


max planck institut  
informatik

**Report 2011**



max planck institut  
informatik



# REPORT

**Report 2011**

### Directors

*Prof. Dr. 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. Pankaj Kumar Agarwal*, Department of  
Computer Science, Duke University, Durham, USA

*Prof. Dr. Douglas L. Brutlag*, Stanford University,  
School of Medicine, USA

*Prof. Dr. Joseph M. Hellerstein*, EECS Computer  
Science Division, University of California, Berkeley, USA

*Prof. Dr. Yannis E. Ioannidis*, Department of Informatics &  
Telecommunications, University of Athens, Greece

*Prof. Dr. Friedhelm Meyer auf der Heide*, Heinz Nixdorf Institute,  
University of Paderborn, Germany

*Prof. Dr. Eugene Myers*, Howard Hughes Medical Institute,  
Ashburn, USA

*Prof. Dr. Frank Pfenning*, Computer Science Department,  
Carnegie Mellon University, Pittsburgh, USA

*Prof. Dr. Claude Puech*, INSITU, Paris, France

*Prof. Dr. Éva Tardos*, Cornell University, Ithaca, USA

*Prof. Dr. Demetri Terzopoulos*, Department of Computer  
Science, University of California, Los Angeles, USA

### Curatorship Board

*Hon. Adv. Professor (Tsinghua) Dr. Reinhold Achatz*,  
Head of Corporate Research and Technologies,  
Siemens AG, Munich

*Dr. Siegfried Dais*, Deputy chairman of the Board of  
Management of Robert Bosch GmbH, Stuttgart

*Christiane Götz-Sobel*, Head of the editorial department  
of Science and Technology, ZDF, Munich

*Dr. Christoph Hartmann*, Minister of Economics and  
Science, Saarland

*Prof. Dr. Joachim Hertel*, CEO of DACOS  
Software GmbH, Saarbrücken

*Prof. Dr. Stefan Jähnichen*, President of the German  
Computer Science Association (GI), Bonn

*Prof. Dr. Henning Kagermann*, President of the National  
Academy of Science and Technology (acatech), Munich

*Prof. Dr. Volker Linneweber*, President of  
Saarland University, Saarbrücken

*Prof. Dr. Wolf-Dieter Lukas*, German Federal Ministry,  
of Education and Research, Bonn

*Dr. Nelson Mattos*, Vice President Engineering,  
EMEA, Google, Zurich, Switzerland

*Prof. Dr. Wolffried Stucky*, Institute of Applied  
Informatics and Formal Description Methods (AIFB),  
University of Karlsruhe

*Prof. Dr. Margret Wintermantel*, President of the  
German Rectors' Conference, Bonn



# C O N T E N T S

7	<b>PREFACE</b>
8	<b>THE MAX PLANCK INSTITUTE FOR INFORMATICS: OVERVIEW</b>
14	<b>DEPARTMENT OVERVIEW</b>
	<b>DEPARTMENTS</b>
14	DEPT . 1 ALGORITHMS AND COMPLEXITY
16	DEPT . 2 COMPUTER VISION AND MULTIMODAL COMPUTING
18	DEPT . 3 COMPUTATIONAL BIOLOGY AND APPLIED ALGORITHMICS
20	DEPT . 4 COMPUTER GRAPHICS
22	DEPT . 5 DATABASES AND INFORMATION SYSTEMS
	<b>RESEARCH GROUPS</b>
24	RG . 1 AUTOMATION OF LOGIC
25	<b>THE MAX PLANCK CENTER</b>
26	<b>THE RESEARCH AREAS</b>
28	UNDERSTANDING IMAGES & VIDEOS
38	BIOINFORMATICS
48	GUARANTEES
56	INFORMATION SEARCH & DIGITAL KNOWLEDGE
66	MULTIMODAL INFORMATION
70	OPTIMIZATION
76	SOFTWARE
82	VISUALIZATION
90	<b>IMPRS-CS</b>
92	<b>NEWS</b>
102	<b>THE INSTITUTE IN FIGURES</b>
104	<b>INFORMATION SERVICES &amp; TECHNOLOGY</b>
108	<b>COOPERATIONS</b>
110	<b>PUBLICATIONS</b>
114	<b>DIRECTIONS</b>



## P R E F A C E

On a regular basis and with this report annually, the Max Planck Institute for Informatics publishes a 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.



The report is structured as follows: After an overview of the institute and its departments and research groups, we present the main areas of recent work. These topics span several departments and will also be the focus of our work in the next years. This part is followed by a 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 an 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 research group “*Automation of Logic*” headed by Christoph Weidenbach, and several research groups, which have been established within the Cluster of Excellence on “*Multimodal Computing and Interaction*” as well as through the cooperation with Stanford University, carry out their work at the institute. They are headed by Mario Albrecht, Jan Baumbach, Tobias Friedrich, Ivo Ihrke, Meinard Mueller, Antti Oulasvirta, Tobias Ritschel, Ralf Schenkel, Carsten Stoll, Martin Theobald, Thorsten Thormählen, Michael Wand, and Tino Weinkauff.

## Research Topics

Algorithms are our central research subject. An algorithm is a general recipe for solving a class of problems. We enlarge the computational universe through the design of new and better algorithms. We prove the correctness of these algorithms and analyze their performance. We implement them and validate them through experiment. We make them available to the world in the form of software libraries and internet services. We study inherent properties of computation and investigate techniques for robust software design. We apply the algorithms to interesting application problems. Improved hardware has led to impressive gains in efficiency; improved algorithms can lead, and have led, to even larger gains. Here is an example: The state of hardware and algorithms in 1970 made it possible to calculate an optimal travelling salesman tour through 120 cities (a classic 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 solely 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 with-



# OVERVIEW

out “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 “*Informatics for Genome Research and Epidemiology*” independent research group, led by Dr. Alice McHardy, develops new methods for the analysis of genomic sequences in order to investigate questions of medical and biotechnological relevance.

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*”, started in 2010 and headed by Prof. Bernt Schiele, 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 multimodal processing. Everyday interpersonal communication is based on numerous differ-

ent 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 four years, and several young scientists have since taken professorships in Germany and abroad.

A substantial part of this past year’s activities revolved around the completion and presentation of the Renewal Proposal for the 2013–2017 period. Final decisions regarding the second program phase of the Excellence Initiative will be announced by the DFG on 15 June 2012.



# O V E R V I E W

## 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), the GISAID EpiFlu database 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 90. Our postdoctoral researchers participate in international collaborations such as the "*Max Planck Center for Visual Computing*" (a cooperation with Stanford University in the area of computer graphics), page 25, or the "*Indo Max Planck Center for Computer Science*" (a cooperation with leading universities in India), page 101, or one of our many EU projects.

Our training effort is successful; about 20 PhDs graduate every year, and more than 100 alumni from the institute are now professors at institutions all over the world.

## 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

Secretary

Ingrid Finkler-Paul | Christina Fries





Phone +49 681 9325-1000

Email [infi@mpi-inf.mpg.de](mailto:infi@mpi-inf.mpg.de)

[chfries@mpi-inf.mpg.de](mailto: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 new algorithms for the provably correct isolation of zeroes, for the distribution of information in data networks, for load balancing in computer networks, for solving the traveling salesman problem, and for real-time scheduling. Our latest analyses of random planar graphs with real network properties, our work on the Complexity Theory for randomized search heuristics, as well as our analyses on how certain optimization processes in nature can be algorithmically modeled provide a fundamental understanding for many algorithmic questions.

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.

Our theoretical work is inspired by our practical work and vice versa. Our theoretical work forms the basis for demonstrators and libraries. For instance, the CompleteSearch Engine is based on new index structures, which are more powerful than known structures, but require no additional space. CGAL algorithms rely on a deep theoretical understanding of algebraic curves and surfaces.

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 from the Swiss National Science Foundation. 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).

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. Part of our training concept is that we continue to employ our students after successful promotion and only after a minimum stay of one year at a research institution abroad. The entirety of these measures ensures that our group members are well-equipped after their stay at the

Max Planck Institute to attain attractive positions in industry not limited to research, or to continue their scientific careers at leading universities or research institutes domestically and abroad. ∴


# Computer Vision and Multimodal Computing

PROF. DR. BERNT SCHIELE

## DEPT. 2



The department was founded in 2010 and currently includes 10 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 cbalzert@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 where object models are learned directly from computer graphic models, or where speech and image processing is combined in a multimodal learning approach.

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. A recently presented approach not only describes people but also entire 3D scenes, representing a further step towards complete image and scene understanding.

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

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

# Computational Biology and Applied Algorithmics

PROF. DR. DR. THOMAS LENGAUER, PH.D.

## DEPT. 3



This department dates back to October 2001 and is directed by Prof. Dr. Dr. Thomas Lengauer.

The department currently comprises about 20 scientists, who perform research exclusively in computational biology.

### CONTACT


**Computational Biology and Applied Algorithmics**

Secretary


Ruth Schnepfen-Christmann

Phone +49 681 9325-3000

Email [ruth@mpi-inf.mpg.de](mailto: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 38). The building blocks of such biochemical networks are DNA, RNA, and proteins. The circuitry is realized by transient binding of the molecules among each other and with small organic ligands. In this manner, proteins catalyze chemical reactions, regulate the expression of genes, and transduce signals in and between cells.

Our research concentrates on two groups of diseases.

On the one hand, we investigate viral infections. The molecular understanding of such diseases requires the elucidation of the function of the relevant viral proteins and their interactions with the molecules of the infected patient as well as with drug molecules. So, analysis must be carried out at the level of the three-dimensional structure of the involved molecules. 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 40), Hepatitis B (*"Analysis of HBV Resistance"*, page 43), Hepatitis C (*"Analysis of Viral Genomes Using Next-generation Sequencing"*, page 41) and EHEC (*"Chasing EHEC with the Computer"*, 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 42). Furthermore, we investigate the pattern of geographical distribution of the virus (*"Phylogenetics of HIV"*, page 45).

Disease like cancer, neurodegenerative diseases, or immunological diseases are based on different principles. Here, the starting point is an interplay between the genomes of the patient and the environment. Therefore, an early diagnosis of a disease like cancer can be based on the analysis of genetic or epigenetic aberrations of the relevant tissue. 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 (*"Deciphering the Second Code – Computational Epigenetics"*, page 44).

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 2011/2012 issue, include the field of epigenetics, the analysis of protein function and protein interaction networks as well as the optimization of AIDS therapies.

The 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, (focusing bioinformatics research for analyzing viral drug resistance), part of the Clinical Research Group 129 of the German

Research Foundation Deutsche Forschungsgemeinschaft (on Hepatitis C), as well as of the National Genome Research Network supported by the German Science Ministry and the European Consortia BLUEPRINT (epigenetics) and PREDEMICS (research on viruses posing substantial epidemiological risks). ...

# 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](mailto: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, telecommunications, 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. 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 algorithms 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 and has been extended in recent years for an additional funding phase. 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. A substantial part of the activities this past year revolved around the completion and presentation of the Renewal Proposal for the 2013–2017 period. Final decisions regarding the second program phase of the Excel-

lence Initiative will be announced by the DFG (German Research Foundation) on 15 June 2012. 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.

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 young scientist prizes and also the Leibniz Prize of the DFG for Professor Hans-Peter Seidel. The group was further strengthened last year with the appointment of Christian Theobalt (Stanford) to a W2 permanent chair at the institute. ...

# Databases and Information Systems

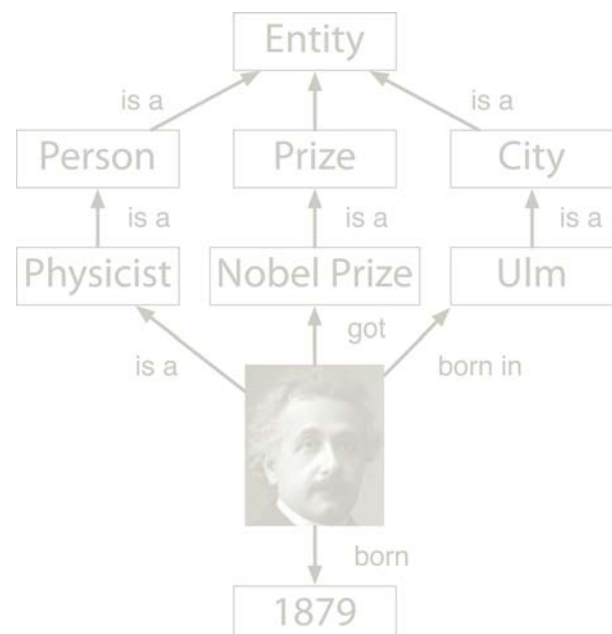
PROF. DR.-ING. GERHARD WEIKUM

## DEPT. 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. Analysis of distributed data, especially in scalable peer-to-peer systems and online communities, for example, in social networks and Web 2.0 media.



## CONTACT

Databases and Information Systems



Secretary

Petra Schaaf

Phone +49 681 9325-5000

Email [schaaf@mpi-inf.mpg.de](mailto:schaaf@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 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.

The group is a worldwide leader in its methods of combining logic-based algorithms for structured data sets and statistically-based methods for unstructured textual data. The former fall into the area of database systems (DB) and data analysis, the latter into the area of information retrieval (IR) and search engines. These two directions were historically divided; their connection has become ever more important with increasingly mixed data forms for digital libraries, social online communities, e-science groups and, last but not least, for companies and in the Web itself. The group is a trendsetter in this strategic direction of DB-IR integration.

The group's research methodology spans the entire spectrum from theoretical foundations to practical systems building for applications and experiments. Many of the prototype systems developed by the group are publicly available as open source software and are used by other research groups around the world. These include, in particular:

1. 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 tools that automatically create and maintain the YAGO knowledge base, as well as the YAGO contents itself;
4. The scalable system *PROSPERA*, which extracts new facts from web pages and natural-language 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”* as well as the DFG Excellence Cluster *“Multimodal Computing and Interaction”*. The group has received a Google Research Award. ...

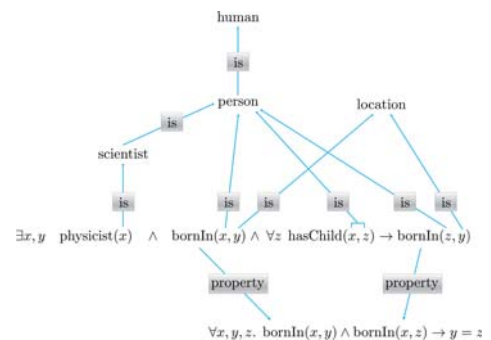
# Automation of Logic

PROF. DR. CHRISTOPH WEIDENBACH

## RG. 1



### RG. 1 Automation of Logic



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 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, protocols, safety properties, software, solutions for decision and optimization problems, and in the computation in linear and non-linear arithmetic. In order to keep up with increasingly complex systems and required properties, the formal procedures used today must significantly increase in productivity. 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](mailto:jmueller@mpi-inf.mpg.de)

# Max Planck Center for Visual Computing and Communication

PROJECT MANAGEMENT: PROF. DR. HANS-PETER SEIDEL

**In 2003, the “Max Planck Center for Visual Computing and Communication” was founded 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 2010, based on the success of the program, funds of up to 7.8 million Euros were approved for an additional funding period.**

The cooperation's research 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, visualization, 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 independently research with 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 in Stanford, where they have the status of a “Visiting Assistant Professor” (*Phase I*), the scientists return to Germany and continue their work as young scientists' 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 offered attractive return perspectives in Germany, 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 the 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 14 young scientists have completed the program. Of these, 12 young scientists (86 %) have meanwhile received professorships, eleven thereof in Germany (79%), and even seven of them being full professorships.

The program'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, which the candidates contribute to developing significantly. 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. Based on their success, these structural elements of the program could possibly be a model for other faculties. ...



## CONTACT

**Prof. Dr. Hans-Peter Seidel**

Secretary Sabine Budde

Phone +49 681 9325-4000

Email [budde@mpi-inf.mpg.de](mailto:budde@mpi-inf.mpg.de)

## Research Areas

### UNDERSTANDING IMAGES & VIDEOS

- 30 People Detection and Pose Estimation  
in Challenging Real-World Scenes
- 31 Scalability of Object Class Recognition
- 32 3D Scene Understanding from Monocular Cameras
- 33 Learning of Visual Representations
- 34 Image-based 3D Scene Analysis
- 35 Markerless Reconstruction of Dynamic Scenes
- 36 Advanced Video Processing

### BIOINFORMATICS

- 40 Protein Structure and Interactions
- 41 Analysis of Viral Genomes Using Next-generation Sequencing
- 42 Bioinformatical Support of HIV Therapy
- 43 Analysis of HBV Resistance
- 44 Deciphering the Second Code – Computational Epigenetics
- 45 Phylodynamics of HIV
- 46 Chasing EHEC with the Computer

### GUARANTEES

- 50 Physarum Can Compute Shortest Paths – a Rigorous Analysis
- 51 Geometric Algorithms – Exact and Efficient
- 52 Dealing with Selfishness in Optimization
- 53 Automated Deduction
- 54 Model Checking for Hybrid Systems
- 55 Quantifier Elimination – Statements Can Also Be Calculated



## INFORMATION SEARCH & DIGITAL KNOWLEDGE

- 58 Random Telephone Chains – Efficient Communication in Data Networks
- 59 Search and Mining in Web Archives
- 60 Information Search in Social Networks
- 61 EnBlogue – What is New and Interesting in Web 2.0 ?
- 62 YAGO – a Collection of Digital Knowledge
- 63 URDF – Efficient Reasoning in Uncertain RDF Knowledge Bases
- 64 PICASSO – Soundtrack Recommendation for Images

## MULTIMODAL INFORMATION

- 68 Recognizing Human Activity
- 69 Knowledge-Based Source Separation for Music Signals

## OPTIMIZATION

- 72 Approximation Algorithms for Hard Optimization Problems
- 73 The Theory of Bio-inspired Algorithms
- 74 The Traveling Salesman Problem
- 75 Rule-based Product Configuration

## SOFTWARE

- 78 TopX 2.0 – Efficient Search in Digital Libraries
- 79 RDF-3X – Fast Searches for Semantic Data
- 80 Decision Procedures for Ontologies

## VISUALIZATION

- 84 Stereo and HDR Imaging: Display Quality Measurement and Enhancement
- 85 Advanced Real-Time Rendering
- 86 Computational Photography and Measurement
- 87 Statistical Models of Humans
- 88 Correspondences and Symmetry in 3D Scenes
- 89 Flow Analysis

# 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. ∴

CONTRIBUTIONS



UNDERSTANDING  
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH  
& DIGITAL KNOWLEDGE

MULTIMODAL  
INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

<b>People Detection and Pose Estimation in Challenging Real-World Scenes</b>	30
<b>Scalability of Object Class Recognition</b>	31
<b>3D Scene Understanding from Monocular Cameras</b>	32
<b>Learning of Visual Representations</b>	33
<b>Image-based 3D Scene Analysis</b>	34
<b>Markerless Reconstruction of Dynamic Scenes</b>	35
<b>Advanced Video Processing</b>	36

## People Detection and Pose Estimation in Challenging Real-World Scenes



Examples of people detection and pose estimation obtained with our approach

Finding and following people is a key technology for many applications such as robotics and automotive safety, human-computer interaction scenarios, or for indexing images and videos from the web or surveillance cameras. At the same time it is one of the most challenging problems in computer vision and remains a scientific challenge for realistic scenes.

