



Association for
Computing Machinery

Advancing Computing as a Science & Profession

Contact: Jim Ormond
212-626-0505
ormond@hq.acm.org

ACM PRIZE IN COMPUTING AWARDED TO ALPHAGO DEVELOPER

David Silver Recognized for Breakthrough Advances in Computer Game-Playing

New York, NY, April 1, 2020 – ACM, the Association for Computing Machinery, today announced that David Silver is the recipient of the 2019 ACM Prize in Computing for breakthrough advances in computer game-playing. Silver is a Professor at University College London and a Principal Research Scientist at DeepMind, a Google-owned artificial intelligence company based in the United Kingdom. Silver is recognized as a central figure in the growing and impactful area of deep reinforcement learning.

Silver’s most highly publicized achievement was leading the team that developed AlphaGo, a computer program that defeated the world champion of the game Go, a popular abstract board game. Silver developed the AlphaGo algorithm by deftly combining ideas from deep-learning, reinforcement-learning, traditional tree-search and large-scale computing. AlphaGo is recognized as a milestone in artificial intelligence (AI) research and was ranked by *New Scientist* magazine as one of the top 10 discoveries of the last decade.

AlphaGo was initialized by training on expert human games followed by reinforcement learning to improve its performance. Subsequently, Silver sought even more principled methods for achieving greater performance and generality. He developed the AlphaZero algorithm that learned entirely by playing games against itself, starting without any human data or prior knowledge except the game rules. AlphaZero achieved superhuman performance in the games of chess, Shogi, and Go, demonstrating unprecedented generality of the game-playing methods.

The ACM Prize in Computing recognizes early-to-mid-career computer scientists whose research contributions have fundamental impact and broad implications. The award carries a prize of \$250,000, from an endowment provided by Infosys Ltd.

Computer Game-Playing and AI

Teaching computer programs to play games, against humans or other computers, has been a central practice in AI research since the 1950s. Game playing, which requires an agent to make a series of decisions toward an objective—winning—is seen as a useful facsimile of human thought processes.

Game-playing also affords researchers results that are easily quantifiable—that is, did the computer follow the rules, score points, and/or win the game?

At the dawn of the field, researchers developed programs to compete with humans at checkers, and over the decades, increasingly sophisticated chess programs were introduced. A watershed moment occurred in 1997, when ACM sponsored a tournament in which IBM's DeepBlue became the first computer to defeat a world chess champion, Gary Kasparov. At the same time, the objective of the researchers was not simply to develop programs to win games, but to use game-playing as a touchstone to develop machines with capacities that simulated human intelligence.

“Few other researchers have generated as much excitement in the AI field as David Silver,” said ACM President Cherri M. Pancake. “Human vs. machine contests have long been a yardstick for AI. Millions of people around the world watched as AlphaGo defeated the Go world champion, Lee Sedol, on television in March 2016. But that was just the beginning of Silver’s impact. His insights into deep reinforcement learning are already being applied in areas such as improving the efficiency of the UK’s power grid, reducing power consumption at Google’s data centers, and planning the trajectories of space probes for the European Space Agency.”

“Infosys congratulates David Silver for his accomplishments in making foundational contributions to deep reinforcement learning and thus rapidly accelerating the state of the art in artificial intelligence,” said Pravin Rao, COO of Infosys. “When computers can defeat world champions at complex board games, it captures the public imagination and attracts young researchers to areas like machine learning. Importantly, the frameworks that Silver and his colleagues have developed will inform all areas of AI, as well as practical applications in business and industry for many years to come. Infosys is proud to provide financial support for the ACM Prize in Computing and to join with ACM in recognizing outstanding young computing professionals.”

Silver is credited with being one of the foremost proponents of a new machine learning tool called deep reinforcement learning, in which the algorithm learns by trial-and-error in an interactive environment. The algorithm continually adjusts its actions based on the information it accumulates while it is running. In deep reinforcement learning, artificial neural networks—computation models which use different layers of mathematical processing—are effectively combined with the reinforcement learning strategies to evaluate the trial-and-error results. Instead of having to perform calculations of every possible outcome, the algorithm makes predictions leading to a more efficient execution of a given task.

Learning Atari from Scratch

At the Neural Information Processing Systems Conference (NeurIPS) in 2013, Silver and his colleagues at DeepMind presented a program that could play 50 Atari games to human-level ability. The program learned to play the games based solely on observing the pixels and scores while playing. Earlier reinforcement learning approaches had not achieved anything close to this level of ability.

