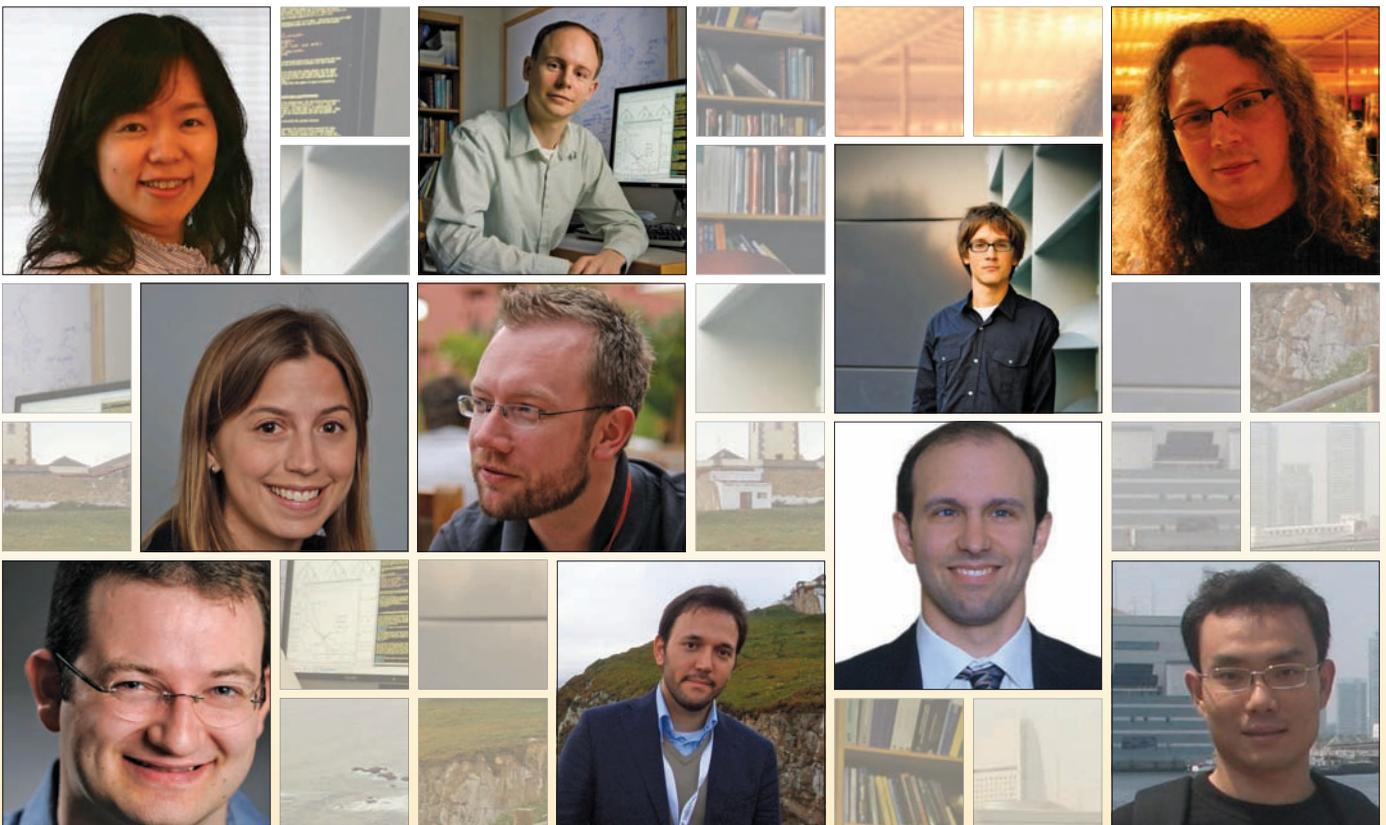


AI's 10 to Watch

It has been a tradition for *IEEE Intelligent Systems* to acknowledge 10 accomplished AI researchers in their early careers every two years as “AI’s 10 to Watch.” These researchers have all completed their doctoral work in the past couple of years. A wide range of senior AI researchers across the world, including primarily academics but also industry researchers, were contacted to nominate young star researchers in all disciplines of AI. A committee comprising members

of our advisory and editorial boards reviewed these nominations and unanimously agreed on the top 10, whom we present to you in this special section. We would like to congratulate these young researchers for winning this special recognition. For the *IEEE Intelligent Systems* readers like you, we are proud to present this glimpse of the future of AI and the practicing AI researchers who promise to be the leaders of the field.