We have developed a new approach for detecting people and estimating their poses in complex street scenes with multiple people and dynamic backgrounds. Our approach does not require multiple synchronized video streams but can operate on images from a moving, monocular and uncalibrated camera. Several examples of people detections and estimated body configurations are shown on the pictures. The important challenges addressed in our approach are frequently full and partial occlusions of people, cluttered and dynamically changing backgrounds, and ambiguities in recovering 3D body poses from monocular data.

### 2D human pose estimation

Several key components contribute to the success of the approach. The first is a novel and generic procedure for people detection and 2D pose estimation that is based on the pictorial structures model and also enables to estimate viewpoints of people from monocular and single images.

Our approach is based on learned local appearance models of human body parts and a kinematic model of the human body. The appearance of the body parts is represented by a set of local image descriptors. We employ a boosting classifier trained on the dataset of annotated human poses in order to learn which of these local features are informative for the presence of the body part at a given image location. Interpreting the output of each classifier as a local likelihood we infer the optimal configuration of the body parts using belief propagation.

### 3D human pose estimation and multi-person tracking

The second key component in our approach enables people tracking and 3D pose estimation. In contrast to prior work, we accumulate evidence in multiple stages thereby reducing the ambiguity of 3D pose estimation effectively in every step.

In particular we propose a novel multi-stage inference procedure for 3D pose estimation. Our approach goes beyond prior work in this area, in which pose likelihoods are often based on simple image features such as silhouettes and edge maps and which often assume conditional independence of evidence in adjacent frames of the sequence. In contrast to these approaches, our 3D pose likelihood is formulated in terms of estimates of the 2D body configurations and viewpoints obtained with a strong discriminative appearance model. In addition, we refine and improve these estimates by tracking them over time, which allows to detect occlusion events and group hypotheses corresponding to the same person. We demonstrate that the combination of these estimates significantly constrains 3D pose estimation and allows to prune many of the local minima that otherwise hinder successful optimization. ...



#### CONTACT

**Mykhaylo Andriluka**

**DEPT. 2 Computer Vision and Multimodal Computing**

Phone +49 681 9325-2119

Email andriluk@mpi-inf.mpg.de



**Bernt Schiele**

**DEPT. 2 Computer Vision and Multimodal Computing**

Phone +49 681 9325-2000

Email schiele@mpi-inf.mpg.de

# Scalability of Object Class Recognition

Because of 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: building reliable object class models requires a sufficiently high number of representative training examples, often in the form of manually annotated images. Since manual annotation is costly, our research aims at reducing the amount of required training examples for building object class models, thereby increasing scalability. We explore three different ways of achieving this goal.

## Re-using object class model components

In this project, we are devising an object class model that represents objects as an assembly of spatially constrained partial shapes. Similar object classes, such as horses and giraffes [Figure 1], share similar representations. Components of one model, e.g., the expected variation based on the localisation of the horses' body parts, can be re-used in order to build another model, e.g., a giraffe model. Re-used model components do not have to be learned. Hence, the number of required training examples is reduced. The focus on object shape rather than on appearance allows model components to be re-used, even for relatively dissimilar classes, such as horses and swans.




Figure 1: Re-using model components, example recognition output

## 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. 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 on par with providing the object class-attribute relations by human subjects.

## Learning from 3D CAD data

In an attempt to abandon training images altogether, this project aims at learning object class models from 3D computer-aided design (CAD) data [Figure 2]. CAD models are typically

employed in product design, computer games, and movie making, and are readily available for many object classes. Since CAD models provide an accurate and detailed representation of object shape, they lend themselves to the learning of object class models. Since they are view-point independent, they allow for arbitrary view-points, that means artificial training examples, through rendering. The key challenge consists in bridging the gap between rendered training images and real-world test images, which we achieve by a shape-based abstraction of object appearance. Our experiments demonstrate state-of-the-art performance for the recognition of cars from multiple viewpoints, even compared to approaches using real-world training images. ...




Figure 2: Automatic identification of reusable attribute models



## CONTACT

**Marcus Rohrbach**

**DEPT. 2** Computer Vision and Multimodal Computing

Phone +49 681 9325-1206

Email [rohrbach@mpi-inf.mpg.de](mailto:rohrbach@mpi-inf.mpg.de)



**Michael Stark**

**DEPT. 2** Computer Vision and Multimodal Computing

Phone +49 681 9325-2120

Email [stark@mpi-inf.mpg.de](mailto:stark@mpi-inf.mpg.de)

Internet <http://www.d2.mpi-inf.mpg.de/nlp4vision>



## 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 such 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 relatively 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** Computer Vision and Multimodal Computing

Phone +49 681 9325-2000

Email [schiele@mpi-inf.mpg.de](mailto:schiele@mpi-inf.mpg.de)



**Christian Wojek**

**DEPT. 2** Computer Vision and Multimodal Computing

Phone +49 681 9325-2000

Email [cwojek@mpi-inf.mpg.de](mailto:cwojek@mpi-inf.mpg.de)

Internet <http://www.d2.mpi-inf.mpg.de/monocular-3d-scene>



# Learning of Visual Representations

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.




Figure 1

Most recently, we have defined a recursive scheme that allows us to form hierarchical representations built on the same principle [see figure 2]. Our studies have shown the importance of holistic

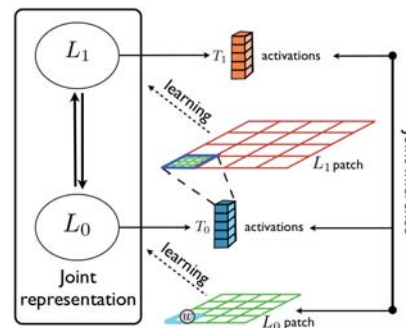


Figure 2

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 3].

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 3

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.



## CONTACT


Mario Fritz

DEPT. 2 Computer Vision and Multimodal Computing

Phone +49 681 9325-1204

Email mfritz@mpi-inf.mpg.de

## Image-based 3D Scene Analysis



**Figure 1:** The estimation of camera motion allows to display the virtual green cube in this stereo image sequence from the correct perspective at any point in time.

views if the calibration pattern is visible in all views. In order to circumvent this problem, our method makes usage of coded patterns, that can be distributed over a larger volume. Now, it is no longer necessary that the calibration pattern is visible in all camera views. It is sufficient if at least one pattern is visible and if a linking of the images via the patterns is possible [Figure 2]. With this practical method an easier calibration of larger scenes is possible. That way, for example, cameras distributed in a room can be calibrated. A specific application emerged in the calibration of an autonomously moving robotic head that can look in all directions and whose cameras could not be calibrated with a single calibration pattern. ...


### Estimation of camera movement from stereo image sequences

Since an ever increasing amount of contemporary movies is filmed in stereoscopic 3D (S3D), there is currently a large demand for tools to generate special effects for stereoscopic image sequences. Thereby the estimation of the camera movement and the 3D reconstruction from image sequences is an essential step that enables the positioning and perceptively correct display of virtual objects [Figure 1]. Until now our algorithms were capable of processing monocular image sequences only. By extending the camera model with parameter that are especial suited for stereo camera systems, results for stereoscopic image sequences can now be generated as well. The temporal dependency of single camera parameters yields a more complex structure in the

algorithms. With our approach we can still achieve similar processing speeds as in the monocular case. We were also able to show that the resulting accuracy of the estimation is higher with our approach than in the case of a separate estimation for both cameras.

### A practical calibration method for cameras in larger scenes

In this project a method was developed that allows the estimation of position and orientation of a single or multiple cameras from a series of images or from a video sequence. For a reliable and accurate estimation of camera parameters usually a calibration pattern is used. However, using a single calibration pattern it is only possible to generate a globally consistent calibration of camera



**Figure 2:** Our approach allows camera calibration from images if the camera views overlap. In every image multiple calibration patterns are visible, whereby in every image at least one of those patterns is visible in the next image.



#### CONTACT

**Thorsten Thormählen**

**DEPT. 4 Computer Graphics**

Phone +49 681 9325-4017

Email [thormae@mpi-inf.mpg.de](mailto:thormae@mpi-inf.mpg.de)

Internet <http://www.mpi-inf.mpg.de/~thormae/>

# 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.

In the past year, 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 inter-




Figure 1: *left to right*: one of 12 camera views; video frames in which individual persons are segmented; reconstructed skeleton and surface model.




Figure 2: *left*: one of 12 camera views; *right*: 3D model with reconstructed high geometry detail.

acting actors in the same scene [Figure 1]. 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 2]. 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). By means of these new algorithms, we come one step closer to the long term goal of becoming able to capture complete dynamic 3D models of arbitrary indoor and outdoor environments, possibly using moving and unsynchronized cameras.

## 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/>





## Advanced Video Processing

Image processing methods are important tools during the postprocessing of photographs. 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. Post-processing 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



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

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. In our research, we have developed one of the first methods in the literature to automatically solve the latter problem, namely inpainting of dynamic scene background [see Figure]. ...



#### CONTACT

**Christian Theobalt**  
**DEPT. 4 Computer Graphics**  
 Phone +49 681 9325-4028  
 Email theobalt@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



<b>Protein Structure and Interactions</b> .....	40
<b>Analysis of Viral Genomes Using Next-generation Sequencing</b> .....	41
<b>Bioinformatical Support of HIV Therapy</b> .....	42
<b>Analysis of HBV Resistance</b> .....	43
<b>Deciphering the Second Code – Computational Epigenetics</b> .....	44
<b>Phylodynamics of HIV</b> .....	45
<b>Chasing EHEC with the Computer</b> .....	46

UNDERSTANDING  
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH  
& DIGITAL KNOWLEDGE

MULTIMODAL  
INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

## Protein Structure and Interactions

### Proteins as biological machines

Proteins are biological molecules that facilitate life's processes: they are molecular engines that synthesize and decompose molecular building blocks such as nucleic and amino acids; convert energy; and facilitate regulation and communication inside cells. They can do so because they are made up of a chain of amino acids that folds into a complicated 3-dimensional structure, thus translating information from 1D to 3D. The specific way in which each individual protein folds depends on the physical and chemical properties of the amino acids that comprise the protein. Structural biology has accumulated extensive data on protein 3D structures that can be analyzed using computational methods.

We extend this paradigm to molecular complexes. Proteins do not act by themselves, but rather are part of the conveyor belt that performs a single action towards a common larger goal. That is why, on many occasions, proteins assemble into complexes to perform their task. These complexes can comprise hundred of proteins. Prediction and analysis of properties of protein complexes is one of the major tasks in structural bioinformatics.

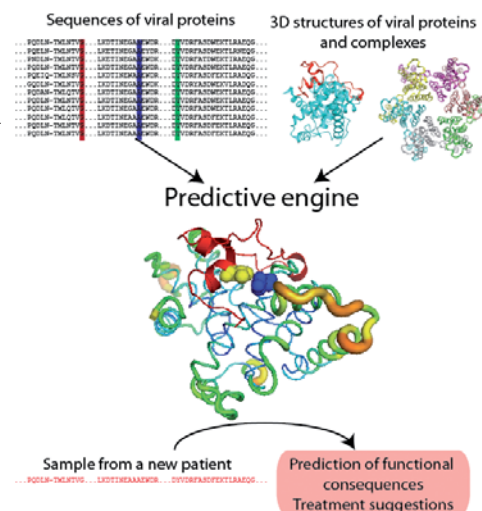
### Structural bioinformatics of viral systems

Viruses invade cells of higher organisms and use them to replicate and multiply. They are small systems, typically containing 10-20 proteins. Using this limited toolkit, they engage in a complex network of interactions with the host cell and recruit host cell proteins to facilitate the virus life cycle. Simultaneously, they adapt very quickly to the persistent pressure exerted by the host immune system and possible drug treatment. Impressive amounts of clinical and laboratory data on the variation of protein sequence in viruses are available. This makes viral systems unique targets for bioinformatics studies.

Protein sequences from virus strains isolated from various patients provide valuable information for patient treatment, as demonstrated by the geno2pheno approach. Including structural information in this analysis will allow us to understand how the variations characteristic of a sequence under consideration influence the function of the protein and its interactions with other molecules in its environment. 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 its affinity to an inhibitor. If a mutation is located in an interaction interface with a host protein, this may destroy the binding, and the mutated virus will not survive.

### A pipeline that combines sequence and structural analysis to predict effect of sequence alterations

We analyze the sequences of viral proteins that have specific properties, e.g. are resistant to certain drugs. We extract the sequence features that are responsible for this phenotype and analyze them in a structural context. We have analyzed multiple sequences of the viral protein that forms the capsid of the human immunodeficiency virus (HIV) and predicted several pairs of mutations that are likely to interact functionally. One of those point mutations restores the viral infectivity that would otherwise be hampered by the other mutation of the pair. Using the structural analysis, we demonstrated that this pair of amino acids could be responsible for keeping individual protein subunits in the ordered capsid structure.



We combine genomic-, modeling- and biophysics-based methods to assess the alterations observed in a given sequence and predict their potential influence on protein function. We also analyze the binding affinity to inhibitors and can suggest possible drug treatment. ...



#### CONTACT


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

# Analysis of Viral Genomes Using Next-generation Sequencing

Research on viral infectious diseases 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 clinical data analysis 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.



Output of the geno2pheno[hcv] web service

## 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 included 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 bringing personalized genomics into clinical practice.

## Screening for oncoviruses in human cancers

Besides causing infectious diseases, viral pathogens are also suspected to be contributing factors to over 20 % of all human cancers worldwide. Several species of viruses are known to carry and integrate into the human genome oncogenic genes, possibly triggering cancerous cell transformations. Several oncogenic viruses such as the human papilloma virus (HPV) and the Merkel cell polyoma virus have been identified so far, enabling the development of vaccines that prevent certain forms of cancer. In collaboration with the Institute of Child Oncology at the University Hospital of Cologne, we recently extended Virana's abilities to detect gene expression of known and unknown viruses within next-generation transcriptome sequencing data of neuroblastoma, an important form of childhood cancer. This viral metagenomics analysis enables researchers to differentiate between human gene transcripts and viral transcripts and can be used to screen cancers samples for viral expression in a high-throughput and interpretable fashion.

...

## CONTACT

Sven-Eric Schellhorn

DEPT. 3 Computational Biology and Applied Algorithmics

Phone +49 681 9325-3028

Email sven@mpi-inf.mpg.de


Internet <http://bioinf.mpi-inf.mpg.de/~sven>



## 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 using methods of genome sequencing. From the genotype, one can identify the viral resistance phenotype with bioinformatics methods. There are two approaches for this purpose. In the first approach experts manually compose a set of rules for deriving the phenotype from the genotype. These rules are collected in computer-based expert systems, which are exercised then put into use by the clinician. 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. Our research over the last 10 years has led to the geno2pheno system, which is freely available on the Internet under [www.geno2pheno.org](http://www.geno2pheno.org) and can be used 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, our research on geno2pheno was honored with the AIDS Research Award of the Heinz-Ansman Foundation.

Several analysis variants offered on the geno2pheno server that are used in routine clinical practice can be described as *virtual phenotypes*. A virtual phenotype is a bioinformatical (i.e., computer-based) procedure that 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 feasible, in that it may be too expensive, may take too much time, or is simply unavailable.




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

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 such virtual phenotypes available, and these can be used as a companion diagnostic for administering certain AIDS drugs.

In addition, geno2pheno offers a second generation of bioinformatical support systems for therapy selection. Not only do these systems estimate viral resistance against individual drugs, they also assess the effectiveness of the combination therapy as a whole. This analysis also entails an estimate of how difficult it will be 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 practice on a broad scale. Reasons for this include the need for additional validation of these

methods. Moreover, the accuracy of the procedures can be improved further by making use of additional information, such as details on the therapy history of the patient or on aspects of his or her genotype, especially those that pertain to the immune system. Finally the predictions are not yet easy enough to interpret. The clinician expects to receive not only the result of a bioinformatical prediction method, but also an argument as to why that result is plausible. Research on improving the second-generation prediction systems is currently a focus of our work.

New sequencing methods make it possible to estimate the composition of essentially the entire viral population in a patient. On the basis of such data, predictions as to whether a therapy will be successful can be made with more accuracy. When the proportion of resistant virus inside the population exceeds a certain threshold, the therapy should not be administered [Figure 2].



**Figure 2: The resistance is estimated for each virus in the population of the patient. If too large a proportion of viruses is resistant, the drug should not be administered.**



### CONTACT


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



# Analysis of HBV Resistance

## Hepatitis B Virus

The hepatitis B virus (HBV) infects human liver cells, which can result in acute or chronic inflammation of the liver (*hepatitis*). The World Health Organization (WHO) estimates that 300 to 420 million people suffer from chronic hepatitis B worldwide and about one million people die per year as a consequence of this disease. In order to prevent long-term consequences such as liver cirrhosis and liver cancer, antiviral therapy must begin early. The goal of treatment is the permanent prevention of viral replication in liver cells. Several drugs are available for this purpose.



Hepatitis B virus particles  
(Image source: Antonio Šiber)

## Reverse Transcriptase Inhibitors

Nucleotide analog reverse transcriptase inhibitors are of prime medical importance for treatment of chronic hepatitis B. These chemical compounds, which originally were introduced for AIDS therapy, prevent the reproduction of viral DNA via a simple, but elegant trick. The building blocks of DNA, the nucleotides, are normally assembled in form of a long chain by the use of a special viral protein, the polymerase. The chemical binding sites relevant for building the chain are found on both sides of the nucleotide. Reverse transcriptase inhibitors used in the treatment of hepatitis B are so similar to the original nucleotides that the viral polymerase incorporates them into the DNA chain. However, one of the chemical groups responsible for chain building is missing, so that chain building is stopped and the


viral DNA cannot be completed. This therapy effectively suppresses the reproduction of the virus, thus preventing long-term consequences of the infection. However, this effect lasts only as long as the virus has not developed resistance to the treatment. Resistance can be effected by suitable mutations in the gene of the viral polymerase. A resistant form of this protein can distinguish reverse transcriptase inhibitors from the original nucleotides and reject them. Resistance is the major problem in the long-term treatment of hepatitis B.

## Resistance Factor

In order to optimize the administration of drugs in the presence of viral resistance, we have set ourselves a challenging goal together with our collaborators in medicine and virology. We aim to completely resolve the resistance profile of the hepatitis B virus harbored by the patient. Through extensive laboratory tests, one can determine a virus' resistance in the form of an easily interpretable numerical value. Ideally, this so-called resistance factor should be available before a drug therapy is selected for a patient. Currently, the resistance factor is not easily derived. The relevant laboratory testing procedures are expensive, not widely available, and they take a long time to produce the result. Therefore, we aim to adapt the computational method that we are successfully employing for treatment of human immunodeficiency virus (HIV) for use in HBV therapy. Our efforts to support HIV therapy are described in the article by Thomas Lengauer on page 42.

