The Future of AI

AI's 10 to Watch

IEEE *Intelligent Systems* once again selected 10 young AI scientists as "AI's 10 to Watch." This biennial celebration of the young stars in the field has been very well-received by the AI community and the IEEE Computer Society. This acknowledgment and celebration not only recognizes these young scientists and makes a positive impact in their academic career but also promotes the community and cutting-edge AI research among next-generation AI researchers, the industry, and the general public alike.

In early 2015, *IEEE Intelligent Systems* solicited nominations for this recognition from a wide range of senior AI researchers from both academia and industry. The nominees all received their PhDs in the past five years. A short list of top candidates was voted on by the selection committee, consisting of members of the *Intelligent Systems* editorial and advisory boards. The final decisions were made by the entire boards. I would like to take this opportunity to thank Raymond Perrault, who served as the Chair of this year's selection committee and did a great job managing the selection process. We also owe our thanks to the members of the selection committee who devoted a lot of time studying the nomination materials and deliberating very carefully about the best our community can offer. In the end, the top 10 surfaced with unanimous support from the advisory and editorial boards. We're particularly pleased about the diversity of the winning group, the breadth of topic coverage, and the global nature of these award-winning works.

IEEE Intelligent Systems presents to its readership and AI researchers around the world the 2015 list of AI's 10 to Watch. We're very proud about these young AI scientists' innovative contributions and impact. We wish the best for their continued excellence and sustained impact in their future careers!

—Daniel Zeng, editor in chief

1541-1672/16/$33.00 © 2016 IEEE
Published by the IEEE Computer Society

Haris Aziz

Data61 and University of New South Wales

Haris Aziz is a senior researcher at Data61 and a conjoint senior lecturer at the University of New South Wales, Sydney. He's also an invited fellow of the Centre for Policy and Market Design, University of Technology Sydney, and was an invited professor at University Paris Dauphine. His research interests lie at the intersection of artificial intelligence and economics, especially computational social choice and algorithmic game theory. Aziz received a PhD in computer science from the University of Warwick, an MSc in mathematics and foundations of computer science from the University of Oxford, and a BSc (Honors) in computer science from the Lahore University of Management Sciences. Contact him at haris.aziz@nicta.com.au.

Collective Decision Making in Multi-Agent Systems

Making collective decisions about allocating resources to agents or agents to teams in a fair, principled, and mutually beneficial manner is a central problem, whether the agents are humans, robots, or some other autonomous entities. Common features of multi-agent settings are that agents express preferences or utilities over allocation outcomes. Based on these preferences as well as other information such as priorities, endowments, and contributions, the allocation is made. For all multi-agent settings, designing formal models and efficient and principled algorithms to make collective decisions is a fundamental goal.

There are many important problems in microeconomics, where economists have largely ignored the computational issues that are necessary to build scalable systems. Artificial intelligence with its toolkit of optimization techniques, tradeoff analysis, and algorithm design is ideal to tackle such problems. In the other direction, concepts from microeconomics, especially social choice theory, mechanism design, fair division, and cooperative game theory, are suitable to reason about optimality, fairness, stability, and incentive properties in multi-agent systems.

A typical challenge is to incentivize agents to report their truthful preferences so that optimization can be done on the correct input. My work is at this very active interface between artificial intelligence and microeconomics. I try to design collective decision-making algorithms and protocols that not only satisfy compelling axiomatic properties but are also computationally efficient. This invariably requires understanding the tradeoffs between the properties.

A fundamental problem in coalition formation and discrete allocation is that of computing Pareto optimal outcomes. My research group came up with a versatile algorithm to solve this problem. A long-standing problem concerning envy-free allocation of divisible goods has been to design a protocol for more than three agents that requires bounded number of queries. A colleague and I solved this problem.

In contrast to work on deterministic voting and allocation rules, randomized mechanisms are relatively less explored. I've been investigating the use of randomization to devise protocols and voting rules that have desirable axiomatic properties. For example, a colleague and I proposed a new randomized voting rule that significantly generalizes a prominent assignment rule and provides an interesting connection between allocation of private and public goods.

As the digital world gives rise to new markets and complex cooperative settings, and multi-agent systems require principled protocols and mechanisms, the interplay between AI and economics will be crucial to come up with innovative solutions. As a researcher, it's an exciting time to be working in this area.