—Fei-Yue Wang





Yiling Chen

Harvard University

Yiling Chen is an assistant professor of computer science at Harvard University's School of Engineering and Applied Sciences. Prior to her appointment at Harvard, she was a postdoctoral research scientist at Yahoo Research, New York. Chen has an MS from the School of Economics and Management at Tsinghua University, Beijing, and a PhD in information sciences and technology from the Pennsylvania State University. Her PhD dissertation, "Markets as an Information Aggregation Mechanism for Decision Support," received the eBRC Doctoral Support Award and the Elwood S. Buffa Doctoral Dissertation Honorable Mention in Decision Science. She won the Outstanding Paper Award at the 2008 ACM Conference on Electronic Commerce and is a recipient of a National Science Foundation Career Award in 2010. Chen is currently the editor in chief of *ACM SIGecom Exchanges*, the newsletter of the ACM's special interest group on e-commerce. Contact her at yiling@eecs.harvard.edu.

A I ' s 1 0 t o W a t c h

Prediction Markets in AI

As computational and social systems are increasingly being developed and used by multiple entities with differing objectives and interests, aligning incentives of participants with an overall system's goals is crucial for success.

My research, situated at the interface between computer science and economics, focuses on analyzing and designing social computing systems according to both computational and economic objectives. My recent work has centered on market-based information elicitation and aggregation mechanisms. Eliciting and aggregating dispersed information is a ubiquitous need for informed decision making. *Prediction markets* have shown great potential for effective information aggregation. A prediction market offers a contract with future payoff that is tied to an event's outcome. Participants express their opinions on the event outcome through trading the contract. The market price hence potentially incorporates the information of all participants and approximately represents a real-time consensus forecast for the event.

This area is naturally amenable to real-world empirical analysis, but my work in this area has focused on

providing strong theoretical and computational foundations for such markets. First, if participants misrepresent their information in the market, the whole endeavor of information aggregation might be called into question. With my collaborators, I have examined game-theoretic equilibria of prediction market play and characterized when participants would play truthfully, bluff, or withhold information. The results answered an important question that had been open for several years and provided a solid ground to consider incentive alignment in designing prediction markets. Second, to obtain and aggregate richer and more refined information, it is highly desirable to provide participants more freedom to express their information. However, the greater expressiveness often comes with an increased computational cost. My research investigated the computational complexity of operating combinatorial prediction markets for various expressive betting

languages and proposed the first tractable market-maker driven combinatorial market. Third, my collaborators and I established a precise connection between a general class of market maker mechanisms and a class of learning algorithms. The results showed that any market maker mechanism of the class could be turned into a learning algorithm, and vice versa. This is the first-known formal connection between markets and learning algorithms. It opens the door for researchers to take advantage of the already vast literature on machine learning to design markets with better properties and would likely benefit machine learning research by introducing the consideration of incentives.

My research also spans other social computing areas. By mining the search query data, my collaborators and I demonstrated that influenza-related searches would rise weeks before an increase in influenza activities. This work showed the possibility of using search term surveillance for early disease outbreak detection. I have also been actively doing research on designing incentives to elicit desirable behaviors from agents or players in various processes. Such work might lead to improved incentives for many online systems that rely on voluntary participation to thrive.

AI and Economic Theory

My research focuses on issues in the intersection of computer science (especially AI) and economics (especially microeconomic theory). Each of these two disciplines can contribute significantly to the other. On the one hand, as computer

systems become more interconnected, multiple parties must interact in the same environment and compete for scarce resources, which necessarily introduces economic phenomena. Here, economic theory can contribute to computer science—although the specific environments faced by computer scientists often require new contributions. On the other hand, in the past, deployed economic mechanisms (such as auctions and exchanges) have been designed to require limited computing and communication resources, even though economic theory allows for much more powerful mechanisms in principle. Computer science can contribute to economics by allowing us to fully operationalize economic theory, resulting in more efficient and entirely novel mechanisms, which requires the design of new algorithms and other contributions. I am devoted to exploring both of these directions.

My work includes the design of new marketplaces and other

mechanisms that let humans and software agents express their preferences naturally and accurately and that generate good outcomes based on these preferences. It also includes the design of software agents that can act strategically in settings where multiple parties all pursue their own interests. This requires the use of concepts from game theory, as well as operationalizing these concepts by finding efficient algorithms for computing the corresponding solutions.

Game theory concerns how to act optimally in environments with other agents who have their own preferences. Some of my best-known work focuses on designing algorithms for computing (or learning) game-theoretically optimal strategies. Such algorithms are now used in several real-world security applications, including the strategic placement of security resources in several US airports as well as with Federal Air Marshals. Our earlier work on

computing Stackelberg strategies has been credited as a launching point for these programs, and we continue to work on new algorithms for similar applications.