## Resistance Landscape

All properties of the virus are encoded in its genome in a complex and concealed fashion. We use statistical methods and clinical data to computationally estimate the resistance factor from the viral genome. In this way, we can replace the complex laboratory testing procedure with a simple computation. The underlying dataset needs to comprise several hundred laboratory measurements of resistance factors for different viral genomes. The judicious selection of these measurements is crucial for our project. First, we can only evaluate mutational patterns that are similar to those we have previously seen in our measurements, and second, we have to try to estimate the resistance factor of many viral variants as precisely as possible on the basis of a limited amount of experimental data. We have developed a system which selects for testing several hundred viral prototypes from several thousand hepatitis B viruses from all over Europe. These prototypes reflect the diversity of the virus, and enable us to develop an algorithm which can support the choice of the most effective individual therapy in the future. ...



Hepatitis B resistance landscape  
(schematic presentation)

## CONTACT

Bastian Beggel

DEPT. 3 Computational Biology and  
Applied Algorithmics

Phone +49 681 9325-3016


Email beggel@mpi-inf.mpg.de

Internet <http://www.genozpheno.org>



## Deciphering the Second Code – Computational Epigenetics

Our DNA encodes the appearance and behavior of our cells. However, there is a multitude of additional layers of control that regulate the ways in which the DNA is packed and processed and thus determine the fate of the cell. Computer scientists at the Max Planck Institute for Informatics contribute to deciphering this second, *epigenetic* code in order to foster the molecular understanding of development and disease and in the long term to improve cancer therapies.



**Methylation of DNA influences its accessibility**

### In cancer, the packing of the DNA becomes highly distorted

The 3 billion letter code known as the *human genome* comprises the blueprint for the molecular building blocks of life. There exists a large variety of different tissues in the human body, all of whose cells share the same genome but run individual genetic programs. The *epigenome* provides a code for the operation of these programs: DNA does not float around naked in the cell's nucleus, but is packed more or less tightly around molecular scaffolds (chromatin). The density of this packing controls how accessible the DNA is to the molecules reading off the genomic information. This packing is controlled by chemical modifications to the DNA itself (methylation) or the histone proteins which comprise the scaffold. Consequently, epigenetics plays an important role in orchestrating the entire organism: Complex tissue structures can arise only if, after countless cell divisions, each cell follows its program, which is configured by the accessible regions in the genome. In cancer cells, this packing is highly deregulated: Programs that contribute to the cell's self-regulation – e.g.,

by programmed cell death (apoptosis) if cellular programs malfunction – are shut down, while programs that induce chaos in the structure of the DNA by reorganizing the genome become active.

### Tackling the layer cake of the epigenome

The first step in finding a remedy is understanding the problem. Following this principle, and in close collaboration with life scientists, computer scientists at the Max Planck Institute for Informatics develop methods for understanding the complex interplay of regulatory layers. Modern experimental techniques allow for charting comprehensive maps of the epigenetic modifications at affordable costs. A major bioinformatics challenge is dealing with and integrating the resulting heterogeneous flood of data. At the Max Planck Institute for Informatics, statistical methods are employed and developed in order to infer the interplay of genomic and epigenomic features. Several web services have been developed supporting the analysis of epigenomic data. One of them, EpiExplorer, facilitates an integrative view and interactive exploration of

genomic and epigenomic features based on technology that has been developed to search for information in the World Wide Web.

### Pinpointing epigenetic determinants of cancer helps to improve therapy

By using computational methods, it is possible to compare epigenomic profiles of tumor and healthy tissues and thus to identify cancer-specific aberrations. Moreover, tumors are not all alike: While one tumor type might grow by being extraordinarily efficient in replicating, another might succeed especially well in evading the human immune response. These characteristics are reflected in the tumors' epigenomes. Different tumor types can be characterized by their epigenetic differences via statistical learning methods. In the future, these findings will have a significant effect on selecting effective therapies on a patient-specific basis. The key phrase here is *personalized medicine*: Clinical assays will be able to identify epigenetic signatures of tumors for individual patients. Using these signatures enables clinicians to prescribe more personalized treatments. ...

#### CONTACT

##### Yassen Assenov

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

Phone +49 681 9325-3009  
Email yassen@mpi-inf.mpg.de



##### Konstantin Halachev

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

Phone +49 681 9325-3112  
Email halachev@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>  
<http://www.computational-epigenetics.de>





## Phyldynamics of HIV

For many types of virus, genetic mutation occurs on the same timescale as disease transmission. A map of ancestral lineages, a phylogeny, recaptures these transmission patterns. Genetic sequencing allows for biologically principled reconstruction of phylogenetic structures. Our group applies and develops phylogenetic methods to infer the transmission dynamics of the most dangerous infectious agent of our age, the Human Immunodeficiency Virus (HIV).

Science and medicine have made great progress in the fight against HIV. Understanding of the viral life cycle has led to an evergrowing arsenal of therapeutic agents. Advances in clinical care have turned a deadly disease into a manageable infection – manageable, but not curable. The disease continues to spread.


### Transmission networks

We, in collaboration with the Institute for Virology at the University of Cologne, have published an investigation of HIV transmission patterns in North Rhine-Westphalia [[http://www.auswaertiges-amt.de/DE/Infoservice/Terminologie/Bundeslaender/Uebersicht\\_node.html](http://www.auswaertiges-amt.de/DE/Infoservice/Terminologie/Bundeslaender/Uebersicht_node.html)]. The study revealed four large and growing transmission networks endogenous to the region. Two of these afflict men having sex with men (MSM), one encompasses heterosexuals, and one plagues intravenous drug users (IVDU). The first three networks have loosely connected structures which each span the entire region. These two factors imply that raising and maintaining broad-based awareness would be an effective control strategy. The IVDU network has a densely connected structure and is strongly tied to Cologne, suggesting that focused intervention in that city could have significant benefit to the entire region.

The study is now being expanded to an EU-wide scale. The first results show that simple distance thresholds yield networks substantially similar to those resulting from more complex ancestor graph reconstructions. The speed, simplicity, and interpretability of threshold methods makes them amenable to partitioning large datasets into epidemiologically relevant subunits. These individual partitions are more tractable for investigation via more sophisticated methods which estimate variation in evolutionary rates or transmission probabilities. They also provide a basis for investigating intra-country migratory flows and country-specific differences in transmission dynamics.

### Modeling disease spread

Infection with HIV is the result of intimate personal contact. This suggests that transmission dynamics should mirror other forms of social exchange. One example is rumor spreading, a topic investigated by the Algorithms and Complexity department at the Max Planck Institute for Informatics. We have joined with them in developing a model of competitive rumor spreading on certain random graphs. In agreement with established theory in network epidemiology, the model shows that an unopposed (false) rumor will rapidly cover the entire network. Our work shows that the best way to fight false information is not via training highly knowledgeable experts, but rather by rapid spread of general and easily comprehended truths.



**Transmission network within the heterosexual community.** Coloring nodes by the year they were sequenced shows that the earliest samples form the core of an expanding network.

### Quasispecies evolution

A viral infection inside a patient manifests itself in thousands of slight variations of one or several basic genomic sequences. It forms a quasispecies. Sequencing technology now allows a glimpse into the composition of the quasispecies. Our research investigates how the population structure relates to the evolutionary course of the infection. Drug resistance, for example, would first emerge as a minority variant of the virus. In addition to work with our local cooperation partners on HIV and Hepatitis C, the research has found application in the study of dengue fever, in cooperation with the Genome Institute of Singapore. :::

### CONTACT

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




## Chasing EHEC with the Computer

Human pathogenic *Escherichia coli*, such as Enterohemorrhagic *E. coli* (EHEC) and Enteroaggregative *E. coli* (EAEC), are globally widespread bacteria. They have gained significance as a serious public health problem. Some strains may cause hemolytic uremic syndromes (HUS). Varying strains cause epidemics all over the world. Last summer, we observed an epidemic outbreak of a multi-resistant EHEC strain in Western Europe, mainly in Germany. The Robert Koch Institute reported more than 4,300 infections and over 50 deaths. Farmers lost several million Euros. Although we have finally arrived in the post-genome era, we still know very little about the function of many genes and their interplay. Here, we contribute to the currently ongoing research with a computer-aided study of EHEC transcriptional regulatory interactions, a network of genetic switches that control, for instance, pathogenicity, survival and reproduction of bacterial cells [see Figure].

Characterizing these genetic switch mechanisms remains expensive and very time-consuming. In consequence, we have very limited knowledge about important post-genome follow-up questions that are crucial for understanding bacterial survival and infection strategies. Particularly striking is the direct juxtaposition of hundreds of genome sequences on a handful of reference databases dedicated to gene regulatory networks. In fact, more than 1,600 microbial whole-genome DNA sequences are available. To date, 37 completely sequenced eukaryotic genomes are available; 480 more projects are in the assembly stage, and 700 more projects are registered to be in progress.

We tackle this problem of gene regulatory network reconstruction by identifying evolutionarily conserved regulatory sites within the DNA of *E. coli* K-12, a harmless strain mainly used for wet lab studies, and 16 human pathogenic *E. coli* bacteria. Just a few genes make EHEC extremely dangerous for humans. If it were not for these genes, EHEC would hardly differ from harmless en-



**The molecular genetic program of EHEC's pathogenicity and toxicity.** Illustrated is the interplay of those genes and proteins that are associated with EHEC's pathogenicity or toxicity capabilities. Elements colored in blue are conserved in *E. coli* K-12, the harmless wet lab strain. For those genes, our EhecRegNet prediction pipeline identified evolutionarily conserved transcriptional gene regulatory interactions, i.e. novel wet lab targets.

teric bacteria. We exploit this similarity to find starting points for drugs against the EHEC pathogen. In order to provide high-potential candidate transcription factors that regulate pathogenicity, we developed the integrated online database and analysis platform EhecRegNet. We utilize 3,489 known regulations from *E. coli* K-12 for predictions of yet – unknown gene regulatory interactions in 16 human pathogens. For these strains we predict 40,913 regulatory interactions. Our computer-aided interspecies network transfer approach identified roughly 68% of the *E. coli* K-12 regulatory interactions as conserved in the 16 pathogens.

We discuss high-potential wet lab targets, i.e. potential transcription factors, for genes known to play an important role in pathogenicity or toxicity [see Figure]. Identifying and characterizing EHEC's genetic control mechanism network on a large scale will allow for a better understanding of its survival and infection strategies. This will support the development of urgently needed new treatments.

All our results as well as further bioinformatics analysis methods are publicly available through the EhecRegNet database and analysis platform at <http://www.ehecregnet.de>.

### CONTACT

Jan Baumbach

DEPT. 3 Computational Biology and Applied Algorithmics

Phone +49 681 302-70880

Email [jbaumbac@mpi-inf.mpg.de](mailto:jbaumbac@mpi-inf.mpg.de)

Internet <http://csb.mpi-inf.mpg.de>





# G U A R A N T E E S

**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 the 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. Or, the developers may not have considered an infrequent but possible border case, and software that has worked reliably for years can suddenly show a wrong behavior. 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"*, 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.

Nature is, in many ways, a model for guaranteed correct and fast calculations. For example, the slime mold *Physarum polycephalum* calculates the shortest path from the entrance to the exit of a maze. The article *"Physarum*

*can Compute Shortest Paths – A Rigorous Analysis"* argues that the behavior observed in the mold can even be proved through the use of appropriate mathematical models of computer science.

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 *"Geometric Algorithms – Exact and Efficient"* 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"*, 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"*. Ultimately, it is not yet possible to give all desired guarantees of an arbitrary complex program within an acceptable period of time. 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"*. ...

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

UNDERSTANDING  
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH  
& DIGITAL KNOWLEDGE

MULTIMODAL  
INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

Physarum Can Compute Shortest Paths – a Rigorous Analysis .....	50
Geometric Algorithms – Exact and Efficient .....	51
Dealing with Selfishness in Optimization .....	52
Automated Deduction .....	53
Model Checking for Hybrid Systems .....	54
Quantifier Elimination – Statements Can Also Be Calculated .....	55

## Physarum Can Compute Shortest Paths – a Rigorous Analysis




Figure 1: (a) shows the maze uniformly covered by Physarum; the yellow color indicates the presence of Physarum. Food (oatmeal) is provided at the locations labelled AG. After a while, the mold retracts to the shortest path connecting the food sources as shown in (b) and (c). (d) shows the underlying abstract graph. From Nature, 407:470, 2000.

*Physarum Polycephalum* is a slime mold that is apparently able to solve shortest path problems. Nakagaki, Yamada and Tóth (Nature, Volume 407, 2000) report on the following experiment [see Figure 1]: They built a maze, covered it with pieces of Physarum (the slime can be cut into pieces that will reunite if brought into each other's vicinity), and then fed the slime with oatmeal at two locations. After a few hours, the slime retracted to a path that follows the shortest path connecting the food sources in the maze. The video shows the mold in action [www.youtube.com/watch?v=tL02n3YMcXw&t=4m43s].

Tero, Kobayashi and Nakagaki (Journal of Theoretical Biology) propose a mathematical model for the behavior of the slime and argue extensively that the model is adequate. Physarum is modeled as an electrical network with time varying resistors [see Figure 2]. The authors implemented the model as a computer simulation and report that the network always converges to the shortest  $s_0$ - $s_1$  path, i.e., the diameters of the edges on the shortest path converge to one, and the diameters on the edges outside the shortest path converge to zero. This holds true for any initial condition and assumes the uniqueness of the shortest path. Miyaji and Ohnishi (International Journal of Pure and Applied Mathematics, 47:353–369, 2008) initiated the analytical investigation of the model. They argued convergence against the shortest path if  $G$  is a planar graph and  $s_0$  and  $s_1$  lie on the same face in some embedding of  $G$ .

Recently a team consisting of three members of the Max Planck Institute for Informatics (Vincenzo Bonifaci, Kurt Mehlhorn and Girish Varma) show convergence for all graphs.

Let  $G = (V, E, s_0, s_1, L)$  be an undirected network with positive length function  $(L_e)_{e \in E}$ . Let  $D_e(0) > 0$  be the diameter of edge  $e$  at time zero. The dynamics (1) are attracted to  $\mathcal{E}^*$ , the set of source-sink flows of value one in the graph of all shortest source-sink paths. If the shortest source-sink path is unique, the dynamics converge to the shortest source-sink path.

A key part of our proof is to show that the function

$$V = \frac{1}{\min_{S \in \mathcal{C}} C_S} \sum_{e \in E} L_e D_e + (C_{\{s_0\}} - 1)^2$$

decreases along all trajectories that start in a non-equilibrium configuration. Here,  $\mathcal{C}$  is the set of all  $s_0$ - $s_1$  cuts, i.e., the set of all  $S \subseteq V$  with  $s_0 \in S$  and  $s_1 \notin S$ ;  $(C_S) = \sum_{e \in \delta(S)} D_e$  is the capacity of the cut  $S$  when the capacity of edge  $e$  is set to  $D_e$ ; and  $\min_{S \in \mathcal{C}} C_S$  (also abbreviated by  $C$ ) is the capacity of the minimum cut. The first term in the definition of  $V$  is the normal-

ized hardware cost; for any edge, the product of its length and its diameter may be interpreted as the hardware cost of the edge; the normalization is by the capacity of the minimum cut.

### Discussion and Open Problems

Physarum may be seen as an example of a natural computer. It can apparently do more than compute shortest paths and solve transportation problems. The computational capabilities of Physarum are applied to network design, and it is shown in lab and computer experiments that Physarum can compute approximate Steiner trees.

Traditionally, the analysis of dynamical systems belonged to the domain of biology, systems theory, and physics. Computer science brings new tools. For example, in our analysis, we crucially use the max-flow min-cut theorem. ...

**The Physarum Model:**  $G = (N, E)$  is an undirected graph with distinguished nodes  $s_0$  and  $s_1$ , which model the food sources. Each edge  $e \in E$  has a positive length  $L_e$  and a positive diameter  $D_e(t)$ ;  $L_e$  is fixed, but  $D_e(t)$  is a function of time. The resistance  $R_e(t)$  of  $e$  is  $R_e(t) = L_e / D_e(t)$ . A current of value 1 is forced from  $s_0$  to  $s_1$ . Let  $Q_e(t)$  be the resulting current over edge  $e = (u, v)$ , where  $(u, v)$  is an arbitrary orientation of  $e$ . The diameter of any edge  $e$  evolves according to the equation

$$\dot{D}_e(t) = |Q_e(t)| - D_e(t). \quad (1)$$

where  $\dot{D}_e$  is the derivative of  $D_e$  with respect to time. In equilibrium, the flow through any edge is equal to its diameter. In non-equilibrium, the diameter grows or shrinks if the absolute value of the flow is larger or smaller than the diameter, respectively.

Figure 2: The mathematical model



### CONTACT

Kurt Mehlhorn

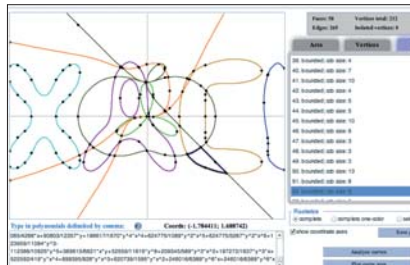
DEPT. 1 Algorithms and Complexity

Phone +49 681 9325-1000

Email melhorn@mpi-inf.mpg.de




## Geometric Algorithms – Exact and Efficient



Geometric algorithms play an important role in solving many everyday problems. A classic example is the so-called piano movers problem, which involves, simply put, using calculations to decide whether or not a certain piano would fit through the narrow staircase the next time you move house. Another important application area is CAD (Computer Aided Design), where methods are developed to design and visualize engineering components. Difficulties frequently occur when geometric algorithms are implemented, arising from wrong decisions within the algorithm due to rounding errors. False results, crashes, or non-terminating programs can be the result.


The Computer Algebra and Geometric Computing Group at the Max Planck Institute for Informatics aims to provide possible solutions for the above-mentioned problems. Here, we primarily deal with an algorithmic framework which eventually constitute the bases for specialized solution methods. Such basic operations mainly include Boolean operations (intersection, union, etc.) and arrangement computations of geometric objects, that is, an explicit description of the decomposition of a plane or space implicitly defined by the given objects.

The algorithms that we develop are designed and implemented in such a way that all possible inputs can be exactly processed. Unfortunately, algorithms with this property often come together with a series of disadvantages. Firstly, the complexity of the algorithm increases and, what is much worse, one must accept a significant increase of computation time in order guarantee completeness and exactness. The reason for this is that



solutions to general polynomial systems must be found and further processed to handle curved geometric objects, and the corresponding operations are computationally intensive. For the latter reason, commercial systems completely rely upon numerical methods and thus show the above mentioned erratic behavior as confirmed by experiments.