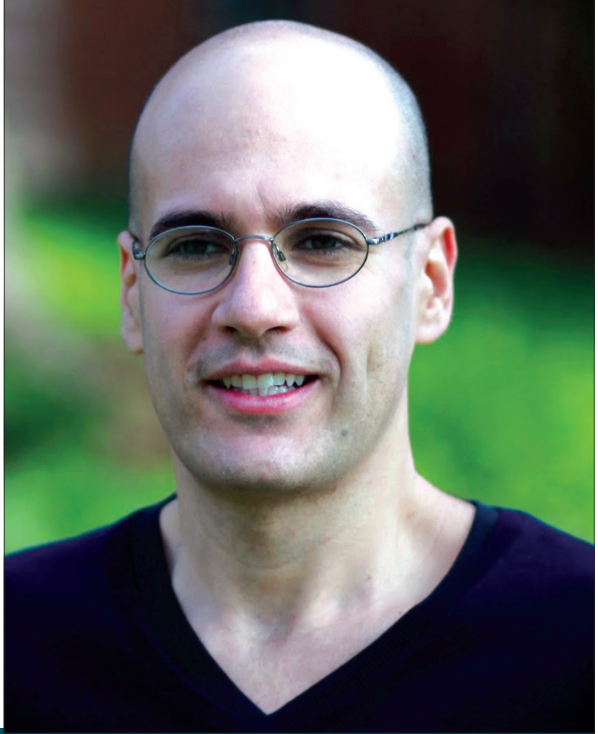

AI'S 10 TO WATCH

Elias Bareinboim

Purdue University

Elias Bareinboim is an assistant professor in the Department of Computer Science at Purdue University, with a courtesy appointment in Statistics. His research focuses on causal and counterfactual inference and their applications to data-driven fields. Bareinboim received a PhD in computer science from the University of California, Los Angeles. His thesis work was the first to propose a general solution to the problem of "data fusion" and provides practical methods for combining datasets generated under different experimental conditions. Bareinboim's recognitions include the Dan David Prize Scholarship, the Edward K. Rice Outstanding Doctoral Student, the Yahoo! Key Scientific Challenges Award, and the 2014 AAAI Outstanding Paper Award. Contact him at eb@purdue.edu.

From Causal Inference and Data Fusion to an Automated Scientist

The words "artificial intelligence" often connote futuristic speculations about how smart machines could become and whether they would eventually take over the planet. But far from the limelight of such extrapolations, a quiet AI revolution has already taken place, one that has profoundly transformed the way scientists look at the world, the language they use to interpret data, and the methods they use to assess cause and effect relationships in many social and medical domains. The advent of graphical methods of causal and counterfactual inference has made it possible to tackle some of the most challenging problems in scientific methodology. These methods have reignited hopes of constructing systems (software, machines, robots) capable of acting like human scientists, and ultimately automating the process of scientific discovery.

With the recent unprecedented accumulation of data, researchers in the empirical fields are becoming increasingly aware of the fact that, to take full advantage of this explosion, traditional statistical techniques, including those based on machine learning, must be enriched with two additional ingredients: the ability to distinguish causal from associational relationships, and the ability to integrate data from multiple, heterogeneous sources.

My research in the past few years has led to a formal theory for handling these two components simultaneously, also known as the "data fusion" problem. Building on the modern language of causation, I developed a theoretical framework for representing and algorithmizing this problem, thus enabling researchers to fuse together data from a heterogeneous mixture of experimental and observational studies as well as generalizing to yet unseen environments.

In the medical and social sciences, this work has resolved several long-standing problems often referred to as "external validity," "selection bias," "transportability," and "experimental generalization," which are pervasive in essentially every nontrivial instance of data analysis. This mathematical framework, in practice, allows scientists to solve various challenges from first principles, which include reducing the cost of data collection and optimizing the design of experiments, predicting in domains with little or no data, and understanding the mechanisms underlying the phenomena being studied. These issues are commonly faced in a wide array of fields, including machine learning, statistics, and the health and social sciences. In the area of robotics, for example, the results of this work can be used to endow intelligent systems with causal-generalization capabilities akin to the work that a human scientist conducts in a laboratory or field study. This means that a robot would be able to probe an environment more effectively and then utilize the knowledge acquired to generalize to a new unexplored setting.

Given the ubiquity of the data fusion problem across empirical disciplines, along with the generality and completeness of current results, I believe that this new framework will be an essential tool for tackling the challenges presented to the next generation of data science research.

