of the therapy consisting of a combination of different drugs. This analysis also entails an estimate of how difficult it is for the virus to attain resistance against the current therapy in the future. The second generation of prediction systems is still in the research stage and has not yet entered clinical routine on a broad scale. Reasons for this include the need for additional validation of these methods. Furthermore, the accuracy of the procedures can be improved further, for instance by resorting to additional information such as on the therapy history of the patient or on aspects of his or her genotype, especially those aspects which pertain to the patient's immune system. Finally, the predictions are not yet easy enough to interpret. The therapist expects from a bioinformatical prediction method not only its result but also an argument as to why this result is plausible. Improving the second-generation prediction systems is currently a focus of our research.

New sequencing methods facilitate the resolution of essentially the whole viral population inside the patient. On the basis of such data, the prediction of therapy success can be made more accurate (see article *"Attacking HIV from New Angles"*, page 47 and *"Analysis of Viral Genomes Using Next-generation Sequencing"*, page 46). :::

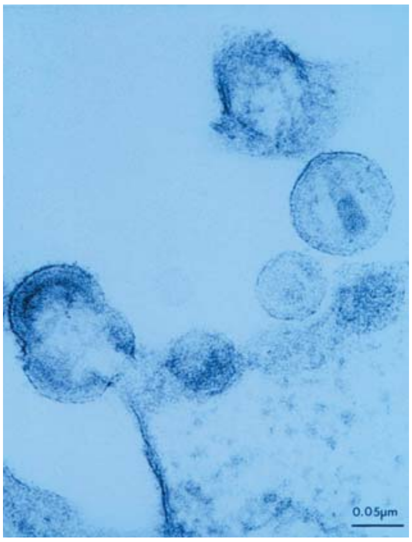

Figure 2: HIV particle budding from an infected immune cell *(lower right), (courtesy Prof. Schneweis, Bonn)*

CONTACT

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

Analysis of Viral Genomes Using Next-generation Sequencing

Research on viral infections is heavily dependent on available genome sequences of both the virus and its human host. These genome sequences provide the basis for understanding the complex molecular interplay between the pathogen and the patient, knowledge that is crucial for both drug development and therapy optimization. The advent of next-generation sequencing technologies resulted in a dramatic drop in cost and increase in throughput of genome sequencing. However, the computational and methodological difficulties regarding storage, analysis, and interpretation of the sequence data are considerable. It has recently become clear that the main bottleneck hindering advances in high-throughput genomic research is not the limited availability of patient samples or sequencing machines, but instead the lack of automated tools and trained bioinformaticians able to analyze this deluge of data.

Standardizing viral sequence analysis with Virana

The situation is further complicated by the fact that viral genomes differ from the genomes of multicellular organisms such as humans in their very high evolutionary rate and intra-species diversity, resulting in an especially complex genotype. These peculiarities compromise the suitability of traditional tools and analysis methods developed for the analysis of the human genome to study infections by viruses. We aim to fill this methodological gap by developing "Virana", a software package and associated web service for viral next-generation sequence analysis. Virana is specifically tailored to the analysis of clinical data and takes the complex viral genotype as well as the longitudinal character of antiviral therapies into account. Analyses performed with the Virana web service can be conducted by non-specialists and are both highly scalable and replicable. Additionally, analysis workflows and results can conveniently be shared over the Internet, encouraging communication and collaboration.

Sequence dynamics of the hepatitis C virus

At the Hospital of the University of Frankfurt, Virana is currently employed in a clinical setting to model the intra-host evolution of the hepatitis C virus. The characterization of the highly complex genotype of this virus is crucial for studying novel resistance mechanisms that the virus evolves to escape antiviral drug regimens. The unique methods offered by Virana enable researchers to statistically infer genotypic markers, such as minority sequence variants and genomic haplotypes, from the sequence data. When tracked over the course of an infection and combined with phenotypic variables such as therapy success, these markers can be used to support further clinical decision-making. Results of such analyses can then be incorporated into our software geno2pheno[hcv], a free Web service that aids medical doctors by predicting drug resistance of viruses derived from patient samples [see figure], thereby translating personalized genomics into clinical practice.

Modification of the natural progression of chronic hepatitis B

Next-generation sequencing is also used to analyze the complex interaction of the hepatitis B virus with the host's immune system on the basis of the viral genome. The natural progression of a chronic hepatitis B infection consists of four phases. The immune-tolerant phase and the inactive phase are characterized by a favorable progression of the infection, whereas patients in the immune-active and the reactive phases are at high risk of developing irreversible liver damage in the medium term. Despite intensive research, relatively little is known about the determinants that cause phase transitions, e.g. from the inactive to the reactive phase. Together with our collaboration partners from the university hospital in Düsseldorf, we analyze the entire viral population within a host of multiple patients at several characteristic points in time. This will eventually provide indicators to characterize and predict phase transitions. By these means it may be possible to prevent deleterious phase transitions in an early stage using antiviral drugs, which would avoid or at least delay fatal progression of a chronic hepatitis B infection using as little medication as possible. :::

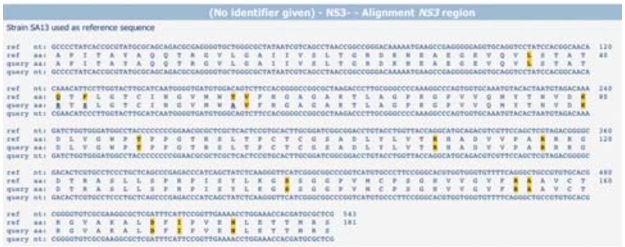

(No identifier given) - NS3- - Alignment NS3 region

Strain SA13 used as reference sequence

Output of the geno2pheno[hcv] web service

CONTACT

Bastian Beggel
DEPT. 3 **Computational Biology and Applied Algorithmics**
Phone +49 681 9325-3016
Email beggel@mpi-inf.mpg.de

Sven-Eric Schelhorn
DEPT. 3 **Computational Biology and Applied Algorithmics**
Phone +49 681 9325-3028
Email sven@mpi-inf.mpg.de
Internet http://hcv.bioinf.mpi-inf.mpg.de/

Attacking HIV from New Angles

Visualization of a prediction result: In addition to the improved prediction method, a new visualization method was developed that shows which amino acids at which positions are most informative for the prediction. Many large red letters mean that the patient should rather not use the entry inhibitor, while green letters symbolize amino acids that can be found in viruses that are susceptible to the treatment. The above sequence is an example of a V3 loop sequence of a virus that should not be treated with the entry inhibitor.

According to the World Health Organization (WHO) 34 million people were living with an HIV infection worldwide in the year 2011. About 1.7 million died due to AIDS-related complications – including 230,000 children. The number of new HIV infections in 2011 was about 2.5 million according to WHO. Preventive measures are an important aspect of tackling these problems. There have been some significant breakthroughs over the past years in HIV vaccine research. For example, Walker et al. were able to extract broadly neutralizing antibodies that can neutralize many different HIV strains from patients with strong immune responses. Since there is no approved method of curing an HIV infection to date, a large part of HIV research is concerned with alleviating symptoms and extending the lifespan of infected patients with the help of antiretroviral drugs. Entry inhibitors are one new class of antiretroviral drugs against HIV that block HIV entry into human cells.

Prediction of HIV coreceptor usage from the genetic sequence

Before prescribing an HIV entry inhibitor it has to be determined whether the drug is effective on the viral population inside the particular patient or not. HI viruses bind to a certain protein on the surface of the host cell, a coreceptor, in order to be able to enter the cell. Different viral variants can use different coreceptors, but the most important coreceptor is the CCR5 coreceptor, which can be blocked by an entry inhibitor. Therefore, only patients with viral populations that exclusively use the CCR5 coreceptor can be treated with this kind of antiretroviral drug. To determine whether the virus harbored by the patient has this property, we follow the approach described in *"Bioinformatical Support of HIV Therapy"*, page 45. In this application scenario, it is general practice to use a part of the sequence of the envelope gene of HIV called the V3 loop, since this part of the corresponding viral envelope protein is mainly responsible for binding to the cellular coreceptor during cell entry. In our research, we could show how to train a statistical model that improves the prediction of coreceptor usage by the virus based on the complex data produced by new-generation sequencers. The better the coreceptor usage prediction model, the better recommendations it will give about whether an entry inhibitor is suitable for the patient. We were also able to craft a statistical model that uses the new-generation sequencing data to improve the prediction of coreceptor usage based on data produced by conventional sequencing machines. This is especially important since many clinics do not have access to the new-generation sequencing techniques. Additionally, we could identify important positions in the V3 loop, which are highly predictive of coreceptor usage. This could provide information on the mechanism of interaction between the V3 loop and the coreceptor and can potentially lead to new measures to attack HIV.

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. :::

CONTACT

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

Networks and Phylogenies: Understanding Disease Spread and Evolution

Network models of disease spread

Networks, where nodes represent individuals or groups, and edges the connections between them, have become the premier mathematical structure for modeling disease transmission because they capture the heterogeneity of human contacts. The structure of the network powerfully influences disease spread. Different network structures can be used, depending on the specifics of the disease under consideration.

The small world is one common model of social contact. This portrays society as a collection of small groups whose members are in close contact, united by nodes called bridges, which connect groups. The key to preventing disease flow in such networks lies in identifying the bridges along which the disease passes from group to group. We have found that a simple morphological measure, the K-shell, identifies bridge nodes critical to disease flow in small world networks. Our work has further shown that interventions such as vaccinations which target these nodes can efficiently slow or stop the disease from spreading through the whole network. For example, if public facilities such as schools and playgrounds are modeled as nodes in such a network, the method could identify which facilities would need to be closed to halt the spread of an outbreak.

A second common model is the scale-free network. This model emphasizes that a few nodes have a large number of connections, while most nodes are peripheral, having only a few connections. Most new or emerging diseases enter networks from this periphery. We have pioneered a metric of how easily an individual node can spread an infection to the rest of the network, based on its local contact structure. Informally, our results show that a node's spreading power is not based on how many connections it has, but rather on how many connections its contacts have.

Phylogenies

Many pathogens evolve rapidly enough that measurable genetic change occurs on the same timescale as disease transmission, making phylogenies an essential tool for studying disease dynamics.

The past five years have seen a growing application of phylodynamics to understand the spread of HIV. Our August 2012 publication in Medical Microbiology and Immunology, *"Endogenous or Exogenous Spreading of HIV-1 in North Rhine-Westphalia, Germany"*, showed that the HIV epidemic has strong endogenous roots in that region. We are currently reconstructing transmission clusters using data from the EURESIST consortium. This data has substantial coverage across Europe, making this the largest pan-European study till date. Preliminary results show evidence for low geographic mixing and small intertransmission intervals.

Deep sequencing technologies give us the ability to inspect the entire viral population circulating in a patient. Often, the data includes multiple time points. This opens unprecedented opportunities to understand intra-host viral evolution. The structure of the data, however, violates many of the base assumptions of traditional phylogenetic methods. We are actively developing appropriate tools, with application to the evolution of drug resistance in HCV.

Phylogenetic methods have applications far beyond understanding rapidly evolving pathogens; they are applied in all aspects of biology. Yet the methods and tools available are still in their infancy. Those which handle large data sets must ignore important subtleties. Bayesian methods address these subtleties, but can only handle small data sets. A hybrid Bayesian/Maximum Likelihood approach, which we are developing, may allow rapid Bayesian estimation of large phylogenies. :::

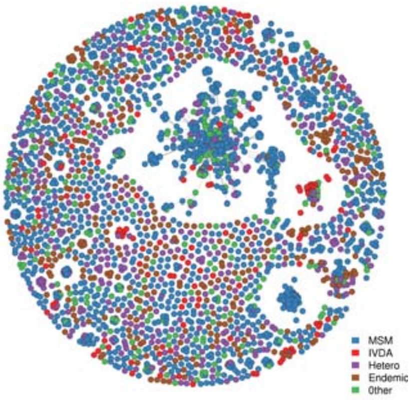

Genetic sequence similarity of HIV-1 samples from patients in North Rhine-Westphalia indicates several large clusters, which implies that the epidemic has strong endogenous roots in the region and remains an important concern to public health.

CONTACT

Glenn Lawyer
DEPT. 3 Computational Biology and Applied Algorithmics
Phone +49 681 9325-3007
Email lawyer@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 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.

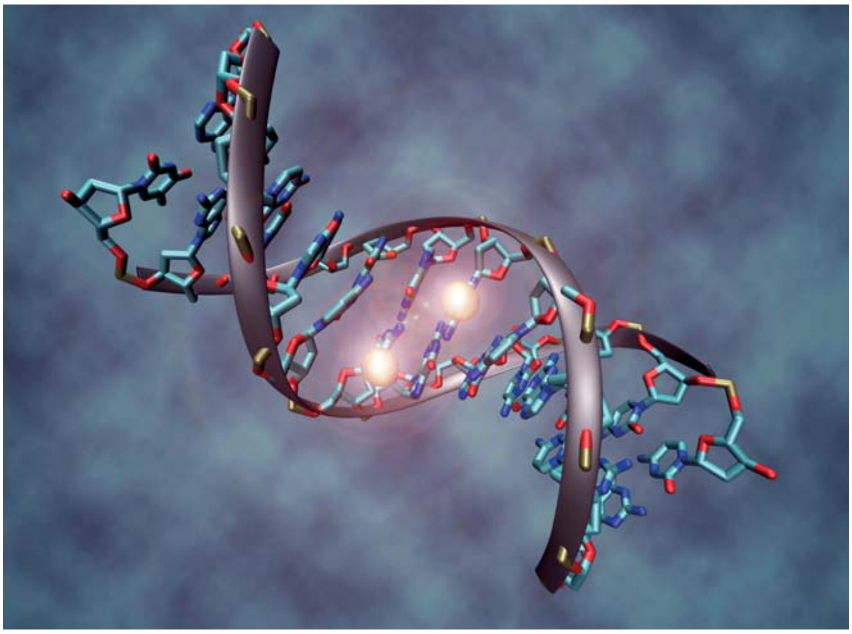

The methylation of the DNA impacts its accessibility by the transcription machinery of the cell.

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 an EU funding volume of 30 million euros, the largest that ever went to a single project in the area of biology, is targeted at charting one hundred epigenomes of cells of the bloodline (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 20 million euros comprises the German contribution to IHEC. Within DEEP, metabolic and immunologic diseases in particular are investigated, and another 70 epigenomes are charted. The Max Planck Institute for Informatics coordinates the bioinformatics part of DEEP. :::

CONTACT

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

GUARANTEES

Software 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 question at the institute.

Today, computers are a ubiquitous part of our lives. We constantly use them, sometimes knowingly, such as a computer on the desk, and sometimes unknowingly, such as the electronic controls in a car, airplane, or washing machine. The more our lives are dependent upon hardware and software, the more we pose the question whether the trust that we place in these products is justified. The answer to this question is especially difficult and complex as software 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 to deliver wrong results. Another common error is ignoring rare but possible inputs, with the result that software suddenly generates a false result after having worked reliably for years. The granting of guarantees is even more complex when we examine not only a single program, but a system of programs, in which parts work independently but also react to each other while pursuing different goals.

In the article *"Dealing with Selfishness in Optimization",* page 56, the latter situation is investigated. We take, for example, our current internet providers who have committed, on the one hand, to forwarding all Internet traffic, although they would ideally prioritize the data transfer of well-paying customers in order to optimize profit.

Guarantees of geometric calculations are especially important today as nearly every larger machine, such as an airplane, for example, is designed on a computer. It would be undesirable if a calculation error caused, for example, the rudder of an airplane to become jammed during its service life. The article *"Exact Solving of Polynomial Systems and Applications in Geometric Computing"*, page 52, shows what is needed for accurate calculations in a geometric context.

In order to provide guarantees of programs in general, deductive methods must be used that enable the desired properties to be derived from a program in the form of proofs. Automatic verification procedures are required for this, which are discussed in the article *"Automated Deduction",* page 53, 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",* page 57. A targeted application of proof methods for verifying the answers of algorithms is described in the article *„How to Trust the Output of Your Program"*, page 54. Ultimately, it is not yet possible to give all desired guarantees of an arbitrary complex program with perfect precision and within an acceptable time period. In such a case, it makes sense to use procedures based on model assumptions, as is presented in the article *"Model Checking for Hybrid Systems",* page 55. :::

CONTRIBUTIONS

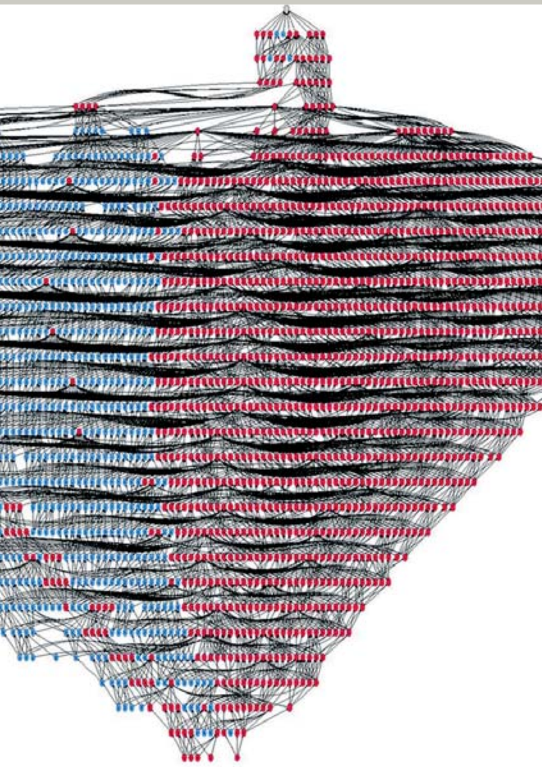

Exact Solving of Polynomial Systems and Applications in Geometric Computing

Solving systems of polynomial equations poses a fundamental problem of algorithmic mathematics. For m given polynomials in n variables, all points in n-dimensional space should be found such that all of the polynomials simultaneously take on the value 0. Such systems of equations occur in nearly all areas of mathematics, physics, and engineering. The research group "Computer Algebra and Geometric Computing" primarily investigates such systems as they typically arise from geometric problems. Most algorithms from computational geometry or computer-aided design require such exact calculations with geometric objects. For example, one might wish to find the intersection points of two curves, or to determine whether multiple curves intersect at the same point. Such problems normally lead to a low-dimensional system of equations, in which some of the polynomials may be of high degree. A common solution method from computer algebra is the implicit projection of all of the solutions onto a line. By using algebraic transformations, one then finds a so-called elimination polynomial, whose zeroes are exactly the projections of the solutions of the system. In many cases, the general problem can thus be reduced to the special case $m=n=1$.

In comparison to pure numeric methods, as are almost exclusively applied in industrial software packages, the above-described approach allows the development of *complete* and *exact* methods. This means that even for particularly difficult inputs, all solutions are found and are provable as such. Unfortunately, there are also downsides: the symbolic transformations generally lead to a very high computational demand, even if the solutions are of a relatively simple geometric nature (for example, solutions that are simple or well separated from one another). This maladaptive behavior results from the fact that the underlying algebra is "ignorant" with respect to such simpler situations. A further problem of the projection approach is that the corresponding elimination polynomial exhibits high complexity, which makes finding its zeroes difficult.

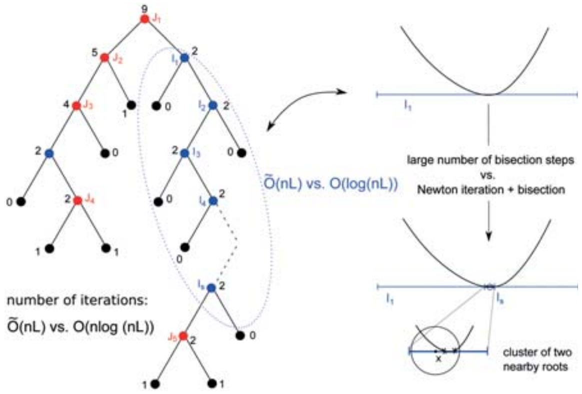

Combination of the Descartes method *(linear convergence)* and Newton's iteration *(quadratic convergence)* produces a simple and nearly optimal process for the determination of the real zeroes of a polynomial.

In previous years, the research group has developed methods to considerably improve the efficiency of elimination procedures. An essential goal is always to replace unnecessary symbolic computations with certified numeric ones, however, without sacrificing exactness or completeness. In addition, we develop efficient methods to carry out the remaining, and also necessary, symbolic steps. For the computation of the elimination polynomial for bivariate systems, as well as that of the greatest common divisor of polynomials, we have developed a new algorithm, whose implementation on graphics hardware takes greatest advantage of the parallelizability of the method. In comparison to the fastest competing implementations, the aforementioned computations can thus be speeded up by a hundredfold.

Through the combination of results from different areas, such as computer algebra, numerical mathematics and algebraic geometry, we have also succeeded in developing procedures that require only those symbolic calculations which can be delegated to graphics hardware. Thus, with regard to speed, we were able to catch up to numerical methods – which deliver no additional guarantees – for the first time. In parallel, we also were able to prove the theoretical efficiency of our algorithms. The corresponding bounds on the number of steps needed to find the solution improve on the best previously known bounds by multiple orders of magnitude. Based on present knowledge, there is evidence that these are actually nearly optimal [see figure].

We also succeeded in finding a similar breakthrough for the classical special case where $m=n=1$. In the 80s and 90s, complex algorithms were first developed to find the zeroes of a polynomial in nearly optimal (theoretical) runtime. However, the practical relevance of these procedures has so far not been confirmed. Taking a different approach, we were able to modify a simple, and in practice very efficient, process such that it could compete with the optimal methods, even in terms of theoretical complexity. :::

CONTACT

Michael Sagraloff

DEPT. 1 Algorithms and Complexity

Phone +49 681 9325-1006
Email msagralo@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/departments/d1/areas/software.html

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 (e.g., 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, e.g., “$x + 0 = x$” only from left to right and not the other way round.