If exactness and completeness should be guaranteed, certain symbolic operations, such as the computation of the greatest common divisor or the resultant of two polynomials, are unavoidable. Hence, for the efficiency of an exact algorithm, it is crucial that such symbolic operations are implemented in an efficient manner, and, in particular, that they are used sparsely. For the implementation of the symbolic operations, we mainly consider modular divide-and-conquer methods from Computer Algebra. Here, the calculations are broken down into many smaller units with significantly lower complexity, and the individual results are ultimately reassembled to form a complete solution. Using massively parallel hardware architecture, as it is found, for example, on graphics cards (GPU), we have been able to accelerate important symbolic calculations by a hundredfold. We further developed novel algorithms for computing arrangements, which use very few symbolic operations. In addition, we have outsourced all of the latter operations to the GPU. This was achieved by combining techniques from



**Arrangements of curves in the plane and on a Dupin-Cyclide. The right figure shows a topological exact triangulation of an algebraic surface.**

different fields such as Computer Algebra, numerical analysis, and algebraic geometry. As a result, we have been able to improve the efficiency of exact methods by several magnitudes compared to existing approaches. In particular, for the first time, we can report that our exact methods can compete with purely numerical approaches that do not come with any additional guarantees.

Finally, we also concentrate on more fundamental problems such as finding the solutions of a polynomial system. In this context, we have developed novel algorithms that are simple and practical and, at the same time, either match or improve the best theoretical complexity bounds known so far for the corresponding problems. We consider this contribution to be fundamental, in particular, since many of the existing theoretically efficient methods from the areas of Computer Algebra and Numerical Analysis have not turned out to be very practically relevant so far. ...

### CONTACT

**Michael Sagraloff**

**DEPT. 1 Algorithms and Complexity**

Phone +49 681 9325-1006

Email [msagralo@mpi-inf.mpg.de](mailto:msagralo@mpi-inf.mpg.de)


Internet <http://www.mpi-inf.mpg.de/departments/d1/areas/software.htm>



## Dealing with Selfishness in Optimization

In my research, I focus on optimization problems that involve selfishness in some way: for example, scheduling problems where the machines are controlled by selfish agents, each of which knows only how fast its own machine is. The agents are only interested in maximizing their own utility, which may be quite different from the overall optimization goal. Hence we need to design mechanisms that ensure that truthtelling is a dominant strategy for the agents.

A problem that I recently studied in this area is maximizing the minimum load on selfish machines, i.e., optimizing the max-min fairness of the job allocation. This is a standard fairness measure in the area of fair division. Such an allocation is relevant, for instance, in the scenario where the jobs represent hard drives of various capacities, and we want to create a backup medium consisting of  $m$  parallel storages of maximum possible capacity; this capacity is then bounded by the size of the smallest storage. The utility of an agent is the payoff that it receives from the mechanism minus the cost of processing the jobs assigned to it. This cost depends on its speed, which only each agent knows which is known only to each individual agent. We want the agents to give their true speeds as input to the mechanism, since otherwise we cannot hope to optimize any objective function. It is known that if the allocation done by the mechanism is monotone, then telling the truth is a dominant strategy for the agents. Monotone means that if one agent increases its bid (i.e. its claimed cost for running one unit of work), the total size of the jobs assigned to it by the mechanism may not increase as a result (the set of assigned jobs does not necessarily have to remain the same, however).



Together with Giorgos Christodoulou and Annamaria Kovacs, I designed the first monotone approximation algorithm with a constant approximation ratio. Very recently, I have further improved this result with Leah Epstein and Asaf Levin, and achieved an arbitrarily good approximation which runs in polynomial time.

Another area that I am working on right now is that of selfish routing, in particular on rings. 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). In general graphs, research has focused on the symmetric setting, where all players want to travel between

the same two endpoints. We consider an asymmetric congestion game, which models the more realistic and difficult setting in which multiple players may have different locations in the network and thus different sets of strategies to choose from.

For the goal function of minimizing the maximum latency of any player, we have shown a tight upper bound of two for the price of anarchy. This means that any Nash equilibrium is within a factor of two of the best possible routing, which could only be achieved by enforcing fixed routes for all the players. See the figure for an example where this bound is tight. Here  $x$  is the number of players on each edge. Note that in the figure on the right, if one player deviates and uses its direct edge, it also has a cost of 2, because the other player also uses this edge. ...



### CONTACT

**Rob van Stee**

**DEPT. 1 Algorithms and Complexity**

Phone +49 681 9325-1005

Email [vanstee@mpi-inf.mpg.de](mailto:vanstee@mpi-inf.mpg.de)

## 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: Obviously 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.

$$x + 0 = x$$

$$x + (-x) = 0$$

$$\frac{x \cdot z}{y \cdot z} = \frac{x}{y}$$

complicated → simple

### 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 to process even very large quantities of data, such as those that occur in the analysis of ontologies (knowledge bases). 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(skf6(U),V),U)* .
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

## 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 ( $b_1, \dots, b_{60}$ ), 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

$$b_5 \text{ and not } b_{32}$$

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:

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

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



## 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  $x_1, \dots, x_{11}$ . 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](mailto: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. ...





<b>Random Telephone Chains – Efficient Communication in Data Networks</b>	58
<b>Search and Mining in Web Archives</b>	59
<b>Information Search in Social Networks</b>	60
<b>EnBlogue – What is New and Interesting in Web 2.0 ?</b>	61
<b>YAGO – a Collection of Digital Knowledge</b>	62
<b>URDF – Efficient Reasoning in Uncertain RDF Knowledge Bases</b>	63
<b>PICASSO – Soundtrack Recommendation for Images</b>	64

UNDERSTANDING  
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH  
& DIGITAL KNOWLEDGE

MULTIMODAL  
INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

## Random Telephone Chains – Efficient Communication in Data Networks

How do you send an urgent message quickly and efficiently to a larger group of people? A classic procedure is the telephone chain. In the simplest case, the sequential telephone chain is based on an organized list of group members. Here, spreading news goes to the first person on this list. Then he calls the second person, who then calls the third, and so forth.

Similar task descriptions appear in informatics if information in computer centers must be distributed. A company with affiliates at several sites can store its customer data locally at each branch office. This has the benefit that in each affiliate one can quickly access this data with little complication, and that the data is also available if the central database is not reachable for a short period of time. When using such replicated databases, there must naturally be changes given in the data inventory of an affiliate, which is slowly communicated to all others.

In two central aspects, this telephone chain problem is different from informatics in the everyday world. First, such database systems can be very large: In a network with several thousand nodes, as in this affiliates example, the information distribution lasts much too long if the nodes are given the news in a fixed series, one after another.

A second aspect is the robustness of the procedure. Robustness means that a process can still function properly even if some partial steps are not error-free. In a telephone chain, the main problem might be, that a participant cannot be reached, or that he is reachable, but for some reason, he does not forward the news. The classic sequential telephone chain appears not to be very robust. If a participant does not pass on the news, all the following participants in the chain remain uninformed.

For this reason, processes based on classic telephone chains are inappropriate for application in informatics. However, there are telephone chain problems in informatics with a surprisingly simple solution. With the *random telephone chain* each caller who has received the news, calls random other participants. To simplify the presentation, one can accept that all calls last the same amount of time. This leads to the information exchange taking place in rounds. In each round, each informed participant calls another randomly chosen participant. At this latest point in time, this chosen participant is also informed.

The random telephone chain is amazingly efficient and robust. If we take an example of a data network with 1,024 nodes which can all communicate with one another, then the random telephone chain needs on average only 18.09 rounds in order to inform all nodes. The same is true for the robustness; even if we accept that every 10th participant is never reachable, it takes only 19.49 rounds on average to inform the remaining members. Therefore, it is not important which of the members fall within this 10%.

The good performance of the randomised telephone chain has stimulated other research efforts. A highly interesting question is: What is the optimal amount of randomness? The first results demonstrate that one sole random decision per participant gives results that are just as good as a completely random telephone chain. Actually, the optimum lies between the two extremes. The following very simple protocol demonstrates this.

We assume that all participants have one (e.g., alphabetically) organized list of all participants. If a participant is informed, he chooses another random position on the list and calls, beginning at this position, all the other participants. If he reaches someone who is already informed, he dials a new, random starting position and continues to call in order. This simple protocol informs 1,024 nodes in only 14.11 rounds on average. Even if (randomly dialed) 10% of the participants are never reachable, only 15.09 rounds are needed on average.

A further benefit of this procedure is that we can give a termination criterion. We can establish, for example, that participants leave the procedure if they have called an informed node a fourth time. This is the first time that a termination criterion has been given for a randomized telephone chain, and it is guaranteed thereafter that all participants will be reached. Furthermore, it provides a very efficient procedure – on average, a participant makes only 5 calls. The recommended termination criterion leads in any case to a minimal increase in running time. In this case, for example, we need only 14.14 rounds in order to inform 1,024 nodes. Apparently, there is a trade-off situation: The earlier the participants stop, the longer the procedure takes. Even if the participants leave immediately after they have called an informed node, the new procedure is slightly better than the fully random telephone chain and needs 17.63 rounds instead of 18.09. ...



### CONTACT

**Benjamin Doerr**

**DEPT. 1 Algorithms and Complexity**

Phone +49 681 9325-1004

Email [doerr@mpi-inf.mpg.de](mailto:doerr@mpi-inf.mpg.de)

## 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 9325-5018

Email kberber@mpi-inf.mpg.de

## Information Search in Social Networks

The advent of Web 2.0 has revolutionized the way users interact with the Web. Instead of simply consuming information as before, now anyone can simply generate and publish contents. Online services such as del.icio.us, Flickr, LibraryThing, YouTube, MySpace, or Facebook offer a simple way to publish pictures, videos, bookmarks and similar data on the Web. In addition, they allow the user to interact with other users, enabling them, for example, to share content. Data published by other users can be explored, commented, assessed, and annotated with so-called tags. These tags are often very good content descriptions, because they can be freely chosen without having to stick to a predefined set. Various users frequently annotate the same image or the same video with different tags, reflecting their different interests. Most services offer easy to use and intuitive interfaces in order to find contents based on annotations, for example, the so-called “tag clouds”.

In addition, users can maintain explicit lists of their friends, and the system often provides additional value to such followers. It is easily possible to look for newly added content, new assessments or, especially lately, to inform friends about your current location. In this way, a social network is built, a dense network of relationships, in which the number

of friends is often taken as an indicator of one's reputation on the network. While these lists of friends are initially filled with friends and other people known from “real” life using the same service, it grows over time to include previously unknown users with similar interests.

This dense user network, in combination with the user-defined annotations, allows to find valuable content recommended by friends (either directly or through transitive friends of friends), exploiting the “wisdom of the crowds”. Such a recommendation can either be explicit (e.g., through judgments and comments) or implicit (e.g., through intensive use of annotations). When searching in such systems, relations and the distance between users should therefore be taken into consideration, as close friends usually are more trusted than remote acquaintances. Some existing social networks already offer such functions, but only in a very limited fashion (for example, only for direct, but not for transitive friends); others do not support neighborhood search at all.

We have developed a search engine coined SENSE (for “*Socially ENhanced Search and Exploration*”) to close this gap. It provides tag-based search for content among the friends of the querying user, weighting annotations based on the

strength of relationship to the friend.

In addition to the distance in the social network, similarity in interests (derived from similarity of annotations) has influence in weights. In comparison to most existing systems that consider plain tag frequency, SENSE retrieves results with significantly better quality for such personalized searches. Our system uses highly efficient and scalable search algorithms to cope with the massive amount of data present in large social networks. The fast growth of these services, and especially the very high rate with which new contents and annotations are generated, requires optimized indexing strategies in order to keep the precomputed transitive user relationships always up to date. SENSE defers updates as long as possible to the point when up-to-date information is needed for query processing; then, updates are performed in a pay-as-you-go manner while processing the query. ...



### CONTACT

**Tom Crecelius**

**DEPT. 5 Databases and Information Systems**

Phone +49 681 9325-5006

Email tcrecel@mpi-inf.mpg.de



**Ralf Schenkel**

**DEPT. 5 Databases and Information Systems**

Phone +49 681 302-70798

Email schenkel@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 volcanic eruption of Eyjafjallajökull in the spring of 2010.



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. ...



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

### CONTACT

Sebastian Michel

DEPT. 5 Databases and Information Systems

Phone +49 681 302 70803

Email smichel@mpi-inf.mpg.de

Internet <http://blogue.mmci.uni-saarland.de>





## YAGO – a Collection of Digital Knowledge

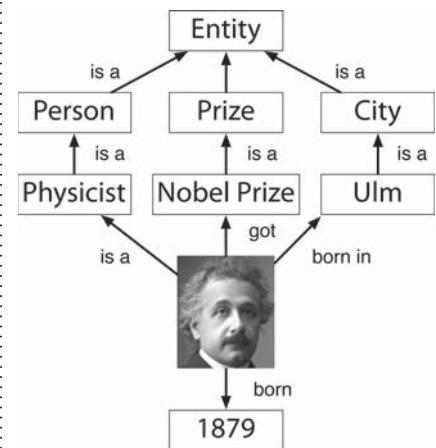
In recent years, the Internet has developed into the most significant source of information we have available today. Train schedules, news, and even entire encyclopedias are available online. Using search engines, we can very efficiently query this information. But, despite efficiency, search engines face limitations in terms of effectiveness, in particular for complicated queries. 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 “*politicians scientists*” 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, for that matter, even understand it. If it would be 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 a 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 major building blocks of an ontology are *entities* and *relations* among 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, into 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. In fact, the great majority of the approximately 900,000 person-entities in YAGO2 are anchored in time by their birth and death dates, allowing us to position them in their historical context. For example, one can ask questions about important historical events during the lifetime of a specific president, emperor or pope, or also ask the question of when the person of interest was actually president.

Most of the approximately 7 million locations in YAGO2 have geographic coordinates which place them on the earth’s surface. Thus, spatial proximity between two locations can be used as search criterion. An exemplary search using the space and time criteria could be: Which 20th century scientists were awarded a Nobel Prize and were born in the vicinity of Stuttgart. In YAGO2, one finds, among others, Albert Einstein, since both his lifetime (1879–1955) and his birthplace Ulm (70 kilometers from Stuttgart) are stored in YAGO2.



In our next article about URDF, we describe how existing knowledge in YAGO2 can also be used to deduce new knowledge and then how to reason with this knowledge under different uncertainty models. ...



### CONTACT

**Johannes Hoffart**

**DEPT. 5 Databases and Information Systems**

Phone +49 681 9325-5028

Email [jhoffart@mpi-inf.mpg.de](mailto:jhoffart@mpi-inf.mpg.de)



**Fabian Suchanek**

**DEPT. 5 Databases and Information Systems**

Phone +49 681 9325-5000

Email [suchanek@mpi-inf.mpg.de](mailto:suchanek@mpi-inf.mpg.de)



**Gerhard Weikum**

**DEPT. 5 Databases and Information Systems**

Phone +49 681 9325-5000

Email [weikum@mpi-inf.mpg.de](mailto: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:

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

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”:

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

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

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

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  $\text{bornIn}(\text{Angela\_Merkel}, \text{Hamburg})$  and  $\text{bornIn}(\text{Angela\_Merkel}, \text{Bremen})$  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 9325-5007

Email [martin.theobald@mpi-inf.mpg.de](mailto:martin.theobald@mpi-inf.mpg.de)

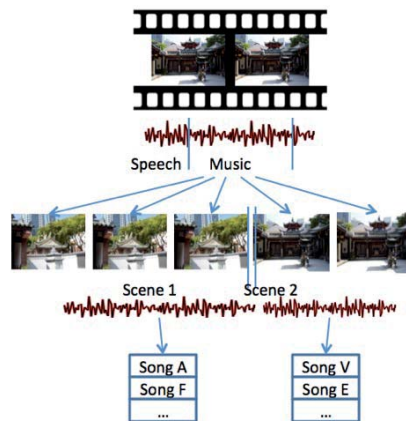
## PICASSO – Soundtrack Recommendation for Images

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.

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. ...



**Figure 2: Screenshot of the iPhone App PicasSound**



### CONTACT

**Sebastian Michel**

**DEPT. 5 Databases and Information Systems**

Phone +49 681 302-70803

Email [smichel@mmci.uni-saarland.de](mailto:smichel@mmci.uni-saarland.de)

**Aleksandar Stupar**

**DEPT. 5 Databases and Information Systems**

Phone +49 681 302-70807

Email [astupar@mmci.uni-saarland.de](mailto:astupar@mmci.uni-saarland.de)

Internet <http://picasso.mmci.uni-saarland.de/demo/>



## MULTIMODAL INFORMATION

**The volume and importance of multimodal data has virtually exploded in recent years. Research on handling such data is thus becoming ever more important. This extends from the modeling and indexing to the understanding of multimodal sensor data. The Max Planck Institute for Informatics accommodates this development and is pursuing some of the most important issues arising from this area.**

Sensors such as cameras, microphones, GPS, and accelerometers are increasingly embedded in equipment and environments and are already useful in a variety of ways. The 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 with access to this sensor information do not fully interpret it and thus cannot truly understand their environments.

There are two work areas that address different problems involved in multimodal sensor processing. The first area is concerned with the automatic processing of music data. A central topic here is the efficient finding of individual audio segments or other music fragments in large music data collections. The processes being researched are not only useful for processing search queries, but also for content-based music analysis. This involves, among other things, the recognition and linking of semantic relationships in various versions of different modalities of a piece of music.


In a second area, we are researching the use of portable sensors already embedded in cell phones or even in clothing. In this case, the focus is on capturing and understanding the context of a person playing an important role in man-computer interaction. Context awareness may enable natural communication, for example, with robots that understand the users goals and offer support at the right time. The work described here concentrates on a specific type of context, namely, the recognition of human activity. ...



# CONTRIBUTIONS

**Recognizing Human Activity** ..... 68

**Knowledge-Based Source Separation for Music Signals** ..... 69



UNDERSTANDING  
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH  
& DIGITAL KNOWLEDGE

MULTIMODAL  
INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

## 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 already 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




Figure 1: Relevant parts for composite activity

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.

### 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




Figure 2: Composite activity for construction

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.



### 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

## Knowledge-Based Source Separation for Music Signals

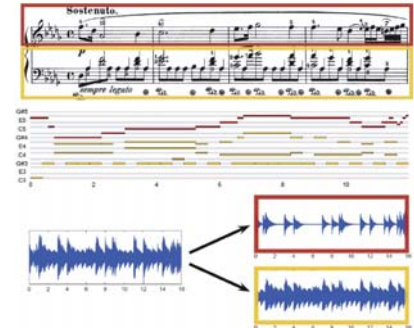
The decomposition of superimposed acoustic sound sources into its constituent components, a task also known as source separation, is one of the central research topics in digital audio signal processing. For example, in speech signal processing, an important task is to separate the voice of a specific speaker from a mixture of conversations with multiple speakers and background noises ("Cocktail party scenario").

