Celebrating the Past, Present, and Future of Computing

Timothy E. Lee (Carnegie Mellon University; timothyelee@cmu.edu)
Justin Svegliato (University of Massachusetts Amherst; jsvegliato@cs.umass.edu)
DOI: 10.1145/3137574.3137581

Abstract

Timothy Lee and Justin Svegliato, two Student SIGAI Scholars, cover *The 50 Years of the ACM Turing Award Celebration*, which convened in San Francisco last June. The semicentennial celebration addressed the past, present, and future advancements of computing, ranging from deep learning and ethics to augmented reality and quantum computing.

As Student Scholars sponsored by SIGAI, we are grateful for the opportunity to be a part of *The 50 Years of the ACM Turing Award Celebration*. For two days in June, hundreds of professors, researchers, and students from across the globe gathered together in San Francisco to celebrate the legacy of the Turing Award (often referred to as the Nobel Prize of computing) and the incredible advances in computing over the last 50 years. The semicentennial celebration also honored this year's Turing Award recipient, Tim Berners-Lee, for inventing the World Wide Web and related networking technologies such as the Semantic Web.

After opening remarks from a Turing laureate each day, we heard from several panels that spanned the field of computing, ranging from deep learning and ethics to augmented reality and quantum computing. Every panel featured a distinguished moderator and several panelists that included Turing laureates, prominent researchers, and rising stars in the field. Interleaved with the panels were several short films. These films featured the life and work of the father of computer science, Alan Turing, and highlighted the Turing laureates' contributions, including those made to the field of artificial intelligence by the AI Turing Award laureates. We were honored by the attendance of several AI Turing Award recipients: Judea Pearl (2011 Turing laureate), Ed Feigenbaum (1994 Turing laureate), and Raj Reddy (1994 Turing laureate).

Figure 1: SIGAI at *The 50 Years of the ACM Turing Award Celebration*. Pictured from left to right: Timothy E. Lee, Yolanda Gil, and Justin Svegliato. The bronze bust of Alan Turing unveiled during the conference is also shown here.

The first panel *Advances in Deep Neural Networks* was particularly relevant to the SIGAI community. Moderated by Judea Pearl, the panel featured Michael Jordan (UC Berkeley), Fei-Fei Li (Stanford), Stuart Russell (UC Berkeley), Ilya Sutskever (OpenAI), and Raquel Urtasun (Toronto). As a popular area not only in AI but also in computing in general, deep learning has emerged as a powerful approach for enabling machine intelligence. Sutskever explained that neural networks are essentially tunable circuits that learn high-dimensional mappings from data. Deep neural networks have emerged from the confluence of several factors: the recent advances in hardware (the "oxygen" of neural networks according to Sutskever), the availability of massive datasets via the Internet, and the accelerated progress of data science.

Despite their promising results, a common theme emerged from the panel. Although effective in particular domains, deep learning in its current form cannot be the fundamental abstraction of machine intelligence sought after by researchers. There are many questions

about its long-term viability as the bedrock of machine intelligence. As Jordan argued, today's neural networks are deep architecturally, but not semantically. Pearl questioned whether these networks could reason about causality, a central theme in his foundational work on Bayesian networks.

Outlining the weaknesses of today's deep neural networks segued into a discussion on the types of intelligent behavior that humans exhibit but these networks currently lack, such as semantic understanding, contextual reasoning, abstraction, and reasoning under uncertainty, all of which are easily handled by humans despite little training data. Russell drew a fitting analogy between Allen Newell, Cliff Shaw, and Herbert Simon's General Problem Solver and the need for exponential computing with deep learning and the need for exponential *data*. In Russell's opinion, hoping to achieve "tabula rasa" machine intelligence with only deep learning may be infeasible in some—or all—domains due to the data demands, and so we must continue to search for better techniques. Li offered a similar anecdote from her work with ImageNet. With enough data, deep neural networks are by far the state of the art in object recognition, but they perform poorly and cannot reason effectively without massive datasets. In the case of robotics, Urtasun noted that being unable to model uncertainty well in deep learning is a considerable drawback in her current work on self-driving vehicles where algorithm robustness is critical.

Still, even with these shortcomings, deep learning performs quite impressively in narrow problems, such as computer vision, image captioning, and object segmentation. In some cases, such as AlphaGo, it enables decision-making capabilities that are superior to human intelligence. Several of the panelists agreed that deep learning has matured enough to be used in industry, but the search for machine intelligence must continue. Ultimately, the panel could be best summarized by Li's comments: we are entering the "end of the beginning" for AI. Deep neural networks may be one of our best existing tools for enabling the development of intelligent agents, but even greater breakthroughs are yet to come.