Yejin Choi

University of Washington

Yejin Choi is an assistant professor in the Department of Computer Science and Engineering at the University of Washington. Her research combines natural language processing with computer vision, exploring image captioning, multimodal knowledge learning, and visual entailment. She has also worked on modeling connotation and writing style to analyze why the text is written (intent) and by whom (identity). Choi received a PhD in computer science from Cornell University. She was a co-recipient of the Marr Prize at ICCV 2013. Contact her at yejin@cs.washington.edu.

Language, Vision, and Social AI

Natural language is at the heart of our everyday lives. We use it to communicate complex ideas, ranging from summaries of what we've seen or experienced to subtle cues about our beliefs, goals, and opinions. Importantly, meaning is conveyed not just by what's literally said but also by what's left to the listener to infer. This ability to reason beyond what is said explicitly is crucial for efficient human-computer communications. We posit that this ability is learnable from large-scale data due to predictable patterns in the way people use and interpret language in various visual and social contexts.

My research group focuses on developing statistical models of language understanding and production under two broad themes: connecting language with visual intelligence, and connecting language with social and emotional intelligence. Common to all these tasks is that the interpretation we seek is frequently latent in text, so the model must learn to recover nonliteral and implied aspects of meaning. This capability to understand communicative intent and participant mental state will be important for enhancing human-computer communication in a wide range of applications.

Language with Visual Intelligence

Online data—including news, books, movies, and social media—presents unprecedented textual and visual records. From this data, my ultimate goal is to attain deep semantic integration between language and vision at a very large scale. One direction of research under this theme is automatic image captioning, in which models to describe images can be learned from the rich spectrum of descriptive language people use in online photo sharing communities. Another direction is learning statistical knowledge about how the world works by mining large collections of text, images, and other modalities of human communications. In our recent work, for example, we've shown that it's possible to infer visual entailment knowledge from online images and text—when a horse is eating, for example, it's likely that the horse is standing.

Language with Social and Emotional Intelligence

We're also developing new models of social and emotional intelligence in automatic text analysis and production. We have, for example, shown that it's often possible to detect an author's deceptive intent, predict whether writing is successful at deception, and forecast the likely emotional states of communication participants. A recent project under this theme is constructing a broad coverage lexicon of connotation—for example, "science" and "surfing" are typically associated with positive judgments or sentiments, whereas "cancer" and "pollution" are associated with negative emotion. We're building a scalable algorithm framework to infer subtle shades of connotation by observing how people use these words in a large-scale online corpus.

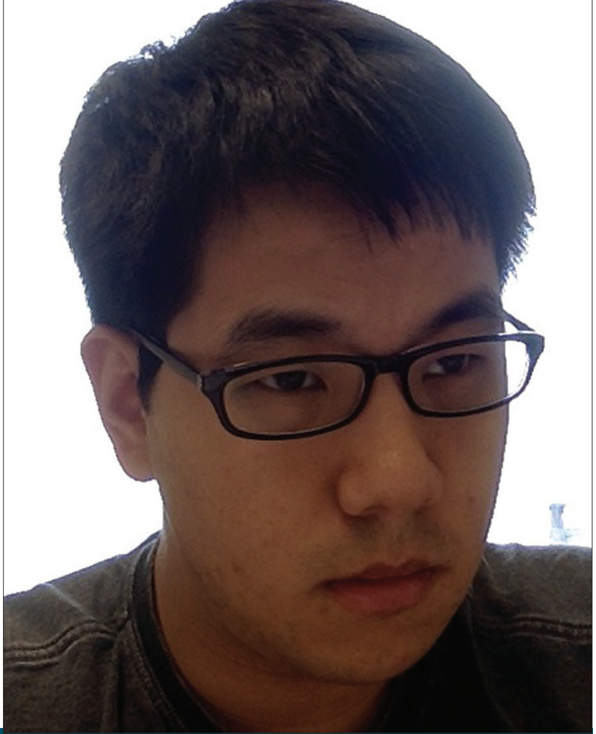

AI'S 10 TO WATCH

Daniel Hsu

Columbia University

Daniel Hsu is an assistant professor in the Department of Computer Science at Columbia University and a member of the Data Science Institute. His research interests focus on algorithmic statistics, machine learning, and privacy. Hsu received a PhD in computer science from the University of California San Diego. Contact him at djhsu@cs.columbia.edu.