Also, in the field of musical signal processing, there are many related issues commonly subsumed under the notion of source separation. In the musical context, sources often correspond to melodies or instrumental voices, which are overlayed to form complex polyphonic sound mixtures. Many source separation methods rely on assumptions such as the existence of multiple channels (e. g., multiple microphone signals recorded from different directions) and some kind of statistical independence of the source signals. In the case of music signals, however, these assumptions are often violated. On the one hand, music recordings are mostly given in mono (one channel) or stereo (two channels). On the other hand, the various musical sources are typically far from being independent. Actually, quite the opposite is true: musicians interact with each other, and instrumental voices follow the same rhythmic patterns, melodic lines, and harmonic progressions. This makes the separation of musical voices from a polyphonic sound mixture an extremely difficult and generally intractable problem.

In decomposing a music signal, one strategy is to exploit music-specific properties and additional musical

knowledge. For example, the separation problem can be alleviated by exploiting the fact that the melody is often the leading voice, characterized by its dominance in dynamics and by its temporal continuity. Or, the track of a bass guitar may be extracted by specifically looking at the lower part of the frequency spectrum. Furthermore, when extracting the drum track, one often relies on the assumption that the other sources are of harmonic nature and can exploit that percussive elements (*vertical spectral structures*) are fundamentally different from harmonic elements (*horizontal spectral structures*). Last but not least, a human singing voice can often be distinguished from other musical sources because of the presence of vibrato and portamento (sliding voice) effects.

In the last years, multimodal, score-informed source separation strategies have also been employed assuming the availability of a score representation along with the music recording. Recent approaches show that the extraction of musical voices from highly overlapping sound sources comes within reach (at least to a certain degree) when exploiting additional musical cues specified by the score. Additionally, the score representation allows a user to conveniently specify the musical voices to be separated. However, to transfer the score information to the audio domain, one requires automated synchronization techniques to establish accurate correspondences between the note events given by the score and time positions in the music recording. The development of efficient and robust synchronization techniques constitutes, in itself, a challenging research direction. ...



**Score-informed source separation strategy.** The score (upper row) is used to specify the voices to be separated (here, the left and right hand of Chopin's Prélude No. 15 "Raindrop"). An audio recording of the same piece is then decomposed into these voices (bottom row). The score-audio synchronization is accomplished by using a MIDI representation as a connective mid-level representation (middle row).



**User interface for specifying note groups to be attenuated, intensified, or modified in corresponding music recordings.**



### CONTACT

**Meinard Müller**

**DEPT. 4 Computer Graphics**

Phone +49 681 9325-4005

Email [meinard@mpi-inf.mpg.de](mailto:meinard@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 optimization 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

<b>Approximation Algorithms for Hard Optimization Problems</b> .....	72
<b>The Theory of Bio-inspired Algorithms</b> .....	73
<b>The Traveling Salesman Problem</b> .....	74
<b>Rule-based Product Configuration</b> .....	75



UNDERSTANDING  
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH  
& DIGITAL KNOWLEDGE

MULTIMODAL  
INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION




## Approximation Algorithms for Hard Optimization Problems

One of the most basic methods for dealing with NP-hard optimization problems is to design polynomial-time algorithms that find a provably near-optimal solution. 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  $x$  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 give an example below.

### Matching of live cell video frames

In joint work with Stefan Canzar, Gunnar W. Klau and Julian Mestre, we have considered a problem that arises naturally in the computational analysis of live cell video data. Studying cell motility using live cell video data helps understand important biological processes such as tissue repair, the analysis of drug performance and immune system responses. Segmentation based methods for cell tracking typically follow a two stage approach: The goal of the first *detection* step is to identify individual cells in each frame of the video independently. In a second step, the linkage of consecutive frames, and thus the tracking of a cell, is achieved by assigning cells identified in one frame to cells identified in the next frame. However, limited contrast and noise in the video sequence often leads to over-segmentation in the first stage: a single cell is comprised of several segments. A major challenge in this application domain is therefore the ability to distinguish biological cell division from over-segmentation.



Example of a feasible matching between two segment sets in two frames

One approach to solve this problem is to match *sets of segments* between neighboring frames, where the segment sets are obtained by a certain *hierarchical* clustering procedure. See the example, which illustrates the hierarchical clustering of two consecutive frames. Matching here means identifying one set of segments at a certain level in the hierarchy in one frame with one set of segments in the other frame. Since segment sets representing different cells in the same frame must be disjoint, no two nodes on any root-to-leaf path in any of the two hierarchies can be matched at the same time. In the figure, we show a feasible matching using dashed lines. Typically, a non-negative weight is assigned to each pair of segment sets (one from the first frame and the other from the second), indicating how likely this pair is to represent the same cell. The objective is to find a feasible matching between the segment sets in the two frames, with maximum total weight.

While the problem was originally formulated in the Computational Biology community, our contribution was to study it from an approximation-algorithms point of view. We first showed that solving the problem exactly is NP-hard; in fact, we showed that there is a constant value  $x > 1$  such that it is even NP-hard to obtain an  $x$ -approximation. On the other hand, we developed a 2-approximation algorithm for that problem. ...



#### CONTACT

Khaled Elbassioni


DEPT. 1 Algorithms and Complexity

Phone +49 681 9325-1007

Email elbassio@mpi-inf.mpg.de

## The Theory of Bio-inspired Algorithms

Bio-inspired algorithms are general search algorithms which are used in many ways in engineering disciplines and in the area of combinatorial optimization. An example of such algorithms are evolutionary algorithms (EA). This classical solution process follows from the evolutionary model and the Darwinian principle of "survival of the fittest." Based on this natural evolution principle, a single solution candidate is called an *individual* and a number of such candidates is called the *population*. A *fitness function*, which depends upon the given problems, assesses the solution candidates. The offspring population can be produced from a parent population according to the biological principle. This occurs through so-called *variation operators*, in which the parents' genetic material is inherited by the children. The most important operators are, in this case, crossover and mutation. Crossover usually produces a child from two adults, while mutation provides the child with new properties [see Figure]. Beginning with a starting population, the goal for a given problem is to retain many possibly-good solution candidates. After having produced the offspring, a new parent population is created by selecting individuals from the parents and offspring. This process is repeated until a good solution is found.



Crossover

Another class of bio-inspired algorithms is ant colony optimization, which is modeled on the pathfinding and communication mechanisms of ants. These algorithms construct many possible solutions for a given problem from various components; components that participate in "good" solutions will (using a material messenger) produce so-called *pheromones* and will be used in later solutions with higher probability. We thereby hope to develop an optimal solution as a combination of many good components.

Bio-inspired algorithms can be especially useful if, for a given (new) problem, no good, problem-specific algorithm is available. Thus, they cannot be expected to surpass a specially-designed solution process. It is therefore not the goal of the research to demonstrate that bio-inspired algorithms surpass problem-specific algorithms. Rather, the focus is to understand the principles by which this new process operates.

### Research focus

While bio-inspired processes have already been successfully employed in many cases, the theoretical understanding of these algorithms in comparison to classical algorithms is still at an early stage.

We are examining whether bio-inspired search processes are capable of solving specific problems, and what structures are advantageous or disadvantageous in their use. The main focus is on how much time the algorithms need

to generate an optimal solution for a given problem. Because bio-inspired algorithms are a special class of randomized algorithms, one can use a large number of classical analytical methods. Further, new analytical methods are being developed in order to respond to the specific characteristics of bio-inspired methods.

We have shown that bio-inspired methods often find good solutions for well-known problems without using problem-specific knowledge. They are able to mimic problem-specific algorithms for a large class of combinatorial problems. For example, it has been formally demonstrated that bio-inspired algorithms can efficiently calculate the shortest path between all nodes in a given graph. Another example is that EAs with recombination and mutation show demonstrable benefits in comparison to EAs that use only mutation as a variation operator.

Other studies show that approaches for multi-criteria optimization of bio-inspired algorithms bring additional benefits to efficient search. Many optimization problems are framed in terms of a target function which must be optimized through many side conditions. These side conditions are perceived as additional, equal target functions in multi-criteria approaches. This provides the search for evolutionary algorithms with new avenues for investigation, which could lead to more efficient processes based on different combinatorial optimization problems. ...



### CONTACT

**Benjamin Doerr**

**DEPT. 1 Algorithms and Complexity**

Phone +49 681 9325-1004

Email doerr@mpi-inf.mpg.de



**Tobias Friedrich**

**DEPT. 1 Algorithms and Complexity**

Phone +49 681 9325-1055


Email tfried@mpi-inf.mpg.de

## The Traveling Salesman Problem

The traveling salesman problem (TSP) is perhaps the most famous problem in combinatorial optimization: given a set of cities and the distances for traveling between each pair of cities, the goal of the problem is to find the shortest tour that visits each city at least once and returns to its starting point. The traveling salesman problem has many applications in logistics and transportation, but also in areas such as chip manufacturing, where the “cities” to be visited are holes that need to be drilled, and the distance between cities is the time needed to move the drill head from one hole to the next.

The traveling salesman problem is an NP-hard problem; this means that the running time of the only known algorithms for optimally solving these problems grows exponentially with the size of the problem (and it is widely believed that this behavior is inherent to the problems). Nevertheless, the past decades have seen great progress in solving real-world instances of the TSP. The starting point for solving real-world TSP instances is to compute a lower bound on the length of the optimal tour, called the “subtour LP” bound. The subtour LP is known to give excellent lower bounds on TSP instances in practice, coming within a percent or two of the length of the optimal tour. However, its theoretical worst case is not well understood.

A famous algorithm called Christofides’ algorithm always finds a tour of length at most  $3/2$  times the value of the subtour LP. This implies that the worst-case ratio of the optimal tour length and the subtour LP bound is at most  $3/2$ . Examples are known that show that the worst-case ratio is at least  $4/3$ , and a long-standing conjecture states that the worst-case ratio, taken over all instances of the problem, is in fact exactly  $4/3$ .



An illustration of the subtour LP bound. The value of the subtour LP is equal to twice the sum of the radii of the red circles, plus twice the width of the blue regions.

Source: <http://www.tsp.gatech.edu/>

Resolving this conjecture would not only satisfy our scientific curiosity (as it is surprising that such a well-studied problem is not better understood) but it could also lead us to the first improvement over the thirty-year-old Christofides’ algorithm.

Together with my co-authors David Williamson and Frans Schalekamp, we recently resolved a related conjecture.

Suppose we allow multiple salesmen, that don’t necessarily start in the same city, but that together cover all the cities. Then, we can prove that the worst-case ratio of the length of the optimal solution to the subtour LP bound is  $10/9$ . In our current research, we are attempting to extend our techniques to answer the main conjecture for the special case of the traveling salesman problem in which all distances are either one or two. ...

### CONTACT

Anke van Zuylen

DEPT. 1 Algorithms and Complexity

Phone +49 681 9325-1035

Email [anke@mpi-inf.mpg.de](mailto:anke@mpi-inf.mpg.de)

Internet <http://www.mpi-inf.mpg.de/departments/d1/areas/optimization.html>



## 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. 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. 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 recording by a rule-based system into 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.

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

## SOFTWARE

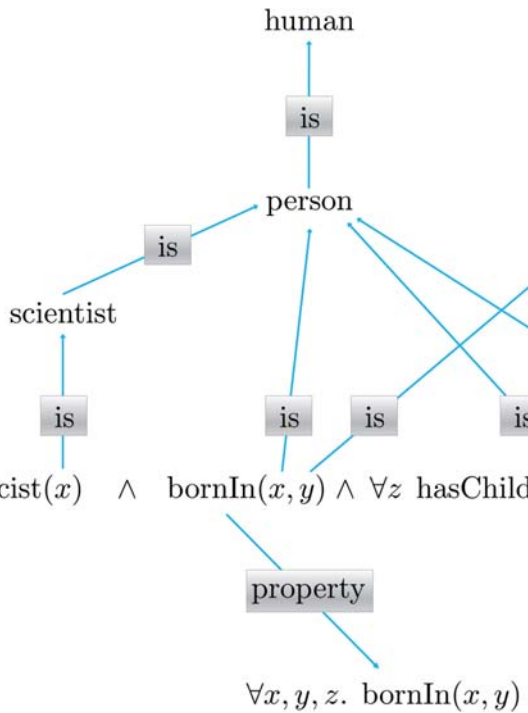
**Informatics is, firstly, a discipline of basic research that deals with universal computation and problem-solving methods and investigates fundamental properties such as correctness and complexity. Secondly, it also resembles an engineering science that supports a great variety of different applications. 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, as mentioned above, 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 engineering if 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. A current research topic is the enhancement of syntactic search possibilities in the World Wide Web (WWW) with semantic information. The articles "*TopX 2.0 – Efficient Search in Digital Libraries*" and "*RDF-3X – Fast Searches for Semantic Data*" describe two software solutions related to this topic. The article "*Decision Procedures for Ontologies*" deals with the semantic structure of everyday knowledge and describes a logic-based software for answering complex search queries in a sound and complete way. ...





**TopX 2.0 – Efficient Search in Digital Libraries** ..... 78

**RDF-3X – Fast Searches for Semantic Data** ..... 79

**Decision Procedures for Ontologies** ..... 80

UNDERSTANDING  
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH  
& DIGITAL KNOWLEDGE

MULTIMODAL  
INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

## TopX 2.0 – Efficient Search in Digital Libraries

### Full text search in semistructured documents

TopX is a search engine for efficient full-text search in semi-structured documents, a data format which is frequently used, for example, in digital libraries. These documents are called semi-structured as they combine rich semantic annotations and structural markup with extensive text passages in a unified format, the popular XML (Extensible Markup Language). In particular semantic annotations are of interest for query evaluation because they can contribute to disambiguating the textual contents and thus to better understanding the textual information. Let us consider, for example, the keyword query “*Who were doctoral students of Max Planck, who themselves became scientists?*”. Assuming that appropriate annotations exist in the collection, this keyword query can be translated into a very specific path query over a corresponding document structure:

```
//article[.//person ftcontains „Max Planck“]
//doctoral_students/scientist
```

(according to XPath 2.0 Full-Text Standard)

The results of this query over a Wikipedia-based, annotated XML corpus, will then be returned by the engine as direct references to “*Gustav Ludwig Hertz*” or “*Erich Kretschmann*” as compact links, for example, to their own encyclopedia entries.

### Exploratory search

The prerequisite for a successful application of this query language is however the user’s detailed knowledge of the document structure, the so-called XML schema. In order to support the user in formulating structured queries over an initially unknown document schema, TopX also allows for an exploratory form of search. Starting with simple keyword queries, the user can employ the search engine to gradually learn more about the document’s schema and then refine the queries step by step via structural conditions. The expressiveness of

the query language (which supports the international XPath 2.0 and XQuery 1.0 Full-Text standards defined by the W3C) allows for a gradual clarification of the requests.

Following the above example, the user may start searching the annotated Wikipedia corpus with a pure keyword query for “*Max Planck*”. This query would provide (amongst others) links to the person Max Planck as well as to the various Max Planck {institutes} among its best matches. Taking advantage of the document structure, a simple refinement of the query might then look like the following:

```
//article[.//person ftcontains „Max Planck“],
```

which would tell the search engine that articles about the *person* “*Max Planck*” would be preferred.

Moreover, different forms of query relaxations and ontological query expansions are supported by TopX. On the one hand, overly specific queries that would not yield any exact matches can be dynamically relaxed at query time by automatically switching from a conjunctive query evaluation mode into a disjunctive evaluation strategy. On the other hand, through ontological query expansions, TopX can also find matches to semantically similar concepts.

### Efficiency and scalability

Much of the implementation of TopX has been continuously improved and in part completely redesigned over

the past years. In particular, the internal index structure used by TopX was transformed from a conventional database system into a completely new, highly optimized index structure, which was specifically designed for TopX. This new index structure supports highly efficient forms of index compression and the distribution of indexes across multiple computers (i.e., compute cluster), which contributes significantly to the scalability and efficiency of the search engine. Our experimental results and the regular, very successful participation in international benchmark competitions support the globally strong position of TopX for information retrieval over semi-structured data that significantly outperforms commercial products.

### Applications and open-source availability

TopX targets both the direct user application as well as the API-based integration with other systems. With the international XQuery 2.0 Full-Text standard, TopX also offers improved support for imperative programming language concepts, such as loops and conditional statements, which increases TopX’s usability with automatically generated queries.

The original prototype of TopX was already made available as an open-source package in 2006. A respective open-source distribution of the current TopX 2.0 version with support for distributed indexing, incremental index updates, and the XQuery 2.0 Full-Text standard is in preparation. ...



#### CONTACT

**Martin Theobald**

**DEPT. 5 Databases and Information Systems**

Phone +49 681 9325-5007

Email martin.theobald@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  $\langle id1, bornOn, 1859 \rangle$ , which express that the subject “*id1*” is associated with the name Arthur Conan Doyle, born in 1859. Further edges, such as  $\langle id1, authorOf, id2 \rangle$ ,  $\langle id1, bornIn, Edinburgh \rangle$  and  $\langle Edinburgh, is\ in, Scotland \rangle$  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 <title> ?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. ...



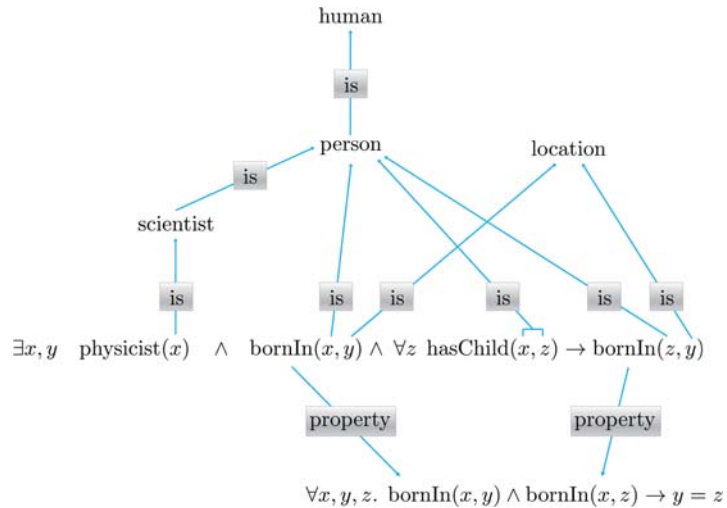
### CONTACT

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

## Decision Procedures for Ontologies

The answers to many questions can already be found with the assistance of search engines on the Internet. However, these search engines perform a purely syntactic search based on keywords. For the query: “Which physicist was born at the place where all his children were born?”, the search engine will extract the keywords “physicist”, “child”, “all” and “born” and return documents where these keywords occur exactly. For the above query, a search engine will return a vast number of documents which contain the extracted keywords, but typically not the answer to the question. A reformulation of the question may lead to the desired result. However, the actual problem is that current search engines do not know the meaning of the information in the documents and in the query.

In order to enable a computer to understand the meaning of words and sentences, additional knowledge about the words must be made explicitly available for the computer. In the above example, this knowledge is the fact that physicists and children are people and they were all born in one specific place. This type of aggregation of general knowledge is called ontology. Certain ontologies can be created automatically. An example of an automatically generated ontology is YAGO (see “YAGO – a Collection of Digital Knowledge”, page 62). Answering quickly all queries that can be answered with the knowledge of an ontology is currently a challenge in research. In order to accomplish this, it is first necessary to convert the ontology into an appropriate language, a logic. Based on the representation in logic, a computer can, with suitable methods, derive everything that is logically implied by the ontology.