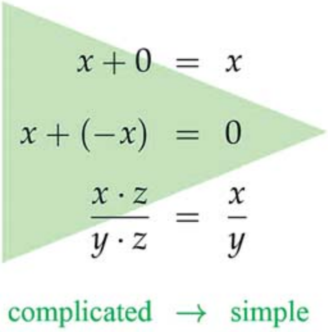

Application of equations

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, we currently focus on refining the general superposition method for special applications. For instance, we are developing techniques for combining the capabilities of various proof procedures (for instance, superposition and arithmetic decision procedures). We are addressing the question of how superposition can be used for business applications (product configuration, authorization checking). We also use superposition to verify network protocols and to analyze probabilistic systems, i.e., systems whose behavior is partially dependent on random decisions :::

```
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

How to Trust the Output of Your Program

Transferring an algorithm (i.e. a solution procedure to computer code is difficult and error prone. For that reason, the code may contain many mistakes, and subtle errors can be very hard to find. For example, the planarity test in LEDA contained an error that was only discovered after years of intensive use.

The field of Software Engineering has developed many tools to improve the quality of code. A standard practice, for example, is testing. This requires a diverse set of inputs with known solutions. Often this is hard to come by, as, intuitively, the program wouldn't be needed if finding the correct answers for example inputs were easy.

Certifying algorithms approach the problem from a different angle. Instead of trying to improve the quality of implementations, certifying algorithms improve the methods to solve the problem such that the implementation is easier to check for correctness. This is achieved by enhancing the algorithms to provide a certificate that proves that their answer is correct; i.e., the algorithm does not provide only a bare yes-or-no answer to the question, but also gives a reason for this answer. The certificate is structured such that a much simpler program, the checker, can verify its correctness. Only the checker needs to be carefully examined for correctness.

An important goal of certifying algorithms is to find certificates that are so simple that the checker can be formally verified, that is, one can give a mathematical proof that the code is correct. Formal verification, unlike mere testing, guarantees the absence of errors, but is impractical for most real-world programs. Through certifying algorithms we can use this powerful technique to prove that our algorithm's answers are always correct if the checker agrees with them.

Developing certifying algorithms is a major goal for problems where the fastest solutions known are complicated. Having a certifying algorithm for such a problem allows verifying the output of every instance, regardless of how complicated the computation of this output is. For this reason, certifying algorithms are now used in software libraries like LEDA for reliability purposes.

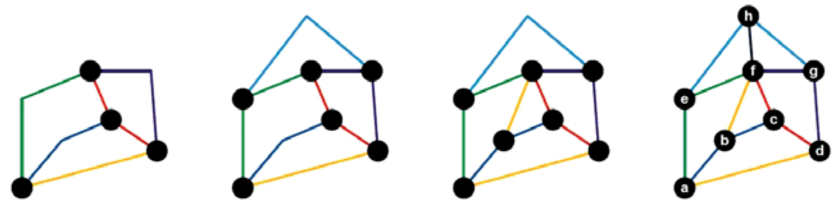

(a) $K_4 = G_4 = smooth(S_4)$ (b) $G_5 = smooth(S_5)$ (c) $G_6 = smooth(S_6)$ (d) $G_7 = G$

A construction sequence of a 3-connected graph G

Current research problems

Prominent examples of problems whose fastest known solutions are complicated include testing a graph for 3-connectivity and for 3-edge-connectivity. These problems are important for various applications. Although linear-time solutions for these problems have been known for decades, none of them describe an easy-to-verify certificate.

A natural certificate to prove 3-connectivity is a so-called *construction sequence*. For a 3-connected graph G, a *construction sequence* is a sequence of operations that constructs G, when applied to K_4.

We proposed the first linear-time certifying algorithm for testing a graph for 3-connectivity. The algorithm has been implemented and made publicly available. By using a well-known reduction from edge-connectivity to vertex-connectivity, we can obtain a linear-time certifying algorithm for 3-edge-connectivity.

A current research problem is to find a direct certificate for 3-edge-connectivity. Future open questions include finding certificates for k-connectivity or k-edge-connectivity for $k \geq 4$.

Verification of Checkers

We developed a framework for verifying checkers. This guarantees instance correctness of all certifying algorithms whose certificates can be checked by this framework. We applied the framework to several examples from the LEDA library in the field of graph theory. Currently, we are evaluating alternative tools for the verification framework and comparing the results in terms of trust and effort needed. :::

CONTACT

Adrian Neumann
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1022
Email aneumann@mpi-inf.mpg.de

Christine Rizkallah
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1112
Email crizkall@mpi-inf.mpg.de

Jens M. Schmidt
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1008
Email jeschmid@mpi-inf.mpg.de

Model Checking for Hybrid Systems

The behavior of many technical devices can be described by specifying their states and state transitions. A washing machine can, for example, be in the starting state "door open, drum empty, machine turned off" or the state "washing." State transitions are changes from one state to another. These can take place automatically, such as the change between washing and drainage, or through other actions, for example when a button is pressed. There are, however, possible states that are obviously undesirable, for example: "water inlet open, door open." If we want to verify that the machine is safe, then we must show that there is no sequence of state transitions leading from the starting state to such an unsafe state.

How can this problem be approached? We can describe states and state transitions using *graphs*. To this end, we symbolize each state by a circle and each possible state transition by an arrow. The result may look as follows:

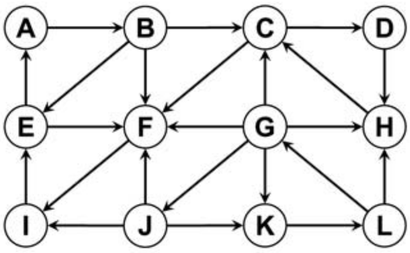

Example

Let us assume that A is the starting state and L is unsafe. What are the predecessor states of L, i.e., the states from which an arrow points to L? Obviously, there is only one such state: K. K is thus also unsafe, since if K can be reached, then L can be reached from there. K's predecessor states, namely G and J, are unsafe for the same reason. If we take a look at the predecessor states of G and J (namely L and G), we can see that these are already known to be unsafe. This means that there is no way of reaching the state set $\{G, J, K, L\}$ from outside.

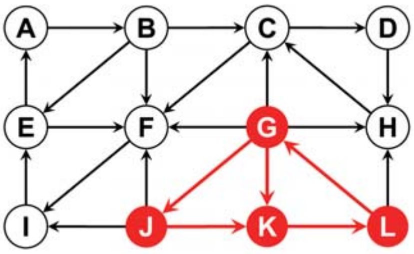

Non-safe states

It is therefore clear that $\{G, J, K, L\}$ are the only states that are unsafe per se, or from which an unsafe state can be reached. All other states, in particular the starting state A, are safe. This technique for verifying the safety of a system is known as *model checking*.

Unfortunately, the number of states can quickly lead to problems in practice. Imagine that the behavior of an electronic device depends on 60 bits ($b_1,\ldots,b_{60}$), i.e., 60 components that can each have the value of true or false. Now, 60 bits are not many, but if all the possible combinations are taken into account, we arrive at 2^{60} states, more than one quintillion. *Symbolic model checking* offers a way out here: We no longer enumerate the unsafe states and permissible state transitions explicitly, but we represent them *symbolically* using logical formulas. For example, the formula

b_5 and not b_{32}

describes all states for which b_5 is true and b_{32} is false – this is a quarter of a quintillion states, which would otherwise have to be enumerated explicitly.

A new problem emerges if we try to verify the safety of a vehicle control system in this way. Such a control system is a hybrid system. This means that one is no longer dealing solely with *discrete* state changes (e.g., acceleration or braking), but also with variables such as speed, which can change continuously. The state formulae must therefore also include numerical variables, for example:

b_5 and not b_{32} and $(x_2 > 50)$.

Discrete state transitions can still be computed in the same way as for purely discrete systems. What about the continuous state transitions? As long as the dynamic behavior is simple from a mathematical point of view (e.g., if the speed changes uniformly), one can use a method known as quantifier elimination. Certain technical tricks are needed, however, to prevent arriving at too large a formula too quickly. In our work group, we are currently focusing on developing methods that can be applied even to more complicated dynamic behaviors. :::

CONTACT

Uwe Waldmann
RG. 1 **Automation of Logic**
Phone +49 681 9325-2905
Email uwe@mpi-inf.mpg.de

Dealing with Selfishness in Optimization

In my research, I focus on optimization problems that involve selfishness in some way. This usually takes the form of selfish agents holding part of the input data and having the option of misrepresenting their data if this is to their own benefit. A prime example of this is given by scheduling problems where the machines are controlled by selfish agents, where only each agent knows how fast their machine is. The agents are only interested in maximizing their own utility, which may be quite different from the overall optimization goal. Our goal in algorithmic mechanism design is to design mechanisms that ensure that reporting their speeds truthfully is a dominant strategy for the agents. Only in this way can we hope to discover the true speeds and optimize the overall objective function successfully.

A mechanism collects bids from the agents that are hopefully true representations of their speeds, allocates jobs to the machines and pays the agents for their work. The utility of an agent is the payout that it receives from the mechanism minus the cost that it has for processing the jobs assigned to it. This cost is also known as the load on its machine, and is equal to the total size of the jobs allocated to it divided by its speed. A strategy (reported speed) of an agent is dominant if this strategy maximizes the utility of this agent regardless of the strategies chosen by the other agents. If telling the truth is a dominant strategy for every agent, we say that our mechanism is truthful.

It is known that in order to get a truthful mechanism, a necessary and sufficient condition is that the job allocation of the mechanism be monotone. Monotone means that if one agent decreases its claimed speed, the total size of the jobs assigned to it by the mechanism as a result may not increase. The set of assigned jobs does not necessarily have to remain the same, however.

Many objective functions for scheduling have been considered in the past, both in a classical setting (without agents) and in the game-theoretical setting described above. Many of these are some function of the loads of the machines, for example, the makespan (maximum load). Together with Leah Epstein and Asaf Levin, I have recently designed a unified monotone approximation scheme for all known objective functions that are functions of the loads, like the makespan, the cover (where the goal is to maximize the minimum load), and the p-norm of the loads. An approximation scheme is a family of approximation algorithms that, in addition to the given machines and jobs, takes a positive value epsilon as input and produces an output for which the value of the objective function is within one plus epsilon of the optimal value for these jobs and machines. This improves and generalizes on previous approximation schemes, which depended on randomization, required many special cases and a difficult analysis, or were only quasi-polynomial-time.

In addition to algorithmic mechanism design, I am also interested in the related question of how bad selfishness is, or put another way, by how much we could improve the performance of a system using centralized coordination compared to letting every participant make their own decisions. A well-known example of a system where everybody makes their own decisions is routing: when you plan your route to work in the morning, you are generally only interested in your own travel time and do not care about the travel times of the other people on the road. We formalize such systems as networks with selfish agents. Every agent has a starting point and a destination, and wants to reach its destination as fast as possible. How long this takes depends on how many other agents use the same links as this agent. Much is known about selfish routing on general graphs, but the ring structure has not yet been studied in much detail (though some initial results have been obtained). This structure is very relevant, however, as the Internet in Europe in fact consists of interlocking rings (e.g. the German network is based on a ring). Together with Xujin Chen, Benjamin Doerr, Xiaodong Hu, Weidong Ma, and Carola Winzen, I have shown that on rings, at most a factor of two can be gained by having a central coordinator, if the goal is to minimize the maximum latency of any player. :::

CONTACT

Rob van Stee
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1005
Email vanstee@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 expanded this to computation with symbolic expressions: A simple example is the calculation of $(x-1) \cdot (x+2)+2$ with the result of $x^2 + x$. That computational work has indeed been performed here, becomes clear, among other things, by the fact that the evaluation of the original expression for a concrete x requires double the amount of arithmetic operations as that of the result. On this basis it became possible, e.g., to automatically derive real functions or even calculate indefinite integrals automatically.

In mathematics, symbolic expressions typically appear in complex statements which can quantify some of the occurring symbols, as for example the following statement about real numbers: For all x there is a y, such that $x^2 + xy + b > 0$ und $x + ay^2 + b \leq 0$. 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 which 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 which 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 integral 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 integral and real numbers, there are many other domains in the natural sciences which are 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 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.

Other interesting applications are 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, multimodality, digital history, and distribution.

- In addition to more than 20 billion web pages, important information sources are: online news streams, blogs, tweets and social networks with several hundred million users, Web2.0 communities for photos, music, books, scientific topics, and last but not least the encyclopedia Wikipedia. The total volume of this data is in the order of exabytes: 10^{18} of bytes – more than one million terabyte disks.

- There is an increasing amount of more expressive data representations: 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.

- 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. This is often accompanied by interpersonal contacts that arise from the Internet, and large online networks are organized.

- 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.

- The quantity and variety of the available information in the Internet has grown to such a scale that search engines are no longer able to maintain all relevant references in a centralized index. Therefore, global information search necessitates distributed algorithms, in which, for example, many local search engines are dynamically coupled for specific tasks. Here, not only local computing and search speed play an important role, but also communication efficiency in the network of networks, the Internet, and in embedded peer-to-peer networks.

This global issue is examined from various perspectives at the Max Planck Institute for Informatics. These include work on efficient search in semistructured XML documents, which play a big role in digital libraries and e-science data. Another research area is the scalable management of graph-structured RDF data which arises in the context of the semantic web but is also important as data representation in computational biology. In other projects, the user-centric view of Web 2.0 communities and multimodal data are the primary issues. 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

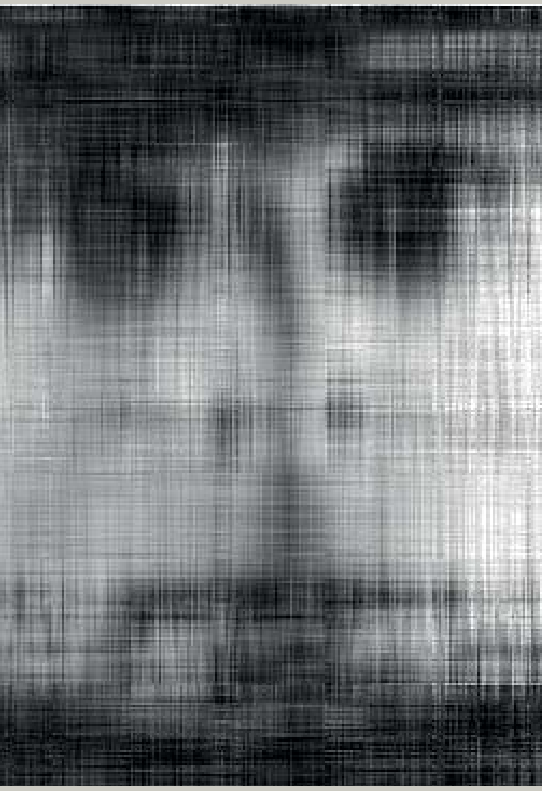

UNDERSTANDING IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH & DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION & VISUALIZATION

OPTIMIZATION

SOFTWARE

Search and Mining in Web Archives

The World Wide Web evolves constantly, and every day contents are added and removed. Some of the new contents are published first and exclusively on the Web and reflect current events. In recent years, there has been a growing awareness that such born-digital contents are part of our cultural heritage and therefore worth preserving. National libraries and organizations such as the Internet Archive [http://www.archive.org] have taken over this task. Other contents were published a long time ago and are now, thanks to improved digitization techniques, for the first time available to a wide public via the Web. Consider, as one concrete example, the archive of the British newspaper The Times that contains articles published from as early as 1785.

Our current research focuses on scalable search and mining techniques for such web archives. Improved search techniques, on the one hand, make it easier for users to access web archives. Mining techniques, on the other hand, help to gain insights about the evolution of language or popular topics. In the following, we describe three aspects of our current work.

Time travel in web archives

Existing search techniques ignore the time dimension inherent to web archives. For instance, it is not possible to restrict a search, so that only documents are retrieved that existed at a specified time in the past.

In our work, we consider time-travel keyword queries that combine a keyword query *(e.g., "bundestag election projection")* with a temporal context such as September 2009. For this specific query, only relevant documents that discuss election projections and which existed back in September 2009 should be retrieved as results.

Our approach builds on an inverted index that keeps a list of occurrences for every word. Depending on the type of query that has to be supported, the inverted index remembers an identifier, how often the word occurs, or the exact positions at which the word can be found in the document for every document in which a word occurs. We extend this information by a valid-time interval to also keep track of when a word was contained in a document and thus to enable time-travel queries.

Consecutive versions of the same document tend to differ only slightly. We exploit this observation to reduce index size drastically. To process time-travel queries more efficiently, we keep multiple lists for every word in the inverted index, each of which is responsible for a specific time interval. This introduces redundancy to the index, increases its size, and thus leads to a trade-off between index size and query-processing performance. Our approach casts this trade-off into optimization problems that can be solved efficiently and determine the lists to be kept for every word in the inverted index.

Temporal information needs

Information needs often have a temporal dimension, as expressed by a temporal phrase contained in the user's query and are best satisfied by documents that refer to a particular time. Existing retrieval models fail for such temporal information needs. For the query *"german painters 15th century"*, a document with detailed information about the life and work of Albrecht Durer (e.g., mentioning 1471 as his year of birth) would not necessarily be considered relevant. This is because existing methods are unaware of the semantics inherent to temporal expressions and thus do not know that 1471 refers to a year in the 15th century.

To capture their semantics, we formally represent temporal expressions as time intervals. We then integrate them into a retrieval approach based on statistical language models that has been shown to improve result quality for temporal information needs.

Mining of characteristic phrases

Mining of web archives is another aspect of our current work. More precisely, we are interested in insights about ad-hoc subsets of the web archive, for instance, all documents that deal with Olympics. Given such an ad-hoc subset, we can identify persons, locations, or in general, phrases that are characteristic for documents published in a particular year. In our Olympics example, these could include Michael Phelps, Beijing and *"bird's nest"* for documents published in 2008.

To identify such characteristic phrases efficiently, one needs frequency statistics for so-called n-grams (i.e., sequences of one or more words). We develop efficient and scalable techniques to compute these n-grams statistics in a distributed environment. One design objective here is to allow for easy scale-out in order to keep up with the growth of web archives in the future. :::

CONTACT

Klaus Berberich

DEPT. 5 **Databases and Information Systems**

Phone +49 681 9325-5005

Email kberberi@mpi-inf.mpg.de

Big Data: 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, as well as advances in gathering scientific data contribute 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 Max Planck Institute for Informatics, we develop efficient, highly scalable methods and systems for the statistical analysis of such big datasets. For example, internet companies such as Amazon, Google, or Netflix provide their users personalized recommendations for products, websites, news, movies, or images. From a user's point a view, such personalized recommendations enable 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: An incomplete matrix

Figure 2: 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 shows such an incomplete matrix; here black dots correspond to unknown ratings. We have developed parallel algorithms that can complete incomplete matrices with millions of users, millions of movies, and billions of entries in 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 simply recommend the movies with the highest predicted ratings. In practice, however, it is important that recommendations are diverse, i.e., that each user's recommendations consist of sufficiently different movies (e.g., different genres or different directors). The availability of licenses or physical media is also important: 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 the ones above.

Apart from personalized recommendations, our group works on methods for interactive exploration of large document collections, for the analysis and automated extraction of knowle dge from natural language text, and for pattern mining and logical reasoning on the resulting knowledge bases. :::

CONTACT

Rainer Gemulla
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5004
Email rgemulla@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 which person (or organization, or place, or film, song, etc.) is mentioned where on the Web, or indeed in any given text, is very useful for a multitude of applications. Where previously search engines could only look for a string of characters, now they can actually understand what exactly the user is looking for, giving much more precise results. Knowledge about the real meaning allows you to actually specify that you are looking for the rock group called *"Bush"* and not a US president. Or imagine a researcher looking for differences in the media reception of Bush Sr. and Bush Jr. She can easily get all articles and how often they were mentioned in each of them without having to look at a single one with her own eyes.

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 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. The 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 suggest otherwise. Take for example the sentence *"Paris had to steal Helen from her husband, the king of Sparta"*. The words in the sentence suggest that this is a person from Greek mythology. This kind of contextual evidence is the second feature used by 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, the 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 of Troy are a better fit than Paris Hilton and Helen, Georgia, a small US city.