Algorithms for Machine Learning

Machine learning is transforming science and technology in ways unimaginable just decades ago. Much of the success is due to the availability of large datasets (whose scale and specificity transcend what was available in the early, formative years of machine learning) and the development of efficient algorithms that process these data to derive statistical models used in applications.

Data and algorithms are, however, also mutually limiting: the effectiveness or applicability of algorithms is limited by the nature of or information contained in data, and the utility of data is limited by how effective algorithms are at extracting useful information. Indeed, insights and otherwise informative signals that can be extracted from unlabeled data are limited by computational barriers to solving hard optimization problems. Unfortunately, most available data are unlabeled and remain unlabeled due to high labeling costs. Thus, except in select domains where existing data and algorithms are already well-matched, much of machine learning's potential remains untapped.

My research tackles these challenges by developing algorithms for learning latent variable models and for learning with minimal supervision. Latent variable models are statistical models with unobserved variables that explain dependencies among observed variables. Understanding these latent variables in applications yields scientific insights and meaningful representations of data; critically, these models can be fit directly to raw, unannotated data. While it's often easy to fit models when data are fully labeled, it's vastly more difficult when only unlabeled data are available. Classical approaches require solving intractable problems and hence can't be used with large, high-dimensional data. I developed new, computationally efficient and statistically consistent algorithms for fitting latent variable models that are ubiquitous in applications. These algorithms have been adopted by both theoreticians and practitioners: they were applied to establish learnability of large Gaussian mixtures and also to analyze chromatin states in epigenetic maps.

In addition, my research addresses challenges in active learning, a model of learning that captures the adaptive process by which labels are obtained. One primary challenge is sampling bias: data that end up being labeled aren't representative of the overall population. Sampling bias can catastrophically derail learning algorithms and, worse yet, can easily go undetected. I developed new active learning frameworks that substantially reduce required labeling effort while simultaneously managing sampling bias for statistical consistency. These frameworks underlie several active learning systems adopted in applications: they've been implemented in popular software packages and have been applied to problems as varied as machine translation and electrocardiogram classification.

AI'S 10 TO WATCH

Shivaram Kalyanakrishnan

Indian Institute of Technology Bombay

Shivaram Kalyanakrishnan is an assistant professor in the Department of Computer Science and Engineering at the Indian Institute of Technology Bombay. His research interests include artificial intelligence and machine learning, spanning topics such as sequential decision making, multi-agent learning, multi-armed bandits, and humanoid robotics. Kalyanakrishnan received a PhD in computer science from the University of Texas at Austin. His contributions to robot soccer have received two Best Student Paper awards at the annual RoboCup competitions. Contact him at shivaram@cse.iitb.ac.in.

Learning Agents

My research is motivated by the goal of creating intelligent agents, especially ones that can learn. In pursuit of this goal, I consider questions from a wide variety of topics. My thesis focused on reinforcement learning (RL), a general paradigm in which an agent, through trial and error, discovers actions that maximize its long-term gain. RL finds application in a variety of domains, including game-playing, stock-trading, medical decision-making, and environmental preservation. Consider the problem of getting a robot to play soccer. The traditional approach is for a programmer to directly specify the robot's behavior in the form of rules such as, "If farther than 21.5 m from the goal, and within 5.6 m of a teammate who makes an angle of more than 23º with any opponent, then pass the ball to that teammate." The attraction of RL is that rather than specifying behavior in this manner, designers merely specify *what is desired* of the behavior. Thus, if the soccer-playing robot is rewarded +1 if the ball gets into the opponent's goal and 0 for every other situation, then by applying an RL algorithm, the robot can eventually learn to play soccer!

Traditional algorithms for RL enjoy good theoretical guarantees when the number of situations in the task is small. However, in practical tasks, the number is typically very large and often even infinite. My thesis demonstrates that when the agent can't perfectly represent the set of situations and their associated long-term rewards, a different class of algorithms (called policy search) can work much better. My colleagues and I applied policy search methods to train a humanoid soccer agent for the RoboCup 3D simulation competition—the resulting agent won the 2011 competitions in Istanbul by a combined score of 136-0 in 24 games.

I routinely encounter theoretical questions while undertaking system-building work. In one such case, I designed sampling algorithms for multi-armed bandits. Suppose we have a set of real-valued random variables, from which we wish to select a small subset of those that have the highest expected values. Our only way to gather any knowledge about the random variables is by sampling them. To identify the top subset with a given level of accuracy, what's the minimum number of samples needed? The theoretical bounds I derived for this problem have significantly advanced the state of the art. Interestingly, much of my contribution to this problem came while I was working at Yahoo—bandits are a natural framework to model online advertising, wherein ads with high click-through rates must be efficiently identified and served.

