

Editorial Introduction to the Special Articles on Successful Research in AI

Reflections on Successful Research in Artificial Intelligence: An Introduction

*Ching-Hua Chen, Jim Hendler,
Sabbir Rashid, Oshani Seneviratne, Daby Sow, Biplav Srivastava*

■ *This editorial introduces the special topic articles on reflections on successful research in artificial intelligence. Consisting of a combination of interviews and full-length articles, the special topic articles examine the meaning of success and metrics of success from a variety of perspectives. Our editorial team is especially excited about this topic, because we are in an era when several of the aspirations of early artificial intelligence researchers and futurists seem to be within reach of the general public. This has spurred us to reflect on, and re-examine, our social and scientific motivations for promoting the use of artificial intelligence in governments, enterprises, and in our lives.*

A twin-win in research is achieved when the work results in a validated theory that can be published and a tested solution that can be widely disseminated (Shneiderman and Hendler 2017). Arguably, examples of the twin-win in artificial intelligence research include the semantic web, *Watson Jeopardy!*, and the self-driving car. In each of these cases, the research was inspired by a concrete, real-world challenge, mobilized long-term investments by one or more teams, and the impact of its success extended beyond its intended scope. We are now in an era when several of the aspirations of early artificial intelligence researchers and futurists seem to be within reach of the general public. This has spurred researchers to reflect on, and re-examine, from a variety of perspectives, our social and scientific motivations for promoting the use of artificial intelligence in governments, enterprises, and our lives.

The first article to be published is an interview with the current president of the Association for the Advancement of Artificial Intelligence, Yolanda Gil, whose long view on successful AI research helps us to see beyond the recent excitement around AI and appreciate how the AI community is evolving and how it has held constant. We are inspired by Gil's approaches to thinking big and measuring progress, and hope that this readership will be as well.

Meanwhile, the field of AI is facing a reproducibility crisis (Hutson 2018). Producing reproducible research is a cornerstone for making scientific progress and is arguably an essential measure of successful research. Computer scientist Odd Erik Gundersen has performed an assessment of this crisis in the AI community based on research presented at top AI conferences and proposes a methodology for addressing this crisis (Gundersen and Kjensmo 2018). His conclusion is that "We are not standing on each other's shoulders. It is more like we are standing on each other's feet. The quality of documentation of empirical AI research must clearly improve." Thankfully, he also includes a discussion of potential barriers to reproducible research, allowing each of us to reflect on the potential role that we can play in overcoming them.

It would be remiss to overlook the role that government plays in enabling successful applications of AI. Therefore, we are also publishing an interview with Arvind Gupta on the Indian government's investment in AI innovations for business and social causes. Gupta has over two decades of experience in leadership, policy, and entrepreneurial roles, in both Silicon Valley and India. His interview provides insights on how the Indian government defines success, and their plans for achieving it.

Additionally, in a subsequent issue, Aaron Mannes, a senior policy advisor to the U.S. Department of Homeland Security, will share his perspective on how governments can use hard and soft governance methods for preventing harmful AI research from being conducted or deployed, while also protecting useful AI research from the blocking effects of unfounded public fear.

For researchers in search of the twin-win, forming partnerships between industry and academe may be part of the answer. In light of the rapid developments in AI technology and the equally dynamic business climate, the barriers between where technologies are developed and where they are deployed need to be lowered. However, forming effective partnerships is far from straightforward. A careful consideration of incentives, sustainability, and market forces are particularly crucial in AI research today when there is a tremendous uptick in the investment of resources in AI research, across academe and industry (Columbus 2019). Based on first-hand knowledge, Lisa Amini and her co-authors Ching-Hua Chen, David Cox, Aude Oliva, and Antonio Torralba, in industry and academe, will reflect on three large, academia-industry initiatives; their article is intended to help spawn a

dialog on the motivations that drive such collaborations, and the execution challenges that shape their designs. Importantly, their experiences reveal the nuanced decision-making around designing organizational structures for promoting successful AI research.

While Aaron Manne's article articulates how political governance can influence the success of AI research endeavors, Jeanna Matthews describes a set of antipatterns of behavior that AI researchers engage in that increase the risk of harm from AI. Both Mannes and Matthews drive home the message that all those involved in creating and deploying AI technologies need to feel more personally accountable for ensuring that the AI is used responsibly.

We find the phrase *be careful what you ask for* to be especially pertinent in AI research, as we invent technologies that are designed to automatically do what we ask. In this issue, we hope that the collection of articles and interviews devoted to this topic are educational, thought-provoking, and even controversial. We don't expect the AI community to ever converge around a universal notion of success; however, we hope to see a continuous and healthy discussion around this topic.

References

- Columbus, L. 2019. 10 Charts That Will Change Your Perspective on Artificial Intelligence's Growth. *Forbes.com*. Accessed on January 12, 2019. www.forbes.com/sites/louiscolumnbus/2018/01/12/10-charts-that-will-change-your-perspective-on-artificial-intelligences-growth/#2e6a0c214758.
- Gundersen, O. E., and Kjensmo, S. 2018. State of the Art: Reproducibility in Artificial Intelligence. *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 1-8. Palo Alto, CA: AAAI Press. www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/viewFile/17248/15864.
- Hutson, M. 2018. Artificial Intelligence Faces Reproducibility Crisis. *Science* 359(6377): 725-6. science.sciencemag.org/content/359/6377/725.
- Shneiderman, B., and Hendler, J. 2017 It's the Partnership, Stupid. *Issues in Science and Technology* 33(4): Summer. issues.org/perspective-its-the-partnership-stupid/.