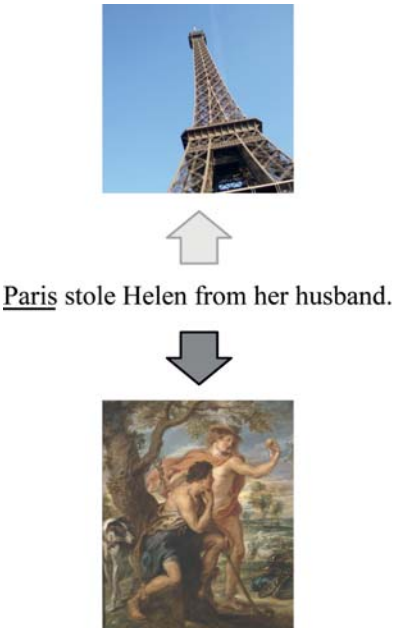

The interplay of the described features can be seen very well in the sentence *"Bush did not handle the aftermath of Katrina in New Orleans very well"*. New Orleans is easy to resolve, as the city is very prominently associated with the name. The first name *"Katrina"* is highly ambiguous, but the way 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 achieves 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 9325-5028
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 mamir@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/yago-naga/aida/

Ontology Matching with PARIS

In the last decade, the World Wide Web has been extraordinarily successful: For many questions of everyday life, the Web has an answer ready. However, computers cannot *"read"* or *"understand"* webpages; they can only display them to us. Therefore, if we are searching for information, we must read through the appropriate webpages ourselves. For example, to search for a French singer's concerts in Germany, one must first find a list of French singers and then study their tour plans. Parallel to the World Wide Web, in recent years the so-called Semantic Web has been developed. It is the computer-readable counterpart to the World Wide Web. There, information is stored in a form that computers can work with directly. Thus, if a computer, at the request of its user, searches for concerts by a French singer in Germany, the computer can first find a list of French singers in the Semantic Web, and then – elsewhere in the Semantic Web – find the singer's concerts in Germany in cities near the user. Development of the Semantic Web has been underway for a few years.

In the Semantic Web, information is stored in so-called ontologies. An ontology is a directed graph in which the nodes are entities (for example the singer Alizée or the country France), and their edges are relationships (for example the relation *"is a citizen of"* between Alizée and France). Furthermore, the entities are grouped in so-called classes. Alizée for instance, is in the class of singers, and France is in the class of European countries. In the Semantic Web there are hundreds of such ontologies. Each one contains entities, classes, and relationships. Altogether, they have multiple milliards of entities and multiple milliards of relationships between them. There are ontologies for musicians, geographic terms, books, medical themes, and (through the state support of the USA and Great Britain) also for public themes like schools, finance, and transport.

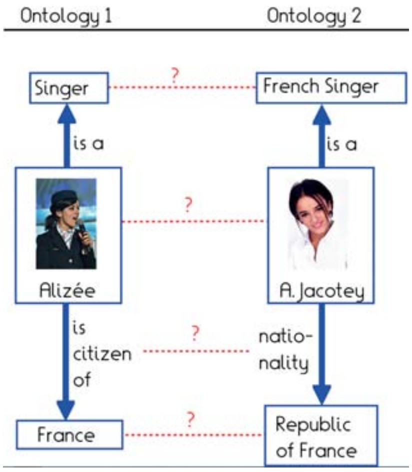

Since individual ontologies exist independently of one another, anyone can contribute ontologies to the Semantic Web. The difficulty lies in how the information in different ontologies can be linked together. If, for instance, one ontology knows that Alizée is a singer (and not, say, a mandolin ensemble), and another ontology contains the concerts of a certain *"A. Jacotey"*, then the entity Alizée in the one ontology should be matched to the entity *"A. Jacotey"* in the other. Otherwise, computers would not be able to integrate them properly in the search for the singer's concerts. This matching process is anything but simple: Even if both ontologies have only 20 entities, there could be be two trillion potential matches between the ontologies. Real ontologies contain multiple millions of entities. Also, it is not necessarily the case that each entity in one ontology is in the other ontology. Furthermore, not only must the entities be matched, but also the names of the relationships. The relation *"is a citizen of"* in one ontology could be called *"has the nationality of"* in another ontology. The situation is similar for the names of classes. Besides, the connections between classes are asymmetric: one ontology could have the special class *"French singers"*, while the another might have only the generic *"singers"*.

Up to now, approaches for aligning ontologies have mainly considered only the matching of classes and relations, or the matching of entities. Our idea was that both of these strategies could complement each other: If we know that *"is a citizen of"* means the same thing as *"has the nationality of"*, that could help us to connect *"Alizée"* to *"A. Jacotey"*. If we have connected these two, then it is simpler to map the classes *"singer"* and *"French singer"* to each other. This in turn simplifies the linking of other singers, whereby further relationships can be inferred. Since these links are never certain, but rather always only incremental, we have developed a probabilistic model for this process, in which the probabilities of links depend on one another.

We have begun this project together with the research center INRIA Saclay, which is based near Paris. Our approach is fittingly named *"PARIS – Probabilistic Alignment of Relations, Instances, and Schema"*. We have implemented this approach and optimized the system so that it can also handle large ontologies. For example, our system can align our YAGO ontology to the DBpedia ontology in a few hours. Both contain multiple millions of entities and facts. Because PARIS takes skillful advantage of the interplay between relationships and entities, it reaches a precision of over 90% in the alignment of the ontologies.

Through this system we have been able to make a contribution to aligning the ontologies in the Semantic Web to one another. In this way the ontologies can complement each other, and computers searching for information can move from one ontology to another – just as hyperlinks allow us to move from one website to another. Thus, the Semantic Web is becoming more and more of a truly connected "Web". :::

CONTACT

Fabian M. Suchanek
RG. 2 **Ontologies**
Phone +49 681 9325-5014
Email suchanek@mpi-inf.mpg.de
Internet http://webdam.inria.fr/paris

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 are not able to relate this data to a given context or even to understand it. If it were possible to make computers understand data as knowledge, this knowledge could be helpful not only for Internet search, 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.

YAGO2, a recently created extension of the original knowledge base, pays particular attention to the organization of entities and facts in space and time – two dimensions that are highly useful when searching in a knowledge base. As an example, the great majority of the approximately 900,000 person-entities in YAGO2 are anchored in time by their birth and death date, 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 also ask the question of when the person in question actually was president.

Most of the approximately 7 million locations in YAGO2 have geographic coordinates which place them on Earth's surface. Thus, spatial proximity between two locations can be used as a search criterion. An example of a 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 YAGO2 one finds, among others, Albert Einstein, as both his lifetime (1879-1955) and his birthplace Ulm (70 kilometers from Stuttgart) are stored in YAGO2. :::

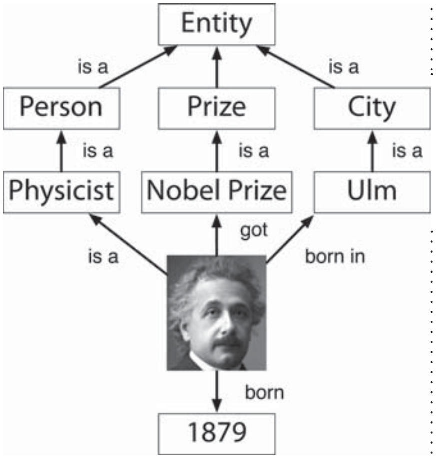

CONTACT

Johannes Hoffart
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5028
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
Internet http://www.mpi-inf.mpg.de/yago-naga/yago/

URDF – Efficient Reasoning in Uncertain RDF Knowledge Bases

Reasoning in uncertain knowledge bases with soft and hard rules

Despite the vast improvements in the field of information extraction in recent years, the very nature of the applied extraction techniques entails that the resulting knowledge bases may exhibit a high degree of uncertainty or even inconsistency. Often, these inconsistencies are not obvious even for a human and can only be uncovered through intricate and potentially expensive inference steps. In large RDF knowledge bases with millions of facts, an automated search for inconsistencies can be accomplished only through complicated, logic-based inference techniques, and an overly eager removal of inconsistencies can even lead to the loss of correct data. Automatic reasoning in large knowledge bases therefore requires a high robustness with respect to incomplete, imprecise, or even inconsistent data, which we will here summarize under the term uncertain data. While traditional query processing techniques for RDF (based on the SPARQL query language) are limited to deterministic data, the research in our latest project, URDF, particularly focuses on efficient, rule-based, and statistical inference techniques for uncertain RDF knowledge bases.

In URDF, we can express, for example, that many people live in the same place as their spouses. Being a rather "soft" form of an inference rule, this rule will also be violated by a number of instances (in this case, people) in the real world, which means that any knowledge derived from such a soft rule will also be uncertain. On the other hand, we can surely exclude that a person could have been born in different geographic locations. This represents a strict constraint, i.e., a "hard" rule, which may not be violated by any instance in our knowledge base. In a logic-based representation, these rules are often formulated as implications, so-called *Horn clauses*, where a conjunctive condition of facts implies a new fact:

$$\begin{aligned}&\mathit{marriedTo}(\mathit{person}_1, \mathit{person}_2) \wedge\\ &\mathit{livesIn}(\mathit{person}_2, \mathit{location}_1)\\ &\longrightarrow \mathit{livesIn}(\mathit{person}_1, \mathit{location}_1)\end{aligned}$$

The above rule is formulated in first-order predicate logic. More generally, our URDF framework allows for identifying common patterns between the instances in the knowledge base, and we can generalize (i.e., learn) first-order rules from these patterns in an inductive way. Conversely, we can then deductively reason about individual instances in our knowledge base by applying these first-order rules. In this case, we can, for example, apply the above first-order rule to the persons "Angela Merkel", "Joachim Sauer", and the place "Berlin-Mitte". This form of uncertain reasoning allows us to infer that "Joachim Sauer" might live in "Berlin-Mitte", given that we know that his wife, "Angela Merkel", also lives in "Berlin-Mitte":

$$\begin{aligned}&\mathit{marriedTo}(\mathit{Joachim\_Sauer}, \mathit{Angela\_Merkel})\\ &\wedge\ \mathit{livesIn}(\mathit{Angela\_Merkel}, \mathit{Berlin\text{-}Mitte})\\ &\longrightarrow \mathit{livesIn}(\mathit{Joachim\_Sauer}, \mathit{Berlin\text{-}Mitte})\end{aligned}$$

Moreover, a special type of Horn clause can be formulated by grouping purely negated facts into a disjunction, as in the following example:

$$\begin{aligned}&\neg\ \mathit{bornIn}(\mathit{Angela\_Merkel}, \mathit{Hamburg}) \vee\\ &\neg\ \mathit{bornIn}(\mathit{Angela\_Merkel}, \mathit{Bremen})\end{aligned}$$

With this Horn clause, we can express that "Angela Merkel" cannot have been born both in "Hamburg" and also in "Bremen". By using negations as in the above example, we now have a formal means to identify inconsistencies. If, in the extraction step, both facts $\mathit{bornIn}(\mathit{Angela\_Merkel}, \mathit{Hamburg})$ and $\mathit{bornIn}(\mathit{Angela\_Merkel}, \mathit{Bremen})$ have been inserted erroneously into the knowledge base, we can formulate that only one of the two facts may be true by using negations of the above form.

Efficient evaluation strategies

URDF supports both logic-based (propositional) and probabilistic evaluation strategies to answer user queries. In propositional reasoning, the user is guaranteed to obtain a consistent overview of a potentially inconsistent knowledge base in response to a query. This problem is reduced to the maximum satisfiability problem (Max-Sat), a classic problem in propositional logics. For URDF, we have developed a highly efficient method for solving a generalization of the Max-Sat problem, which is specifically tailored to the combination of soft and hard rules described above. In probabilistic reasoning, on the other hand, URDF does not only assign binary true/false values to facts, but also confidence weights, which correspond to a probabilistic interpretation of the derivation of the facts via the rules.

This enormous expressivity of rule-based and probabilistic inference techniques certainly poses major challenges to the development of new and efficient evaluation strategies for user queries. Both logic-based and probabilistic evaluation strategies underlie a much higher combinatorial complexity than traditional query evaluation strategies in relational databases. Therefore, exact evaluation algorithms cannot be applied to large volumes of data, which we obtain through information extraction from the World Wide Web. Our research concentrates on efficient approximation algorithms with good approximation guarantees, which are specifically tailored to these evaluation strategies. :::

CONTACT

Martin Theobald
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5015
Email martin.theobald@mpi-inf.mpg.de

EnBlogue – What is New and Interesting in Web 2.0?

The success of the Web 2.0 has drastically changed the way information is generated on the Web, with millions of users actively contributing content in form of blog entries or micro-news (e.g., Twitter). In order to maintain an overview of this continuously flowing stream of information, we need methods that extract the essence of current events and display them in a suitable way to users.

With enBlogue, the Max Planck Institute has developed an approach which continuously scans Web 2.0 streams for interesting events, which have received attention in blogs, Twitter, or other media. The name enBlogue is supposed to sound like *"in vogue"*, meaning it is popular.

EnBlogue processes web contents, annotated with the time of creation and semantic notes to detect surprising changes in the behavior of topics, which can be interpreted as interesting events. Not only individual semantic annotations (so-called *tags*) are assessed for their popularity, but all correlations between different tags are considered. Two tags are strongly correlated if the relative number of documents that report about two tags at the same time is high. Therefore, the dynamic increase in correlations is particularly interesting.

For example, normally the tags *"Iceland"* and *"flight operation"* are not very strongly correlated. This changed abruptly with the vulcanic eruption of Eyjafjallajokull in the spring of 2010.

en Blogue
emergent topic tracking
Pronunciation: ['en blog]
Function: adjective
Etymology: derived from "en vogue"
Date: early 2010
Synonyms: interesting, highly discussed

Web 2.0 streams, like blog postings, micro-blogging tweets, or RSS feeds from online communities, offer a wealth of latest news about real-world events and societal discussion.

WHAT DO WE DO?

We have devised an approach to continuously monitor emergent topics over these streams. This website gives a first glimpse of it.

SEE WHAT'S EN BLOGUE

In this example we continuously feed our topic detection engine with the New York Times archive. Please wait a couple of seconds to receive the first update.

Last updated on 16 Jul 1996 22:00:00 GMT

Most interesting **emergent topics**

BIOGRAPHICAL-INFORMATION ✻ PRESIDENTIAL-

The enBlogue website

The choice of tags to be considered is a priori unlimited but is primarily based on annotations generated by users, for example, when commenting on a blog entry or a news paper article. Methods can also be used for classifying information into topic areas as well as for discovering names of persons or places. The enBlogue system then also enables a mixture of peoples, places, and regular topics.

A model that takes into account the quantity of change as well as time-lines is used to extract interesting emergent topics and to present them in a clear and user-friendly manner.

EnBlogue is the starting point to current information on the web and it helps users being linked to actual contents. Personalization, i.e., re-ranking of topics depending on user preferences and interests, and additionally also considering the locations of events and users (if this information is available), is currently being researched. :::

OLYMPIC-GAMES-(1996) ✻ SUMMER-GAMES-(OLYMPICS)
COMPUTERS-AND-INFORMATION-SYSTEMS ✻ INTERNET-AND-WORLD-WIDE-WEB
TERRORISM ✻ UNITED-STATES-INTERNATIONAL-RELATIONS
BOMBS-AND-BOMB-PLOTS ✻ SAUDI-ARABIA-INTERNATIONAL-RELATIONS-US
ARSON ✻ BLACKS
MILITARY-BASES-AND-INSTALLATIONS ✻ UNITED-STATES-ARMAMENT-AND-DEFENSE
COMPANY-AND-ORGANIZATION-PROFILES ✻ MERGERS,-ACQUISITIONS-AND-DIVESTITURES
ASYLUM(POLITICAL) ✻ BOXING
ADVERTISING ✻ AWARDS,-DECORATIONS-AND-HONORS
CORRECTION-STORIES ✻ DEMONSTRATIONS-AND-RIOTS
DOLPHINS-AND-PORPOISES ✻ LAW-AND-LEGISLATION
BASEBALL ✻ LEFT-HANDEDNESS

Events identified by enBlogue on July 8, 1996 ***(with the assistance of the New York Times Archives)***

CONTACT

Foteini Alvanaki
DEPT. 5 Databases and Information Systems
Phone +49 681 302 70801
Email alvanaki@mmci.uni-saarland.de

Sebastian Michel
DEPT. 5 Databases and Information Systems
Phone +49 681 302 70803
Email smichel@mmci.uni-saarland.de
Internet http://qid3.mmci.uni-saarland.de/

MULTIMODAL INFORMATION & VISUALIZATION

Modern computer systems feature a variety of output modalities. These range from visual and auditory output devices to haptic output devices. Visual output modalities, ranging from standard monitors to complex virtual reality setups, are of particular importance. The development of algorithms to correctly display visual information, therefore, plays an important role at the Max Planck Institute for Informatics.

Nowadays, computer systems also feature a variety of novel sensors in addition to traditional input devices. This comprises cameras, haptic sensors, GPS receivers, and orientation sensors, to name a few. We develop algorithms to optimally process and combine information from these multimodal input and output devices.

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

By using such novel reconstruction techniques, we can also learn generalized statistical models of the static and dynamic real world. By means of such models, structural similarities across and within 3D models can be discovered, scene reconstruction can be made more robust, scene models can be modified more easily, and basic building principles of 3D shapes can be inferred, which enables the generation of new shape variants.

In order to show lifelike models of virtual worlds, we have also developed new simulation methods for light propagation in scenes, so-called global illumination.

The images created through image synthesis or image construction 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.

Rendering approaches also play an important role in the visualization of complex high-dimensional data. Such data frequently occurs, for example, as an output of weather simulations, flow simulations, or the analysis of certain effects in human society. We develop new approaches to visualize this complex data such that correlations and important features can be effectively deduced.

Nowadays, the optimal processing of optical input data is equally important. 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. Future computer systems will use sensors to perceive the state of their environment, for instance, in order to understand human activities. Correctly fusing information from multimodal input and output devices paves the way for more efficient and more intuitive human-computer interaction. :::

CONTRIBUTIONS

UNDERSTANDING IMAGES & VIDEOS
BIOINFORMATICS
GUARANTEES
INFORMATION SEARCH & DIGITAL KNOWLEDGE
MULTIMODAL INFORMATION & VISUALIZATION
OPTIMIZATION
SOFTWARE

Stereo and HDR Imaging: Display Quality Measurement and Enhancement

Computer models of human perception

We understand human visual perception as a final and mandatory component in a visual processing pipeline. An *"as complete and as correct as possible model"* of human visual perception is therefore crucial when aiming to improve the visual pipeline as a whole. While psychology provides a range of theoretic models of human perceptual performance, they are often not applicable for two reasons. First, too many simplifying assumptions are made. Second, the models are often only passive descriptions of findings and not algorithms. The challenge in our work is to deliver concrete models and algorithms. These should reliably and efficiently predict human perception of 2D or 3D content, and, at the same time, improve the content using these models.

We have developed an application to predict the perceived difference between two film sequences. Building on perceptual models, it uses existing physiological data to analyze the size, strength, and change of visual patterns in the input sequences and from this produces a third sequence that visualizes the perceived difference of the input sequences. It can be surprising how different the perceived difference is from what naïve numerical differences would suggest. Applications of our approach include video compression and steering image synthesis techniques.

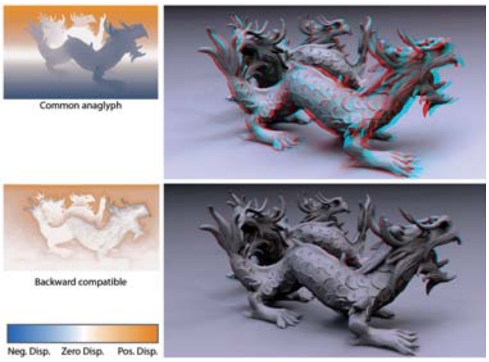

A conventional anaglyph stereo image *(top)* looks unpleasant, when observed without proper anaglyph glasses, due to color artifacts. Using our backward compatible solution *(bottom)*, we can reduce the artifacts and, using anaglyph glasses, still provide a sense of depth when seeing the image.

In a different track, we have developed an approach to compute the perceived difference between two stereo images. Here, our work goes beyond what is known to the field of visual perception: We first performed studies to acquire a model of human stereo perception that describes how a sinusoidal corrugation with a certain amplitude and frequency is perceived. From this model, we devised an approach that reports the perceived difference of two complex stereo stimuli. This finds applications in the compression and manipulation of stereo images. Further, the display of stereo content on various devices can benefit. Stereo images created for a movie theater but displayed on the screen of a small cell phone are unpleasant and uncomfortable to watch. Our model helps to "re-target" content between different devices.

Beyond the description of human perception, our models can be used to improve the user experience when depicting content. 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. Along similar lines, our approaches add depth impression to an image by using only minimal cues or by letting colors appear brighter than the monitor can actually display.

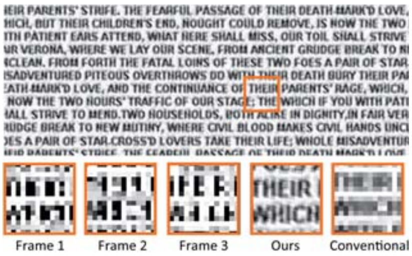

Apparent resolution enhancement for moving images: We optimize individual frames *(1-3, from bottom left)*, so that they appear more detailed *(ours, frame 4)* on a human retina than conventional solutions *(bottom right)*.

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. And, this provides us with an improved way of depicting content for all users on all devices. :::

CONTACT

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

Tobias Ritschel
DEPT. 4 Computer Graphics
Phone +49 681 9325-4041
Email ritschel@mpi-inf.mpg.de

Advanced Real-Time Rendering

Efficient simulation of indirect illumination

The creation of photorealistic images using computers is a basic technology required in many types of visual media. Although initially only used for creating special effects, nowadays, the usage of computers to render entire films is common. However, the required realism is only possible with high computational effort, and as a result, the synthesis of a single frame of motion picture can typically take several hours. At the same time, interactive real-timer rendering has become part of our everyday lives: In computer games, in geo-visualization on a cell phone, or in interactive kitchen planers on desktop computers. The required images are produced instantaneously from a userís 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.