I recently embarked on an academic career in India, and I look forward to applying AI and machine learning to problems in the local ecosystem.

Reshef Meir

Technion-Israel Institute of Technology

Reshef Meir is a senior lecturer in the Department of Industrial Engineering and Management at the Technion-Israel Institute of Technology. His research interests focus on computational game theory, artificial intelligence, mechanism design, decision making, and behavioral game theory. Meir received a PhD in computer science from the Hebrew University. He was awarded the 2013 Michael B. Maschler Prize to an outstanding research student by the Israeli Chapter of the Game Theory Society. Contact him at reshefm@ie.technion.ac.il.

Strategy and Bounded Rationality

Adam Smith's "invisible hand" principle asserts that when multiple self-interested parties each pursue individual goals without central intervention, parties will converge to a socially desirable state. However, in practice, competition often quells cooperation, and outcomes can be undesirable and even disastrous for society.

In my work, I apply a game-theoretic approach to understand and characterize strategic behavior in multi-agent interactions, where agents are self-interested. The challenges involved include both economic factors and a significant computational aspect, such as computing optimal strategies or an equilibrium point. While part of my research deals with algorithmic problems arising from playing and designing games, I'm more interested in the other direction, of how tools from computer science and artificial intelligence can assist us in solving fundamental economic and game-theoretic questions.

One of the biggest obstacles in applying game theory to real-life problems is that the classic model of a rational agent does not describe accurately enough the strategic behavior of people in most interactions. In fact, any single rule is probably not enough to capture the variety of behaviors demonstrated in a society.

My present research is concerned with making realistic assumptions about the knowledge and capabilities of agents of different types, as well as comparing theoretical predictions with empirical and experimental data. Within this broad field, known by the title *behavioral game theory*, I focus on the perception of uncertainty and how it affects strategic behavior. More specifically, I consider nonprobabilistic ways of representing uncertainty and combine them with game-theoretic notions of equilibrium, to explain the imperfect strategies people adopt in various situations such as voting or commuting.

One example from my research shows how to define an alternative notion of equilibrium in routing games that takes into account arbitrary behavioral biases, and that some of these biases, such as moderate risk aversion under uncertainty, lead to reduced congestion and higher social welfare. Another work that's in more preliminary stages is the construction of a public database for strategic behavior under uncertainty. Data from our strategic voting experiments can already be downloaded from the project's website, www.votelib.org.

Beyond deepening our understanding of games in general, my research is aimed at designing better mechanisms founded on more accurate behavioral models and applying them in the domains of voting, group scheduling, resource allocation, crowdsourcing, routing in networks, and more. Such mechanisms will allow the market's "invisible hand" to work properly and efficiently, for the benefit of everyone.

Suchi Saria

Johns Hopkins University

Suchi Saria is an assistant professor of computer science and health policy at Johns Hopkins University. Her research interests are in statistical machine learning and computational healthcare. Saria received a PhD in computer science from Stanford University. Her work has received recognition in the form of two cover articles in Science Translational Medicine (2010, 2015), a best student paper award by the Association for Uncertainty in Artificial Intelligence (2007), a best student paper finalist by the American Medical Informatics Association (2011), an Annual Scientific Award by the Society of Critical Care Medicine (2014), and competitive awards from the National Science Foundation (2011), the Gordon and Betty Moore Foundation (2013), and Google Research (2014). Contact her at ssaria@cs.jhu.edu.

A Reasoning Engine for Tailoring Healthcare to the Individual

Healthcare is in the early stages of a digital revolution. My research focuses on a question that's fundamental to transforming the delivery of healthcare: How can we design intelligent computational systems that reason about health?

Such systems must flexibly integrate and even solicit many sorts of data about an individual, over time, understanding how the individual fits into patterns observed elsewhere in the population. What they discover can help shape minute-by-minute clinical decisions as well as institutional best practices and new scientific questions.

Many chronic diseases are difficult to treat because of the remarkable degree of variation among affected individuals. Patients exhibit highly variable symptoms, complications, and outcomes—the degree of variability in scleroderma, for example, makes it difficult to choose the optimal therapy for any given patient.