Ching-Hua Chen is a research staff member at the T.J. Watson Research Center in Yorktown Heights, New York. She manages the Health Behavior and Decision Science group within the Center for Computational Health. She graduated from Penn State University with a dual-title PhD in Business Administration and Operations Research.

Jim Hendler is the Tetherless World Professor of Computer, Web and Cognitive Sciences at Rensselaer Polytechnic Institute. He was the recipient of the 2017 Association for Advancement of Artificial Intelligence Distinguished Service award and is a fellow of the Association for Advancement of Artificial Intelligence, Association for Computing Machinery, Institute of Electrical and Electronics Engineers, and the National Academy of Public Administration.

Sabbir Rashid is a graduate student working with Deborah McGuinness at Rensselaer Polytechnic Institute on research related to the semantic web. Rashid has contributed to technologies involving data annotation and harmonization,



Call for Nominations for the 2020 AAAI Executive Council Election

The 2020 Nominating Committee is seeking nominations from the AAAI membership for the positions of AAAI President-Elect and Executive Councilor. In 2020, AAAI members will elect one individual to serve a two-year term as president-elect, followed by two years as president, and finally, two years as immediate past president. In addition, members will elect four new councilors to serve three-year terms on the AAAI Executive Council.

All elected officers and councilors are required to attend all council meetings each year (usually 1–2 in person and 2–3 via telecon), and actively participate in AAAI activities. Nominees must be current members of AAAI. The Nominating Committee encourages all regular AAAI members in good standing to place an individual's name before them for consideration. (Student and institutional members are not eligible to submit candidates' names.) The Nominating Committee, in turn, will nominate two candidates for president-elect and eight candidates for councilor in early spring. In addition to members' recommendations, the committee will actively recruit individuals in order to provide a balanced slate of candidates. AAAI regular members will vote in late spring, and the new members of the Executive Council will be installed in the summer of 2020.

To submit a candidate's name for consideration, please send the following information to Carol Hamilton, Executive Director, AAAI, 2275 East Bayshore Road, Suite 160, Palo Alto, CA 94303; by fax to 650/321-4457; or by email to hamilton@aaai.org:

Name
Affiliation
City, State or Province, Country
Email address
URL
Year of membership in AAAI
Approximate number of AAAI publications
At least two sentences describing the candidate and why the nominee would be a good candidate

Please include any additional information or recommendations that would be helpful to the Nominating Committee. Nominators should contact candidates prior to submitting their names to verify that they are willing to serve, should they be elected.

The deadline for nominations is March 1, 2020.

using semantic data dictionaries to semantically represent and integrate several publicly available datasets. His current work includes the application of deductive and abductive reasoning techniques over linked health data. This research is being applied as part of the Health Empowerment by Analytics, Learning, and Semantics project to help explain the actions of physicians, as well as adverse drug reactions of patients, in the context of chronic diseases such as diabetes.

Oshani Seneviratne is the director of health data research at the Institute for Data Exploration and Applications at the Rensselaer Polytechnic Institute. Seneviratne's research interests lie at the intersection of decentralized systems and health applications. At Rensselaer Institute for Data Exploration and Applications, Seneviratne is involved in the Health Empowerment by Analytics, Learning, and Semantics project, and leads the Smart Contracts Augmented with Analytics Learning and Semantics project. Seneviratne obtained her PhD in computer science from the Massachusetts Institute of Technology under the supervision of Sir Tim Berners-Lee. Before Rensselaer, Seneviratne worked at Oracle specializing in distributed systems, provenance, and healthcare-related research and applications.

Daby Sow is a principal research staff member at IBM Research. Since August 2017, he has managed the Biomedical Analytics and Modeling group, part of the IBM Research

Center for Computational Health. In this role, he is leading a team of AI scientists developing novel AI and machine learning solutions for various open healthcare research problems. These problems range from modeling the progression of complex chronic conditions, to pharmacovigilance with the development of signal detection algorithms for early adverse drug reaction detection from real-world evidence data (electronic health records, claims, spontaneous reporting systems) and the generation of time-varying treatment strategies using data collected during clinical practice. Sow is an alumni of Columbia University, where he received a PhD degree in electrical engineering in 2000.

Biplav Srivastava is a distinguished data scientist and master inventor at IBM's Chief Analytics Office. With over two decades of research experience in AI, services computing and sustainability, most of which was at IBM Research, Biplav is also an Association for Computing Machinery Distinguished Scientist and Distinguished Speaker, and an Institute of Electrical and Electronics Engineers Senior Member. Srivastava is exploring new approaches for goal-oriented, ethical, human-machine collaboration via natural interfaces using domain and user models, learning, and planning. He is leading efforts for adoption of AI technologies in a large-scale global business context and understanding their impact on workforce. Srivastava received his MS and PhD from Arizona State University, and a BTech from Indian Institute of Technology, India, all in computer science.