In particular, our approaches were the first to allow interactive simulation of indirect illumination in dynamic scenes. Most previous interactive techniques assume direct illumination, where the light is emitted from point light sources and reflected once inside the scene before arriving at the observer. This simplification, however, is often not true in nature. Many visual effects such as the appearance of materials or the understanding of spatial configuration are known to be affected by indirect light. To achieve efficient indirect illumination, our approaches make assumptions that are different from classic numerical approaches (Finite-Element Radiosity, Raytracing), first, by fitting the demands of parallel graphics hardware and, second, by including knowledge of the human visual system.

Interactive simulation of a deforming piece of cloth.

In order to efficiently make use of existing graphics hardware, our approaches employ existing and optimized hardware functionality such as the drawing of points or the application of local image filters. Such operations are executed in ten thousands of parallel cores on modern hardware; by succeeding to keep those cores busy, quality or performance can be improved by orders of magnitude. In particular, the resulting approaches allow the simulation of dynamic scenes in real-time.

Our approach can compute results similar to a reference *(left, hours of computation time)* orders of magnitude faster *(right, seconds of computation time).*

In order to only compute what a human observer perceives, our work includes, and extends in independent perceptual studies, the knowledge about the human visual system; for example, it was not known until now how humans perceive indirect shadows. Our studies of this matter have shown that such shadows are perceived as largely smooth, which allows us to make strong simplifications and leads to improved performance.

Recently, the (long-standing) idea of performing expensive computations in a computational *"cloud"*, which then streams its results to client machines such as cell phones, has received renewed interest. In our work, we consider streaming approaches that fit best to modern graphics hardware and human perception. First, 3D information is encoded on the server side with human perception in mind, e.g., by faithfully encoding luminance and depth edges. Next, existing graphics hardware functions are used to efficiently extrapolate information on the client side. :::

CONTACT

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

Tobias Ritschel
DEPT. 4 Computer Graphics
Phone +49 681 9325-4041
Email ritschel@mpi-inf.mpg.de

PICASSO – Soundtrack Recommendation for Images

Most of us have encountered the following situation: we have taken multiple images during a vacation trip and, when back at home, want to present these images to our friends. What is missing is a background music accompaniment that underpins the emotions felt during the vacation. Our approach, named PICASSO, recommends the most appropriate soundtrack, fully automatically, from a user-provided music collection.

There are two worlds to be dealt with here: the world of music and the world of images. Considered separately, both of these worlds are relatively well understood. For each of these worlds, there exist methods that determine which objects are most similar to a given query object. For example, a user can find the 10 images most similar to a given vacation image. Features describing colors, brightness, and edges are used to calculate this similarity. However, different features are used in the world of music, for example, timbre and rhythm, which are crucial for precise search results.

Up to now, we have not been able to make a connection between the world of music and the world of images. With PICASSO, we have developed an approach that is able to learn and apply this connection.

The key principle is to look at very many examples of semantically related images and music pieces, aiming at covering a wide range of both media. The idea behind PICASSO is to analyze movies, which offer exactly the data needed for our task: (moving) images with matching soundtrack music, manually assigned by experienced movie directors. We had PICASSO analyze more than 50 movies, hence, obtain a large collection of pairs of the form (image, movie soundtrack).

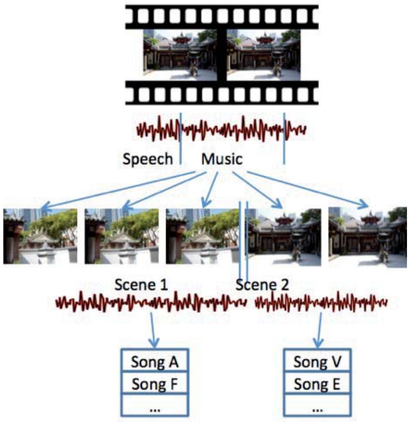

Figure 1: Process of Generating the PICASSO Database

In the recommendation process, PICASSO retrieves the most similar movie scenes for a given query image. In doing this, we move from the world of pictures to the world of music, using the already collected image-music pairs. Then, we can select the most similar songs out of a user-provided collection, with regard to the soundtracks in the movie scenes we found. For efficiency reasons, we have precomputed the last step of this process, i.e., we can jump directly from a movie scene to the user-provided songs. Figure 1 illustrates this process.

This recommendation procedure can also be applied to a given collection of the images (i.e., a slide show). First, multiple soundtracks are recommended for each individual image. Subsequently, these results are combined to recommend a soundtrack that matches all images. When doing so, we need to avoid choosing a soundtrack that is totally unfitting for some images, even if it might be a great fit for others. PICASSO is also applicable to videos, in which case the video is considered as a series of images, for which we recommend a soundtrack.

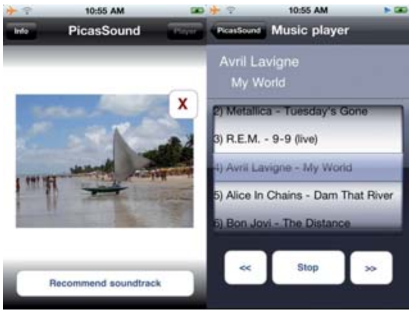

Figure 2: Screenshot of the iPhone App PicasSound

We have released an application, named PicasSound, which queries our PICASSO database via iPhone or Android-based smartphone. PicasSound is available free-of-charge. It enables users to capture an image or to use an already existing one [Figure 2], for which PICASSO recommends a list of songs out of the songs available on smartphones. :::

CONTACT

Sebastian Michel
DEPT. 5 Databases and Information Systems
Phone +49 681 302-70803
Email smichel@mmci.uni-saarland.de

Aleksandar Stupar
DEPT. 5 Databases and Information Systems
Phone +49 681 302-70807
Email astupar@mmci.uni-saarland.de
Internet http://picasso.mmci.uni-saarland.de/demo/

Computational Photography

Computational photography

Our group investigates innovative algorithms in order to correct common weak points in ordinary cameras and develop new applications, such as 3D reconstruction.

High-speed capture of human actors

One example of limited camera capabilities is its limited frame rate, or temporal resolution, in video mode. It is, therefore, difficult to record very fast motions that are, e.g., performed by human actors in action scenes. The resulting imagery is motion-blurred a well-known effect that is deliberately utilized to illustrate fast motions, e.g., in comic books. In modern visual effects, such motion blur can be undesirable, for example, if slow motion shots are to be shown. Moreover, modern techniques allow for the capture of virtual humans, which enables a virtual change of viewpoint after the scene has been shot. Motion blur can be a major handicap to these techniques. In cooperation with the Graphics, Vision and Video group, we have developed a motion capture system that can counteract motion blur and, therefore, facilitate the markerless motion capture of high-speed human motions. The developed system employs standard camera hardware in a surround imaging setup. As an example, we have demonstrated capturing at 400 frames per second with a set of 20 cameras that can physically record only at 20 frames per second. Future systems exploiting our ideas could enable the inexpensive surround capture of high-speed motions.

A view inside the measurement device. The object to be scanned is a bust of Max Planck *(partial view)*.

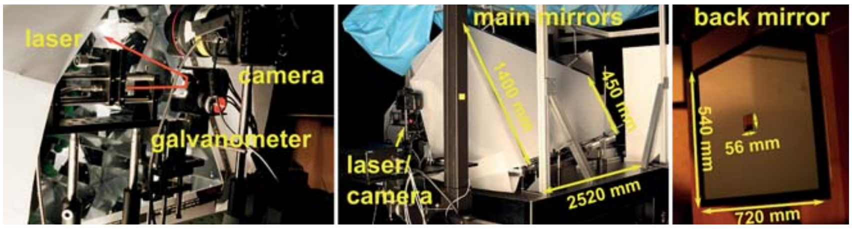

Device for the surround scanning of an object's geometry and reflectance. A large mirror system multiplies the number of available view points. A three-colored laser is used to illuminate the object and simultaneously scan its geometry.

Theoretical analysis of multiplex camera systems

Cameras are often used in multiplex mode, i.e., various forms of information, such as color or directional information, are digitized by suitable optical filters in different pixels of a single sensor. This process allows for less expensive production in comparison to the use of separate sensor chips for each information channel such as the colors red, green, and blue, for example. These filters are known as color filter arrays (CFAs). Recent analysis shows that, in addition to color information, the directional distribution of light, the so-called light field, can be recorded in different channels. The recording of light fields makes it possible, among other things, to change the focus of an image after the recording. Our group has established that many new types of cameras are special variations of one single overarching model. This knowledge can be used to create improved reconstruction algorithms, e.g., for color and directional image reproduction and for the analysis of camera performance. The application of the new model enables a comparison of different camera prototypes in terms of noise and resolution.

A surround imaging and projection system for reflectance measurements

The three-dimensional digitization of objects is another important application area for new camera technologies. Digitization, also called 3D scanning, is used more and more commonly in measurement, quality control, in the digital preservation of cultural artifacts, but also in the entertainment industry. A fundamental problem with 3D scanning systems is the inability of conventional cameras to have an all-around view; an object can only be viewed from one perspective at a time. In previous work we have shown how kaleidoscopic mirror systems can be employed to circumvent this restriction. For capturing a life-like appearance of real objects, however, it is not only necessary to image the object from all possible view points, but also under all possible illumination conditions. Established technology requires a tremendous hardware effort: Typically, 200 cameras and 200 light sources are employed in such devices. Extending our idea that a single sensor can image hundreds of views by repeated mirroring of the viewing rays, we have developed a prototype system that can generate 200 viewing and lighting directions with a single projector and a single camera. This capability, in turn, enables the acquisition of an object's reflectance in conjunction with its geometry, using a completely static setup. Our prototype does not yet match the accuracy of commercial scanners, but it provides an exciting outlook at the potential capabilities of future scanning devices. :::

CONTACT

Ivo Ihrke
DEPT. 4 Computer Graphics
Phone +49 681 302-70756
Email ihrke@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/departments/d4/areas/giana/

Calibrated Display Management

The amount of light available when the sun shines is about 10 orders of magnitude more than under starlight. Yet, humans are over time able to adapt to this enormous range of light. At any one time, the dynamic range of human vision is much smaller at about 4 orders of magnitude. Nonetheless, this is still 2 orders of magnitude (a factor of 100) more than current video, film and photo cameras can capture, and conventional monitors, televisions and projectors can display.

High dynamic range imaging is a field of research that studies hardware and software solutions to the dynamic range mismatch between current technology and human vision on the one hand, and current technology and the range of light available in scenes on the other hand. This has given rise to a set of capture techniques that allow images and video to be recorded with a dynamic range that meets or exceeds that of human vision.

An important aspect of high dynamic range imaging is that of the display of such images on conventional display devices. In other words, there is a question of how to display images with a dynamic range much greater than 2 orders of magnitude on a display which can only distinguish about 100 useable levels. Algorithms that solve this problem are commonly known as tone reproduction operators. They scale the image to be within the display range, and additionally aim to preserve the visual quality of the image so that the resulting image looks natural, as shown in Figure 1. Tone reproduction is certainly successful in that sense. Nonetheless, tone reproduction algorithms tend to negelect aspects of human vision, in particular human colour vision, which means that the results are not quite as natural as they could be.

Human colour vision involves similar adaptive processes as mentioned above, albeit affecting the perception of colour. For instance, if we view a scene under coloured illumination we would to a certain extent ignore the colour of the illuminant to the extent that surfaces in the scene that have a white reflectance would indeed appear white. In addition, the colourfulness of a scene depends strongly on the strength of illumination. If we had a dimmer switch connected to a light source and slowly increase the amount of light emitted, colours in the scene would increase in colourfulness. In general, the study of how coloured patches are perceived under specific illumination falls under colour appearance modeling.

Figure 1: A high dynamic range image tonemapped for display with a conventional operator.

Figure 2: An image processed using our combined tone reproduction and colour management algorithm.

This has significant implications for tone reproduction that are often ignored. In particular, if we reduce the dynamic range of an HDR image, we should also reduce the encoded colourfulness of the image to match what would happen in reality if we turned down our dimmer switch. Additionally, if the image was captured under a specific colour of illumination, and it is displayed on a monitor that is located in a room with different illumination, we would have to apply a correction to the image.

The goal of this project is to combine colour appearance effects into tone reproduction, effectively unifying parts of high dynamic range imaging with colour appearance modeling. Our algorithm carefully models all relevant aspects of human colour vision, and as a result is capable of producing images that are perceptually correct as measured against several existing psychophysical datasets and shown in Figure 2. :::

CONTACT

Erik Reinhard
DEPT. 4 Computer Graphics
Phone +49 681 9325-4020
Email reinhard@mpi-inf.mpg.de

Recognizing Human Activity

Recognizing human activity from wearable sensors

Sensing and understanding the context of the user plays an essential role in human-computer-interaction. It may enable natural communication, for instance, with robots that understand the user's goals and offer support at the right time. Our work focuses on a particular type of context, namely human activity recognition.

While impressive progress has been made in recognizing short and atomic activities (such as shaking hands and walking), research on more complex human activity lasting for minutes or hours (such as a morning routine or an assembly task) is far less mature. We thus focus on several important aspects of human activity recognition related to complex and long-term activities.

Collecting human activity data

Wearable sensors attached to the body have great potential for sensing what the user is doing, at any given time or place, from a first-person perspective. Given the advances in micro-technology, inexpensive sensors are al-ready today becoming widely available in watches, cell phones, and even clothing. Motion data can then be collected and analyzed for activity understanding, using machine learning techniques.

Identifying and combining relevant activity events

It is not usually the sequence of atomic activities that is interesting, but rather the higher level goal at which these activities are directed. There are several ways to infer the higher level goal from observing atomic activities. Since composite activities can contain large amounts of unrelated activity, using the *complete* observation can be suboptimal and therefore confuse the recognition. For many composite activities, it is sufficient to spot only a few underlying activity events to allow their recognition [Figure 1]. For example having lunch can be characterized by walking at a certain time of day, without even observing the actual eating activity. In a discriminative analysis, we observe that a surprisingly small fraction of relevant parts can be sufficient to recognize the higher level composite activity and allow efficient recognition algorithms.

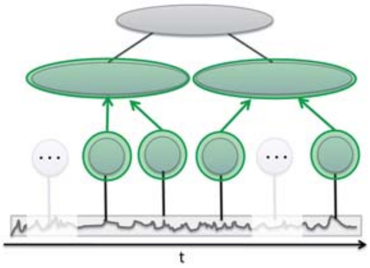

Figure 1: Relevant parts for composite activity

Hierarchical model for composite activities

Preserving the structure of a hierarchical activity offers several benefits. Considering a construction manual for a mirror in figure 2, one of several tasks is to fix the frame to the panel. This seemingly simple task consists of various steps, and it becomes obvious that composite activities add significant variation. Composite activities can be interrupted, the duration can vary strongly across different users, or the underlying activities can happen in different order. Using the same algorithms recognizing atomic activities can be suboptimal, as these require prohibitive amounts of training data. Therefore, we propose a hierarchical model that observes relevant activity events and combines them to recognize composite activities, similar to the way in which letters create words. Experiments show indeed superior performance compared to the standard approaches usually used for activity recognition.

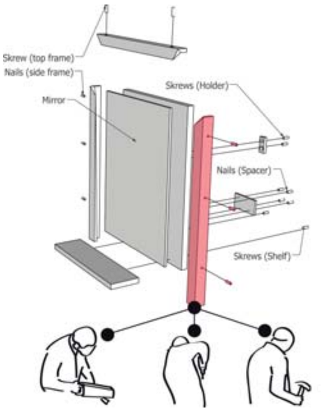

Figure 2: Composite activity for construction

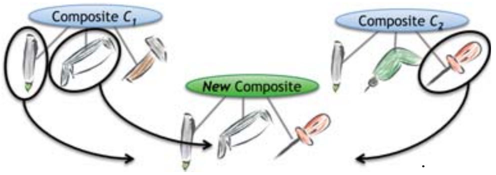

Transferring knowledge

Parts that are similar in different composite activities can be shared, much like vocabulary. Instead of re-learning composite activities from scratch, transferring shared parts reduces the training effort for new composite activities significantly. :::

CONTACT

Ulf Blanke
DEPT. 2 **Computer Vision and Multimodal Computing**
Phone +49 681 9325-2000
Email blanke@mpi-inf.mpg.de

Bernt Schiele
DEPT. 2 **Computer Vision and Multimodal Computing**
Phone +49 681 9325-2000
Email schiele@mpi-inf.mpg.de

Correspondences and Symmetry in 3D Scenes

Correspondence problems and shape understanding

An important challenge in modern computer science is the *understanding* of *structures* in data. Humans have remarkable capabilities in understanding structures in the world from their sensory input. Obviously, computers are not able to perform tasks at even a vaguely comparable level of cognitive skill. Nevertheless, machine processing of data has significant advantages: Computers can work themselves through enormous quantities of data at minimal costs. Not only do they unburden the human of tedious tasks but could potentially detect patterns in data that are only visible in large data bases.

Our research area deals with the *machine understanding* of *shapes*. Our goal is to enable a computer algorithm to understand to a certain extent how geometric shapes are structured. In the last few years, we have mostly worked on an important low-level problem in the domain of shape understanding: establishing *correspondences*. The goal is to determine whether two shapes are essentially the same, and if so, which points are corresponding, i.e., matching equivalent geometry. We have addressed two aspects of this problem. The first is the primary problem of actually establishing correspondences. The second problem is the higher-level problem of using correspondences among data to analyze the structure of the shapes.

Correspondence problems and symmetry

There are a large number of correspondence problems: The first aspect is *how* correspondences are established. The simplest option is to match geometry only under rigid motions, i.e., rotations, mirroring, and translations [Figure 1]. However, more complex models are possible; for example, shapes can undergo deformations [Figure 2], or we might seek to establish correspondences according to semantic categories, learned from a small set of training examples provided by the user [Figure 3] using machine learning algorithms.

Furthermore, not only is it possible to establish correspondences between pairs of separate shapes [Figure 2], but we can also look for reoccurring building blocks within one and the same object [Figure 1]. This process is called *partial symmetry detection*.

Applications

Correspondence information has a number of obvious applications; for example, we can fully automatically register multiple scans of a building (rigid matching) or of a scanned person (deformable matching), or automatically perform semantic tagging, for example, in large scale 3D scans of cities. However, interesting and non-obvious applications arise when we study the structure of the correspondences. By analyzing the structure of the partial symmetries of an object, we can find building blocks, out of which we can fully- or semi-automatically assemble variants of the shape [Figure 4]. An analysis can also reveal structural regularity such as arrangements of windows in regular grids [Figure 1]. This regularity can be used as an invariant for shape editing, thereby allowing a quick creation of plausible variations of the original shape. Such automated editing tools are useful, for example, for reducing the costs of content creation in computer games. :::

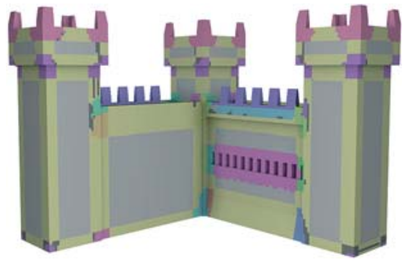
Figure 1: Decomposition of a 3D Model into elementary building blocks. The color indicates matching types.

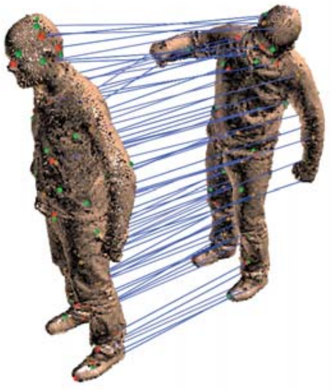
Figure 2: Correspondences between two 3D scans of a person in different poses *(computed automatically).*

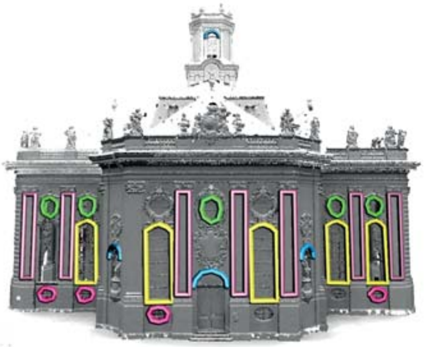
Figure 3: Correspondence between semantically related parts of a building. The matching model is learned from a few user examples. *Dataset: LKVK Saarland.*


Figure 4: Building blocks *(as in Figure 1)* can be reassembled automatically to create plausible shape variations from a single example.

CONTACT

Michael Wand
DEPT. 4 Computer Graphics
Phone +49 681 9325-4008
Email mwand@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/statgeo/

Flow Analysis

Flows of fluids or gases are of great importance in many technical processes. For example, designing a fuel efficient car involves reducing its aerodynamic drag. To this end, the flow around the car is simulated with a computer program. A subsequent flow analysis reveals vortex structures, which usually have a negative influence on the aerodynamic drag. Modern analysis methods aim at raising the physical understanding of a flow in order to allow for conclusions about necessary changes to the car's design.

Real-world flow visualization

The properties of a flow can be explained by considering the motion of particles. A particle moves with the flow on its so-called *trajectory*, which is a line consisting of all point locations visited by the particle over time. Unfortunately, a trajectory is hard to visualize in a real-world flow experiment since a simple photo shows only the current location of the particle. Therefore, other *characteristic curves* are used for visualization: a *streak line* consists of a large number of particles, which have been injected into the flow one after another from the same location. A common way to achieve this effect in a flow lab is to constantly release smoke from a nozzle. It moves with the flow and thereby forms the streak line. A *time line* is another means of visualizing a flow. It is created by an instantaneous release of smoke from a slit: the initially straight line of smoke is transported by the flow and rolls up in vortices. This reveals interesting flow patterns.

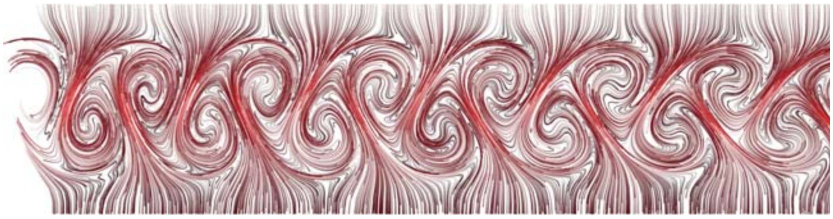

Figure 1: The time lines of the flow behind a cylinder reveal – in contrast to particle trajectories – the swirling motion in the flow.

Computer-based flow visualization

While streak and time lines are easier to visualize than particle trajectories in real-world experiments, it is exactly the other way around in computer-based visualizations. It is well-known that particle trajectories can be expressed mathematically using *ordinary differential equations*, which can be solved using standard methods. Furthermore, these equations allow conclusions about important properties of particle trajectories. For example, it is possible to determine their curvature without actually computing the trajectories themselves. Many important methods of computer-based flow analysis rely on the simple, yet powerful, representation of particle trajectories using ordinary differential equations.