In addition to the deep learning panel, the opening day of the celebration also featured four other panels with many prominent researchers from industry and academia, along with a talk by 2008 Turing laureate Barbara Liskov that explored the history of computing. First, in *Restoring Personal Privacy without Compromising National Security*, Whitfield Diffie, 2015 Turing laureate, along with several leaders in security, cryptography, and networking discussed how governments could obtain useful information using backdoors and other intentional vulnerabilities to aid criminal investigations without jeopardizing the privacy of society. Following a short film on Alan Turing's life, we then turned to Vint Cerf, 2004 Turing laureate, and several other distinguished researchers in *Preserving Our Past for the Future*. They considered the problem of how to store data for centuries to come and whether corporations or governments should fund such an endeavor.

Later that day, *Moore's Law Is Really Dead: What's Next?* headlined the 1992 Turing laureate Butler Lampson. The panel explored the ways in which the field can continue the trend of exponential technological growth despite that Moores Law has continued to slow down. During the panel, a common theme emerged: researchers will eventually leverage special-purpose hardware and quantum computing to push the boundaries of computing forward.

At the end of the day, we heard from Raj Reddy in *Challenges in Ethics and Computing*. Given the increasing relevance of AI and machine learning, Reddy believes that ethical questions in computing have become more important than ever. Noel Sharkey added several important questions in light of recent progress in self-driving cars and machine learning. How can self-driving cars make decisions that were once reserved for humans in life and death situations? And how do we ensure that data-driven algorithms escape bias against minorities in the justice system?

On the final day of the conference, there were two panels on some of the most rapidly growing fields in computing following a talk from Donald Knuth, 1974 Turing laureate. *Quantum Computing: Far Away? Around the Corner? Or Maybe Both at the Same Time?* that featured the 2000 Turing laureate Andrew Yao

investigated the current state of quantum computing and how it might drive software development in the next 50 years. Like the deep learning panel, John Martinis cautioned that quantum computing is only a powerful tool in certain combinatorial problems but useless in others. However, in areas like AI, machine learning, and cryptography, it has the potential to revolutionize the field.

The celebration culminated with a panel on an area of computing that has recently seen rapid progress: *Augmented Reality: From Gaming to Cognitive Aids and Beyond*. Fred Books, 1999 Turing laureate, and Ivan Sutherland, 1988 Turing laureate, reminisced about the early work of augmented reality while they were aspiring researchers. Peter Lee discussed the impact of augmented reality on the gaming industry. Other panelists considered Google Glass, Pokemon Go, and Oculus Rift and explored the inevitable future of augmented reality in the home and at the workplace.

For all attendees across the spectrum of computing, the advancements of the last 50 years honored during the conference will undoubtedly shape our own contributions for the next 50 years to come. And, for the SIGAI community, the experts in our field gave insight into the ongoing search for machine intelligence and how deep learning might play a role. Given the recent groundbreaking advances in AI, it was only fitting that the celebration of computing's greatest achievements was in honor of who many call the grandfather of AI, Alan Turing.

Acknowledgements

We gratefully acknowledge the support of SIGAI for the opportunity to attend *The 50 Years of the ACM Turing Award Celebration*. We also thank Yolanda Gil for her help with this article.

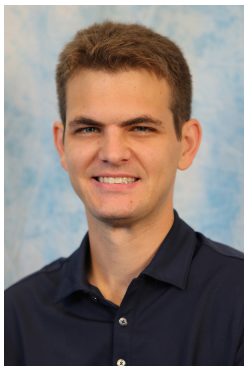

Timothy E. Lee is a M.S. student in Robotics at Carnegie Mellon University. As a member of the Robust Adaptive Systems Lab under Professor Nathan Michael, Timothy's research focuses on improving mobile robot intelligence to enable the automation of challenging tasks in real-world settings. He is currently investigating robust, vision-based navigation of a submersible robot to automate the precision inspection of underwater infrastructure.

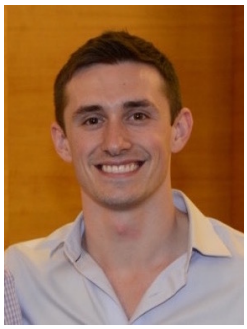

Justin Svegliato is a second year Ph.D. student in Computer Science at the University of Massachusetts Amherst. In the Resource-Bounded Reasoning Lab under Professor Shlomo Zilberstein, Justin's research focuses on bounded rationality, real-time decision making, and autonomous agent architectures. He is currently developing metareasoning techniques that monitor and control algorithms that trade decision quality with computation time.