In order to effectively answer questions in this way, the ontology is compiled: It is converted into a compact representation of all logical consequences, which is free of contradictions. Based on the compiled representation of an ontology, our procedures efficiently answer all questions that are implied by the original ontology. However, the operations required for the compilation are computationally intensive. The previously known methods cannot compile ontologies consisting of tens of millions of entries in an acceptable time period. The challenge is to adapt the algorithm to the structure of the ontologies in such a way that we obtain an efficient compilation method even for huge ontologies. Currently, our methods compile ontologies with up to ten million entries in less than one hour.

The above-mentioned query contains a complex structure; a so-called quantifier alternation: “there is (a physicist) – for all (children)”. Queries with such a structure generally require a lot of effort to answer. With our new approach, basically any nesting of such quantifier alternations can be processed efficiently. For example, the question “Does every German physicist have at least one ancestor who was also a scientist?” is a “for all – there is” query. Another example of a quantifier alternation is this question about the first German chancellor: “Is there a German chancellor, for whom all German chancellors are not a predecessor?” Our method already answers all these queries using a compiled ontology in less than one second.



### CONTACT

**Christoph Weidenbach**

**RG. 1 Automation of Logic**

Phone +49 681 9325-2900

Email weidenbach@mpi-inf.mpg.de



**Patrick Wischniewski**

**RG. 1 Automation of Logic**

Phone +49 681 9325-2908

Email wischniew@mpi-inf.mpg.de





## VISUALIZATION

**Pictures are the fastest access to human consciousness. Algorithms for appropriate visualization of digital information therefore play a significant role in informatics. The requirements for visualization algorithms have dramatically increased: Artificial and natural worlds must be presented in an ever more realistic and fast manner: in flight simulators, surgical operation planning systems, computer games or for the illustration of large amounts of data in the natural and engineering sciences.**

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*”).

Another subject of our research is the reconstruction of statistical models of specific object categories such as faces or human bodies from 3D scans. These models make it possible for us to simulate a movement of the face or body, and also the changing face or body form (for example, after gaining weight). The solutions of many problems in image processing, computer animation, and 3D movement measurement are dramatically simplified with these models. One can discover underlying similarities and symmetries between different 3D models as well as within individual 3D models using statistical analysis of 3D geometry. In this way, the basic building blocks of 3D shapes can be seen, and rules can automatically be derived as new versions of these forms are created.

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. For this, we focus on the development of real-time algorithms.

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.

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

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

CONTRIBUTIONS



UNDERSTANDING  
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH  
& DIGITAL KNOWLEDGE

MULTIMODAL  
INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION


<b>Stereo and HDR Imaging:</b>	
<b>Display Quality Measurement and Enhancement</b> .....	84
<b>Advanced Real-Time Rendering</b> .....	85
<b>Computational Photography and Measurement</b> .....	86
<b>Statistical Models of Humans</b> .....	87
<b>Correspondences and Symmetry in 3D Scenes</b> .....	88
<b>Flow Analysis</b> .....	89

## 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.

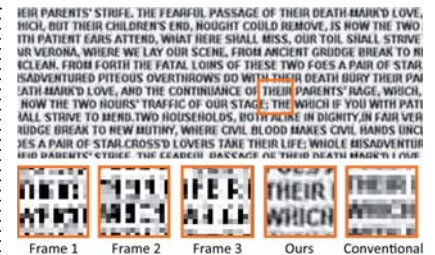
We have developed an application to predict the perceived difference between two film sequences. Building on perceptual models, it uses existing physiological data to analyze the size, strength, and change of visual patterns in the input sequences and from this produces a third sequence that visualizes 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

Internet <http://www.mpi-inf.mpg.de/~karol>



**Tobias Ritschel**

DEPT. 4 Computer Graphics

Phone +49 681 9325-4157

Email ritschel@mpi-inf.mpg.de


Internet <http://www.mpi-inf.mpg.de/~ritschel>

## 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-time 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 user's input. In order to achieve this performance, severe simplifications had to be made, which led to the development of highly specialized graphics hardware (GPUs). In our work, we aim to fill the gap between highly realistic offline-rendering and fast interactive rendering.

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

Internet <http://www.mpi-inf.mpg.de/~karol>



**Tobias Ritschel**

**DEPT. 4 Computer Graphics**

Phone +49 681 9325-4157

Email ritschel@mpi-inf.mpg.de

Internet <http://www.mpi-inf.mpg.de/~ritschel>

## Computational Photography and Measurement


Digital cameras are everywhere – yet they are still used in the same way as film-based cameras were in the old days. There are two key differentiating aspects with respect to analogue technology: their computational capabilities, and the direct on-board digitization of intensity measurements. The latter aspect enables the use of digital cameras as measurement devices, whereas the former aspect enables computation on the camera itself. Our group investigates novel algorithms for the analysis of light transport processes by means of digital cameras as well as novel applications based on camera images such as 3D reconstruction.

### Image-based measurement of material reflectance under general conditions

Subtle appearance differences due to changes in lighting and/or viewpoint are very important for realistic image synthesis. Material reflectance is one of the major causes for these variations. In recent years, it has been widely recognized that the parameters of traditional mathematical reflectance models are too difficult to control manually, resulting in an artificial look to rendered objects. Based on this observation, the dominating trend has become to *measure* these parameters. Our work in this area has significantly extended the scope of these techniques:

traditional techniques estimate one reflectance function for each of the three color channels – red, green, and blue.

In reality, many colors, especially very intensive ones, are produced by processes that transfer energy between different wavelengths of light; this effect is, for example, used in laundry detergents. We have developed an extended bispectral reflectance model for this purpose; and we were the first to measure this effect to provide data for analysis. In a similar vein, we have measured the dependence of the reflectance function on the refractive index of the surrounding medium. This helps in simulating effects associated with different appearances of the same object in different media like, for example, air, water, and oil.



**A measurement system for immersed material samples:** A flat material sample is immersed in a medium of refractive index different from air. The material sample is illuminated by a laser source, and the response is recorded on the screen. Performing the measurement for different refractive indices allows for simulations in different media.

### Kaleidoscopic cameras for generating hemispherical viewpoint distributions

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, above all in film production. A fundamental problem with 3D scanning systems is the restricted visible volume covered by conventional cameras: an object can only be viewed from one perspective at a time. This implies that all-round digitization of *dynamic* objects can only be achieved at considerable expense: by employing multi-camera systems. In addition, there is a direct correlation between the number of cameras and the quality of the digitization: as in computed tomography technology, a higher number of viewpoints of the object, that is, more measurements, provide more robust results and yield an increased reconstruction quality. We have developed a method in our group that uses a mirror system to generate a high number of virtual viewpoints (> 200). These can be recorded *at the same time* with only a *single* camera. ...

### 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/>





## Statistical Models of Humans



**Figure 1: left to right: Subject in the light stage, acquisitions of the face under different illumination conditions, reconstructed direction of the surface normals (color-coded)**

### Statistical face models


In order to accurately capture the 3D geometry of the face, a so-called light stage is used at the institute. A light stage is a sphere, on which approximately 160 light sources are arranged. The face of a subject is located in the center of the light stage [see Figure 1]. This apparatus allows pictures to be taken of the face under different illumination conditions in order to acquire its properties very precisely. For example, single pores are visible in 3D reconstructions of the face. Furthermore, the diffuse and specular component of the reflected light, the gloss of the skin, or the amount of sub-surface scattering in the skin can be acquired.

We are currently working on creating a large database of faces that have been captured with the light stage. In order to capture facial expressions of subjects, a single measurement in the light stage should only take split seconds because especially tense facial expressions are hard to keep for a longer time span. Such a measurement places high requirements on the used hardware components. The goal of this project is to significantly improve the resolution and the degree of realism of our current statistical face model with the new database. The resulting face models and their facial expressions should at a first glance not be distinguishable from real faces. There are many applications for such facial animations, e.g., in Internet applications where virtual assistants can guide visitors on a website. The usage of a statistical face model can also improve the automatic tracking of faces in videos in order to control a computer or a robot by facial expressions, for example.

### Statistical body models

A detailed model of the human body is of great importance in many applications of image analysis and computer graphics, for example, in markerless motion capture, human-computer interaction, or computer animation. In general current models only describe the geometry and movement of a certain subject in a very accurate way. The necessity of creating a subject-specific model, however, prevents the application of image processing algorithms to arbitrary persons. Hence, we have developed a new statistical body model of the human, which is described by two low dimensional parameter spaces. The first parameter space determines the joint angles of the human skeleton and hence the movement of the body. The main innovation is the second space, which allows the control of body shape and constitution by a set of few intuitive parameters. Examples for such parameters are: weight, height, leg length, waist measurement, muscularity, etc. The model was created by means of a machine learning technique on laser scans of more than 120 persons of different sex and age.

Hence, it does not only model the freedom of movement of a human, but the range of body shapes of the population as well. We can fit this model to image and video data and measure the movement and constitution of a particular person much more accurately. By means of the newly developed MovieReshape technique, we can also track movements in 2D video streams, which are normal movies and videos, from a single camera perspective. After capturing the movement of a person, the body shape parameters can be changed and transferred onto the entire video sequence. Meanwhile, it is possible to fit the body model to the data of a depth camera, which allows a more precise measurement of the movement in real-time [Figure 2].



**Figure 2: The motion of an actor in a video can be captured with a statistical body model. Afterwards, for example, the body height can be changed in the entire video.**



#### CONTACT

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



**Thorsten Thormählen**  
**DEPT. 4 Computer Graphics**  
 Phone +49 681 9325-4017  
 Email thormae@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 1,3]. 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: Partial symmetry detection: a point cloud from a 3D laser scan is decomposed automatically into repeating elements that form building blocks of the shape. Data set: IKG Hannover.




Figure 2: Correspondences between deformed surfaces are computed automatically. Data set: TOSCA data base.



Figure 3: Correspondences between building blocks can be learned from user input. This enables semi-automatic semantic tagging of geometric data. Data set: IKG Hannover.



Figure 4: Given partial symmetries of an object, the resulting building blocks can be reassembled to form new shapes. Data set: IKG Hannover.

### 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/departments/d4/areas/statgeoprocessing>



## 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 streak 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 streak lines in Figure 1 were computed within a minute using our new method. The classic algorithm requires more than two hours for this. ...




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

### CONTACT

Tino Weinkauff

DEPT. 4 Computer Graphics

Phone +49 681 9325-4020

Email [weinkauff@mpi-inf.mpg.de](mailto:weinkauff@mpi-inf.mpg.de)

Internet <http://feature.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](mailto:jgerling@mpi-inf.mpg.de)  
Internet <http://www.imprs-cs.de>



## N E W S

**This section contains an overview of 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 2011 included two international conferences, the roofing ceremony for the new Max Planck Institute for Software Systems building that will also host one floor for the Informatics institute, and the assessment of the institute by the scientific advisory board.**

## CONTRIBUTIONS

ARFUL PASSAGE OF THEIR DEATH  
 END, NOUGHT COULD REMOVE, IS  
 THAT HERE SHALL MISS, OUR TOIL  
 OUR SCENE, FROM ANCIENT GRUD  
 TAL LOINS OF THESE TWO FOES /  
 RTHROWS DO WITH THEIR DEATH  
 CONTINUANCE OF THEIR PARENTS  
 FIC OF OUR STAGE; THE WHICH IF  
 HOUSEHOLDS, BOTH ALIKE IN DIGN  
 Y, WHERE CIVIL BLOOD MAKES CI  
 LOVERS TAKE THEIR LIFE; WHOL  
 ADETH PASSAGE OF THEIR DEATH



Frame 3



Ours

Honors and Awards .....	94
Fellowships .....	94
Personal Data .....	94
Cooperations .....	95
CeBIT 2011 .....	96
Conference Reports .....	97
Cluster of Excellence on Multimodal Computing and Interaction .....	98
Scientific Review of Max Planck Institute for Informatics .....	99
Research Days 2011 .....	100
Topping-out ceremony, new building for Max Planck Institute for Software Systems .....	100
IMPECS .....	101
IOI Training 2011 .....	101

## Honors and Awards

### Paris Kanellakis Award

Kurt Mehlhorn won one of the most prestigious awards ever received by members of our Institute. In July, he was presented, as the first German ever, with the Paris Kanellakis Theory and Practice Award for 2010: *"For contributions to algorithm engineering by creating the LEDA library for algorithmic problem solving."* The world's largest association of computer scientists, the Association for Computing Machinery (ACM), bestows this prize annually for *"specific theoretical accomplishments that have had a significant and demonstrable effect on the practice of computing."* The reputation of this award is accordingly very high. LEDA itself is easy to use, efficient and accurate. To achieve these characteristics, significant theoretical advances were required. For example, a revolutionary concept is the fact that the LEDA routines not only deliver the required results in a short time, but at the same time also provide the reasoning for the result. ...



### ACM-SIGMOD Award

Gerhard Weikum, Director at the Max Planck Institute for Informatics, was honored with this year's SIGMOD award by the Special Interest Group on Management of Data within the Association for Computing Machinery. He received the award *"for his significant service to the database community, through research leadership and academic mentoring..."* He was praised as *"... a long term leader and influential organizer in innovating our collaboration and publication culture ..."* in the reasoning for the award. Special mention was made of this pioneer work in the field of self-tuning databases and his strict requirements for publishing. ...

## Fellowships

### Google Europe Scholarship



With its Google Doctoral Fellowship Programme, the Internet giant Google supports outstanding doctoral students in computer science or associated fields.

It is not a requirement that the research topics are relevant for Google.

Carola Winzen, who is a doctoral student under Professor Kurt Mehlhorn and is researching complexity theory for randomized algorithms, applied for and received one of only thirteen fellowships for 2011. ...

### Marie Curie Intra-European Fellowship



In the simplest case, a society can be described as the amount of interpersonal relationships in which individuals interact in pairs, each individual is considered a node and the relationship to another individual is considered as an edge. Within social networks, e.g. Facebook, opinions and views spread, for example the reputation of certain products, by word-of-mouth, i.e. from one individual to a neighboring individual. This opinion propagation can be simulated by means of a graphs method. Nikolaos Fountoulakis successfully applied for a Marie Curie Intra-European Fellowship for the research on development of models which are based on similar graphs structures. Furthermore, seven doctoral students sponsored by the Microsoft Research PhD Scholarship studied at the Max Planck Institute for Informatics in 2011. ...

## Personal Data

### Appointments

Alice McHardy, head of the research group *"Computational Genomics and Epidemiology"* at the Max Planck Institute for Informatics since 2007, has received an appointment to the full professorship for Computational Biology at the Heinrich-Heine University in Düsseldorf. Additionally, she continues to direct the research project at Max Planck Institute for Informatics. ...

Viorica Sofronie-Stokkermans was appointed as professor at the University of Koblenz-Landau. ...

Hendrik Lensch, former doctoral student under Hans-Peter Seidel, was appointed to the prestigious Chair for Computer Graphics, formerly held by Prof. Strasser, at the University of Tübingen. ...

### Senior Researchers

In 2011, the following scientists were appointed as Senior Researches due to their excellent performance: First, Rainer Gemulla, Dept. 5, whose research areas include *Probability & Uncertainty in Data Management* and *Large-Scale Statistical Computing*. Second, Tino Weinkauf, Dept. 4, who deals with the subject *Topological Methods and Feature Extraction*. And third, Michael Sagraloff, who researches the areas of *Algebraic Geometry* and *Nonlinear Computational Geometry* in Dept. 1. ...

## Cooperations

### Siemens cooperation (Support of doctoral students)

In 2010, a cooperation agreement was concluded with Siemens AG, in which Siemens sponsors the training of doctoral students at the Max Planck Institute by granting fellowships. In 2011, three young scientists were already supported:



*Claudio Magni*, a doctoral student under Professor Kurt Mehlhorn, receives support for his thesis on Improved Approaches in Production Planning.

Among other things, the thesis examines the complex process of the production of LEDs at Siemens. Here, the analysis of varied and multiply-connected process parameters for the individual production steps is required. To accelerate the optimization process, Claudio Magni analyzes the current management processes and develops a more holistic approach.



As a research assistant of Professor Gerhard Weikum, *Mohamed Yahya* tries to enable automated replies by the means of knowledge-based techniques, even

for complex issues. Siemens needs and wants to offer its clients detailed and in-depth support for its high-tech products. The existing knowledge data bases are often confusingly large and interrelated on many levels. The collaboration with Mohamed Yahya includes the development of interfaces which understand natural language structures. The Siemens managers are thereby hoping for e.g. easier and more intuitive access to their knowledge-based systems.



*Ching Hoo Tang*, under the guidance of Professor Christoph Weidenbach, completes the trio. His research fits into Siemens' quest to keep product

development efficient and manageable, despite increasing complexity. For example, rule-based systems enable a variable implementation of product configuration. Ching Hoo Tang's goal is to find a suitable formalization for systems in which products are automatically derived based on user decisions and rules. This would enable the investigation of the attributes of such systems by means of automated procedures. This especially includes the critical issue of potential inconsistencies in the regulations caused by conflicting rules. ...

### Intel Visual Computing Institute – IVCI

Two years ago the IVCI was founded as the largest Intel research facility in Europe. Based on the globally unique concentration of computer scientists in the Saarland, Intel decided to establish its institute in Saarbrücken. The IVCI is managed by a steering committee, on which Intel and four Saarland partners (both Max Planck Institutes, namely MPI-INF and MPI-SWS, DFKI, and the University of Saarland) are equally represented. Of the three directors, only one is from Intel, a fact that demonstrates their confidence in the quality of the location. The high-tech company has made twelve million euros available for the next five years. The projects are not dictated by Intel, but are assessed according to their scientific value. Prof. Christian Theobalt is the member of the steering committee for the Max Planck Institute of Informatics. One of his projects on "Markerless motion capture" is also a subject of research at the Intel Institute. ...



### SoftwareCampus

The Max Planck Institute for Informatics is a partner in the EIT ICT Labs, a scientific and innovative community of the European Institute of Innovation and Technology. Within this framework our institute supports the combination of computer science in a scientific environment with IT as a strategic industry field by participating in the Initiative SoftwareCampus. The initiative was preceded by the recognition that there is a lack of appropriate management trainees in the high-tech industry who can look back on solid training in computer science or related fields. SoftwareCampus offers support by German software companies, parallel to an academic training, to extraordinary Master's or Doctoral Students. Students are in contact with the top management of the industrial partners. For example, René Obermann, CEO of Deutsche Telekom AG, will spend a weekend with the young scientists. The first young scientists were introduced at this year's IT summit on December 12, 2011. ...