For a long time, only comparatively complex algorithms existed for the computation of streak and time lines, and they did not allow any conclusions about the inherent properties of these lines. Together with our collaboration partner, Prof. Dr. Holger Theisel of Magdeburg University, we have succeeded in developing a novel mathematical approach that allows streak and time lines to be described using ordinary differential equations. This work was awarded the *Best Paper Award* of the annual IEEE Visualization conference.

New possibilities

This allows for a wide range of new methods for analyzing flows since the inherent properties of streak and time lines can now be expressed in a compact mathematical form. Figure 2 shows a so-called vortex core line in the center of swirling streak lines. For the first time, such vortex cores can be computed using our novel mathematical approach. They are an important basis for the identification of vortices in flows.

Furthermore, streak and time lines can often be computed significantly faster using the new approach. The 5000 time lines in Figure 1 were computed within 14 seconds using our new method. The classic algorithm requires 45 minutes for this. :::

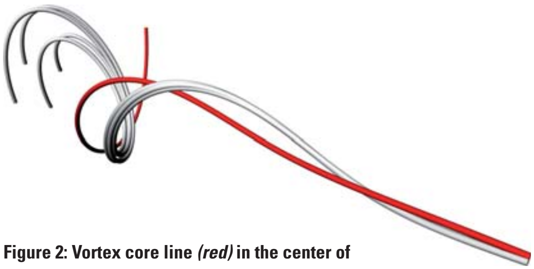

Figure 2: Vortex core line *(red)* in the center of swirling streak lines.

CONTACT

Tino Weinkauf
DEPT. 4 Computer Graphics
Phone +49 681 9325-4005
Email weinkauf@mpi-inf.mpg.de
Internet http://feature.mpi-inf.mpg.de/

Optimizing User Interfaces for Human Performance

Despite decades of research and enormous investments by the industry, our most commonly used user interfaces are age-old. For example, the QWERTY keyboard was invented in the 19th century, the menu in the 1950s, and the mouse and touchscreen in the 1960s. In software engineering projects, the user interface is among the most difficult things to get right. We believe that these problems are due to the fact that the space of alternative designs is too enormous to be explored by trial and error. A designer can entertain only the minutest fraction of all designs. Let us consider the case of designing a simple menu, one of the most commonly used user interface. The number of possible designs for a menu with 20 items is 20! = 2432902008176640000 – more than there are stars in the observable universe (10^{24}).

We believe that computational methods can be exploited for interface design. The automatization of recurring problems will allow a designer to focus on truly novel aspects. Instead of generating and trying out one or only a few instances at a time, the designer defines optimization goals, assumptions about the user and use, and sets constraints, and the computer explores the best designs. Our interface optimizer allows a designer to answer two critical questions:

1) How close to the optimal is a given design?
2) What is the optimal design?

The four general steps of this approach are:

1) representing the design space in terms of continous/discrete variables that are free/fixed,

2) operationalizing the desirable effects/outcomes of interaction as optimization goal(s),

3) constructing a predictive model that can take as an input any instance of the design space, and

4) identifying a suitable optimization method.

We explore this approach in a number of cases ranging from classic problems like text entry and menu selection to those novel multimodal interfaces such as computer vision based full-body control.

In our first case, we developed an optimized keyboard layout two-thumb text entry on tablet devices. To derive a predictive model, we devised a bimanual tapping task. The derived model accurately predicts the time to coordinate thumb movement for any given layout. The model was used in an optimizer that considered millions of alternative layouts. The rate with the predicted keyboard is the fastest reported for two-thumb text entry on a touchscreen device and improved the users' performance by 34 %.

The group presently works to optimize keyboards for different languages, challenging the hegomony of QWERTY as the universal keyboard. Over a century after its invention, the qwerty layout is ubiquitous: in smartphones, PCs, and large displays alike, despite the fact that the new form factors and input technologies may require radically different movement ranges of fingers. The optimization approach to design critically relies on the availability of valid predictive models that reliably estimate user performance for any given design from the design space. For simple sensorimotor tasks, such as reaching a finger to press a button, there are a few models that are robust enough to make such predictions, for example, for typing. The hard problem has been to acquire new models to extend the approach to other user interfaces. The group consults the behavioral sciences and biomechanics to derive such models.

For instance, in another case, we study full-body control, as in the now popular Kinect games. Here, human movement can be mapped to virtual movement in almost infinite ways, and each mapping is associated with different performance vs. fatigue properties. The light pen, for example, was touted as a serious alternative to the mouse but was never adopted because it could not be used by information workers. We have developed a novel method that allows searching for optimal mappings by combining biomechanics simulations with analyses of speed-accuracy-tradeoff in a single experiment. The method allows designers to identify optimal gestures for games and applications. :::

CONTACT

Antti Oulasvirta
DEPT. 4 Computer Graphics
Phone +49 681 302 71927
Email oantti@mpi-inf.mpg.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 difficult 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 efficient design of planning processes. Its importance will continue to grow in the future. :::

CONTRIBUTIONS

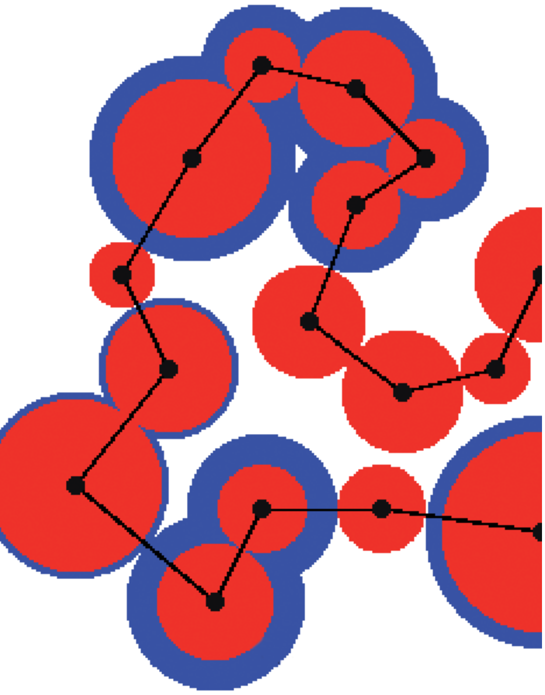

Computation of Equilibrium Prices

The balance between supply and demand takes place through pricing. For a particular good, the question of the equilibrium price is easy to resolve. Demand is a decreasing function of the price and supply is an increasing function of the price. Therefore, there is always an equilibrium price at which supply and demand balance out. But what happens in an economy with many goods, many suppliers, and many buyers?

The market models of Fisher and Walras

Already in the 19th century, Fisher (1890) and Walras (1875) introduced mathematical models for markets. Fisher's model is as follows: There exist buyers and goods. Buyers have a certain budget and preference with respect to the goods; goods are available in a certain quantity. The preference states how much utility a buyer obtains from a given quantity of a particular good. In the simplest case of a linear preference, the utility is a linear function of the quantity. We further assume that goods have prices: a unit of good X has a price $p(X)$. Then for each buyer and each good, there is an amount of utility per unit of money. Given a price of 2 Euro for a bottle of champagne and 1 Euro for a bottle of beer, for many buyers the utility per unit of money is higher for champagne than for beer. Fisher then postulates that each buyer buys only those goods that would, for them, maximize their utility per unit of money. Fisher asks whether there are always prices such that all of the buyers spend their entire budget and all of the goods are sold.

Walras' model is a bit different: He does not assume that buyers have a certain amount of money right from the start. Rather, buyers are also sellers. They own goods and receive money only through the sale of goods. The question is again the same: Are there prices such that all of the buyers spend all of the money they have received by selling their goods, and are all of the goods sold?

In Fisher's model, money has an intrinsic value; in Walras' model it serves only to compare goods. Fisher's model is a special case of Walras'.

Existence of equilibrium prices

It took until 1954 before the existence of equilibrium prices was mathematically stringently proven by Arrow and Debreu. Arrow and Debreu were awarded the Nobel Prize in Economics for that and for other achievements. Their proof is nevertheless purely a proof of existence and does not produce a method for calculating equilibrium prices.

Calculation of equilibrium prices

Can one efficiently calculate equilibrium prices? It took multiple decades before the answer was found for Walras' model. First in 2007, K. Jain found an algorithm with polynomial runtime, although it was still not satisfactory because of their use of the ellipsoid method. In 2012, Ran Duan and Kurt Mehlhorn found a relatively simple combinatorial algorithm. :::

CONTACT

Ran Duan
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1009
Email duanran@mpi-inf.mpg.de

Kurt Mehlhorn
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1000
Email mehlhorn@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 on a common base.

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 e.g. of a temporal or technical nature. 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 total weight. The rules themselves describe dependencies between the components and contexts: A powerful engine requires specific tire combinations based on the potential achievement of high speeds, whereas trailer hitching devices exclude weak engines. The interactions represent so-called if-then rules, which become active through decisions by the user specifying 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 is not the development of the concrete rules, but rather ensuring important global properties such as consistency. The above-described rule approach is particularly characterized by its flexibility. Rules can easily be adapted to the type of products being configured and the associated components and the respective area of use. The desired level of detail can be chosen freely: Thus, a product can be considered purely from a sales perspective, or down to the last, single 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, e.g. applications for software products that are composed of individual modules are also possible. Production lines, which consist of a sequence of individual steps, or the composition of business services from individual activities, can also be naturally described 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 representation of a rule-based system by a language of logic, is still the subject of 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 applications on 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 calculation of properties. :::

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

Approximation Algorithms for Profit-maximizing Pricing Problems

One of the most basic methods for dealing with NP-hard optimization problems is to design polynomial-time algorithms that find a solution which is provably not far from the optimum. Formally, an algorithm is said to be an x-approximation algorithm for a maximization (or minimization) problem if it runs in polynomial time in the size of the input instance and produces a solution whose objective value is within a factor of x of the optimum solution. Under this general framework, we have been working on problems from different areas, including graph problems, computational economics, computational geometry and mathematical programming. We will give an example below.

Item pricing

Profit (or revenue) maximization is a classic and fundamental economic goal. The design of computationally-efficient item-pricing schemes for various profit-maximization problems has received much attention recently, due in large part to the emergence of e-commerce applications on the Internet. With the availability of vastly emerging online markets, it is possible now for retailers to display their goods, such as books, flight tickets and hotel rooms, for purchase to interested consumers sitting at home. The added flexibility in such a market allows consumers to submit their budgets for every subset of items they are interested in to the retailer who considers this input from all the customers, and aims at coming up with an item-pricing strategy that maximizes the revenue. A typical example would be how to set the nightly rate for rooms in a hotel so as to maximize the revenue achievable from customers who submit their budget for the entire duration of intended stay. There is an obvious tradeoff on how to price the items: pricing cheaply implies a large volume of sales as many people will be able to afford their preferred subsets of items, but the revenue from each item sold is small; pricing high increases the revenue from each sold item, but reduces the volume. The optimization problem that arises from such tradeoff is quite complex and inherently non-linear. In the case when the supply of each item is limited, an added complication is sometimes mandated by the requirement to have a *fair allocation* of items to customers, in the sense that it might be necessary to raise the prices of certain items to ensure that there is enough supply, so that customers who can afford to purchase their preferred subsets at the announced prices are granted these subsets.

The complexity of these kinds of problems translates into hardness of approximation; that is, under plausible complexity assumptions, there exist in general no polynomial-time algorithms with good approximation ratios for the optimal profit-maximizing pricing. In joint work with Parinya Chalermsook, Danupon Nanongkai, and He Sun, we made an attempt at overcoming this inherent difficulty, in one basic variant of the problem, by focusing on the more realistic situation in which the consideration sets that the customers are interested in satisfy a certain geometric restriction. In particular, instead of looking at a full generality where each consumer C considers any set of items S_C, it is reasonable to assume that each consumer has some criterion in mind for each attribute of the item, and her consideration set consists of any item that passes all her criteria. This natural assumption has been a model of study in other fields such as marketing research, healthcare economics and urban planning. It is referred to as the attribute-based screening process. Assuming there are d-relevant attributes, consumers and items are represented by points in the d-dimensional Euclidean space, and the consideration set S_C of consumer C is the set of item-points that dominate the point corresponding to C in every attribute; see the figure below. Examples of attributes for a customer buying, say, a car, are horsepower, engine size, number of doors, etc. Each consumer C declares her budget B_C, and once the prices are declared, among all the items in the consideration set S_C, the cheapest one with price at most B_C is purchased.

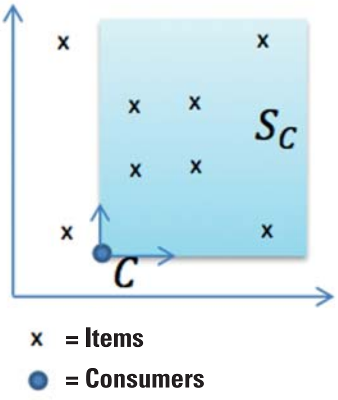

Consideration set

We showed that the *low dimensionality* of the consideration sets arising in this pricing problem does lead to improved approximation ratios, by presenting sublinear-approximation algorithms for the above variant of the problem. Our algorithm is obtained by combining algorithmic pricing and geometric techniques. These results suggest that considering geometric aspects might be a promising research direction in obtaining improved approximation algorithms for such pricing problems. :::

CONTACT

Khaled Elbassioni
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1007
Email elbassio@mpi-inf.mpg.de
Internet http://www.mpi-inf.mpg.de/departments/d1/areas/optimization.html

Query Complexity: From Theory of Evolutionary Algorithms to *Mastermind*

How much do I need to know about a problem to be able to solve it optimally? This question is not so much of an algorithmic nature, but rather pertains to the area of query complexity, which (together with communication complexity and coding theory) considers the fundamental problems of the handling of information.

However, almost every child knows a simple problem of query complexity, namely the game *Mastermind*. In this game, a player tries to guess the secret color code that the other players have arranged out of four pegs with six possible colors. With each guess, the first player learns the number of positions where the guess agrees with the secret code, but not which positions those are. The challenge in Mastermind is to find the secret code, using this limited information, in as few guesses as possible.

Computer science is interested in problems of query complexity for various reasons. One is that query complexity is clearly a fundamental measure of the difficulty of obtaining sufficient information about an unknown object. It is thus a central concept in theoretical computer science.

Mastermind: an everyday example from the area of query complexity

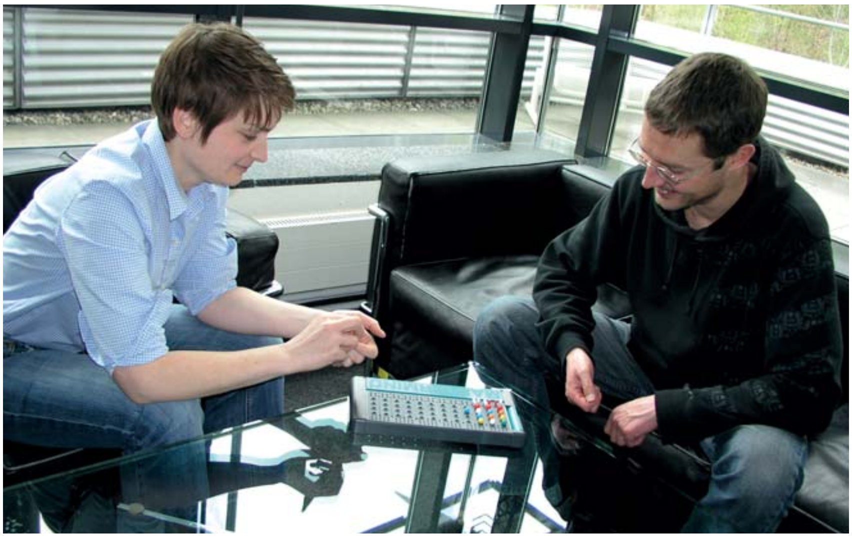

In addition to a foundationally-based interest, there are also practical grounds for the investigation of query complexities. Evolutionary algorithms and other randomized heuristics attempt to solve problems using generic methods. They have no access to an explicit problem statement; instead, they obtain information about the underlying problem only by generating and assessing possible solutions. Therefore, an evolutionary algorithm for a problem cannot be better than the corresponding query complexity.

Confidentiality issues can also lead to query complexity problems, the focus here naturally being that the unknown information cannot be discovered through the queries. An example of such a problem is the comparison of genetic data. Nobody wants their genetic information to be accessible to others. Therefore, in the comparison of two genome sequences, for example, only vague information about the similarity of the two sequences is disclosed. Just as in the Mastermind game, the question is again posed, how much of this vague information would suffice for the gene sequence to be unlocked.

Aside from fundamental results, our most recent efforts have also yielded an entertaining result, namely the current best strategy for the generalized Mastermind game with n colors and n positions. While up to now different strategies were known for guessing the code in approximately n log(n) questions, in our new strategy, approximately n log(log(n)) questions suffice. Even though this does not yet definitively solve the problem (the best lower bounds imply that one needs at least n questions), our result is nonetheless a great advance on a problem that has fascinated mathematicians and computer scientists for more than 30 years. :::

CONTACT

Benjamin Doerr

DEPT. 1 Algorithms and Complexity

Phone +49 681 9325-1004

Email doerr@mpi-inf.mpg.de

Optimality in Matching Problems

Matching is a fundamental concept in computer science. Besides having numerous applications, its study has led to many important advances in polyhedral combinatorics, graph theory, and algorithm design.

In a matching problem, we are given a graph with a set of vertices and edges. We want a matching, which is a subset of the given edges where none of these edges share an endpoint. Moreover, the matching should satisfy certain optimality criteria (often depending on the applications). We discuss a few common optimality criteria and our obtained results below.

Matching under preferences: stability and popularity

Suppose that a set of boys and girls are given, each with a preference over the members of the other sex. Ideally, we would like to have a "good" matching – a matching in which the boys and girls would be happy with the current situation. One possible optimality condition to capture the "goodness" of a matching is that of the *stability*: a matching is stable if no two persons strictly prefer each other to their assigned girlfriend/boyfriend. How to find a stable matching in this context is often called the stable marriage problem. This problem has many important applications in real world, the most famous one being the NRMP program in the United States, where the medical residents are assigned to the hospitals.

Another possible optimality condition is that of the *popularity*. A matching is more popular than another if more people prefer the former to the latter. A matching is popular if there is no matching more popular than it. Roughly speaking, popularity can be regarded as a sort of "global" stability: that is, we want a situation in which the *majority* of the people prefer the matching.

We remark that both popular matchings and stable matchings are very active research topics. Every year many research papers are published discussing various aspects of them. In particular, just recently, Alvin Roth and Lloyd Shapley have been awarded the Nobel prize in economics for their work on stable matching and market design.

We show that finding a popular matching of the maximum size can be done in polynomial time and a generalization of the stable matching problem, called *classified* stable matching, can be solved in polynomial time as well.

Maximum weight matching: approximate approach

In the maximum weight matching problem, we want the sum of weights of edges in the matching to be maximum. The most traditional and important motivation of maximum/minimum weight matching is the assignment problem, one of the fundamental combinatorial optimization problems. In this problem, there are a number of workers and the same number of tasks. Any worker can be assigned to perform some task, with different costs/benefits. It is required to perform all tasks by assigning exactly one worker to each task in such a way that the total cost of the assignment is minimized or the total benefit is maximized.

If we represent every worker and task by a vertex in the graph, and the weight of the edge between a worker and a task is the cost/benefit of assigning the worker to the task, it is easy to see that the assignment problem is equivalent to the minimum/maximum weight matching problem. Note that such a graph is bipartite, but there are also motivations to find maximum matchings in general graphs, which will greatly complicate the algorithm.

Although the maximum weight matching problem has been studied for decades, the computational complexity of finding an optimal matching remains quite open. We gave the first linear-time algorithm for computing the approximate maximum matching in general graphs, which can achieve an arbitrarily small approximation ratio. :::

CONTACT

Ran Duan
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1009
Email duanran@mpi-inf.mpg.de

Chien-Chung Huang
DEPT. 1 Algorithms and Complexity
Phone +49 681 9325-1016
Email villars@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.

At the Max Planck Institute for Informatics, this philosophy has been implemented with great success since the founding of the institute. All departments and groups work to implement the results of their basic research efforts into practical software systems and make them available to science and industry. There is a considerable number of software systems, developed in all departments and groups at the institute, that have found their way into the scientific community and that are used in research in many places around the world. Most of our software is distributed as free open source software. In some cases, however, start-up companies have been founded to further develop, distribute, and sell the software.

The clever implementation of mathematical models and algorithms into executable software is an important research subject. Algorithms that show nice mathematically analyzable run time and storage space worst case properties (so-called asymptotic complexity dimensions) are not necessarily efficient when implemented on modern computers, particularly on current distributed computer hardware. Properties of communication, processors, storage and magnetic disks as well as the characteristics of real data must be appropriately considered in the software's design if complex software systems are to be successfully applied in practice. When we succeed in building a software system for a newly developed method, the system in turn provides valuable information on relevant special cases, meaningful generalizations of the problem investigated, or the methods used.