Fortunately, the HITECH Act in 2009 and the Affordable Care Act in 2010 accelerated investment in digital technologies. Detailed data on health measurements, treatments, costs, and outcomes are finally being collected in electronic form. By analyzing such data for an entire population over a long period, we can uncover disease subtypes—subpopulations whose symptoms are more homogeneous. For example, by using probabilistic methods developed in our lab, we've identified several new coherent—and some surprising—subtypes of scleroderma with distinct patterns of progression. We can use these subtypes within a Bayesian framework to obtain individualized estimates of a scleroderma patient's future disease course.

A challenge of fitting these models is that the data we observe are influenced by many latent phenomena unrelated to the target disease of interest, including the patient's family history, clinician and nurse interventions, and even actions motivated by financial incentives. For example, information on whether a patient experienced a condition could be less frequently recorded for conditions that aren't reimbursed. Thus, we need higher-quality models that can account for the bias due to the "healthcare process" that's generating the data.

Based on our work, a prototype tool for obtaining individualized estimates of a patient's lung disease course is being piloted at the Johns Hopkins Scleroderma Center, one of the largest scleroderma treatment centers. By using such tools, clinicians can target only the sickest individuals with therapies that have strong side effects. The scleroderma subtypes we've found have also sparked an exciting collaboration with cell biologists to study whether different collections of autoantibody markers can explain these. The study of complex diseases via the lens of computational subtyping is an emerging area of research with the potential to redefine the way these diseases are classified. Beyond scleroderma, using similar techniques, we're asking questions such as can we detect life-threatening adverse events in hospitalized patients earlier? Can we identify which patients would benefit from being transferred to a higher acuity setting? Realizing these applications will require advances in representation, inference, and learning.

With the amount of health data exploding, we desperately need more accurate and scalable tools for integrative analysis. Imagine a reasoning engine that layers one patient's health data on that from 7 billion others to monitor, track, and make recommendations that are individualized! This sounds like a pipe dream, but in 1998, organizing all of the world's documents into a searchable graph seemed equally unlikely.

AI'S 10 TO WATCH

Gerardo I. Simari

Universidad Nacional del Sur in Bahia Blanca and CONICET

Gerardo I. Simari is an assistant professor at Universidad Nacional del Sur in Bahía Blanca and a researcher at CONICET, Argentina. His research focuses on topics at the intersection of AI and databases, and reasoning under uncertainty. Simari received a PhD in computer science from the University of Maryland, College Park, and then joined the Information Systems group at the University of Oxford before returning to Argentina in 2014. He received an honorable mention for the IJAR Young Research Award 2011 and the Best Student Paper prize at ICLP 2009. Simari is a former Fulford Junior Research Fellow of Somerville College, University of Oxford. Contact him at gis@cs.uns.edu.ar.

Pushing the Limits of Knowledge Representation and Reasoning

Reasoning tasks that deal with large amounts of data quickly run into two main problems: a representational/algorithmic one (essentially, what are the best ways to represent and manipulate knowledge) and a computational one (finding the most efficient way to do so). R&D in the former has traditionally been the realm of knowledge representation and reasoning (KR), whereas the most important practical advances in the latter have come largely from the databases community.

I have been fortunate enough to work in research projects that addressed key questions in both areas, mainly due to great people like V.S. Subrahmanian, Thomas Lukasiewicz, and Georg Gottlob. My first steps were in the area of autonomous intelligent agents, in which I developed novel comparisons and relations between prescriptive and descriptive architectures. For my doctoral work, I researched models and algorithms for reasoning with probabilistic logic programming-based knowledge bases. As a member of Maryland's Laboratory for Computational Cultural Dynamics, I helped develop such models for characterizing the behavior of agents or groups of interest, with applications in real-world domains. Also at Maryland, I started working on problems related to reasoning with inconsistent information, focusing on answering queries to relational databases. When I moved to Oxford, I was welcomed into a research group that fostered my growth in both areas, by working on the development of novel probabilistic extensions to ontological languages such as Datalog+/- and lightweight description logics, as well as novel approaches to querying inconsistent ontologies. Also at Oxford, I became interested in reasoning with preferences, especially in combination with the probabilistic and inconsistency management branches of my work.

These research lines have proved to be fertile grounds for developing a keen interest in areas straddling both theory and practice. Although seemingly disparate, these areas are tied by problems that neither databases nor KR have successfully addressed on their own. My main goal is to continue to push the limits of both theory and practice, in search of powerful and tractable tools for effective reasoning and knowledge engineering in an ever-changing information landscape that's currently at the heart of a revolution in the way data is produced and consumed by virtually everyone on the planet.