### Max Planck Science Gallery

The Max Planck Society has established the Science Gallery as a public attraction in the center of Berlin, in the WissenschaftsForum (Science Forum) at the Gendarmenmarkt and is an absolute crowd puller. The Science Gallery is a permanent "open door" to the world of international top research, which is targeted both at passers-by and visitors. The exhibitions all follow a digital multimedia interactive concept. Exhibitions can therefore be changed within minutes. Results from the scientific work by Prof. Dr. Dr. Thomas Lengauer on the topic of medication for HIV patients and material from Prof. Dr. Alice McHardy's research on the evolution of the H1N1 flu virus were utilized for the first exhibition. ...

## CeBIT 2011

The world renowned CeBIT computer trade fair is one of the most attractive platforms for networking outside of the research environment for introducing oneself to the press and to industrial partners and for inspiration in terms of industry-relevant research topics.

Traditionally, the Max Planck Institute for Informatics presents its latest developments to the public at a joint stand. The combination of partners, which in addition to the Max Planck Institute for Informatics includes the University of the Saarland, the German Research Institute for Artificial Intelligence and the Cluster of Excellence on Multimodal Computing and Interaction, guarantees exhibits of the highest scientific quality.

This year, Michael Wand and Art Tevs presented a method for the processing of point clouds which originate e.g. from laser scans of the façades of an entire city. This presents two significant challenges: Firstly, such 3D scans result in huge amounts of data ranging into several terabytes. However, the data need to be visualized and processed in real time. New algorithms and data structures which allow processing in real time were developed for this purpose, whereby the scene size is restricted only by the available disk space. Secondly, manual processing of such large amounts of data is impossible due to the workload involved.



The Max Planck Institute at CeBIT 2011

This problem was solved by newly developed algorithms, that automatically identify similar elements (e.g. doors or windows in building façades), which enables further automated data processing.

In addition, the software – developed by Thorsten Thormählen and his team – that suggests make-up tips received a great response from visitors and the press. Based on a database of reference examples, a make-up style is suggested to a test subject, and is then applied to a 3D model of her face. This combination of direct everyday relevance and the latest computer graphics tech-

niques was so appealing to visitors and journalists alike that the program was mentioned in various media and enquired for by cosmetic companies and beauty salons. Numerous well-known national and international agencies reported on it: AFP, The Independent, BBC, Discovery News, Bayern 3, RBB, D-News. The demonstration video on YouTube was accessed several thousand times. Even 40 weeks after the CeBIT, the news magazine “Der Spiegel” published an article about it in its 50/2011 edition. ::



## Conference Reports

Conferences are a great opportunity for researchers and parties interested in a particular field to spend a few days focusing exclusively on a given topic. Both established researchers and beginners benefit from this concentration of specialists and gain many new ideas. Organizers have creative leeway within a program, and a successfully completed event earns a good reputation.

In the late summer of 2011, Saarbrücken was at the center of European algorithmics for two weeks. Two internationally renowned events on algorithms were held in succession: From August 29 to September 02, the ADFOCS (Max Planck Advanced Course on the Foundations of Computer Science) summer school was held, and a week later, from September 05-09, the ALGO, the European algorithm conference took place. The FroCos2011: 8th International Symposium on Frontiers of Combining Systems (October 07-09) was the third major conference that the Max Planck Institute for Informatics organized this year.

### ADFOCS

Established computer science researchers, together with young scientists from around the world, focus on current topics which are abstract and thus difficult to penetrate. The freshly acquired knowledge is directly tested in specific tasks. Algorithms thus learned and developed can later be applied in computer science, e.g. to come to a conclusion more quickly. They also serve to accurately perform computations, which previously could only be solved approximately. They also help in simultaneously achieving interrelated results, which previously could only be solved by mutual approximation.

More than 100 course participants, mostly post graduates in computer science and mathematics, but young postdocs as well, learn effective solution approaches and a deeper understanding of the inner correlations of the theory from the professors. The requirements for students are high: *"They require a broad knowledge of mathematics, scientific curiosity and excellent skills"*, says Anke van Zuylen, former researcher at the Tsinghua University in Beijing, China, who is also one of this year's summer school organizers. Participants included young mathematicians and computer scientists from Germany and several other European countries, as well as from the USA, Japan, Iran and India. The summer course was held for the twelfth time in annual succession.

### ALGO

More than 300 participants, a quarter of whom are non-Europeans, presented their findings of the previous year. Issues such as *"How successful is strategic market behavior?"* are always practically relevant and usually are the direct result of research activities.

The ALGO was comprised of six different workshops devoted to the different applications of algorithms, for example bioinformatics. It was investigated how one can compute even more quickly which active substance will most likely contain a flu epidemic. Another workshop dealt with the issue of transport problems, for example at what intervals, speed, and stopping times a railway line can be best utilized for both passenger and freight transport.

In eight by invitation only lectures (so-called key notes) leading researchers summarized the current state of scientific knowledge in their respective fields. In addition, more than 170 individual presentations dealt with current research topics. Many of the new issues presented at the conference are later found in sensor networks and mobile applications, in faster or more accurate computations within the Internet or as well as in the development of energy networks.

### FroCoS 2011

The various fields of computer science (e.g. logic, artificial intelligence, program development) require specialized and thus at times very different formalisms and inference methods.

For practical applications, it is necessary to combine these special structures and to integrate them into systems which can be applied across fields. Currently, the development of general methods has begun in many ways where it is demonstrated how special formal systems should be combined and integrated or how complex systems can be modularized and analyzed.

FroCoS (Frontiers of Combining Systems) focuses exactly on this research subject and aims for the further development thereof. The FroCos 2011 was an ideal platform for this: Approximately 30 scientists from 8 countries and 20 further participants from the Max Planck Institute for Informatics and the University of Saarland exchanged information in the field of combination, modularization and integration of (especially logic-based) systems and their practical application. ...

## Cluster of Excellence on Multimodal Computing and Interaction

The institute plays an important role in the Excellence Cluster “*Multimodal Computing and Interaction*”, which was established by the German Research Foundation (DFG) in 2007. All institute directors are members of the cluster as principal investigators; Prof. Hans-Peter Seidel is the scientific coordinator of the cluster.

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 arises 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 type 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. We call this multimodal processing. In addition, daily interpersonal communication is based on a multitude of different modalities. 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 gestures as well as answer in appropriate modalities.

The proposed research program builds on existing strengths. The cluster includes 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 participating institutions have agreed to a joint, long-term research program as the foundation for their work. 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 es-

tablishment of scientific junior 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.

Last year, the completion and presentation of the Renewal Proposal for the 2013–2017 period was a significant part of the cluster’s activities. Final decisions for the second program phase of the Excellence Initiative will be announced by the DFG (German Research Foundation) on June 15, 2012. ...

## Scientific Review of Max Planck Institute for Informatics

The scientific achievements of the Max Planck Institute for Informatics are reviewed by its advisory board every two years. The advisory board consists of world-renowned scientists from the field of computer science [<http://www.mpi-inf.mpg.de/about/index.html>]. Members are selected based on their expertise in terms of the research conducted at the institute and their independence from the institute. The inspection by the advisory board is based on a biennial report, a written status report which provides information about the concluded, current and planned scientific research and projects of the institute. For 2009–2011 the report contained almost 700 pages; it can be downloaded from the institute's website [<http://mpi-inf.mpg.de/publications/index.html>]. The inspection by the advisory board takes two days; in 2011 it took place on the last two days of May.



The advisory board report and the inspection are crucial for the evaluation of the institute, which is subsequently performed by the advisory board. The inspection itself is divided into different stations: To begin with, the department heads give presentations both on the



development of the respective departments/large groups and on the research highlights achieved in the assessment period. For example, one of these highlights concerns the use of randomness in the distribution of news: The department “*Algorithms and Complexity*” was able to accurately mathematically prove that the parallel distribution of a message in a group of  $n$  people who randomly choose their communication partners lasts a maximum of  $\log(n)$  rounds. After the presentations, the advisory board visits all senior researchers/group leaders of the institute in individual meetings, to discuss their findings and to gain a comprehensive picture of all groups. On the second day of the inspection, the foyer of the main building is turned into a landscape of posters, where all researchers, from graduate students to professors, present their work.


The advisory board takes this opportunity to learn more about individual research achievements.

Thus, the institute is assessed in all aspects from the overall structure to the individual scientists. The result of the assessment is binding and applies to the institute as a whole: from the global structure to the individual research achievements. ...

## Research Days 2011

Every year, the Max-Planck Institute for Informatics organizes the “Computer Science Research Days” for gifted high school students, in cooperation with the Computer Science Department of Saarland University and the DFKI.

This year, from July 18–20, 2011, successful participants from the second round of the Federal Competition for Informatics, and the best two high school graduates from Rhineland-Palatinate schools completing an advanced computer science course, were invited. The students were given the opportunity to get a taste of current research topics and to actively participate in research activities.



The “tasting” consisted of attending presentations, for example on the topic “Slime Mold Computer”, or visiting a current computer science lecture on the topic of “Optimization”. In various workshops, the participants could e.g. build a 3D scanner and thereafter test it themselves, or learn how the wireless high-performance digital transmission of data is set up within the building. The group work was designed in such a way that all areas were covered: from theoretical foundations 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 “Challenge and promote”. One objective of the Max-Planck Institute is to attract young talented people to informatics, to identify and to promote new talent. ...

## Topping-out Ceremony, New Building for Max Planck Institute for Software Systems

After nearly one and a half years of construction, the topping-out ceremony for the new building was celebrated in February 2011. The new building architecturally completes the “*Square of Informatics*” towards the east. The building was erected for our sister institution Max Planck Institute for Software Systems and following the completion of the interior, will also be home to the joint administration of the MPI-INF and MPI-SWS. The move into the building is planned for Spring 2012. ...





## IMPECS – Indo-German Max Planck Center for Computer Science

On February 3, 2010, the “*Indo-German Max Planck Center for Computer Science*” (IMPECS) was opened by the Federal President Horst Köhler at the Indian Institute of Technology in Delhi (IIT Delhi). The goal of the Center is the promotion of first-class basic research in computer science by means of close cooperation between Indian and German scientists. This cooperation is to be promoted by joint research, the exchange of doctoral students and postdocs, a Max Planck guest professorship and several workshops and schools. For the first five years, IMPECS will be funded by MPG, the Federal Ministry of Education and Research (BMBF) and the Indian Ministry for Science and Technology (DST). On the MPG side, the governing body consists of Kurt Mehlhorn (MPI-INF) and Rupak Majumdar (MPI-SWS),

and on the Indian side of Naveen Garg (IIT Delhi) and Manindra Agrawal (IIT Kanpur).

IMPECS is funding approximately ten German-Indian research groups. On the Indian side, any university or institution may participate; on the German side, however, a partner must be a member of the MPI for Informatics or MPI for Software Systems. Currently, six binational research groups are being sponsored; several other groups are being called for.

The IMPECS has allowed for an active exchange of postdocs and students in both directions. In 2010, Prof. Pankaj Agarwal from Duke University received a Max Planck guest professorship. Prof. Manindra Agarwal of IIT Kanpur, winner

of the Humboldt Research Award, visited the Max Planck Institute for Informatics for two months in the summer of 2011. A joint workshop was held in Delhi in April 2010 as a kick-off event.

In addition, IMPECS co-financed the “*School on Parameterized and Exact Computation*” at the IMSc Chennai in December 2010, a school and a workshop on “*Geometric Computing*” at the IIT Delhi in October and November 2010 and, together with Microsoft Research, a “*School on Approximability*” at the IISc Bangalore in January 2011. More information is available at [\[www.impecs.org\]](http://www.impecs.org). ...

## IOI Training 2011

From June 14-16, Prof. Dr. Markus Bläser, Fabian Grundlach, Stephan Seufert, Dr. Wolfgang Pohl, Prof. Dr. Christoph Weidenbach and Dr. Anke van Zuylen trained the German Olympic team for the Computer Science Olympiad 2011 in Thailand. The training took place in the guesthouse of the Saarland Turnerbund in Braunshausen. Students Sebastian Gießel, Tobias Lenz, Johannes Kirschner, Patrick Klitzke and Aaron Montag were recruited from among the winners of national computer science competitions, in particular the BWINF (“*Bundeswettbewerb Informatik*”). They are the best in their age group in Germany. The Computer Science Olympiad competition consists of solving tasks with the aid of the computer. It requires a level of skill which would be expected from a computer science student with a bachelor’s degree. The aim of the training was both to prepare the students for the competition and to provide them with a broader insight into informatics as a science. In this context, we dealt with

the scientific issues related to the solution of Hilbert’s tenth problem, solved by the scientists Martin Davis, Hilary Putnam, Julia Robinson and Yuri Matijasevitch. Beyond that, we discussed associated problems with regard to the famous computer science  $P=NP$  question. Eventually, a presentation on the history of the Chaos Computer Club,

with subsequent discussions on IT security in general, was held. The obligatory soccer competition “*Trainer against Trainees*” was won by the trainer team. However, this came at a high price: On the next day clear deficiencies presented themselves in the general fitness of the trainers, observable as a persistent inability to move. ...






# The Institute in Figures


Budget without third-party funds from 2008 to 2011



Staff Current as of 1/1/2012

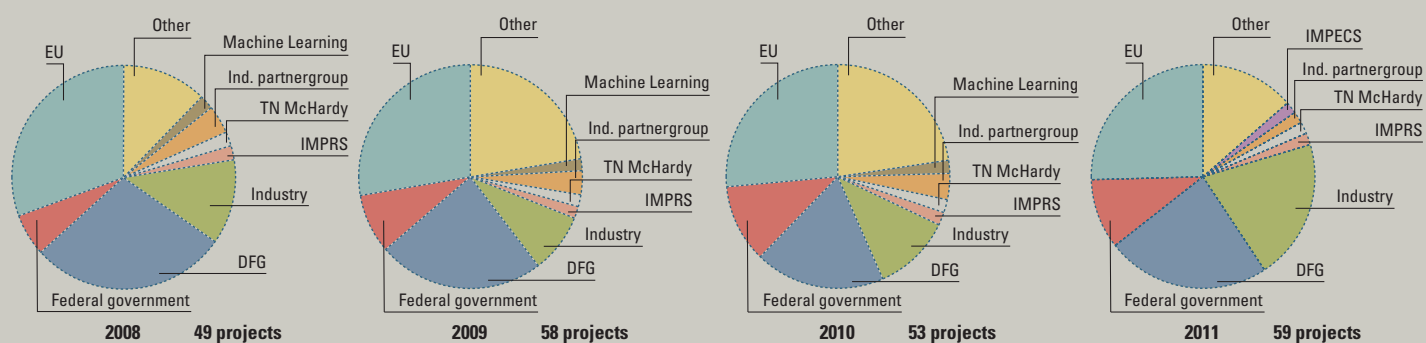


Ratio of scientific to non-scientific staff members Current as of 1/1/2012




## Third-party funding from 2008 to 2011

Numbers and distribution



## Third-party funding from 2008 to 2011

Revenues



## 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 reaction times. The advantages for operating secu-



rity, total expenses and time until larger changes are ready, however, clearly outweigh these disadvantages.

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 sys-

tems 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 auto-



matic 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 cur-

rently 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. 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-klb.mpg.de](mailto:jh@mpi-klb.mpg.de)

Internet <http://www.mpi-klb.mpg.de/gvw/ist.html>

# Selected Cooperations

## UNDERSTANDING IMAGES & VIDEOS

- ... Intel Visual Computing Institute, Saarbrücken, Germany
- ... Katholieke Universiteit Leuven, Leuven, The Netherlands
- ... Microsoft Research, Cambridge, UK
- ... Saarland University, Saarbrücken, Germany
- ... Stanford University, Stanford, USA
- ... Technical University of Darmstadt, Darmstadt, Germany
- ... Toyota, Brussels, Belgium
- ... University of British Columbia, Vancouver, Canada
- ... University of California, Berkeley, USA
- ... University of Passau, Passau, Germany

## BIOINFORMATICS


- ... BioSolveIT GmbH, Sankt Augustin, Germany
- ... Heinrich-Heine-University, Düsseldorf, Germany
- ... Informa s.r.l, Rome, Italy
- ... Karolinska Institute, Stockholm, Sweden
- ... Saarland University, Saarbrücken, Germany
- ... Stanford University, Stanford, USA
- ... University of British Columbia, Vancouver, Canada
- ... University of Cologne, Cologne, Germany
- ... University of Freiburg, Freiburg, Germany
- ... University of Heidelberg, Heidelberg, Germany
- ... University Tor Vergata, Rome, Italy
- ... Wolfgang-Goethe-University, Frankfurt, Germany

## GUARANTEES

- ... Carnegie Mellon University, Pittsburg, USA
- ... Chinese Academy of Sciences, Beijing, China
- ... New York University, New York, USA
- ... Technion, Haifa, Israel
- ... University of Canberra, Canberra, Australia
- ... University of Freiburg, Freiburg, Germany
- ... University of Haifa, Haifa, Israel
- ... University of Nancy, Nancy, France
- ... University of Oldenburg, Oldenburg, Germany
- ... University of Pittsburgh, Pittsburgh, USA

## INFORMATION SEARCH & DIGITAL KNOWLEDGE

- ... DERI, Galway, Ireland
- ... EPFL Lausanne, Lausanne, Switzerland
- ... European Archive Foundation, Amsterdam, The Netherlands
- ... FluidOps AG, Walldorf, Germany
- ... German-Israeli Foundation for Scientific Research and Development, Israel
- ... Google, Zurich, Switzerland
- ... Gottfried Wilhelm Leibniz University Hanover, Hanover, Germany
- ... Hebrew University of Jerusalem, Jerusalem, Israel
- ... IIT Bombay, Bombay, India
- ... Jülich Supercomputing Centre, Jülich, Germany
- ... Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
- ... Microsoft Research, Cambridge, UK



::: Microsoft Research, *Redmond, USA*  
 ::: RWTH Aachen, *Aachen, Germany*  
 ::: Siemens AG, *Munich, Germany*  
 ::: Stanford University, *Stanford, USA*  
 ::: Stichting European Archive,  
*Stichting, The Netherlands*  
 ::: Technical University Ilmenau,  
*Ilmenau, Germany*  
 ::: University of Duisburg-Essen,  
*Duisburg, Germany*  
 ::: Yahoo! Research, *Barcelona, Spain*

#### **MULTIMODAL INFORMATION**

::: Saarland University, Excellence  
 Cluster Multimodal Computing and  
 Interaction, *Saarbrücken, Germany*  
 ::: Saarland University of Music,  
*Saarbrücken, Germany*  
 ::: Technical University of Darmstadt,  
*Darmstadt, Germany*  
 ::: University of Bonn, *Bonn, Germany*

#### **OPTIMIZATION**

::: Aarhus University, *Aarhus, Denmark*  
 ::: Chennai Mathematical Institute,  
*Chennai, India*  
 ::: Danish Technical University,  
*Kopenhagen, Denmark*  
 ::: Durham University, *Durham, UK*  
 ::: ETH Zurich, *Zurich, Switzerland*  
 ::: Indian Institute of Technology Delhi,  
*New Delhi, India*  
 ::: Indian Institute of Technology  
 Kanpur, *Kanpur, India*  
 ::: INRIA Saclay - Ile de France, *Paris, France*  
 ::: Prostep AG, *Darmstadt, Germany*  
 ::: Saarland University, *Saarbrücken, Germany*