The articles with a research focus on "Software" are a cross-section of our activities. The article *"RDF-3X – Fast Searches for Semantic Data"*, page 91, describes the pattern matching search for information in the WWW (World Wide Web). The software system goes far beyond the mere textual search used today. The article *"EpiExplorer and RnBeads: Integrative Analysis of Epigenomic Data"*, page 90, presents two software systems from bioinformatics that enable the study of DNS characteristics. The article *"Markerless Performance Capture"*, page 92, describes a system that captures movements, without there being explicit markings on the people, and then produces a movement model that can be utilized for further visualization or analysis. What all three software systems have in common is that they are able to filter and process relevant information for individual applications from huge amounts of in part "noisy" data. :::

CONTRIBUTIONS

UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION
& VISUALIZATION

OPTIMIZATION

SOFTWARE

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 thus are 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 searching information in a fast and intuitive fashion. *RnBeads* is a comprehensive software package supporting the analysis of experimentally obtained DNA methylation data.

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 that are of particular interest.

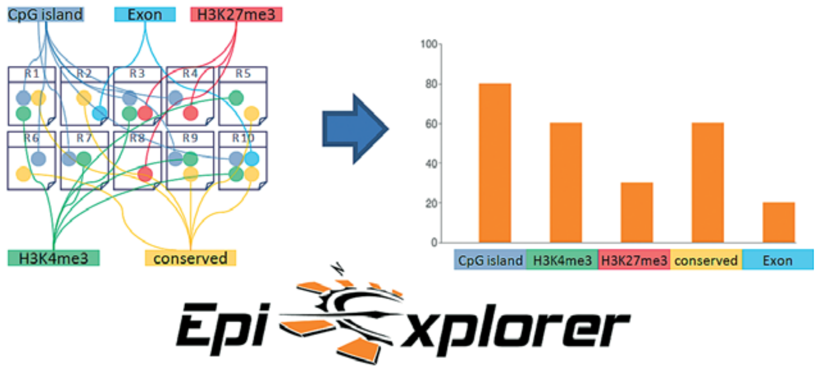

EpiExplorer facilitates fast, interactive searches in genomic and epigenomic maps and visualizes them in an intuitive fashion.

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 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 RnBeads' default workflow 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 exchanging 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

RDF-3X – Fast Searches for Semantic Data

Semantic data

RDF-3X is a database system to manage and search for semantic data, that is, data which 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 web 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 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, born in 1859. Further edges, such as *<id1, authorOf, id2>*, *<id1, bornIn, Edinburgh>* and *<Edinburgh, is in, Scotland>* demonstrate that the object *"id2"* is written by id1, that id1 is born in Edinburgh, and that Edinburgh is in Scotland.

Search in RDF graphs

Also, complex relationships can be formulated relatively simply with this kind of notation. 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:

?author <bornIn> ?city
?city <locatedIn> Scotland
?author <authorOf> ?book
?book <titel> ?titel.

The parts beginning with question marks are individual variables whose values must be determined by database systems. 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 triples that fulfill the patterns of the query. Because the data graphs are very large, there can be many millions of candidate triples, making efficient searches very difficult.

The database system we have developed, RDF-3X (RDF Triples Express), approaches these problems on several levels. First, the data itself is stored appropriately, so that individual triple patterns can be efficiently evaluated. We have compressed the data and indexed it using search trees, so that any triple pattern can be evaluated very quickly. But this is not enough, however; 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.

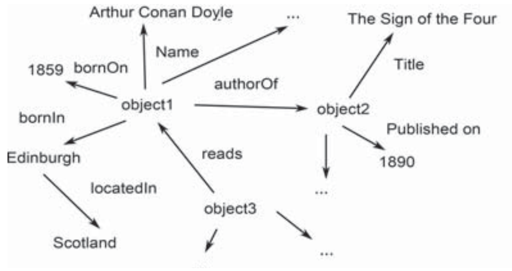

Example RDF graph

Efficiency and scalability

Searching in semantic data is already difficult for even relatively small graphs because the graphs frequently show no known structure (i.e., they follow no schema). Despite this, it was important for us to use RDF-3X not only for small data graphs, but also to make sure that it scales efficiently to very large data sets with billions of edges. One big challenge is updating such large data sets, i.e., inserting new data and, if necessary, removing old data, all without having to stop the whole database. Over time, we have therefore incorporated many techniques, such as *Sideways Information Passing* and *Triple Versioning* in RDF-3X in order to scale to such data sizes. In experimental comparison with other systems, RDF-3X usually performs extremely well. Even for complex queries with many triple patterns, RDF-3X is frequently much faster than other systems. :::

CONTACT

Thomas Neumann
DEPT. 5 Databases and Information Systems
Phone +49 681 9325-5007
Email neumann@mpi-inf.mpg.de

Markerless Performance Capture

The ability to capture detailed motion data from real people or from technical processes is of critical importance in many application areas of commerce and science. In medicine or biomechanics, motion capture is used to measure the success of a treatment or the efficiency of a sequence of movements. Motion capture is also applied in product design and in ergonomics research, for example to draw conclusions about the usability of technical devices. Today it is even used in onboard systems in vehicles, for example to recognize pedestrians and their movements in the vicinity of an automobile. Moreover, motion capture processes are used in the film and computer game industries. Virtual actors (avatars) are fundamental components of special effects and virtual 3D worlds.

With respect to commercially-available motion capture methods, marker-based optical methods are, as a general rule, the state of the art in industry. These processes are invasive, since certain markers need to be brought into the scene. The use of these systems is costly and time-consuming. In many cases these methods cannot be applied, for example when it is not possible to actively intervene in the scene.

With support from the Federal Ministry of Economics and Technology (BMWi) through its EXIST Transfer of Research (EXIST-Forschungstransfer) program, our team, consisting of Dr. Nils Hasler, Dr. Carsten Stoll, MsC Michal Richter, and our mentor, Prof. Dr. Christian Theobalt, develops software for the markerless reconstruction of human movements from normal video images. First, the skeletal movement is computed using a new, very robust and also fast approach. This represents the video images and the tracked person through implicit functions, and then attempts to reconcile these. In contrast to industrial methods, this approach requires neither markers nor explicit silhouettes. Thus, it is possible for us to capture the movements of people wearing everyday clothing. This is especially important for motion capture of athletes or virtual actors.

The markerless tracking method we have developed is qualitatively sophisticated and comparable to marker-based approaches. It is, however, clearly more flexible and works without as many restrictive constraints. Thus, the new method functions without a green screen, and without synchronization of the cameras, while at the same time being faster than all of the markerless tracking processes from the literature. In addition, the characteristics that (with a greatly increased runtime) gave earlier approaches an accuracy advantage, for example feature tracking and implicit stereo reconstruction, can be integrated in the new method. In this way, in addition to the skeletal movement, a 3D surface reconstruction of the person can be captured (so-called *performance capture*). Here a complex optimization consisting of a combination of multi-view stereo and shape-from-shading is carried out. This enables the temporal reconstruction of the movements of loose clothing or folds. :::

CONTACT

Nils Hasler
DEPT. 4 Computer Graphics
Phone +49 681 9325-4055
Email hasler@mpi-inf.mpg.de

Carsten Stoll
DEPT. 4 Computer Graphics
Phone +49 681 9325-4004
Email stoll@mpi-inf.mpg.de

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 scholarship which covers fees, living expenses, and health insurance coverage for themselves as well as for their spouses and children who accompany them. 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
Internet http://www.imprs-cs.de

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 in 2012 included the successful evaluation of the International Max Planck Research School for Computer Science (IMPRS-CS) and the visit of the Indian Consul General, who chose our institute for his inaugural visit to Germany and was impressed by our sustainable Indo-German Max Planck Center for Computer Science (IMPECS). The VERIDIS agreement, which was signed in August, is the first INRIA project with joint French-German leadership. The year ended with a symposium held on the occasion of the 60th birthday of Prof. Dr. Lengauer, Director of the Department of Bioinformatics.

CONTRIBUTIONS

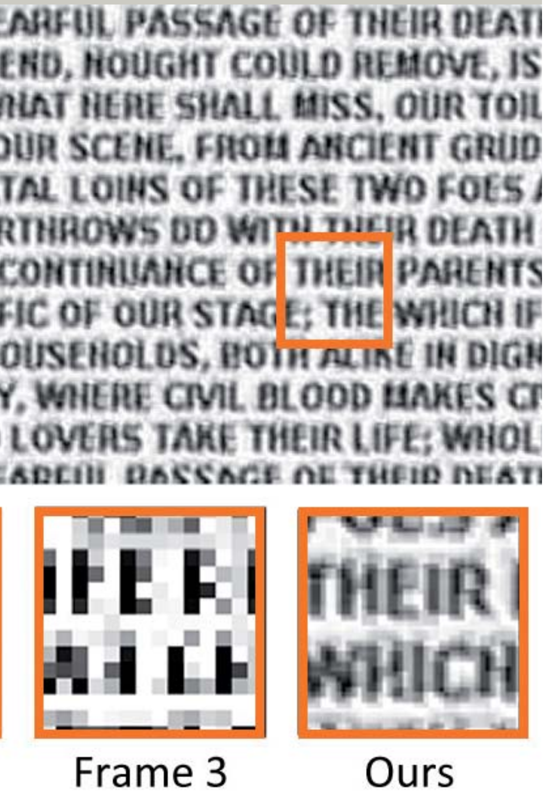

Prizes, Honors, Awards

Professor Hans-Peter Seidel, Head of DEPT. 4 for Computer Graphics has received the Eurographics Distinguished Career Award 2012 from the European Association for Computer Graphics for his pioneering scientific contribution to the field of computer graphics as well as his prominent role in the training and support of young university lecturers in Germany and Europe.

The Eurographics Distinguished Career Award, which is presented every two years, is the most prestigious computer graphics research award in Europe. Seidel was the first German scientist to receive this award. :::

Professor Hans-Peter Seidel was appointed as member of The Academy of Europe (*Academia Europaea*). The aim of Academia Europaea is the promotion and interaction of science and research throughout Europe. :::

Professor Kurt Mehlhorn, the founding director of the MPI-INF and Head of DEPT. 1 for Algorithms and Complexity, has been appointed as Corresponding Member by the Bavarian Academy of Sciences and Humanities. The academy, which is one of the oldest in Germany, does not specialize in particular fields or subjects but is committed to broad interdisciplinary communication.With this appointment, Mehlhorn, who moved to Saarland in 1974, gained another connection to his birthplace. :::

WISSENSCHAFTSRAT

Professor Gerhard Weikum has been reappointed as member of the German Council of Science and Humanities for three years. According to its own understanding, the German Council of Science and Humanities 'provides advice to the German federal government and the state (Länder) governments on the structure and development of higher education and research.'
24 scientists and 8 public representatives are appointed by the German Federal President and serve on two commissions that constitute the German Council of Science and Humanities. The formation of the council dates back to 1957. Ever since, it has provided impetuses for the development of higher education in the Federal Republic of Germany. In addition to its advisory role, the council is a decision-making body on topical issues such as the accreditation of private universities and participation in the Excellence Initiative. :::

The Academy of Sciences and Literature Mainz has included **Professor Gerhard Weikum** as full member of the Class of Mathematics and Natural Sciences, following a decision of its meeting on November 9, 2012. The academy, which was founded in 1949 by former members of the Prussian Academy of Sciences, in its position as responsible body for research projects, is committed to the promotion of science and preservation of culture. :::

Christian Theobalt, a professor at Saarland University and head of a research group in the Department for Computer Graphics at the MPI-INF was awarded with the German Pattern Recognition Prize 2012, a major award for early-stage researchers. The German Working Group for Pattern Recognition honored Theobalt for his excellent and ground-breaking work in the field of 3D reconstruction based on camera pictures and video sequences as well as for trend-setting papers on the theme of motion capturing without markers. The award is bestowed annually and includes prize money of 5,000 euros. :::

Personal Data

Google Research Award

Hannah Bast, former PhD student under the supervision of Kurt Mehlhorn and postdoc in DEPT. 1 has been honored with this year's Google Research Award. Up to now, only 42 candidates have been awarded this prize worldwide; eight of whom were European. With almost 1 million U.S. dollar, it is one of the highest awards presented by Google. During her time at the MPI-INF, Hannah Bast spent one and a half years as a researcher at Google, where she developed a procedure for accelerating the route planning of public transport networks. Since 2009, she has been a professor of computer science at the University of Freiburg, where she was been offered the chair of Algorithms and Data Structures. :::

Lise Meitner Award

This year, the Lise Meitner Award, which MPI-INF presents to early-stage female researchers in the field of computer science, has been given to **Anna Adamaszek** of the University of Warwick. The award includes a tax-free two-year research stay at the MPI-INF and research funds. She is continuing her studies on algorithms under Kurt Mehlhorn (DEPT. 1). :::

Dr. Eduard Martin Award

Gerard de Melo, DEPT. 5, received the Dr. Eduard Martin Award 2012 for his dissertation on *"Graph-based Methods for Large-Scale Multilingual Knowledge Integration"*. The award is presented by the Saarland University and bears the name of the long-standing director of the incorporated association "Friends of Saarland University" ("Vereinigung der Freunde der Universität des Saarlandes"). Each year, the award is given to the best early-stage researchers at Saarland University. :::

Otto Hahn Medal

Matthias Hullin, DEPT. 4, was awarded the Otto Hahn Medal 2011 for the *"Development of New Procedures for the Measurement and Reproduction of the 3D Geometry and Reflectance of Real Objects"*. With this award, the Max Planck Society honours young scientists and researchers for outstanding scientific achievementswithin the scope of their doctoral thesis. :::

Eurographics PhD Award

Tunç O. Aydin, a PhD candidate under the supervision of Prof. Seidel in DEPT. 4, until 2011, was honored with the Eurographics PhD Award 2012 for his disseration on *"Human Visual System Models in Computer Graphics"*. :::

Bertram Somieski

Professorships offered

Over the past year, professorships were offered to a number of early-stage researchers.

Panagiotou, Konstantinos | DEPT. 1
University of Munich

Baumbach, Jan | DEPT. 3
Southern Denmark University

Doerr, Benjamin | DEPT. 1
École Polytechnique, Paris

Theobald, Martin | DEPT. 5
University of Antwerp

Hose, Katja | DEPT. 5
University of Aalborg

Elbassioni, Khaled | DEPT. 1
Masdar Institute of Science
and Technology in Abu Dhabi

Friedrich, Tobias | DEPT. 1
University of Jena

Megow, Nicole | DEPT. 1
TU Berlin

Hermelin, Danny | DEPT. 1
Ben Gurion University

Spöhel, Reto | DEPT. 1
Bern University of Applied Sciences

van Zuylen, Anke | DEPT. 1
College of William and Mary
Williamsburg (VA)

Lasowski , Ruxandra | DEPT. 4
Furtwangen University

Thormählen, Thorsten | DEPT. 4
University of Marburg

Müller, Meinard | DEPT. 4
University of Erlangen

Sauerwald, Thomas | DEPT. 1
University of Helsinki

Elbassuoni, Shady | DEPT. 5
American University of Beirut

Wahlström, Magnus | DEPT. 1
Royal Holloway University
of London

Personal Data

New Senior Researchers in 2012

Five young scientists were appointed as Senior Researchers due to their excellent work:

Olga Kalinina | DEPT. 3
Structural Bioinformatics
of Protein Interactions

Nicole Megow | DEPT. 1
Combinatorial Optimization &
Graph Algorithms

Antti Oulasvirta | DEPT. 4
Human-Computer Interaction

Thomas Sauerwald | DEPT. 1
Randomization and Load Balancing

Fabian Suchanek | RG. 2
Otto Hahn Group, *Ontologies*

Fellowships

Fellowships are used in many ways to directly support researchers or scientists and have become an integral part of contemporary academic education. Depending on the sponsor and purpose, fellowships provide minimum support or meet the entire financial needs of the beneficiaries, and sometimes even of a work group including PhD candidates and postdocs. The duration of fellowships and the requirements to be met by beneficiaries vary widely. Some fellowships are held in such high esteem that they significantly advance the recipients careers. At the MPI-INF, six young scientists are currently supported by such excellent fellowships.

In 2012, **Ivo Ihrke** of DEPT. 4 and the MMCI, as well as **Nicole Megow** and **Konstantinos Panagiotou**, both of DEPT. 1, have been admitted to the Emmy Noether Program for the support of outstanding early-stage researchers. Each of them will be supported by the program for five years. They are given the possibility to employ their own PhD candidates, for example. The beneficiaries had applied for the program with the following themes: *“Plenoptic Image Acquisition and Projection: Theoretical Developments and Applications”* by Ivo Ihrke; *“Models, Algorithms and Complexity for Scheduling Under Uncertainty: On the Tradeoffs between Performance and Adaptivity”* by Nicole Megow. The title of Konstantinos Panagiotou’s project is *“Real-world Networks and Random Graph Models”*. Ivo Ihrke is supported with a total of 1.2 million euros.

Artur Jez, postdoc of DEPT. 1, successfully applied at the Alexander von Humboldt Foundation. He was admitted to a Humboldt Research Study, which he will pursue for of 22 months from February 1, 2013.

Since 2011, **Carola Winzen**, DEPT. 1, has been supported with a *Google European Doctoral Fellowship* by the Google Foundation.

Over the year 2012, the following three PhD candidates at MPI-INF were on fellowships sponsored by the Siemens AG: **Claudio Magni** of DEPT. 1, **Ching Hoo Tang** of RG. 1 and **Mohamed Yahya** of DEPT. 5.

Three further PhD candidates have earned doctoral fellowships from Microsoft: **Martin Suda** of RG. 1, **Elena Tretyak** of DEPT. 2 and **Xiaokun Wu** of DEPT.1. :::

Bertram Somieski

Non-academic Activities

Patents

At times, the basic research at the Max Planck Institute for Informatics 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 a 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 2012. :::

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	Ritschel / Eisemann / Myszkowski / Seidel
Method and Apparatus for Encoding High Dynamic Range Video	Patent extension	Mantiuk / Krawczyk / Myszkoski / Seidel
A Luminance-Contrast-Aware Disparity Model and Applications	Invention reported	Didyk / Eisemann / Ritschel / Myszkowski / Seidel / Matusik
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

Spin-offs

The work on the analysis and editing of motion pictures done by the group around Prof. Christian Theobalt, DEPT. 4, has been widely covered by the media. Moreover, it has attracted the attention of the movie industry, which is interested in using this technology of capturing and quickly editing video sequences for commercial purposes. Establishing a business enterprise outside the Max Planck Institute appeared to be the best way to transfer this technology to the market place.

Dr. Nils Hasler, Dr. Carsten Stoll and Prof. Christian Theobalt have successfully applied for initial public funding by *“EXIST-Forschungstransfer”*, a support program of the German Federal Ministry of Economics and Technology, for the purpose of founding a company. For details, refer to page 92: *“Markerless Performance Capture”*. :::

Bertram Somieski

Cooperations

Intel Visual Computing Institute

Three 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. Moreover, Dr. Michael Wand, DEPT. 4, with his *"Symmetry-based Shape Processing and Analysis"* project and Dr. Thorsten Thormählen, also DEPT. 4, with his project *"3D Reconstruction Using Active Illumination"* are all heads of project groups at the IVCI. [www.intel-vci.uni-saarland.de] :::

Bertram Somieski

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. In summer 2011, the Max Planck Institute for Informatics was visited by Prof. Manindra Agarwal of IIT Kanpur, Humboldt Research Award laureate, over a period of several months. In fall 2011, Prof. Helmut Seidl of TU Munich spent two months at the IISc Bangalore.

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

Science Tunnel

The Max Planck Science Tunnel addresses the megatrends of the 21st century. Which areas of research are currently being explored? In what fields does science address forthcoming challenges and how? The name *"Science Tunnel"* is derived from the exhibition concept that, unlike classic exhibitions, confronts the visitor with an environment with vertical limits. The Science Tunnel is a traveling exhibition, which is open in its current form to the public in Paderborn. The Heinz Nixdorf Museums-Forum, which is dedicated to the German computer pioneer, Heinz Nixdorf, has been chosen as the first location to hold the exhibition. On October 18th, the third version of the Science Tunnel was launched and in March 2013, the next stop will be Moscow.

Eight modules reflect on scientific (universe, matter), social (energy, society, life sciences, health) and cognitive (brain, complexity) themes.

Two MPI-INF themes have been added to the *"complexity"* module of Science Tunnel 3.0.Under the motto *"Learning from nature"*, systems intelligent and capable of learning will be exhibited and analyzed. An exhibit titled *"Language made from stone"* presents the work of Art Tevs and Michael Wand, both of DEPT. 4, on automatic pattern recognition inside scanned 3D objects. The algorithms developed by computer graphics experts are capable of recording and analyzing huge 3D scans of entire cities and creating new models of cities on the basis of the recognized patterns, due to the learning capacity of pattern recognition systems. For this purpose, the algorithms analyze characteristics (analogous to words of a language) and typical structures (analogous to the grammar of a language) of the patterns. They learn, so to speak, the *"language of the patterns"*.

"Virtual Mirror Cabinet – The Real time Reshaping of Humans" by Michal Richter and his colleagues, also of DEPT. 4, uses video sequences to show new methods of real-time video editing: Different characteristics of the human body can be changed using a simple slider. The visitors can see directly on the computer screen how their legs are becoming longer, their muscles stronger, or their hips leaner, while moving in front of the camera.

On the occasion of the inauguration of the Science Tunnel, Prof. Thomas Lengauer, head of DEPT. 3, was asked to deliver one of the introductory lectures at the site of the new exhibit. On November 29, Thomas Lengauer lectured on *"Hunting the virus – How using the computer to fight AIDS"*. [www.sciencetunnel3.de]. :::

Bertram Somieski

Software Campus

Software Campus, an initiative committed to supporting future computer science graduates in becoming leaders in the German IT industry, was launched in 2011 and fully established in 2012.

The topic *"Robust 3D Reconstruction from Visual Data"*, which was suggested by Christian Kurz, has gained acceptance in the evaluation commission and has won Bosch in Stuttgart Schwieberdingen as an appropriate industrial partner. The aim of the project is to develop a market-suitable interactive application toperform 3D reconstruction of everyday's objects based on images, videos, and other visual data. To achieve this, Christian Kurz will be funded with up to 100,000 euros over the next 12 to 15 months. Maintaining close contact with the scientific partners and mentors at the industrial partner is equallz as important important for the success of the project. In order to prepare them for a potential leading position, the participants of Software Campus are provided with intense training and access to people in senior management positions. :::

Bertram Somieski

New Max Planck Institute for Software Systems Building Inaugurated

“Computer Science Plaza” completed

On October 15, 2012 the new institute building of the Max Planck Institute for Software Systems was ceremonially inaugurated. Now all researchers of the MPI-SWS based in Saarbrücken are under the same roof and the years of being dispersed across the region are over.

Rupak Majumdar, Managing Director of the MPI-SWS, thanked the state government of the Saarland, the Max Planck Society and, last but not least, the construction crews for the new six-story research building. Not only the scientists of the Max Planck Institute for Software Systems but also parts of the staff of the Joint Administration of MPI-INF and MPI-SWS have offices in the new building.

Managing Director of the MPI-SWS, Rupak Majumdar

Photo: Johannes-Maria Schlorke