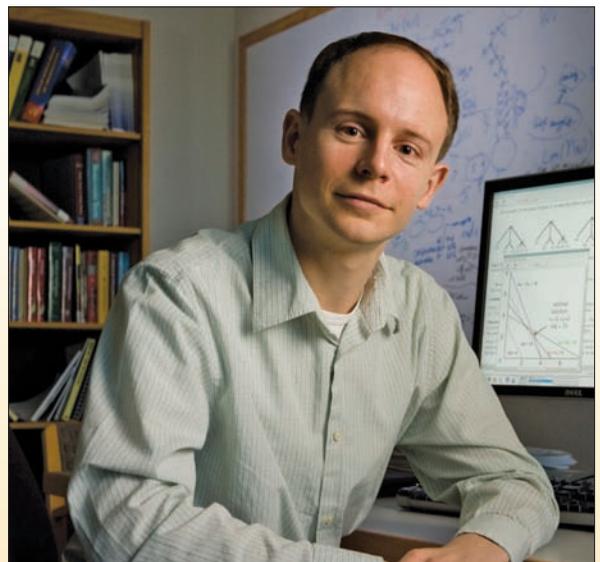
Mechanism design concerns settings where we can design the game—for example, designing an auction mechanism to allocate scarce resources or an election mechanism for choosing a joint plan of action. Among other work, we outlined and developed the general agenda of automated mechanism design, where we let a computer search through the space of possible mechanisms in an intelligent way. Using these techniques, we have obtained new results in economic theory, some of which were recently published in the leading game theory journal. More generally, my group performs research in computational social choice, a rapidly growing subarea in the AI community.

I am generally interested in addressing settings in computer science where multiple parties each act in their own self-interest. In particular, we have been working on the design of mechanisms that are robust in highly anonymous environments such as the Internet, where it is easy for an agent to participate under multiple identities.

AI's To Watch

Vincent Conitzer Duke University

Vincent Conitzer is an assistant professor of computer science and economics at Duke University. He has an AB in applied mathematics from Harvard University and an MS and PhD in computer science from Carnegie Mellon University. Conitzer has received a National Science Foundation Career Award in 2010, an Outstanding Paper Award for AAAI 2008, and an Alfred P. Sloan Research Fellowship. His dissertation, “Computational Aspects of Preference Aggregation,” received the 2006 IFAAMAS Victor Lesser Distinguished Dissertation Award and a honorable mention for the 2007 ACM Doctoral Dissertation Award. Contact him at conitzer@cs.duke.edu.



Semantic Technology Applications

The key theme underlying my research activities concerns investigating and extending the possibilities of concrete applications of semantic technologies. In particular, this theme is pursued by exploring new architectures and epistemological

paradigms for semantic technologies investigating the use of semantic technologies in novel domains, and combining different reasoning methods (both semantic and non) in innovative ways. During my PhD research at the LORIA Laboratory in Nancy, France, I worked with experts in cancer treatment to create new approaches to support the decision-making process in this knowledge intensive area. This required the use, extension, and combination of a variety of techniques, including description logic reasoning, case-based reasoning, and multi-viewpoint representations.

Since then, I have been working at the Open University's Knowledge Media Institute, at the forefront of research related to the Semantic Web. Here I have been leading activities around new infrastructures, methods, and approaches making it possible for intelligent applications to efficiently exploit the knowledge that is now available online and is

distributed and modeled according to heterogeneous ontologies. This research has led in particular to the development of the Watson Semantic Web search engine (<http://watson.kmi.open.ac.uk>), which provides a platform enabling applications to dynamically exploit online ontologies.

Among the various environments, domains, and applications in which I have explored this approach to exploiting the Semantic Web, I am now conducting research on monitoring and making sense of the activity of a Web user, using semantic technology and Semantic Web knowledge to interpret and make sense of online, personal information exchange. This approach can help support the understanding and management of personal information exchanges, or *semantic weblifelogging*.

Also at the Open University, I am director of the Lucero project, which involves making university-wide

resources available to everyone in an open, linked data approach. We are building the technical and organizational infrastructure for institutional repositories and research projects to expose their data on the Web as linked data (<http://data.open.ac.uk>) and pioneering novel methods to deriving knowledge from the aggregation of interlinked, distributed and heterogeneous data sources.

Another key aspect of my research is its focus on large-scale empirical studies of the distributed and heterogeneous knowledge available online, which is very different in nature from classic work in knowledge representation that typically lacks an empirical element. This type of research has only now become possible because of the wide availability of large amounts of formalized knowledge, distributed through the Semantic Web, and of the realization of platforms such as Watson to exploit such online knowledge. For example, it has allowed me to investigate the notions of agreement, disagreement, and consensus and controversies in online knowledge. I am now focusing on extracting interesting patterns in the way Web knowledge bases are modularized or the way they evolve.