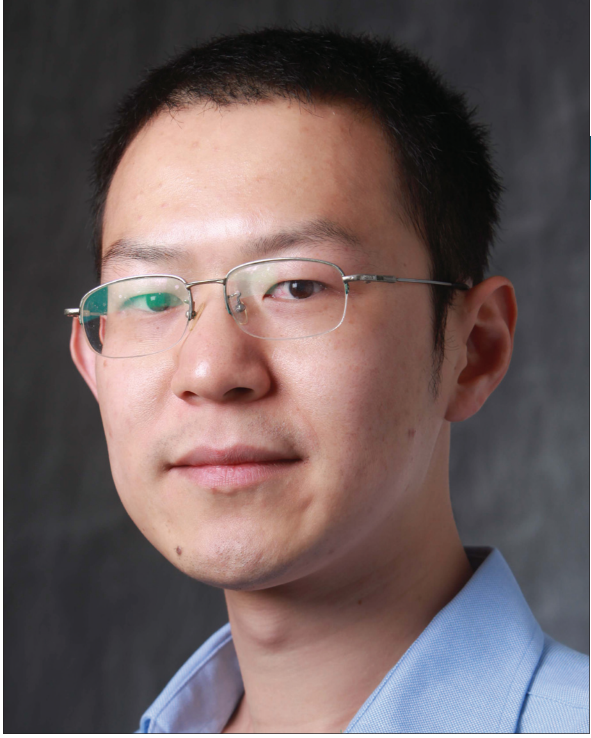

Lirong Xia

Rensselaer Polytechnic Institute

Lirong Xia is an assistant professor in the Department of Computer Science at Rensselaer Polytechnic Institute (RPI). Prior to joining RPI in 2013, he was a CRCS fellow and National Science Foundation Computing Innovation Fellow at the Center for Research on Computation and Society at Harvard University. Xia received a PhD in computer science from Duke University. He's an associate editor of *Mathematical Social Sciences* and is on the editorial board of the *Journal of Artificial Intelligence Research*. Xia received a National Science Foundation CAREER award and a Simons-Berkeley Research Fellowship for the Economics and Computation Program. Contact him at xial@cs.rpi.edu.

Better Group Decision-Making

My research aims to combine principles and techniques of economics, statistics, and AI to help a group of agents resolve disputes and make a joint decision. For example, the Internet Movie Database (IMDb) wants to rank movies by aggregating user ratings, which represent subjective preferences and can be viewed as data generated by a statistical model—how should IMDb design its ranking mechanism? Similar group decision-making problems arise in a wide range of settings beyond product ranking: How should a metasearch engine merge results from multiple search engines? How should a task publisher answer a question by aggregating online workers' noisy answers in crowdsourcing?

Classical social choice theory and statistics can't solve these problems. The former uses desirable economic properties, called axioms, to evaluate group decision-making mechanisms. Unfortunately, various impossibility theorems, including Arrow's celebrated theorem, tell us that no mechanism can satisfy the combination of even a few natural axioms, thus tradeoffs among them are inevitable. On the other hand, little work in statistics explores tradeoffs among the satisfaction of social choice axioms and statistical properties. A traditional social choice paradigm is to manually design mechanisms to explore such tradeoffs, but this doesn't solve the group decision-making problems that arise from the vast variety of statistical models and desirable axioms. I'm pursuing a new paradigm to efficiently explore these tradeoffs: once the decision-maker chooses the axiomatic, statistical, and computational properties she deems important, an automated designer comes up with a mechanism to explore the best tradeoffs among these properties. More precisely, I want to build *principled and systematic frameworks* for researchers and engineers to choose and discover *application-specific* group decision-making mechanisms.

Accordingly, my collaborators and I have built frameworks based on statistical decision theory and analyzed axiomatic and computational properties of Bayesian estimators; used social choice theory to reverse-engineer existing manually designed mechanisms; and applied machine learning to efficiently learn mechanisms that satisfy a customizable set of axioms. We're currently exploring the idea of converting satisfaction of an axiom into multiple training data and then using our frameworks to learn a mechanism.

There are a lot of exciting directions and challenges here. The new paradigm and proposed frameworks work best for applications in which complicated mechanisms are acceptable to agents as a black box. How to incorporate simplicity and usability into the automated framework is a key challenge for future research. Can state-of-the-art AI techniques such as constraint programming and deep learning be leveraged to help learning new mechanisms? How well do AI systems and statistical estimators work in societal contexts where data providers are humans pursuing their own goals? Can we characterize AI systems and statistical estimators using sensible axioms?