The new, large lecture hall is equipped with 178 seats and state-of-the-art media technology. It is used by both institutes. For example, the *“Joint MPI-INF/MPI-SWS Lecture Series”* is held there on the first Wednesday of each month. At these events, senior researchers from both institutes hold lectures to a broad audience of experts on the latest state of research. The MPI-INF research group *“Automation of Logic”* has moved to the top floor of the new building.

The building in Saarbrücken, had cost around 17 million euros. It provides office space for 185 employees. The structure of the MPI-SWS building is characterized by a high level of functionality and is in line with the concept of concept to foster communication and cooperation. Open and interconnected structures allow for targeted and casual encounters, while diagonals and tapers eliminate strict right angles. For the time being, this will be the last building constructed on the *“Computer Science Plaza”* on the campus of Saarland University and is designed to complement the architecture and contribute to the completeness of the existing buildings. The construction of the building was made possible by the Saarland’s gorvernment. :::

Bertram Somieski

Consul General of the Republic of India

Inaugural visit to Saarbrücken

On January 19, 2012, Shri Taranjit Singh Sandhu, consul general of the Republic of India and based in Frankfurt/Main, paid his inaugural visit to Saarbrücken. Besides issues regarding the economic collaboration between India and the Saarland, he was especially interested in the living conditions of the Indian citizens of the Saarland. After a short visit to the German Research Center for Artificial Intelligence, he was welcomed at both Max Planck Institutes. There he was informed about scientific collaboration and working conditions. All Indian scientists and students of the two institutes were invited to a joint discussion. :::

Bertram Somieski

New Scientific Advisory Board

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, which is why new members were also appointed for the Max Planck Institute for Informatics for the term beginning in 2013.

We thank the following for their service in the last term:

Prof. Pankaj Kumar Agarwal of Duke University, **Prof. Douglas L. Brutlag** of Stanford University, **Prof. Joseph M. Hellerstein** of Berkeley University, **Prof. Friedhelm Meyer auf der Heide** of the Heinz Nixdorf Institute and the Computer Science Department of Paderborn University, **Prof. Eugene Myers** of the Howard Hughes Medical Institute, **Prof. Frank Pfenning** (Chairman of the former Scientific Advisory Board) of Carnegie Mellon University, and **Prof. Claude Puech** of INRIA Paris and Paris University.


The Scientific Advisory Board's visit in 2011. The next visit will be on May 14th and 15th, 2013.

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

Prof. Travor Darell of Berkeley University, **Prof. Nir Friedman** of the Hebrew University of Jerusalem, **Prof. Pascal Fua** of the EPFL of Lausanne, **Prof. Jürgen Giesl** of RWTH of Aachen, **Prof. Alon Halevy** of Google Research of Mountain View, California, **Prof. Yves Moreau** of Leuven University, **Prof. Nicole Schweikardt** of Frankfurt University, **Prof. François Sillion** of INRIA of Grenoble Rhône-Alpes and **Prof. Emo Welzl** of ETH Zurich. They will complement the existing members **Prof. Yannis E. Ioannidis** of Athens University, **Prof. Éva Tardos** of Cornell University and **Prof. Demetri Terzopoulos** of Los Angeles University. :::

Christoph Weidenbach

Journalists' Journey to the World of Computer Science

European journalists visit campus of Saarland University

From March 28 to 30, more than a dozen European journalists informed themselves on the latest research and development in the field of computer science and adjoining research fields in Saarland. The journey *"How Computer Science and Materials Science are improving the world"* took them to all institutes relevant to the field of computer on the campus of Saarland University science. Keeping to a tight schedule, representatives of all participating bodies presented their organizations and research themes. Thomas Lengauer, who presented the topic *"Bioinformatics Support of HIV Therapy – Spearheading Personalized Medicine"*, and Bernt Schiele, who talked about the topic *"Cars learn to see and predict"* familiarized the journalists with some important research interests at the MPI-INF. Thorsten Thormählen and Christian Theobalt complemented the selection of topics of our institute with the talks *"Hollywood-style 3D Animations for Everyone"* and *"Optical Performance Capture"*.

It was for the second time since 2008 that representatives from the field of computer science in Saarland had invited journalists to the journey through computer science. Journalists from all over Europe may participate if they are members of the EUSJA umbrella organization (European Union of Science Journalists' Association).This year's participants represented the whole spectrum of the media: daily newspapers, radio stations and magazines of science and technology as well as online media.

After visiting the Computer Science section of the campus, the journalists were invited to stay overnight in Dagstuhl Castle – Leibniz Center for Informatics where they were informed on-site about the activities of the Material Engineering Center Saarland. :::

Bertram Somieski

INRIA Research Project VERIDIS

Framework agreement on Franco-German project signed

On August 28, 2012 Prof. Karl Tombre and Prof. Stephan Merz, on behalf of the French side, and Prof. Kurt Mehlhorn and Prof. Christoph Weidenbach, on behalf of the German side, signed the framework agreement on the VERIDIS research project of INRIA (Modeling and Verification of Distributed Algorithms and Systems).

Binational and multinational projects are not uncommon in the scientific community. However, VERIDIS is the first project in the history of the French state-owned computer science research organization INRIA (Institut National de Recherche en Informatique et en Automatique) where one of the partner groups involved is not from France. The group around Prof. Stephan Merz from the INRIA location in Nancy and the group around Prof. Christoph Weidenbach of the Max Planck Institute for Informatics intend to do research on the verification of distributed algorithms and systems in a joint effort. While the French group contributes great expertise in the modeling and interactive verification of such systems, the German group is one of the world's leaders in automated verification procedures. The intention behind combining the two approaches is to make it possible in the coming years to verify the properties of distributed systems more quickly and with considerably less manual effort. :::

Christoph Weidenbach

Guests from ENS Cachan

40 French students attend lectures and workshops

The ENS (École Normale Supèrieure) schools in France constitute the elite universities for France's young scientists. As part of their training, the students spend one semester at another university or do their doctorate at an external institution such as the Max Planck Institute for Informatics.

For more than 20 years, we have maintained contact with ENS in Cachan near Paris, and students of ENS have regularly been guests at the MPI-INF and in the Computer Science Department in Saarbrücken. On January 26, 2012 we again received 40 students of ENS Cachan under the guidance of Dr. Sylvain Schmitz. Prof. Kurt Mehlhorn provided an overview of the computer science hub in Saarbrücken; afterwards, the students participated in six different workshops and took the opportunity for discussion with the lecturers. :::

Christoph Weidenbach

Workshops:

Prof. Dr. Holger Hermanns:
„Dependable Systems and Software“

Dr. Khaled Elbassioni:
„Combinatorial Optimization“

Prof. Dr. Markus Bläser:
„Computational Complexity“

Prof. Dr. Bernd Finkbeiner:
„Reactive Systems“

Dr. Sebastian Altmeyer:
„Programming Languages and Compiler Construction“

Prof. Dr. Christoph Weidenbach:
„Automation of Logic“

IOI Training 2012

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, refer to http:// www.bwinf.de/]. We are especially committed to preparing successful students in a national competition so as to qualify them for international competitions. This is why we, together with Dr. Pohl, manager of the BMWINF, also organized the final selection and training for the 2012 International Olympiad in Informatics (IOI).

From January 12th to 15th , Pawel Gawrychowski, former trainer of the Polish IOI-team and current researcher at the MPI-INF, and Patrick Klitzke, former successful Olympiad participant and now student of computer science in Saarbrücken, trained the six selected students. 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 Schloss Dagstuhl 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 required to have background knowledge that extends even to topics only taught in graduate courses, for example, Tarjan's linear algorithm for finding the biconnected components of a graph. The training incorporates not only finding solutions to such problems, but also subsequent discussion of the participants' solutions and, thus, the acquisition of new knowledge and the learning of new methods.

The concept of the workshop aims not only to help the participants be successful at the Olympiad, but 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 *"Better Algorithms by Chance?"* by Prof. Markus Bläser, *"What is Quantifier Elimination?"* by Dr. Thomas Sturm, and *"How do Modern SAT Solvers Work?"* by Prof. Christoph Weidenbach.

The Olympiad was held from September 23rd to 30th, 2012 on Lake Garda in Sirmione-Montichiari, Italy. Tobias Lenz earned a silver medal, and Julian Labeit received a bronze medal. :::

Christoph Weidenbach

MPC-VCC Review

Scientific Advisory Board impressed by excellent scientific work

The Max Planck Center for Visual Computing and Communication, see page 26, was established by the German Federal Ministry of Education and Research (BMBF) to address the particular career-development needs of young scientists and is located at both the Max Planck Institute for Informatics, Saarbrücken, and the School of Engineering of Stanford University. During the support period, the junior group leaders have to develop their research themes, supervise PhD candidates and become independent scientists. Each year, the Max Planck Center meets with its Scientific Advisory Board.

For the annual report to the Scientific Advisory Board, all university instructors working as mentors for junior group leaders as well as all junior group leaders with their PhD candidates and scientific advisors meet and discuss the state of the support program and its research results.

This year, Saarbrücken was the venue for the annual review meeting. For almost a whole day, the junior group leaders held 13 excellent speeches, in which they explained the state and outlook of their work. This was followed by a presentation by the PhD candidates of approximately 30.

The members of the Scientific Advisory Board were impressed by the variety of themes and by how the members of the Max Planck Center had cooperated with each other. They especially praised the outstanding excellence of the achieved scientific work. :::

Bertram Somieski

Saarland Company Race

The “Planck Quanten” place in top quarter

In 2012, enthusiastic runners from the Max Planck Institute again participated in the Saarland Company Race, which attracts thousands of runners each year. Within the scope of this big event, which took place on June 28, the runners were to run a 5 km-long circuit through the town center of Dillingen. Seven teams of four runners each from the Max Planck Institute of Informatics had registered for the challenge: The “Fire-Breathing Dragons” and the “Planck Quanten” teams with the numbers 1 to 6 ran in the mixed-team category, respectively. The “Planck Quanten 6” team took 241st place out of 1129 in that category. :::

MPI runners before leaving for the company race

Bertram Somieski

CeBIT 2012

Great media interest in the field of computer science in the Saarland

Each year, the CeBIT attracts all kinds of visitors, from ordinary men and women to scientific experts. It is the event where the latest products and developments in the field of information technology are shown, and national as well as international companies present themselves and their products. Students, pupils, and housewives are as common on the CeBIT as are peers, business partners, and competitors.

As in the previous years, the Max Planck Institute for Informatics shared the *"Saarland"* stand with other presenters. Together with staff of the German Research Center for Artificial Intelligence, the Computer Science Department of Saarland University, and the Center for IT-Security, Privacy and Accountability as well as other institutions from the Saarland University campus location, MPI researchers presented the results of their scientific work to interested members of the public. In Exhibition Hall 26, which had the motto *"The Laboratory"*, the Saarland stand showed a mix of exhibits from diverse fields of basic research. As is tradition, representativers from the state government of the Saarland, led by its Prime Minister Annegret Kramp-Karrenbauer, toured the CeBIT itsself and observed the excellent work on display. Two exhibits shown at the CeBIT 2012 presented research results from the MPI-INF. *"Markerless Motion Capture"* by Dr. Nils Hasler and Prof. Christian Theobalt, DEPT. 4, showed how human movements can be reconstructed on the basis of video data.

The work group developed a procedure to reconstruct movements in real time without using any markers or special backgrounds (greenscreens). *"3D animation à la dream factory Hollywood for everyone"* was the title of the exhibit shown by the group of scientists around Dr. Thorsten Thormählen.The presented software uses freely available or low-priced databases that contain computer-readable representations of three-dimensional objects.They can be changed and extended easily, and transferred into each other in a step-by-step approach. Thus, they allow for defining new 3D characters. Ambitious home users canbe successful in the animation of characters à la Toy Story with this software. As in the previous years, the exhibits of the Saarland stand attracted great media interest. :::

Bertram Somieski

Meeting of the Curatorship Board

Members of the Curatorship Board informed themselves about state of institute

For the Max Planck institutes, supervisory boards offer a means to present the scientists work to representatives of the general public. The Curatorship Board of the Max Planck Institute for Informatics is composed of high-level representatives from executive, industrial, research, academic, and media organizations. The members of the Board are responsible for conveying the intention of the institute and fostering the trust of the public in autonomously organized and funded research. Due to their many links within society, they are able to give advice on how to strengthen the location and on the strategic orientation of the organization.

On February 29, 2012 the Curatorship Board of the MPI-INF had its annual meeting. Led by Chair Prof. Wolffried Stucky, the Curatorship Board informed Peter Jacobi, Saarland Minister for Economic Affairs and Science, and Prof. Margret Wintermantel, Chair of the German Rector's Conference (HRK), among about the state of the institute. In his report, Prof. Bernt Schiele, Managing Director of the MPI-INF, presented the scientific achievements and personnel developments of the last year. He emphasized that the institute receives continuous feedback on its acceptance within the society and the interest of the public due to the voluntary activity of the members of the Board. Prof. Kurt Mehlhorn reported on the efforts and progress made by the work group which deals with the establishment of spin-offs.

The young professionals, Rainer Gemulla, Michael Sagraloff, Fabian Suchanek und Tino Weinkauf displayed their latest results as examples of cutting-edge research. Theywere awarded certificates by Prof. Stucky appointing them as senior researchers. At the subsequent poster session, young researchers of the MPI-INF showed the latest research results in condensed form. :::

Bertram Somieski

Evaluation of IMPRS

Great result ensures further funding

In July 2012, the International Max Planck Research School for Computer Science (IMPRS-CS), also see the IMPRS-CS article on page 94, was evaluated by a team of international experts. Within the scope of a large poster session, the PhD candidates of the IMPRS-CS answered the reviewers' questions. These ultimately testified to the excellence of the IMPRS-CS: The quality of the PhD candidates and their research work as well as the individual counseling were praised. The reviewers especially underlined the exemplary cooperation between the Max Planck Institute for Informatics and Saarland University. Due to very favorable judgment, the Max Planck Society decided to continue funding the IMPRS-CS for six more years. :::

Gerhard Weikum

60th Birthday of Thomas Lengauer

MPI-INF holds colloquium in honor of Thomas Lengauer

On the occasion of Prof. Thomas Lengauer's 60th birthday, the Max Planck Institute for Informatics and the Center for Bioinformatics Saar invited to a celebratory colloquium, which was held in the Günther Hotz Lecture Hall on November 16, 2012. Professors Kurt Mehlhorn of the MPI-INF, Burkhard Rost of TU Munich, and Volkhard Helms of Saarland University/Bioinformatics were full present to praise the scientific work of Thomas Lengauer. They pointed out that Prof. Thomas Lengauer is capable of getting things moving and taking things forward more quickly than most others. They said that he had started a new subject in science every ten years and had made contributions at a global level.

Thomas Lengauer with former Bioinformatics PhD students

Six guest speakers outlined the field of bioinformatics in their speeches. Three of them, Niko Beerenwinkel, Christoph Bock and Matthias Rarey, who now themselves hold chairs, were formerly PhD Students under Thomas Lengauer's guidance. Three further experts from the field of computer science, Eugene Myers, Burkhard Rost, and Martin Vingron, completed the group of speakers. In condensed form, the audience was lectured on a wide selection of topics from Bioinformatics on an international level. :::

Bertram Somieski

ADFOCS 2012

Summer School for early-stage and leading scientists

Each year in summer, young algorithm experts meet in Saarbrücken for the *Max Planck Advanced Course On the Foundations of Computer Science* (ADFOCS) at the MPI-INF. From August 13 to 17, more than 50 Master's and PhD students gathered to improve their knowledge of theoretical computer science. Guided by He Sun of DEPT. 1, the young scientists were to listen carefully and practice diligently over a period of five days. In 2012, the organizers were again able to attract internationally renowned lecturers:
Luca Trevisan of Stanford University, Berthold Vöcking of RWTH Aachen, and Avi Wigderson of Princeton who delivered outstanding lectures followed and seminars.

Prof. Kurt Mehlhorn launched this annual educational event in 2000 to bring young scientists into close contact with leading scientists outside the normal teaching environment. It is the lecturers' task to include current research topics in the lectures and seminars. The ADFOCS has now taken place 13 times, which reflects the sustained interest in this summer school. [www.mpi-inf.mpg.de/conference/adfocs/index.html] :::

Bertram Somieski

VTSA 2012

World-class researchers teach at school for PhD candidates

The universities of Luxembourg and Liège, the INRIA research institutes in Nancy, and the Max Planck Institute for Informatics take turns organizing the VTSA, the Summer School for *"Verification Technology Systems and Applications"*,at their locations every year. The idea behind the school for PhD candidates [http://www.mpi-inf.mpg.de/vtsa12/] is to give students the possibility to extend their knowledge and exchange views with leading researchers in the area of verification.

The Max Planck Institute for Informatics hosted the Summer School in 2012, from September 3rd to 7th, which makes it more of a "fall school" than a "summer school". The organizers were once again successful in convincing world-class researchers to hold courses for the summer school: **Prof. Carsten Schürmann** of the University of Copenhagen gave a lecture on *"Logical Frameworks"*, **Prof. Armin Biere** of Linz University talked about *"Understanding Modern SAT Solvers"* and **Prof. David Monniaux** of the Verimag research institute, which is located at the Grenoble University, explained the basic ideas of *"Abstract Interpretation"*. **Prof. Jürgen Giesel** of RWTH Aachen expounded the current approaches to *"Automated Termination Analysis"*, and **Prof. Ahmed Bouajjani** of Université Paris Diderot gave an introduction to the *"Verification of Concurrent Programs under Weak Memory Models"*.

The lectures were accompanied by talks by PhD candidates on the topics of their doctoral theses, leaving time afterwards for them to discuss their work with the experienced researchers. :::

Christoph Weidenbach

Computer-Science Research Days 2012

Challenging and supporting gifted pupils of computer science

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

In 2012, the winners of the German Federal Competition in Computer Science (*"Bundeswettbewerb Informatik"*) as well as two students from each school in Rhineland-Palatinate who had the best Abitur (final exams taken at the end of secondary school in Germany) results and computer science as one of their main subjects were invited to the event, which took place from July 19th to 20th. The students were given the opportunity to become acquainted with, and actively participate in, current research themes within the Max Planck Institute, the Computer Science Department, and the DFKI.

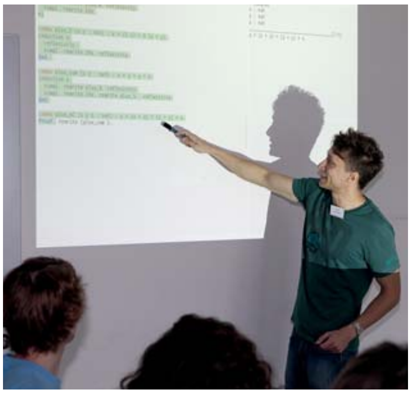

The participants spent the first evening getting to know each other while bowling. They then attended lectures on specific themes, such as *"Time analysis"*, but also current lectures from the curriculum of computer science, like *"Programming II"*. At different workshops, the participants were given the opportunity to build and subsequently test, e.g., 3D scanners, or learn how to convert a simple webcam into a game controller. The group work covered all areas of computer science, from the fundamentals of computer science, to programming and networks, to graphics and computer vision; there was something for everyone.

The motto of the Computer Science Research Days is: *"Challenging and supporting"*. With this event, the Max Planck Institute for Informatics, the Computer Science Department of Saarland University, and the DFKI aim to fill young people with enthusiasm for the subject of computer science as well as discover and support the development of new talent. :::

Jennifer Müller

The Institute in Figures

Budget without third-party funds 2009 to 2012

Staff Current as of 1/1/2013

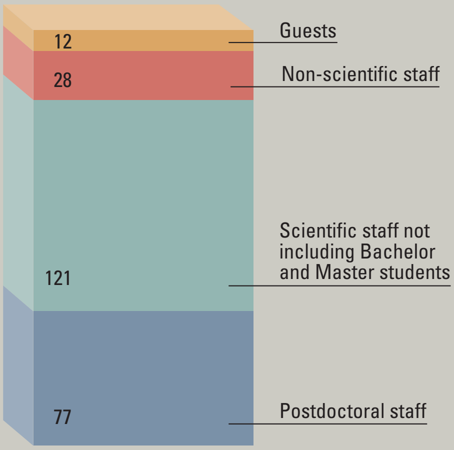

Ratio of scientific to non-scientific staff members Current as of 1/1/2013

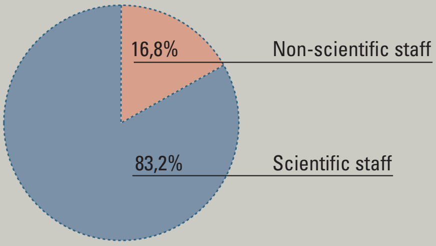

Third-party funding from 2009 to 2012 Numbers and distribution

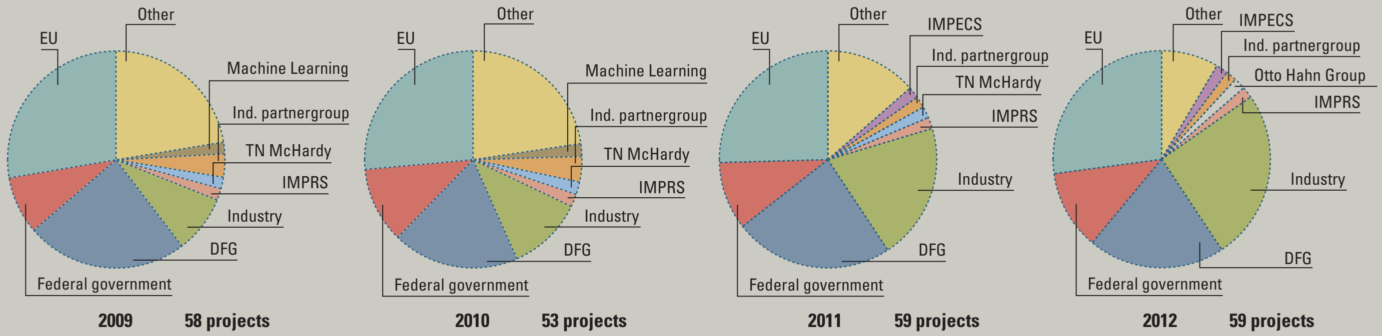

Third-party funding from 2009 to 2012

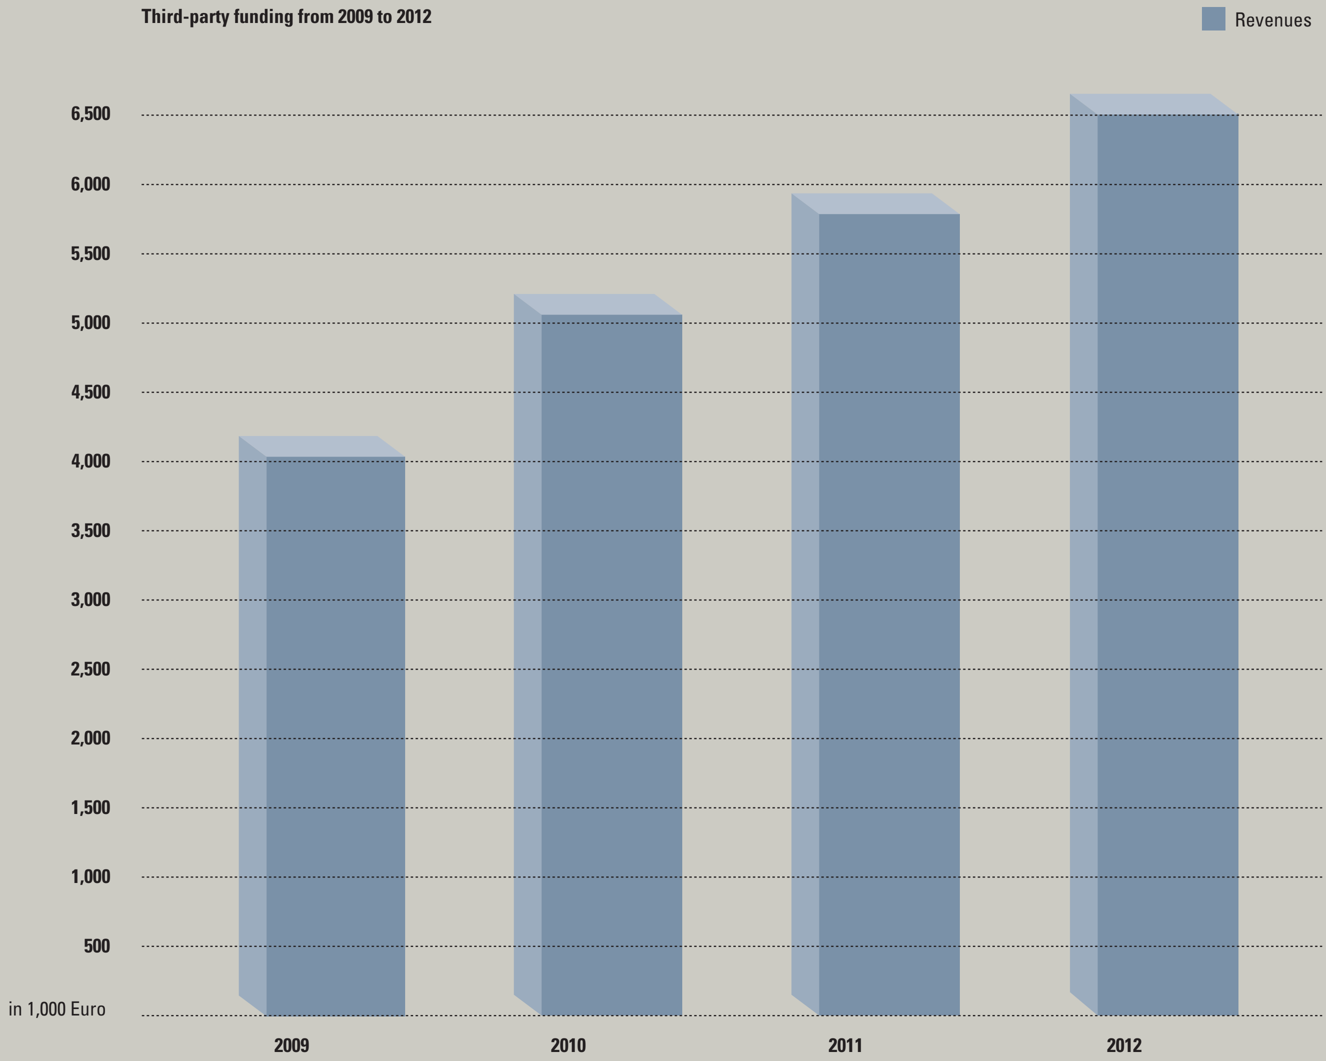

CONTACT

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

Information Services and Technology

Unhindered global cooperation and communication in a motivating environment form the basis for an institute whose aim is to produce first-class research. The flexibility, quality and reliability of the equipment as well as its usability make a decisive contribution to this research.

This aim can be transferred to our IT infrastructure: We operate a versatile system that can adapt to rapidly evolving requirements and provide the user with consistency and reliability. It does not neglect security, despite the openness necessary 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 deploy hardware from various manufacturers, as well as all major operating systems: MacOS for clients, Solaris for file servers and Linux and Windows for client and server systems. User data and the most important software packages are available independently of platform and OS. This homogeneity eases the use of these heterogeneous systems.