A l s o t o W a t c h

Mathieu d'Aquin

The Open University

Mathieu d'Aquin is a research fellow in the Knowledge Media Institute at the Open University, UK. He is also the director of the Lucero JISC Funded project (<http://lucero-project.info>), exposing and using linked data at the Open University (<http://data.open.ac.uk>). He is currently "tutoring" a group of 25 distance learners, using online communication and face-to-face tutorials, and teaching AI in the broad sense, including symbolic methods, nature-inspired techniques (such as ant colony optimization), neural networks, and evolutionary computing. D'Aquin has a BSc and MSc in computing and a PhD from LORIA, Nancy University, France. He received the Best Paper Award at the Third Asian Semantic Web Conference (ASWC 2008). Contact him at m.daquin@open.ac.uk.

Kristen Grauman

University of Texas at Austin

Kristen Grauman is the Clare Boothe Luce Assistant Professor in the Department of Computer Science at the University of Texas at Austin. Her research focuses on visual search and object recognition. Before joining UT-Austin in 2007, she received her PhD in computer science from the Massachusetts Institute of Technology, and a BA in computer science from Boston College. Grauman's work on large-scale visual search for learned metrics with Prateek Jain and Brian Kulis received the Best Student Paper Award at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) in 2008, and the pyramid match kernel approach developed in her PhD thesis is often used within object recognition and image search systems. She is a Microsoft Research New Faculty Fellow, and a recipient of an NSF Career award and the Howes Scholar Award in Computational Science. Contact her at grauman@cs.utexas.edu.



AI's 10 to Watch

Visual Search and Analysis

Images and videos are easy to capture and store: handheld digital video cameras record memories, magnetic resonance imaging scanners display fine structures within the body, and cameras on satellites view the Earth's glaciers. Strikingly, however, our

ability to collect large amounts of interesting visual data has outpaced our ability to analyze it. Today we rely heavily on direct human interpretation and are accustomed to primitive search mechanisms based on tags or keywords. My research aims to remove this disparity and transform how we access visual information.

To this end, my group is developing scalable methods to recognize objects, actions, and scenes, and to automatically search large collections of images and videos based on their content. Our current work focuses on two key questions:

- How can we design algorithms that ensure the most effective transfer of human insight into computer vision systems?
- How can we provide large-scale search for meaningful similarity

metrics, which are often coupled with complex visual representations?

The first component entails coordinating the cooperation between machine and human vision systems. Success in recognition problems depends largely on how learning algorithms solicit and exploit human knowledge about visual content. Existing approaches rely on carefully prepared hand-labeled data, and once trained, the models are fixed. In contrast, we are developing visual learners that build their models actively and can discover patterns with minimal and imperfect instruction. Our cost-sensitive active visual learning strategies and unsupervised discovery methods will enable evolving systems that continuously refine their models of objects and activities, drawing on human assistance only where it is most needed.

In the second major component of our research, we are developing fast and accurate metrics to search images or videos. How one measures “nearness” in the space of visual descriptors is crucial to the success of many tasks, whether retrieving similar exemplars, mining for noisy instantiations of a theme, or constructing classifiers between object categories. However, the computational complexity of good visual similarity measures often precludes their use in large-scale problems. Thus, we are exploring efficient data structures and embeddings to accommodate preferred metrics. Among our contributions thus far are linear-time correspondence measures and sub-linear time search strategies for families of learned metrics and arbitrary kernel functions. In practice, this means that by examining only a small fraction of millions of images, our algorithms can provide nearly the same response as if they were to scan every one. Such fast visual search capabilities lay the groundwork for emerging data-driven approaches to computer vision problems.



Tom Heath

Talis

Tom Heath is lead researcher at Talis. He has a BSc in psychology from the University of Liverpool and a PhD in computer science from the Knowledge Media Institute at the Open University. Heath's dissertation, "Information-Seeking on the Web with Trusted Social Networks: From Theory to Systems," received the PhD of the Year Award from Semantic Technologies Institute International in 2009, and he was the winner of the Semantic Web Challenge at the International Semantic Web Conference in 2007. Contact him at tom.heath@talis.com.

A I ' s 1 0 t o W a t c h

Linked Data Research

As lead researcher at Talis, a Birmingham, UK, software company that focuses on developing and commercially exploiting Linked Data and Semantic Web technologies, I have been responsible for building the company's research capacity. I have

established an internship program and the funding of PhD studentships at leading universities. I have also worked to define the direction of the company's research program, which spans the breadth of key issues related to realization of the Semantic Web vision, from methods to increase the density and diversity of links between data sets in the Semantic Web, to techniques for enhancing the utility of data through mining and collective intelligence, to questions of how the transition from a Web of documents to a Web of data may change the nature of human interaction with the Web.