::: Siemens AG, *Vienna, Austria*  
 ::: TATA Institute of Fundamental  
 Research Mumbai, *Mumbai, India*  
 ::: University of Adelaide, *Adelaide, Australia*  
 ::: University of Birmingham,  
*Birmingham, UK*  
 ::: University of Kiel, *Kiel, Germany*  
 ::: University of Liverpool, *Liverpool, UK*  
 ::: University of Paderborn, *Paderborn, Germany*

#### **SOFTWARE**

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


#### **VISUALIZATION**

::: Indian Institute of Technology Delhi,  
*New Delhi, India*  
 ::: Media Lab, Massachusetts Institute  
 of Technology, *Massachusetts, USA*  
 ::: Saarland University, Excellence  
 Cluster Multimodal Computing and  
 Interaction, *Saarbrücken, Germany*  
 ::: Stanford University, *Stanford, USA*  
 ::: Technical University of Denmark,  
*Kopenhagen, Denmark*  
 ::: Telecom ParisTech, *Paris, France*

::: University College London,  
*London, UK*  
 ::: University of Bristol, *Bristol, UK*  
 ::: University of British Columbia,  
*Vancouver, Canada*  
 ::: University of Magdeburg,  
*Magdeburg, Germany*

# Selected Publications

- [1] F. ALVANAKI, S. MICHEL, K. RAMAMRITHAM AND G. WEIKUM. EnBlogue – Emergent Topic Detection in Web 2.0 Streams. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, Athens, Greece, 2011*. ACM.
- [2] A. ANAND, S. BEDATHUR, K. BERBERICH AND R. SCHENKEL. Temporal Index Sharding for Space-Time Efficiency in Archive Search. In W.-Y. Ma, J.-Y. Nie, R. A. Baeza-Yates, T.-S. Chua and W. B. Croft, eds., *34th ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2011)*, Beijing, China, 2011, pp. 545–554. ACM.
- [3] M. ANDRILUKA, S. ROTH AND B. SCHIELE. Discriminative appearance models for pictorial structures, 2011.
- [4] B. ARANDA, H. BLANKENBURG, S. KERRIEN, F. S. BRINKMAN, A. CEOL, E. CHAUTARD, J. M. DANA, J. DE LAS RIVAS, M. DUMOUSSEAU, E. GALEOTA, A. GAULTON, J. GOLL, R. E. W. HANCOCK, R. ISSERLIN, R. C. JIMENEZ, J. KERSEMAKERS, J. KHADAKE, D. J. LYNN, M. MICHAUT, G. O'KELLY, K. ONO, S. ORCHARD, C. PRIETO, S. RAZICK, O. RIGINA, L. SALWINSKI, M. SIMONOVIC, S. VELANKAR, A. WINTER, G. WU, G. D. BADER, G. CESARENI, I. M. DONALDSON, D. EISENBERG, G. J. KLEYWEGT, J. OVERINGTON, S. RICARD-BLUM, M. TYERS, M. ALBRECHT AND H. HERMJAOK. Psicquic and psiscore: accessing and scoring molecular interactions. *Nature Methods*, 8(7):528–529, 2011.
- [5] P. BAUMGARTNER AND U. WALDMANN. A combined superposition and model evolution calculus. *Journal of Automated Reasoning*, 47(2):191–227, August 2011.
- [6] E. BERBERICH, P. EMELIYANENKO, A. KOBEL AND M. SAGRALO. Arrangement computation for planar algebraic curves. In M. Moreno Maza, ed., *Proceedings of the 4th Internal Workshop on Symbolic-Numeric Computation*, San Jose, USA, June 2011, pp. 88–99. ACM. An extended version has been submitted to Theoretical Computer Science.
- [7] K. BERBERICH, A.-C. KÖNIG, D. LYMBERPOULOS AND P. ZHAO. Improving Local Search through External Log. In *34th ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2011)*, Beijing, China, 2011, pp. 785–794. ACM.
- [8] U. BLANKE, R. REHNER AND B. SCHIELE. South by south-east or sitting at the desk: Can orientation be a place? In *Wearable Computers (ISWC), 2011 15th Annual International Symposium on*, San Francisco, USA, July 2011, pp. 43–46. IEEE.
- [9] M. BOKELOH, M. WAND, V. KOLTUN AND H.-P. SEIDEL. Pattern-aware deformation using sliding dockers. *ACM Transactions on Graphics (Proc. Siggraph Asia)*, 30(6):123:1–123:10, December 2011.
- [10] T.-H. H. CHAN AND K. M. ELBASSIONI. A QPTAS for TSP with fat weakly disjoint neighborhoods in doubling metrics. *Discrete & Computational Geometry*, 46(4):704–723, 2011.
- [11] F. Y. L. CHIN, Z. GUO AND H. SUN. Minimum Manhattan Network is NP-complete. *Discrete & Computational Geometry*, 45(4): 701–722, 2011.
- [12] W. DAMM, C. IHLEMANN AND V. SOFRONIE-STOKKERMANS. Decidability and complexity for the verification of safety properties of reasonable linear hybrid automata. In E. Frazzoli and R. Grosu, eds., *Hybrid Systems: Computation and Control 2011*, Chicago, 2011, pp. 73–82. ACM.



- [13] P. DIDYK, T. RITSCHER, E. EISEMANN, K. MYSKOWSKI AND H.-P. SEIDEL. A perceptual model for disparity. *ACM Transactions on Graphics (Proceedings SIGGRAPH 2011)*, 30(4), 2011.
- [14] B. DOERR, M. FOZ AND T. FRIEDRICH. Social networks spread rumors in sublogarithmic time. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011*, 2011, pp. 21–30. ACM.
- [15] B. DOERR, M. FOZ AND T. FRIEDRICH. Social networks spread rumors in sublogarithmic time. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC 2011)*, 2011, pp. 21–30. ACM Press.
- [16] P. DOLLAR, C. WOJEK, B. SCHIELE AND P. PERONA. Pedestrian detection: An evaluation of the state of the art. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99):1, August 2011.
- [17] M. Y. ELTABAKH, Y. TIAN, F. ÖZCAN, R. GEMULLA, A. KRETTEK AND J. MCPHERSON. CoHadoop: flexible data placement and its exploitation in Hadoop. In *Proceedings of the VLDB Endowment*, Seattle, USA, June 2011, vol. 4, pp. 575–585. VLDB Endowment.
- [18] D. EMIG AND M. ALBRECHT. Tissue-specific proteins and functional implications. *Journal of Proteome Research*, 10(4):1893–1903, 2011.
- [19] L. EPSTEIN, A. LEVIN AND R. VAN STEE. Max-min online allocations with a reordering buffer. *SIAM J. Discrete Math.*, 25(3):1230–1250, 2011.
- [20] L. FEUERBACH, R. B. LYGSO, T. LENGAUER AND J. HEIN. Reconstructing the ancestral germline methylation state of young repeats. *Molecular Biology and Evolution*, 28(6):1777–1784, January 2011.
- [21] N. FOUNTOLAKIS, M. KHOSLA AND K. PANAGIOTOU. The multiple-orientability thresholds for random hypergraphs. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, 2011, pp. 1222–1236. SIAM.
- [22] T. FRIEDRICH AND L. LEVINE. Fast simulation of large-scale growth models. In *Proceedings of the 15th International Workshop on Randomization and Computation (RANDOM 2011)*, 2011, LNCS 6845, pp. 555–566. Springer.
- [23] F. GASSE AND V. SOFRONIE-STOKKERMANS. Efficient TBox subsumption checking in combinations of  $\mathcal{CL}_0$  and (fragments of)  $\mathcal{CL}_0$ . In R. Rosati, S. Rudolph and M. Zakharyashev, eds., *Proceedings of the 24th International Workshop on Description Logics (DL-2011)*, Barcelona, Spain, June 2011, CEUR Workshop Proceedings, vol. 745, pp. 125–135. CEUR Workshop Proceedings.
- [24] A. GEIGER, C. WOJEK AND R. URTASUN. Joint 3D estimation of objects and scene layout. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira and K. Weinberger, eds., *Neural Information Processing Systems (NIPS)*, Granada, Spain, December 2011, pp. 1467–1475.
- [25] R. GEMULLA, E. NIJKAMP, P. J. HAAS AND Y. SISMANIS. Large-Scale Matrix Factorization with Distributed Stochastic Gradient Descent. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, San Diego, USA, 2011, pp. 69–77. ACM New York.
- [26] R. HARREN, K. JANSEN, L. PRÄDEL AND R. VAN STEE. A  $(5/3 + \epsilon)$ -approximation for strip packing. In *WADS*, 2011, pp. 475–487.
- [27] J. HOART, M. A. YOSEF, I. BORDINO, H. FÜRSTENAU, M. PINKAL, M. SPANIOL, B. TANEVA, S. THATER AND G. WEIKUM. Robust Disambiguation of Named Entities in Text. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, Edinburgh, Scotland, UK, July 2011, pp. 782–792. The Association for Computational Linguistics.
- [28] S. KARAYEV, M. FRITZ, S. FIDLER AND T. DARRELL. A probabilistic model for recursive factorized image features. In *IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, CO, USA, June 2011. IEEE.
- [29] M. KERBER AND M. SAGRALO. Efficient real root approximation. In *International Symposium on Symbolic and Algebraic Computation (ISSAC)*, San Jose, USA, 2011, pp. 209–216. ACM. An extended version has been submitted to Journal of Symbolic Computation.
- [30] M. H. KIM, T. RITSCHER AND J. KAUTZ. Edge-aware color appearance (SIGGRAPH 2011). *ACM Trans. Graph.*, 30(2):13, 2011.
- [31] T. LU, S. MERZ AND C. WEIDENBACH. Towards verification of the Pastry routing protocol using TLA+. In R. Bruni and J. Dingel, eds., *FMOODS/FORTE 2011*, Reykjavik, Iceland, 2011, LNCS 6722, pp. 244–258. Springer. Accepted.
- [32] P. LUTSIK, L. FEUERBACH, J. ARAND, T. LENGAUER, J. WALTER AND C. BOCK. BiQ analyzer HT: locus-specific analysis of DNA methylation by high-throughput bisulfite sequencing. *Nucleic Acids Research*, 39(suppl 2):W551–W556, May 2011.
- [33] P. MIETTINEN AND J. VREEKEN. Model Order Selection for Boolean Matrix Factorization. In *Proceedings of the 17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD-2011)*, San Diego, CA, USA, 2011. ACM.



- [34] M. MÜLLER, D. P. W. ELLIS, A. Klapuri AND G. RICHARD. Signal processing for music analysis. *IEEE Journal on Selected Topics in Signal Processing*, 5(6):1088–1110, 2011.
- [35] N. NAKASHOLE, M. THEOBALD AND G. WEIKUM. Scalable Knowledge Harvesting with High Precision and High Recall. In *4th ACM International Conference on Web Search and Data Mining (WSDM)*, Hong Kong, 2011, pp. 227–236. ACM.
- [36] L. PISHCHULIN, A. JAIN, C. WOJEK, M. ANDRILUKA, T. THORMAEHLEN AND B. SCHIELE. Learning people detection models from few training samples. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, Colorado Springs, USA, 2011, pp. 1473–1480. IEEE.
- [37] S. REISS, I. REBHAN, P. BACKES, I. ROMERO-BREY, H. ERE, P. MATULA, L. KADERALI, M. POENISCH, H. BLANKENBURG, M.-S. HIET, T. LONGERICH, S. DIEHL, F. RAMREZ, T. BALLA, K. ROHR, A. KAUL, S. BÜHLER, R. PEPPERKOK, T. LENGAUER, M. ALBRECHT, R. EILS, P. SCHIRMACHER, V. LOHMANN AND R. BARTENSCHLAGER. Recruitment and activation of a lipid kinase by Hepatitis C virus NS5A is essential for integrity of the membranous replication compartment. *Cell Host & Microbe*, 9(1):32–45, 2011.
- [38] I. RESHETOUSKI, A. MANAKOV, H.-P. SEIDEL AND I. IHRKE. Three-dimensional kaleidoscopic imaging. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 353–360.
- [39] M. ROHRBACH, M. STARK AND B. SCHIELE. Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In *2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Colorado Springs, USA, 2011, pp. 1641–1648. IEEE.
- [40] H. SAIGO, A. ALTMANN, J. BOGOJESKA, F. MÜLLER, S. NOWOZIN AND T. LENGAUER. Learning from past treatments and their outcome improves prediction of in vivo response to anti-HIV therapy. *Statistical Applications in Genetics and Molecular Biology*, 10:Article 6, 2011.
- [41] S.-E. SCHELHORN, J. MESTRE, E. ZOTENKO AND M. ALBRECHT. Inferring physical protein contacts from large-scale purification data of protein complexes. *Molecular & Cellular Proteomics*, 10(6):M110.004929, 2011.
- [42] K. SCHERBAUM, T. RITSCHER, M. HULLIN, T. THORMÄHLEN, V. BLANZ AND H.-P. SEIDEL. Computer-suggested facial makeup. *Computer Graphics Forum (Proceedings EUROGRAPHICS 2011)*, 30(2):485–492, April 2011.
- [43] M. STIKIC, D. LARLUS, S. EBERT AND B. SCHIELE. Weakly supervised recognition of daily life activities with wearable sensors, 2011.
- [44] 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.
- [45] V. SVICHER, V. CENTO, R. SALPINI, F. MERCURIO, M. FRAUNE, B. BEGGEL, Y. HAN, C. GORI, L. WITTKOP, A. BERTOLI, V. MICHELI, G. GUBERTINI, R. LONGO, S. ROMANO, M. VISCA, V. GALLINARO, N. MARINO, F. MAZZOTTA, G. M. DE SANCTIS, H. FLEURY, P. TRIMOULET, M. ANGELICO, G. CAPPIELLO, X. X. ZHANG, J. VERHEYEN, F. CECCHERINI-SILBERSTEIN AND C. F. PERNO. Role of Hepatitis B virus genetic barrier in drug-resistance and immune-escape development. *Digestive and Liver Disease*, 43(12):975–983, December 2011.
- [46] N. TANDON, G. DE MELO AND G. WEIKUM. Deriving a web-scale common sense fact database. In W. Burgard and D. Roth, eds., *AAAI*, 2011. AAAI Press.
- [47] A. TEVS, A. BERNER, M. WAND, I. IHRKE AND H.-P. SEIDEL. Intrinsic Shape Matching by Planned Landmark Sampling. In O. Deussen and M. Chen, eds., *Computer Graphics Forum (Proc. EUROGRAPHICS)*, Llandudno, UK, 2011, pp. 543–552. Blackwell.
- [48] L. TOLOSI AND T. LENGAUER. Classification with correlated features: unreliability of feature ranking and solutions. *Bioinformatics*, 27(14):1986–1994, May 2011.
- [49] T. TUYTELAARS, M. FRITZ, K. SAENKO AND T. DARRELL. Kernelized nbnn. In *IEEE International Conference on Computer Vision*, Barcelona, Spain, 2011. IEEE.
- [50] T. WEINKAUF, H. THEISEL, A. V. GELDER AND A. PANG. Stable Feature Flow Fields. *IEEE Transactions on Visualization and Computer Graphics*, 17(6):770–780, June 2011.
- [51] T. WITTKOP, D. EMIG, A. TRUSS, M. ALBRECHT, S. BÖCKER AND J. BAUMBACH. Comprehensive cluster analysis with Transitivity Clustering. *Nature Protocols*, 6(3):285–295, 2011.
- [52] C. WOJEK, S. WALK, S. ROTH AND B. SCHIELE. Monocular 3D scene understanding with explicit occlusion reasoning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, Colorado Springs, 2011, pp. 1193–2000. IEEE.
- [53] C. WU, B. WILBURN, Y. MATSUSHITA AND C. THEOBALT. High-quality shape from multi-view stereo and shading under general illumination.

In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Colorado Springs, 2011, pp. 969–976. IEEE.

[54] F. XU, Y. LIU, C. STOLL, G. BHARAJ, H.-P. SEIDEL, J. KAUTZ AND C. THEOBALT. Video-based characters - creating new human performances from a multi-view video database. *ACM Transactions on Graphics (Proceedings of ACM Siggraph 2011)*, 30(4):1–20, 2011.

[55] M. A. YOSEF, J. HOART, M. SPANIOL AND G. WEIKUM. AIDA: An Online Tool for Accurate Disambiguation of Named Entities in Text and Tables. In *37th International Conference on Very Large Data Bases*, Seattle, USA, 2011. VLDB Endowment.

[56] M. J. ZILLER, F. MÜLLER, J. LIAO, Y. ZHANG, H. GU, C. BOCK, P. BOYLE, C. B. EPSTEIN, B. E. BERNSTEIN, T. LENGAUER, A. GNIRKE AND A. MEISSNER. Genomic distribution and inter-sample variation of non-CpG methylation across human cell types. *PLoS Genetics*, 7(12):e1002389, December 2011.

# 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 Campus”

## from Frankfurt or Mannheim on Autobahn A6

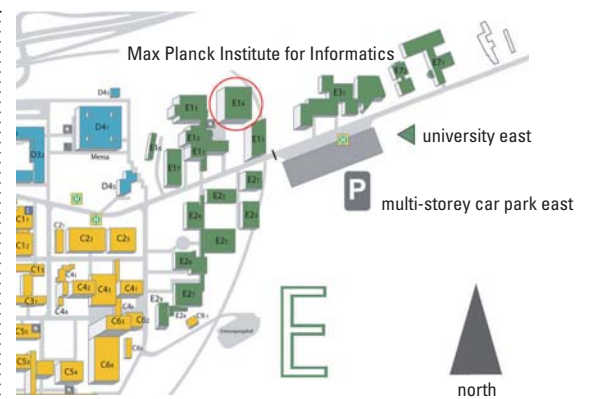
exit “St.Ingbert-West”

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



## I M P R I N T

**Publisher**

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

**Editors & Coordination**

Volker Maria Geiß  
Kamila Kolesniczenko  
Manuel Lamotte-Schubert  
Thomas Lengauer  
Kurt Mehlhorn  
Jennifer Müller  
Marcus Rohrbach  
Thomas Sauerwald  
Bernt Schiele  
Jens Schmidt  
Hans-Peter Seidel  
Bertram Somieski  
Ingolf Sommer  
Martin Theobald  
Christian Theobalt  
Uwe Waldmann  
Christoph Weidenbach  
Gerhard Weikum

**Proofreading**

Krista Ames  
Margaret Delap  
Kamila Kolesniczenko

**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](mailto:info@mpi-inf.mpg.de)  
Internet <http://www.mpi-inf.mpg.de>

**Report Period**

1<sup>st</sup> of January, 2011 to 31<sup>st</sup> of December, 2011

**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](mailto:info@mpi-inf.mpg.de)  
Internet <http://www.mpi-inf.mpg.de>



max planck institut  
informatik