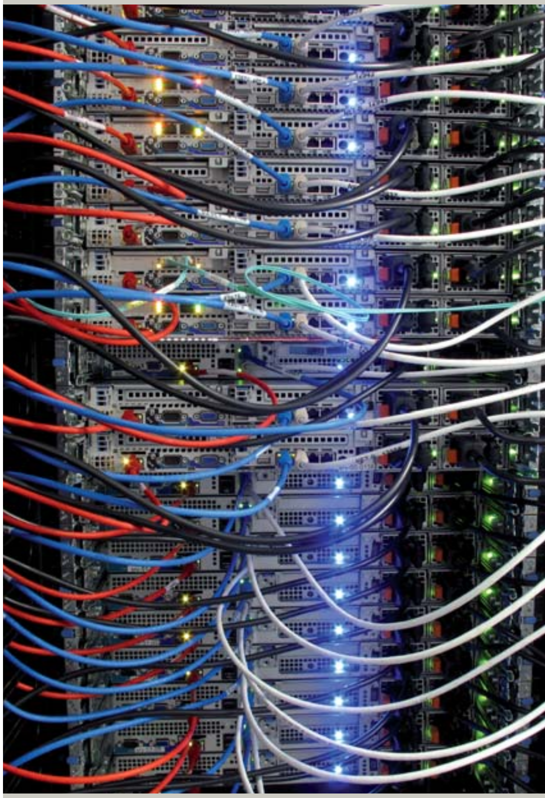

Automation as a guarantee for reliability

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 system that is designed to take individual requirements 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. Short-lived special installations are done manually.

Security and protection

Fixed mechanisms which protect against sabotage and espionage are not possible in open systems. They limit usability to too great a degree. 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 combated through keeping our local software installations and virus scanners continuously 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 big outages.

A monitoring system based on Nagios (open source) provides information on critical states of the server systems and the network, and on malfunctions of complex processes via e-mail and text messaging.

Cooperation and communication

Our network is divided into various areas according to organizational and security-relevant points. Externally accessible, but anonymous services (DNS (Internet address book), WWW, FTP (data transfer), SMTP (e-mail), etc.) 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. 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 e-mail, important databases and other services. Cooperation in software development is supported by protected access to a revision control database (software repository: Subversion).

This logical structure is spread over a 10 GB backbone and 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 2.6 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 1 GB or 10 GB 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 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, with 128 systems with 8 cores and 48 GB main memory apiece, 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 may 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.

File service and data security

The roughly 150 TB data of the institute is made available via NFS and CIFS (SMB). About 20 file servers distribute the volumes of about 60 RAID systems with more than 3000 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 if a data block is correct and if not, which mirrored half should be used to correct it. This implies as well that time consuming file system checks prior to the access of the data may be skipped. 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 apiece share a view of the disk status and can substitute for one another 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 secures the data directly using Tivoli Storage Manager (TSM) on a tape robot (StorageTek), 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 615 TB of gross data can be reduced to about 36 TB in the different backup runs. In order to combine the benefits of disk and tape technologies, we will combine this system with the tape robot in the future.

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. In the next construction phase, the system will be extended to a second robot that can then operate at a remote site to protect the data against disaster situations. This is at least currently achieved offline through the storage of copies in a special fire-safe data safe or at a second site.

Special systems

For special research tasks, especially in the area of computer graphics, diverse special systems are required. Available systems include a video editing system, several 3D scanners, multi-video 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, sometimes even in international collaborations such as the 'CADE ATP System Competition' [http://www.cs.miami.edu/~tptp/CASC/]. In order to meet the various resulting requirements, IST offers specifically tailored degrees of support that extends 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 with self-administered systems, however, IST only provides consultation.

Responsibilities

IT procurement, installation, administration, operation, applications support and continuation of the described systems and techniques are the duty of IST (Information Services and Technology), a subdivision of the Joint Administration of the Institutes for Software Systems and Informatics. Due to the cooperation with several departments and institutes of the University, IST is also responsible for the campus library (‘Campusbibliothek für Informatik + Mathematik’) and the Cluster of Excellence ‘Multimodal Computing and Interaction (MMCI)’.

IST is divided into one core and several institute-specific support groups. Presented in a simplified fashion, the core group is responsible for services that are identical for both institutes or are even operated together. The support groups accordingly cover the specific needs of the institutes and the service for their researchers.

Staff structure

In addition to management and procurement (two positions), the Max Planck Institute for Informatics provides six scientific staff members and one technician, three of them to the core group. The core group is supplemented by three scientific staff members from the Max Planck Institute for Software Systems. There is also one employee apiece for additional administrative tasks to support the joint library and MMCI.

For their service desk, the institute’s support group is assisted by a team of students. They are reachable either by e-mail or 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 Department
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*

::: Katholieke Universiteit Leuven, *Leuven, The Netherlands*

::: Max Planck Institute for Intelligent Systems, *Tübingen, Germany*

::: Microsoft Research, *Cambridge, UK*

::: RWTH Aachen, *Aachen, Germany*

::: Saarland University, *Saarbrücken, Germany*

::: Stanford University, *Stanford, USA*

::: Swiss Federal Institute of Technology, *Zurich, Switzerland*

::: Technicolor, *Paris, France*

::: Technische Universität Darmstadt, *Darmstadt, Germany*

::: Tsinghua University, *Bejing, China*

::: University of British Columbia, *Vancouver, Canada*

::: University College London, *London, UK*

::: University of California, *Berkeley, USA*

BIOINFORMATICS

::: BioSolveIT GmbH, *Sankt Augustin, Germany*

::: CeMM Research Center for Molecular Medicine, *Vienna, Austria*

::: Goethe University, *Frankfurt am Main, Germany*

::: Heinrich Heine University, *Düsseldorf, Germany*

::: Informa s.r.l, *Rome, Italy*

::: Karolinska Institutet, *Stockholm, Sweden*

::: Saarland University, *Saarbrücken, Germany*

::: Stanford University, *Stanford, USA*

::: Siemens Healthcare Diagnostics, *Berkeley, USA*

::: University of Freiburg, *Freiburg, Germany*

::: University of Heidelberg, *Heidelberg, Germany*

::: University of Cologne, *Cologne, Germany*

GUARANTEES

::: Chinese Academy of Sciences, *Bejing, China*

::: Google, *Zurich, Switzerland*

::: Microsoft Research, *Cambridge, UK*

::: New York University, *New York, USA*

::: TU Munich, *Munich, Germany*

::: University of Freiburg, *Freiburg, Germany*

::: University of Haifa, *Haifa, Israel*

::: University of Melbourne, *Melbourne, Australia*

::: University of Oldenburg, *Oldenburg, Germany*

INFORMATION SEARCH & DIGITAL KNOWLEDGE

::: DERI, *Galway, Ireland*

::: Google, Mountain View, *USA*

::: Hasso Plattner Institute, *Potsdam, Germany*

::: Hebrew University of Jerusalem, *Jerusalem, Israel*

::: Hungarian Academy of Science, *Budapest, Hungary*

::: IBM Almaden Research Center, *San Jose, USA*

::: INRIA Saclay, *Paris, France*

::: Internet Memory, *Paris, France*

::: Microsoft Research, *Redmond, USA*

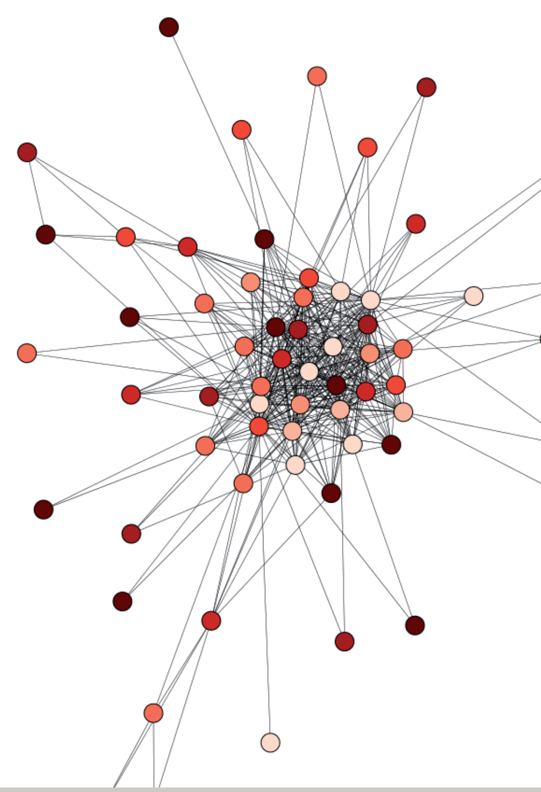

- ::: Paris Institute of Technology, *Paris, France*
- ::: Saarland University, *Saarbrücken, Germany*
- ::: Siemens AG, *Munich, Germany*
- ::: University of Antwerpen, *Antwerpen, Belgium*
- ::: University of Patras, *Patras, Greece*
- ::: University of Versailles, *Versailles, France*
- ::: Yahoo! Research, *Barcelona, Spain*

MULTIMODAL INFORMATION & VISUALIZATION

- ::: Delft University of Technology, *Delft, The Netherlands*
- ::: Inria Rhône Alpes, *Grenoble, France*
- ::: Massachusetts Institute of Technology, *Cambridge, USA*
- ::: Microsoft Research, *Cambridge, UK*
- ::: MIT Media Lab, *Cambridge, USA*
- ::: Stanford University, *Stanford, USA*
- ::: Swiss Federal Institute of Technology, *Zurich, Switzerland*
- ::: Telecom ParisTech, *Paris, France*
- ::: Tsinghua University, *Bejing, China*
- ::: University of Bonn, Bonn, Germany
- ::: University of British Columbia, *Vancouver, Canada*
- ::: University College London, *London, UK*
- ::: University of Glasgow, *Glasgow, UK*
- ::: University of Helsinki, *Helsinki, Finland*
- ::: University of Magdeburg, *Magdeburg, Germany*
- ::: University of St. Andrews, *St. Andrews, UK*

OPTIMIZATION

- ::: Aarhus University, *Aarhus, Denmark*
- ::: Centre National de la Recherche Scientifique, *Paris, France*
- ::: CRS, Rome, Italy
- ::: Danish Technical University, *Kopenhagen, Denmark*
- ::: Indian Institute of Technology Delhi, *New Delhi, India*
- ::: Indian Institute of Technology Kanpur, *Kanpur, India*
- ::: INRIA Nancy, *Nancy, France*
- ::: INRIA Saclay – Ile de France, *Paris, France*
- ::: Kiel University, *Kiel, Germany*
- ::: Logic4Business GmbH, *Saarbrücken, Germany*
- ::: Max Planck Institute for Biological Cybernetics, *Tübingen, Germany*
- ::: Saarland University, *Saarbrücken, Germany*
- ::: Siemens AG, *Vienna, Austria*
- ::: Swiss Federal Institute of Technology, *Zurich, Switzerland*
- ::: TATA Institute of Fundamental Research Mumbai, *Mumbai, India*
- ::: University of Adelaide, *Adelaide, Australia*
- ::: University of Birmingham, *Birmingham, UK*
- ::: University of California, *Berkeley, USA*
- ::: University of Jena, *Jena, Germany*
- ::: University of Liverpool, *Liverpool, UK*
- ::: University of Sheffield, *Sheffield, UK*

SOFTWARE

- ::: FGAN, *Wachtberg, Germany*
- ::: Linköping University, *Linköping, Sweden*
- ::: Los Alamos National Laboratory, *Los Alamos, USA*
- ::: TU Dortmund, *Dortmund, Germany*
- ::: University of Bonn, *Bonn, Germany*
- ::: University of California, *Davis, USA*
- ::: University of Hanover, *Hanover, Germany*
- ::: University of Siegen, *Siegen, Germany*

Selected Publications

[1] F. Abed and C.-C. Huang. Preemptive coordination mechanisms for unrelated machines. In *ESA*, 2012, pp. 12–23.

[2] A. Anand, S. Bedathur, K. Berberich and R. Schenkel. Index Maintenance for Time-Travel Text Search. In *Proceedings of the ACM SIGIR International Conference on Research and Development on Information Retrieval (SIGIR 2012),* Portland, Oregon, USA, 2012, pp. 235–244. Association for Computing Machinery (ACM).

[3] M. Andriluka, S. Roth and B. Schiele. Discriminative appearance models for pictorial structures, 2012.

[4] R. Awadallah, M. Ramanath and G. Weikum. Harmony and Dissonance: Organizing the People's Voices on Political Controversies. In *Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM 2012)*, Seattle, Washington, USA, 2012, pp. 523–532. Association for Computing Machinery (ACM).

[5] B. Beggel, M. Neumann-Fraune, M. Doering, G. Lawyer, R. Kaiser, J. Verheyen and T. Lengauer. Genotyping hepatitis b virus dual infections using population-based sequence data. *Journal of General Virology,* 9(93):1899–1907, September 2012.

[6] J. C. Blanchette, A. Popescu, D. Wand and C. Weidenbach. More spass with isabelle superposition with hard sorts and configurable simplification. In L. Beringer and A. Felty, eds., *Third International Conference on Interactive Theorem Proving*, Princeton, NJ, USA, 2012. Springer.

[7] C. Bock and T. Lengauer. Managing drug resistance in cancer: lessons from hiv therapy. *Nature Reviews Cancer,* 12(7):494–501, 2012.

[8] J. Bogojeska and T. Lengauer. Hierarchical bayes model for predicting effectiveness of hiv combination therapies. *Statistical Applications in Genetics and Molecular Biology*, 11(3):1–19, April 2012.

[9] M. Bokeloh, M. Wand, H.-P. Seidel and V. Koltun. An algebraic model for parameterized shape editing. ACM Transactions on Graphics (Proc. SIGGRAPH), 31(4):78:1–78:10, 2012.

[10] V. Bonifaci, K. Mehlhorn and G. Varma. Physarum can compute shortest paths. In *SODA*, 2012, pp. 233–240. full version to appear in Journal of Theoretical Biology.

[11] K. Bozek, M. Eckhardt, S. Sierra, R. Kaiser, H.-G. Kräusslich, B. Müller and T. Lengauer. An expanded model of hiv cell entry phenotype based on multi-parameter single-cell data. *Retrovirology,* 9(1):60, 2012.

[12] Coja-Oghlan and K. Panagiotou. Catching the k-naesat threshold. In *Proceedings of the 44th ACM Symposium on Theory of Computing (STOC'12)*, 2012, pp. 899–908.

[13] W. Damm, H. Dierks, S. Disch, W. Hagemann, F. Pigorsch, C. Scholl, U. Waldmann and B. Wirtz. Exact and fully symbolic verification of linear hybrid automata with large discrete state spaces. *Science of Computer Programming*, 77(10-11):1122–1150, September 2012.

[14] P. Didyk, T. Ritschel, E. Eisemann, K. Myszkowski, H.-P. Seidel and W. Matusik. A luminance-contrast-aware disparity model and applications. *ACM Transactions on Graphics (Proceedings SIGGRAPH Asia 2012)*, 31(6), 2012.

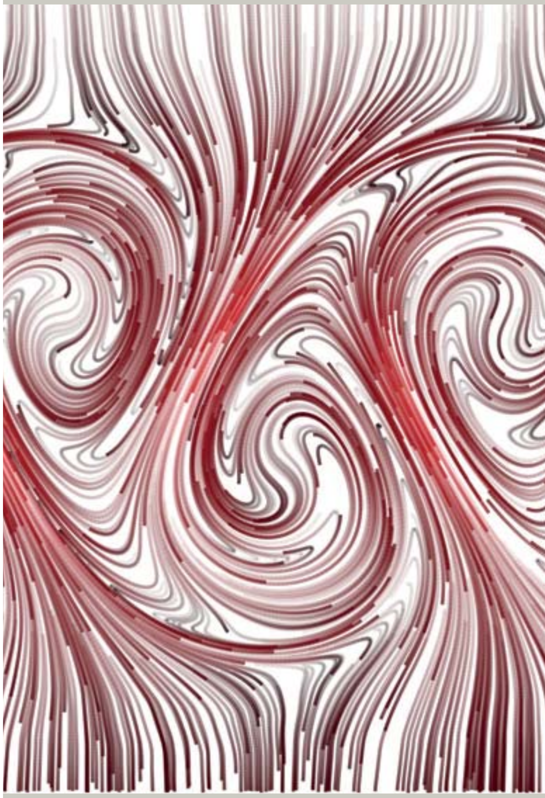

[15] M. Dietzen, E. Zotenko, A. Hildebrandt and T. Lengauer. On the applicability of elastic network normal modes in small-molecule docking. *Journal of Chemical Information and Modeling,* 52(3):844–856, February 2012.

[16] B. Doerr, M. Fouz and T. Friedrich. Why rumors spread so quickly in social networks. *Commun. ACM*, 55(6):70–75, June 2012.

[17] P. Dollár, C. Wojek, B. Schiele and P. Perona. Pedestrian detection: An evaluation of the state of the art. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PP(99):1, August 2012.

[18] N. T. Doncheva, Y. Assenov, F. S. Domingues and M. Albrecht. Topological analysis and interactive visualization of biological networks and protein structures. *Nature Protocols*, 7(4):670–685, 2012.

[19] L. Epstein, A. Levin and R. van Stee. Approximation schemes for packing splittable items with cardinality constraints. *Algorithmica*, 62(1-2):102–129, 2012.

[20] A. Fietzke, E. Kruglov and C. Weidenbach. Automatic generation of invariants for circular derivations in SUP(LA). In N. Bjørner and A. Voronkov, eds., *18th International Conference on Logic for Programming, Artificial Intelligence, and Reason-ing (LPAR-18)*, Mérida, Venezuela, 2012, LNCS 7180. Springer.

[21] P. Fontaine, S. Merz and C. Weidenbach. Combination of disjoint theories: Beyond decidability. In B. Gramlich, D. Miller and U. Sattler, eds., Automated Reasoning – *6th International Joint Conference, IJCAR 2012, Manchester, UK, June 26-29, 2012. Proceedings*, 2012, LNCS 7364, pp. 256–270. Springer.

[22] M. Granados, K. I. Kim, J. Tompkin, J. Kautz and C. Theobalt. Background inpainting for videos with dynamic objects and a free-moving camera. In *Proc. ECCV*, Berlin, Heidelberg, 2012, ECCV'12, pp. 682–695. Springer-Verlag.

[23] K. Halachev, H. Bast, T. Lengauer and C. Bock. Interactive exploration of large genome and epigenome datasets in real time. *Genome Biology*, 13(10):r96, 2012.

[24] J. Hoffart, S. Seufert, D. B. Nguyen, M. Theobald and G. Weikum. KORE: Keyphrase Overlap Relatedness for Entity Disambiguation. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM 2012)*, Maui, Hawaii, USA, 2012, pp. 545–554. Association for Computing Machinery (ACM).

[25] J. Hoffart, F. M. Suchanek, K. Berberich and G. Weikum. Yago2: a spatially and temporally enhanced knowledge base from wikipedia. *Artificial Intelligence Journal, Special Issue on Artificial Intelligence, Wikipedia and Semi-Structured Resources*, 2012.

[26] A. Jacobson, T. Weinkauf and O. Sorkine. Smooth shape-aware functions with controlled extrema. *Computer Graphics Forum (Proc. SGP)*, 31(5):1577–1586, July 2012.

[27] D. M. Kane, K. Mehlhorn, T. Sauerwald and H. Sun. Counting arbitrary subgraphs in data streams. In *Proceedings of ICALP 2012*, 2012, pp. 598–609.

[28] S. Karayev, T. Baumgarnter, M. Fritz and T. Darrell. Timely object recognition. In *Advances in Neural Information Processing Systems 25 (NIPS)*, Lake Tahoe, USA, 2012. MIT Press.

[29] S. Kratsch and M. Wahlstrom. Representative sets and irrelevant vertices: New tools for kernelization. In *IEEE 53nd Annual Symposium on Foundations of Computer Science, FOCS 2012, Los Alamitos, CA, USA, 2012*, 2012, pp. 450–459.

[30] N. Megow, K. Mehlhorn and P. Schweitzer. Online Graph Exploration: New Results on Old and New Algorithms. *Theoretical Computer Science*, pp. 62–72, 2012. preliminary version appeared in ICALP 2011, LNCS 6756, 478 – 489.

[31] P. Miettinen. On Finding Joint Subspace Boolean Matrix Factorizations. In *SIAM International Conference on Data Mining (SDM 2012)*, Annaheim, California, USA, 2012, pp. 954–965. Society for Industrial and Applied Mathematics (SIAM).

[32] S. Miller, J. van den Berg, M. Fritz, T. Darrell, K. Goldberg and P. Abbeel. A geometric approach to robotic laundry folding. *International Journal of Robotics Research*, 31(2):249–267, 2012.

[33] N. Nakashole, G.Weikum and F. M. Suchanek. PATTY: A Taxonomy of Relational Patterns with Semantic Types. In *Proceedings of the ACL Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, (EMNLP-CoNLL 2012)*, Jeju Island, Korea, 2012, pp. 1135–1145. Association for Computational Linguistics (ACL).

[34] N. Nakashole, G. Weikum and F. M. Suchanek. PATTY: A Taxonomy of Relational Patterns with Semantic Types. In *EMNLP,* 2012.

[35] B. Pepik, P. Gehler, M. Stark and B. Schiele. 3d^{2}pm-3d deformable part models. In *Computer Vision – ECCV 2012, 12th European Conference on Computer Vision*, Firenze, Italy, 2012. Springer.

[36] N. Pfeifer and T. Lengauer. Improving hiv coreceptor usage prediction in the clinic using hints from next-generation sequencing data, September 2012.

[37] L. Pishchulin, A. Jain, M. Andriluka, T. Thormaehlen and B. Schiele. Articulated people detection and pose estimation: Reshaping the future. In S. Belongie, A. Blake, J. Luo and A. Yuille, eds., *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Providence, United States, 2012, pp. 3178–3185. IEEE Computer Society.

[38] N. Preda, F. M. Suchanek, W. Yuan and G. Weikum. SUSIE: Search Using Services and Information Extraction. In *IEEE International Conference on Data Engineering (ICDE)*, 2013.

[39] F. Ramírez, G. Lawyer and M. Albrecht. Novel search method for the discovery of functional relationships. *Bioinformatics*, 28(2):269–276, 2012.

[40] E. Reinhard, T. Pouli, T. Kunkel, B. Long, A. Ballestad and G. Damberg. Calibrated image appearance reproduction. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 31(6):article 201, 2012.

[41] M. Rohrbach, M. Regneri, M. Andriluka, S. Amin, M. Pinkal and B. Schiele. Script data for attribute-based recognition of composite activities. In *Computer Vision - ECCV 2012, 12th European Conference on Computer Vision*, Firenze, Italy, October 2012, LNCS 2012. Springer.

[42] F. Schalekamp, D. P. Williamson and A. van Zuylen. A proof of the boyd-carr conjecture. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, 2012, SODA '12, pp. 1477–1486. SIAM.

[43] J. M. Schmidt. Certifying 3-connectivity in linear time. In *Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP'12)*, 2012, pp. 786–797.

[44] F. Simmer, A. B. Brinkman, Y. Assenov, F. Matarese, A. Kaan, L. Sabatino, A. Villanueva, D. Huertas, M. Esteller, T. Lengauer, C. Bock, V. Colantuoni, L. Altucci and H. G. Stunnenberg. Comparative genome-wide dna methylation analysis of colorectal tumor and matched normal tissues. *Epigenetics*, 7(12):1355–1367, December 2012.

[45] H. O. Song, S. Zickler, T. Althoff, G. R. B., M. Fritz, C. Geyer, P. F. Felzenszwalb and T. Darrell. Sparselet models for efficient multiclass object detection. In *12th European Conference on Computer Vision (ECCV)*, Florence, Italy, 2012. Springer.

[46] 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*, 33(12):2521–2537, 2011.

[47] A. Stupar and S. Michel. Being Picky: Processing Top-k Queries with Set-Defined Selections. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM 2012)*, Maui, Hawaii, USA, 2012, pp. 912–921. Association for Computing Machinery (ACM).

[48] T. Sturm and A. Tiwari. Verification and synthesis using real quantifier elimination. In E. Schost and I. Z. Emiris, eds., *Proceedings of the 36th international symposium on Symbolic and algebraic computation (ISSAC 2011)*, San Jose, CA, June 2011, pp. 329–336. ACM.

[49] F. M. Suchanek, S. Abiteboul and P. Senellart. PARIS: Probabilistic Alignment of Relations, Instances, and Schema. *Proceedings of the VLDB Endowment (PVLDB)*, 5(3):157–168, Nov. 2011. Very Large Data Base Endowment Inc. (VLDB Endowment).

[50] F. M. Suchanek, S. Abiteboul and P. Senellart. PARIS: Probabilistic Alignment of Relations, Instances, and Schema. *PVLDB*, 5(3):157–168, 2011.

[51] F. M. Suchanek, J. Hoffart, E. Kuzey and E. Lewis-Kelham. YAGO2s: Modular High-Quality Information Extraction with an Application to Flight Planning. In *German Database Symposium (BTW 2013)*, 2013.

[52] C. Teflioudi, F. Makari and R. Gemulla. Distributed Matrix Completion. In *Proceedings of the IEEE International Conference on Data Mining (ICDM 2012)*, Brussels, Belgium, 2012. Institute of Electrical and Electronics Engineers (IEEE).

[53] K. Templin, P. Didyk, T. Ritschel, K. Myszkowski and H.-P. Seidel. Highlight microdisparity for improved gloss depiction. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 31(4), 2012.

[54] A. Tevs, A. Berner, M. Wand, I. Ihrke, M. Bokeloh, J. Kerber and H.-P. Seidel. Animation cartography – intrinsic reconstruction of shape and motion. *ACM Trans. Graph.*, 31(2):12:1–12:15, Apr. 2012.

[55] J. Tompkin, K. I. Kim, J. Kautz and C. Theobalt. Videoscapes: exploring sparse, unstructured video collections. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 31(4):68:1–68:12, July 2012.

[56] L. Valgaerts, C. Wu, A. Bruhn, H.-P. Seidel and C. Theobalt. Lightweight binocular facial performance capture under uncontrolled lighting. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 31(6):187:1–187:11, Nov. 2012.

[57] G. Weikum, J. Hoffart, N. Nakashole, M. Spaniol, F. M. Suchanek and M. A. Yosef. Big Data Methods for Computational Linguistics. *IEEE Data Engineering Bulletin*, 35(3):46{64, 2012. Institute of Electrical and Electronics Engineers (IEEE).

[58] T. Weinkauf, H.-C. Hege and H. Theisel. Advected tangent curves: A general scheme for characteristic curves of flow fields. *Computer Graphics Forum (Proc. Eurographics)*, 31(2):825–834, May 2012.

[59] C. Wojek, S. Walk, S. Roth, K. Schindler and B. Schiele. Monocular visual scene understanding: Understanding multi-object traffic scenes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99), 2012.

[60] M. Yahya, K. Berberich, S. Elbassuoni, M. Ramanath, V. Tresp and G. Weikum. Natural Language Questions for the Web of Data. In *Proceedings of the ACL Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL 2012)*, Jeju Island, Korea, 2012, pp. 379–390. Association for Computational Linguistics (ACL).

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.

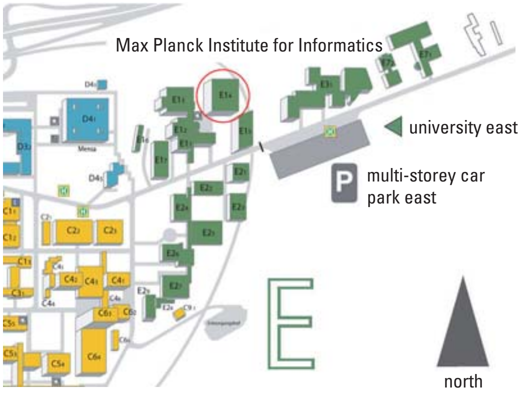

Saarbrücken has its own airport (Saarbrücken-Ensheim) and is well-connected to the Frankfurt and Luxemburg airports by shuttle bus, train, or car. There are good hourly train connections from Saarbrücken to cities all over Germany, and trains to Metz, Nancy and Paris leave at regular intervals. The express 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
Rainer Gemulla
Volker Maria Geiß
Olga Kalinina
Kamila Kolesniczenko
Manuel Lamotte-Schubert
Thomas Lengauer
Kurt Mehlhorn
Jennifer Müller
Marcus Rohrbach
Bernt Schiele
Jens Schmidt
Hans-Peter Seidel
Bertram Somieski
Fabian M. Suchanek
Christian Theobalt
Uwe Waldmann
Christoph Weidenbach
Gerhard Weikum
Andreas Wiese

Proofreading
Krista Ames
Margaret Delap

Photos & Photo Editing
Manuel Lamotte-Schubert

Contact
Max Planck Institut for Informatics

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, 2012 to 31st of December, 2012

Design
Behr Design | Saarbrücken

Printing
Kubbli GmbH | Blieskastel

max planck institut
informatik

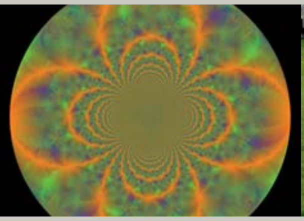
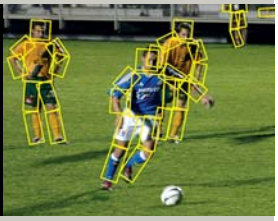

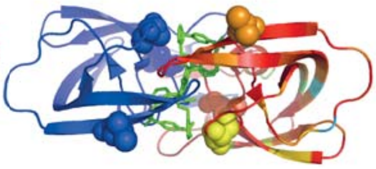


Max Planck Institute for Informatics
Campus E1 4
D-66123 Saarbrücken

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