These topics are key to the future of AI. They not only acknowledge the significance of the Web as a source of background knowledge but are explicitly predicated on its existence as a medium for connecting and sharing data, upon which novel, intelligent applications can be built. Understanding this potential, and how it can be translated into applications and services that benefit Talis customers and the world at large, are fundamental components of the research culture we are building in the company.

Prior to joining Talis, my doctoral research took an interdisciplinary approach to investigating key problems in computing, such as how to realize

a more personally relevant search experience on the Web. Specifically, it focused on using the Semantic Web to support recommendation-seeking in social networks. The aim was to understand how people choose word of mouth information sources from among members of their social networks. I also looked at the factors that affect these source-selection decisions, modeled these processes as computational algorithms so they can be replicated in an online environment, and built Semantic Web-based systems that support people in information seeking through their trusted social networks. This was achieved by exploring how trust relationships influence the choice of information sources when seeking recommendations from social networks, before developing a computational framework to evaluate the potential for these factors to enhance or even redefine Web-based information seeking.

Large-Scale Web Data Analysis

The Web is a vast information resource that captures the pulse of humanity: what we are thinking, what we are doing, and what we know. By studying the Web, we can study human activity and thought. At the same time, by studying how it

is used, we can improve Web technology so it better serves humanity. Thus, the Web has become arguably the most important and largest information resource of all time. Through the emergence of online social networks, social media applications, and social gaming applications, the daily activities of millions of people are migrating to the Web and leaving massive digital traces. With such data, we can create networks of interactions that can be studied and analyzed. This provides enormous potential, letting us study whole societies, human activity, and thought.

To address these challenges and harness the opportunities the Web puts forward, my research consists of analyzing models of network structure and evolution; developing models and algorithms to capture, model, and predict the pulse

of communities and society; and working with massive data sets of terabyte scale as certain behaviors and patterns are observable only when the amount of data is large enough.

My principal research interest is in large-scale data mining and machine learning, focusing on the analysis and modeling of large real-world networks as the study of phenomena across the social, technological, and natural worlds. Such graphs frame numerous research problems and high-impact applications. For example, social networks on the Internet generate revenue of multiple billions of dollars, detection of virus outbreaks can save lives, and anomaly detection in computer-traffic networks is vital for security. My long-term research goal is to harness large-scale social and information networks to understand, predict,

and ultimately, enhance social and technological systems. I would like to create explanatory and predictive models of actions of large groups of people and societies, and large technological systems.

Through my work, I have addressed a number of important questions regarding the properties and patterns of large evolving networks by revealing how local behavior and structure lead to large-scale phenomenon and use full applications. What does a “normal” network look like? How will new links be created over time? Is the network or a community “healthy”? How does information spread? How does information mutate as it spreads? Which other nodes in the network does a given node trust or distrust? Who are friends and foes of a node? How can we identify and find influential nodes or select nodes to immunize in networks? Answers to such questions are vital to a range of application areas from the identification of illegal money-laundering rings, misconfigured routers in the Internet, viral marketing, and protein-protein interactions to disease outbreak detection.

A I ' s T O t o W a t c h

Jure Leskovec

Stanford University

Jure Leskovec is an assistant professor in the Department of Computer Science at Stanford University. He has a BSc in computer science from the University of Ljubljana, Slovenia, and a PhD in computational and statistical learning from Carnegie Mellon University. Leskovec received the ACM SIGKDD Dissertation Award in 2009, the Best Research Paper Award at ASCE Journal of Water Resources Planning and Management Engineering in 2009, and the Best Student Paper Award at the ACM 13th International ACM SIGKDD Conference on Knowledge Discovery and Data Mining in 2007. Contact him at jure@cs.stanford.edu.



Event Detection and Prediction

With the critical importance of addressing global policy problems such as disease, crime, and terrorism and the continuously increasing mass of available data, machine learning has become essential for the development of new, practical

information technologies that can be applied for the public good. One major focus of my research is event detection, with the goal of developing new statistical and computational methods for early and accurate detection of emerging events and other patterns in massive, complex real-world data sets. My research group is currently applying these methods to three main areas: public health (to achieve early detection of emerging outbreaks of disease), law enforcement (to predict and prevent crime hot spots), and health care (to detect anomalous patterns of patient care and discover new best practices). Other applications include monitoring of water quality and food safety, network intrusion detection, fraud detection, and scientific discovery.

