

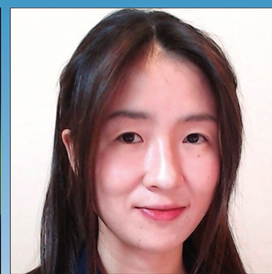





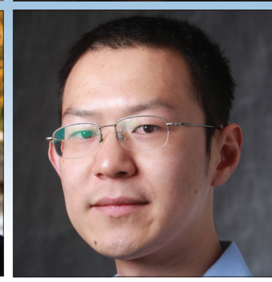


AI's 10 to Watch

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

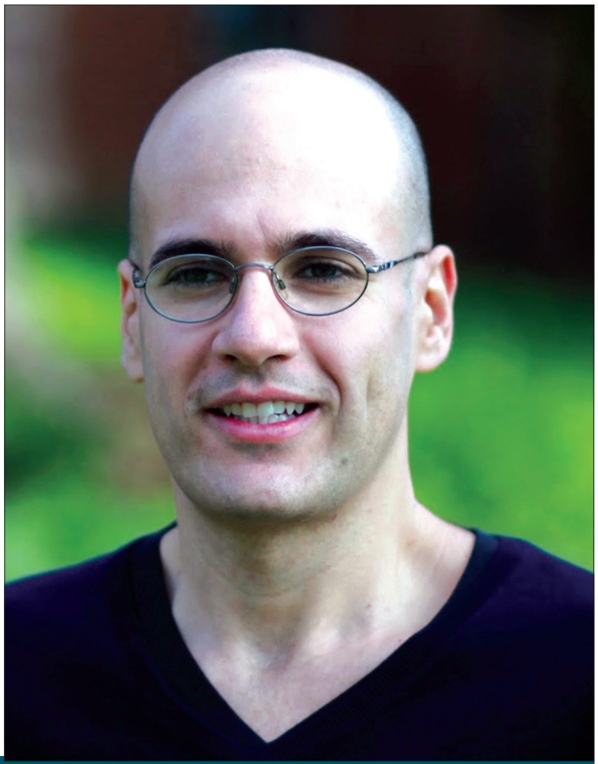

Elias Bareinboim

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.