Silver and his colleagues published their method of combining reinforcement learning with artificial neural networks in a seminal 2015 paper, "[Human Level Control Through Deep Reinforcement Learning](#)," which was published in *Nature*. The paper has been cited nearly 10,000 times and has had an immense impact on the field. Subsequently, Silver and his colleagues continued to refine these deep reinforcement learning algorithms with novel techniques, and these algorithms remain among the most widely-used tools in machine learning.

AlphaGo

The game of Go was invented in China 2,500 years ago and has remained popular, especially in Asia. Go is regarded as far more complex than chess, as there are vastly more potential moves a player can make, as well as many more ways a game can play out. Silver first began exploring the possibility of developing a computer program that could master Go when he was a PhD student at the University of Alberta, and it remained a continuing research interest.

Silver's key insight in developing AlphaGo was to combine deep neural networks with an algorithm used in computer game-playing called Monte Carlo Tree Search. One strength of Monte Carlo Tree Search is that, while pursuing the perceived best strategy in a game, the algorithm is also continually investigating other alternatives. AlphaGo's defeat of world Go champion Lee Sedol in March 2016 was hailed as a milestone moment in AI. Silver and his colleagues published the foundational technology underpinning AlphaGo in the paper "[Mastering the Game of Go with Deep Neural Networks and Tree Search](#)" that was published in *Nature* in 2016.

AlphaGo Zero, AlphaZero and AlphaStar

Silver and his team at DeepMind have continued to develop new algorithms that have significantly advanced the state of the art in computer game-playing and achieved results many in the field thought were not yet possible for AI systems. In developing the AlphaGo Zero algorithm, Silver and his collaborators demonstrated that it is possible for a program to master Go without any access to human expert games. The algorithm learns entirely by playing itself without any human data or prior knowledge, except the rules of the game and, in a further iteration, without even knowing the rules.

Later, the DeepMind team's AlphaZero also achieved superhuman performance in chess, Shogi, and Go. In chess, AlphaZero easily defeated world computer chess champion Stockfish, a high-performance program designed by grandmasters and chess programming experts. Just last year, the DeepMind team, led by Silver, developed AlphaStar, which mastered the multiple-player video game StarCraft II, which had been regarded as a stunningly hard challenge for AI learning systems.

The DeepMind team continues to advance these technologies and find applications for them. Among other initiatives, Google is exploring how to use deep reinforcement learning approaches to manage robotic machinery at factories.

Biographical Background

David Silver is Lead of the Reinforcement Learning Research Group at DeepMind, and a Professor of Computer Science at University College London. DeepMind, a subsidiary of Google, seeks to combine the best techniques from machine learning and systems neuroscience to build powerful general-purpose learning algorithms.

Silver earned Bachelor's and Master's degrees from Cambridge University in 1997 and 2000, respectively. In 1998 he co-founded the video games company Elixir Studios, where he served as Chief Technology Officer and Lead Programmer. Silver returned to academia and earned a PhD in Computer Science from the University of Alberta in 2009. Silver's numerous honors include the Marvin Minsky Medal (2018) for outstanding achievements in artificial intelligence, the Royal Academy of Engineering Silver Medal (2017) for outstanding contribution to UK engineering, and the Mensa Foundation Prize (2017) for best scientific discovery in the field of artificial intelligence.

About the ACM Prize in Computing

[The ACM Prize in Computing](#) recognizes an early to mid-career fundamental innovative contribution in computing that, through its depth, impact and broad implications, exemplifies the greatest achievements in the discipline. The award carries a prize of \$250,000. Financial support is provided by an endowment from Infosys Ltd. The ACM Prize in Computing was previously known as the ACM-Infosys Foundation Award in the Computing Sciences from 2007 through 2015. ACM Prize recipients are invited to participate in the Heidelberg Laureate Forum, an annual networking event that brings together young researchers from around the world with recipients of the ACM A.M. Turing Award, the Abel Prize, the Fields Medal, and the Nevanlinna Prize.

About ACM

ACM, [the Association for Computing Machinery](#), is the world's largest educational and scientific computing society, uniting computing educators, researchers and professionals to inspire dialogue, share resources and address the field's challenges. ACM strengthens the computing profession's collective voice through strong leadership, promotion of the highest standards, and recognition of technical excellence. ACM supports the professional growth of its members by providing opportunities for life-long learning, career development, and professional networking.

About Infosys

[Infosys](#) is a global leader in next-generation digital services and consulting. We enable clients in 46 countries to navigate their digital transformation. With over three decades of experience in managing the systems and workings of global enterprises, we expertly steer our clients through their digital journey. We do it by enabling the enterprise with an AI-powered core that helps prioritize the execution of change. We also empower the business with agile digital at scale to deliver unprecedented levels of performance and customer delight. Our always-on learning agenda drives their continuous improvement through building and transferring digital skills, expertise, and ideas from our innovation ecosystem.

###