All these applications require integration of noisy, heterogeneous data from many sources. For example, in the early stages of a disease outbreak, we can observe the many subtle ways

that the illness changes the behavior of affected individuals, even before they are sick enough to visit a doctor or hospital. These minor behavior changes, such as increased consumption of over-the-counter medications and changes in mobility patterns, might not be sufficiently out of the ordinary when we look at a single individual. However, they are often visible when we view the problem on a societal scale and when we combine traditional public health data sources (such as hospital visits and medication sales) with population-wide data from emerging technologies such as location-enabled mobile devices, online social networks, and user-generated Web content.

We have developed novel machine learning methods that can achieve much more timely detection of emerging events, accurately characterize the type of event and pinpoint the affected subset of the population, and significantly reduce false

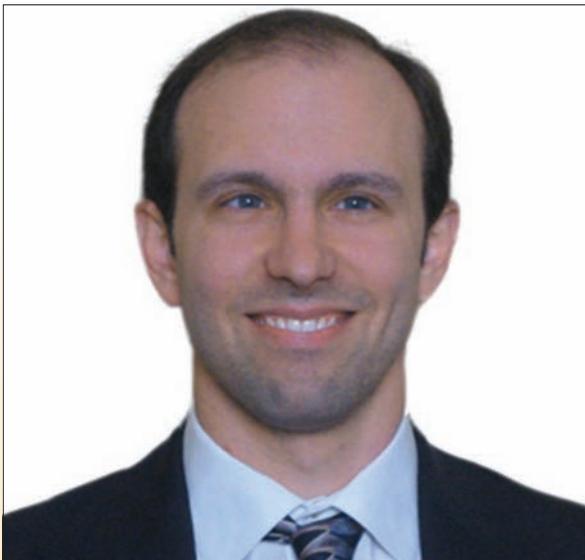
positives. These methods address the fundamental statistical challenges of detecting anomalous patterns from multiple data streams and learning which anomalies are actually relevant to human users. In addition, we address the fundamental computational challenges of scaling up to huge and high-dimensional data. We have developed new, fast optimization algorithms that can solve previously impossible detection problems in a matter of seconds, exactly and efficiently identifying the most relevant subsets of a large, complex data set. Finally, we are working directly with public health practitioners and law enforcement agencies to develop and deploy systems for event detection and prediction. Our detection methods have been incorporated into deployed disease surveillance systems in the US, Canada, India, and Sri Lanka, and our CrimeScan software is in day-to-day operational use by the Chicago Police Department to predict and prevent emerging hot-spots of violent crime. These collaborations will facilitate rapid and effective responses to disease outbreaks, crime waves, and other emerging events, directly contributing to the public good by improving health, safety, and security.

Also to Watch

Daniel B. Neill

Carnegie Mellon University

Daniel B. Neill is an assistant professor of information systems at Carnegie Mellon University's Heinz College, where he directs the Event and Pattern Detection Laboratory and the PhD program in Machine Learning and Public Policy. He also holds courtesy appointments in Carnegie Mellon's Machine Learning Department and Robotics Institute and is an adjunct professor in the University of Pittsburgh's Department of Biomedical Informatics. Neill has an MPhil from Cambridge University and an MS and PhD in computer science from Carnegie Mellon. He is the recipient of a National Science Foundation Career Award and an NSF Graduate Research Fellowship, and he received the Best Research Presentation Award at the 2005 National Syndromic Surveillance Conference. Contact him at neill@cs.cmu.edu.



Verifying Cyberphysical Systems

Many intelligent systems are embedded in cyberphysical systems (CPS), including parts of intelligent driver assistance technology for cars and autopilots in aircraft. How do we know that their designs will work as intended? Most initial designs do not, and some deployed systems still don't. Ensuring the correct functioning of CPS is a central challenge in computer science, mathematics, and engineering because it is the key to designing smart and reliable control. Scientists and engineers need analytic tools to understand and predict the behavior of their systems. Yet, there are fundamental mismatches between the actual dynamics of CPS applications and the limits imposed by previous modeling and analysis techniques.

By closing these fundamental analytic gaps, my research addresses an intellectual grand challenge that has substantial scientific, economical, societal, and educational impact arising from the benefits of improved CPS analysis and design. My research focuses on the development of the logical foundations and formal verification techniques for complex intelligent CPS, especially hybrid systems with nontrivial interacting discrete and continuous dynamics. I am particularly interested in designing and analyzing verification techniques and verification logics and

in using these to develop automatic verification tools that help produce reliable complex systems, such as in aeronautical, railway, automotive, robotic, and biomedical applications.