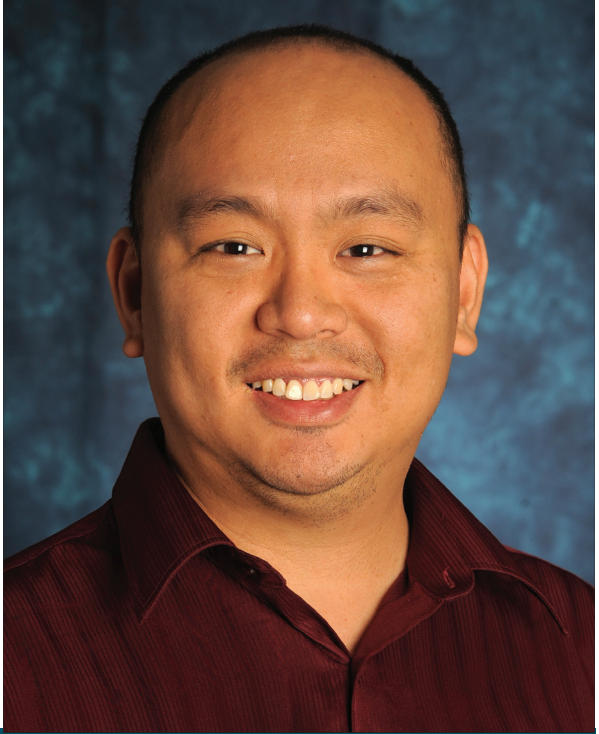

AI'S 10 TO WATCH

William Yeoh

New Mexico State University

William Yeoh is an assistant professor in the Department of Computer Science at New Mexico State University. His research interests lie in the area of multi-agent systems, specifically in distributed constraint optimization and decentralized planning under uncertainty. Yeoh has a PhD in computer science from the University of Southern California and currently serves on the editorial board of the Journal of Artificial Intelligence Research. Contact him at wyeoh@cs.nmsu.edu.

Distributed Constraint Optimization

Through the proliferation of smart devices in our daily lives, the multi-agent system design paradigm is rapidly growing in applicability and popularity. In a multi-agent system, autonomous agents, often with limited computation and communication capabilities, need to coordinate with one another to perform aggregated tasks and optimize aggregated objective functions. For example, agents representing smart devices within a home can coordinate with each other on their device settings and schedules with the objective of improved comfort of occupants, improved energy efficiency, and reduced operational costs.

The distributed constraint optimization problem (DCOP) formulation is well suited for modeling multi-agent coordination problems. In a DCOP, agent interactions are represented by constraints whose utilities depend on the values taken on by the agents. The goal is for the agents to coordinate their value assignments to maximize the sum of the resulting constraint utilities. DCOPs have been used to model a variety of problems including distributed coordination in disaster evacuation scenarios, sensor networks, and logistics operations.

My research is primarily aimed at improving the scalability and efficiency of DCOP algorithms in an effort to deploy these algorithms in the field. As part of my PhD work, I introduced several algorithms that use classical heuristic search techniques, such as choosing appropriate search strategies and appropriate caching heuristics, to speed up prior state-of-the-art offerings by an order of magnitude. Together with colleagues, I've also investigated the use of sampling strategies to efficiently find solutions with probabilistic error bounds as well as the use of general-purpose graphics processing units to speed up parallel computations within DCOP algorithms.

An orthogonal research goal of mine is to enrich the representative power of DCOPs. The default vanilla model doesn't allow for uncertainty and dynamicity, and many real-world applications have elements of both. As such, I've begun to investigate the synergies between DCOPs and decentralized Markov decision processes, which are commonly used to model sequential planning problems under uncertainty, with preliminary results that use reinforcement learning methods.

One of my long-term goals is to transition some of my research into real-world settings. An application of interest is the smart grid, the vision being that we'll be able to monitor, control, and optimize our energy usage, trade excess energy produced by our solar panels with our neighbors, and better react to energy disruptions and contain their impact. The AI community has a lot to offer to realize this vision. The multi-agent system paradigm is particularly well suited for modeling many inherently distributed smart grid problems. I've begun to investigate this direction through collaborations with power engineers and am working toward the deployment of DCOP algorithms in a microgrid testbed at my institution.