I have developed fully compositional logic-based verification techniques for hybrid systems models—CPS with interacting discrete and continuous dynamics. These techniques have been successfully used to verify collision freedom for intelligent control systems in the railway and air-traffic control domains. I have proven the first completeness result for CPS, which shows that all true properties of hybrid systems can be proven by decomposition from elementary properties of differential equations. Decomposition is a crucial divide-and-conquer technique here because it makes the verification problem tractable.

This result was instrumental in the successes I had with verifying real case studies for intelligent systems from the European Train Control System and from roundabout collision avoidance maneuvers in air-traffic control using KeYmaera, our hybrid verification tool

for hybrid systems that combines deductive, real algebraic, and computer algebraic prover technologies. With my verification techniques, I discovered and fixed a bug in an aircraft collision avoidance maneuver that caused collisions instead of preventing them. Other achievements include contributions to Bayesian statistical model checking, the foundations of the hybrid systems image computation problem, and deductive verification of object-oriented programs.

In a recent breakthrough, I have further developed the first verification approach for distributed hybrid systems such as distributed car control. These systems have been considered for more than two decades, but no previous verification technique has ever been capable of actually verifying their complicated dynamics.

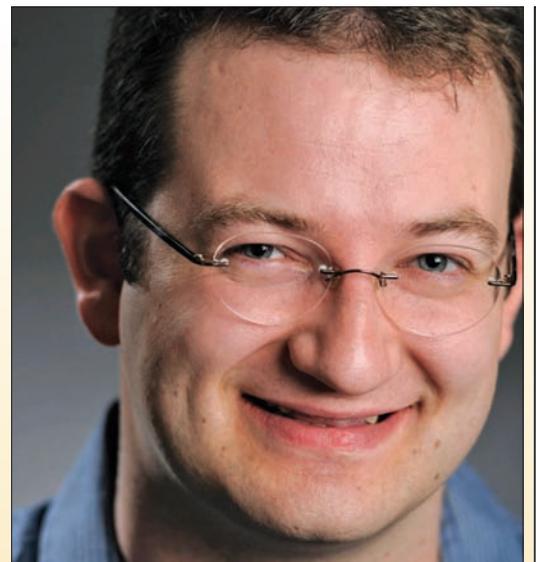
Logic and foundations have played an invaluable role in the development of classical computer science, and its outcome had a substantial impact on practical applications, including hardware design, with all its subsequent societal impact. I am convinced that an investigation of the logical foundations of CPS will undoubtedly have a similar impact on complex physical system design. They will probably have an even bigger societal impact, given the ubiquitous nature of CPS technology and our increasing reliance on this technology and its safety in modern life.

AI's 10 to Watch

André Platzer

Carnegie Mellon University

André Platzer is an assistant professor in the Computer Science Department at Carnegie Mellon University. He has a MSc in computer science from University of Karlsruhe, Germany, and a PhD in computer science from University of Oldenburg, Germany. Platzer received *Popular Science Magazine's* Brilliant 10 Award in 2009 and the ACM Doctoral Dissertation Honorable Mention Award in 2009. He is the author of *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics* (Springer, 2010). Contact him at aplatzer@cs.cmu.edu.





Talal Rahwan

University of Southampton

Talal Rahwan is a senior research fellow in the Intelligence, Agents, Multimedia (IAM) group at the University of Southampton. He has a BEng in informatics from the University of Aleppo, Syria, where he was awarded the Syrian Ministry of Higher Education's Presidential Award for Academic Achievement. He has a PhD in computer science from the University of Southampton. His thesis, "Algorithms for Coalition Formation in Multiagent Systems," earned him the British Computer Society's Distinguished Dissertation Award, which recognizes the best PhD in the UK in computer science. Contact him at tr@ecs.soton.ac.uk.

A | ' | s | T | O | t | o | W | a | t | c | h

Coalition Formation in Multiagent Systems

Multiagent systems represent an important and rapidly growing design paradigm for understanding and building distributed systems. This paradigm is based on modeling the computational components as agents—autonomous entities (such

as software programs or robots) that can decide for themselves what they need to do to satisfy their objectives. Typically, these agents need to interact with one another to improve their performance and compensate for each others' deficiencies. In this context, one of the most fundamental types of interactions is *coalition formation*, which involves creating goal-directed and short-lived coalitions (such as teams); they are formed with a purpose in mind and dissolve when that purpose no longer exists, or when they cease to suit their design purpose. Potential applications include multisensor networks, grid computing, and distributed vehicle routing.

Forming effective coalitions is a key research challenge in the field of multiagent systems. Central to this

endeavor is the coalition structure generation problem, which involves determining which of the many possible coalitions to form to optimize the performance of the system as a whole (an NP-complete problem with an exponentially growing search space).

During the past seven years, I have developed a completely new approach to this problem. Specifically, my multidisciplinary approach draws on research from computer science, game theory, operations research, and combinatorial optimization. It builds on novel representations of the search space and combines different search techniques (such as linear programming, dynamic programming, divide and conquer, branch and bound, and depth-first search). In so doing, I developed a number of novel

algorithms that were several orders of magnitude faster than extant approaches. These advances provided new impetus to the field and have sparked significant interest from contemporary researchers.

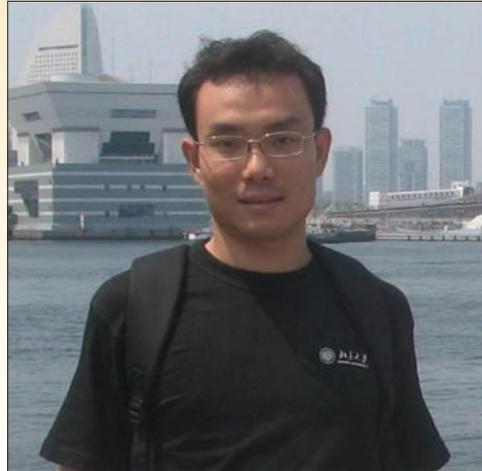
I also have recently focused on another fundamental question: If the space of possible solutions is too large to be fully searched, then can we search through only a subset of this space and be guaranteed to find a solution that is within a specified bound (that is, a percentage, say 50 percent) from the optimal? If so, what is the smallest such subset? This is a long-standing open question in the field that I have recently solved. In particular, I am now able to identify the smallest subset that must be searched given "any" specified bound.

My ambition is to make coalition formation technology applicable to whole new classes of applications and on a totally different scale than hitherto possible. In so doing, I will come closer to realizing my long-term goal of contributing the underlying theoretical foundations to the area of agile teaming problems.

Liwei Wang

Peking University

Liwei Wang is an associate professor in the Department of Machine Intelligence at Peking University, China. He has a BS and MS in electronics engineering from Tsinghua University and a PhD in applied mathematics from Peking University. His honors include the *Pattern Recognition Letters* Top Cited Article 2005–2010 and a coauthored paper that won the MIRU outstanding paper award. Outside of work, Wang enjoys playing GO, badminton and volleyball. Contact him at wanglw@cis.pku.edu.cn.



A I ' s T O t o W a t c h

Machine Learning Theory

Making computers learn as well as people learn is one of the major goals of AI. To achieve this goal, two problems need to be solved: developing practical learning algorithms, and understanding the power and limitations of learning.

The latter is what machine learning theory studies.

Assume that we want a computer to learn a rule that decides whether an email is a spam, given an amount of training data consisting of a set of emails and their corresponding labels (made by humans) indicating spam or nonspam. Before trying to solve it, we might want to ask if this task is learnable at all. Learning theory aims to answer this fundamental question by characterizing the conditions under which a task is learnable, and the best possible performance for learnable problems. Classical learning theory pointed out that the capacity of the space from which the learned rule is chosen is crucial to learnability, and there is a subtle trade-off between the capacity and the number of training data to obtain good performance. This result sheds light not only on the nature of

machine learning, but human learning as well.

My research focuses on establishing theory on new learning paradigms that generalize those studied in traditional learning theories. Previously, machine learning systems always assumed that data are feature representations of objects. Although such a representation provides many analytical tools and makes it convenient to develop theories and algorithms, extracting features that are useful for learning is often difficult. On the other hand, in some applications (computer vision, for example) defining dissimilarity between a pair of objects is relatively easy. Thus, an alternative approach is to learn from data that are pairwise dissimilarities of objects, rather than features. In this setting, we proposed sufficient conditions under which a problem is learnable and find interesting applications in computer vision.

Another learning paradigm that is of great interest to me is active learning. In contrast to traditional supervised learning in which the learner (computer) receives training data passively from the teacher (human), in active learning the learner can actively select which data to request. The hope is that active learning can achieve the same performance as passive learning with fewer training data. However it can be shown that no active learning algorithm is always better than supervised learning—there is no free lunch in learning. Hence, the most important question for active learning is whether there is any reasonable condition under which active learning is superior. I proved that boundary smoothness guarantees that active learning reduces the amount of training data. It is a polynomial reduction for finite smoothness and an exponential reduction for infinite smoothness.

I believe there are many sufficient conditions under which problems are learnable, and each condition leads to a powerful learning algorithm. The challenge for me is to discover these conditions and then discover the learning algorithms.

This article was featured in

computing **now**

ACCESS | DISCOVER | ENGAGE

For access to more content from the IEEE Computer Society,
see computingnow.computer.org.



IEEE

IEEE  **computer society**

Top articles, podcasts, and more.



computingnow.computer.org