

What does AI's success playing complex board games tell brain scientists?

Dale Purves^{a,1}

In research, it sometimes happens that advances in one field unexpectedly inform another in a fundamental way. A case in point may be what computer-generated gameplay suggests about how brains operate. Spearheaded by Demis Hassabis, David Silver, and their colleagues at the artificial intelligence (AI) company Google DeepMind, a series of articles published over the past several years has reported ongoing improvement of a computational strategy (1–3) that now beats all players in complex board games (4). Beginning in 2015 with a computer program that beat human opponents at all 49 games in the Atari 2600 suite (challenges that include Pong, Space Invaders, and Pac-Man) (1), the group progressed to a program called AlphaGo that

beat the European and world champions at Go (2, 3), a territorial board game far more complicated than chess. The latest version in this effort, called AlphaZero (4), now beats the best players—human or machine—in chess and shogi (Japanese chess) as well as Go.

That machines can beat expert human players in board games is not news. The achievement of this long-standing goal in computer science first attracted widespread attention in 1997 when IBM's program "Deep Blue" beat Garry Kasparov, then the world champion chess player (5). Human players have since been shown to be weak opponents in such games compared with a variety of machine programs. What is different about Google DeepMind's AlphaZero is




Fig. 1. The search space of the territorial board game Go is intractable using logical algorithms. Image credit: Shutterstock/Saran Poroong.

^aDuke Institute for Brain Sciences, Duke University, Durham, NC 27708

The author declares no conflict of interest.

Published under the [PNAS license](#).

Any opinions, findings, conclusions, or recommendations expressed in this work are those of the author and have not been endorsed by the National Academy of Sciences.

¹Email: purves@neuro.duke.edu.

that the system operates by learning on a wholly empirical (trial and error) basis.

Previous game-playing machines, including earlier versions in the AlphaGo series, used a combination of brute force logic, tree search, successful moves made by experts in response to specific board positions, and other “hand-crafted” expedients. In contrast, AlphaZero uses only the empirical results of self-play. The starting point is simply a randomly initialized “deep” neural network (an artificial neural network with multiple hidden layers) that gradually improves by playing against itself millions of times using only the rules of the game and reinforcement learning.

What, then, is the key to this success, and what are the implications for understanding of how human brains work? The lesson seems to be that neural success hinges on connectivity generated by trial-and-error learning over evolutionary and individual time rather than algorithms that compute behavioral answers.

The Underlying Strategy

To understand the power of this approach in playing board games, consider the search spaces involved. The “game tree complexity” of tic-tac-toe—i.e., an estimate of the number of possible positions that must be evaluated to determine the worth of an initial position—is about 2×10^4 . This number of possibilities is easily searched by an algorithm, as my children learned at a theme park when they were beaten by a chicken pecking one of the tic-tac-toe squares in what amounted to a cleverly designed Skinner box.

The take-home message may be that understanding neural circuitry in terms of steps specified by algorithms is not a path to success much beyond the challenge of tic-tac-toe.

There are several different ways to measure the complexity of more difficult board games, leading to different estimates. But no matter the method used, as the complexity of board games grows, the relevant search spaces become unimaginably large. The game tree complexity is estimated to be about 10^{20} for checkers, about 10^{120} for chess, and an astonishing 10^{320} for Go (6). For comparison, the estimated number of atoms in the universe is on the order of 10^{87} . Unlike tic-tac-toe, fully searching spaces of this magnitude is not feasible. In the strategy used by Google DeepMind, success follows not from a logical exploration of possible moves but from the instantiation of trial and error play in the network’s connectivity that beats a competing machine playing in the same empirical way.

The relevance of AlphaZero to neuroscience is the size of the search spaces organisms—and ultimately animals with brains—must contend with as they play “the game of life” by making “moves” that are rewarded by survival and reproductive success. Although the complexity of this biological “game tree” for humans

is beyond reckoning, it clearly dwarfs the complexity of board games such as Go.

“Intelligent” Brains?

In thinking how these facts bear on understanding brains, some caution is in order when using phrases such as “artificial intelligence (AI).” Although we use the word “intelligence”—artificial or otherwise—as if we knew what it meant, the term simply refers to the ability to solve some kind of problem. The presumption is that AI solves problems the way humans do, ignoring the fact that the way we solve problems is largely a mystery.

With respect to humans, “intelligence” usually refers to skills most of us lack, such as the ability to do higher mathematics. But this usage is misleading: is the accomplished athlete solving complex problems on the playing field any less “intelligent” than the math prodigy? This definitional deficiency has led to endless arguments, not least the value of IQ tests and other dubious measures of a vague concept whose determination is subject to social bias and other influences.

We researchers accord a high status to the ability to think through logical problems and imagine that is what brains must be doing. Psychologists showed long ago, however, that systematic thinking is awkward for humans with consequences that can be embarrassing (7). For instance, when high school students were briefly presented with the sequences

$$1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8$$

and

$$8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1$$

and asked to estimate the products, the median answers given for the first series was 512 and for the second 2,250. This result shows that when it comes to logic we are often flummoxed, in this case by failing to recognize that the sequences are multiplications of the same numbers and must, therefore, generate the same product (which is actually 40,320). Because rationality and logic play little part in survival and reproduction, a generally feeble approach to problems such as this is what one might expect. These strategies are far too ponderous for generating answers to most biological problems, and problems whose solution demands reasoning and logic are relatively rare outside the classroom.

Implications

AI in the guise of mimicking computations that brains might be carrying out dates from a conference of mathematicians, computer scientists, information theorists, and others in 1956. The consensus was famously that this challenge should be relatively easy to solve. More than 60 years on, however, the goal has not been met. The central problem now as then is that no one knows what the operating principle of a brain actually is.

This may be where success in gameplay points the way. A wholly empirical approach to winning complex board games revives a trial-and-error (neural network) approach that was introduced in the 1940s (8) but has had a checkered history over the subsequent decades (9). Solving problems by unsupervised learning

using neural networks has often been overshadowed by the immense success of algorithmic computation (executing a series of specified steps), a strategy that has repeatedly proven its worth in applications as diverse as manufacturing, genomics, and drug discovery.

But the success of reinforcement learning in tackling otherwise intractable games suggests that making behavioral responses based on what worked for a species and its members over accumulated past experience is the most efficient way to win in life. Because the brain of any animal is the executor of behavioral choices, it is hard to escape the implication that contending with extremely large “search spaces” must be predicated on neural (or machine) connectivity that reflects accumulated learning by trial and error. This is, after all, the strategy used in evolution.

If this empirical mode of operation is what brains such as ours are doing, then this game-related evidence will have to be taken into account by brain scientists, despite the daunting challenge of understanding nervous systems in this way. Although AlphaZero uses a variety of algorithms to specify the rules of the game and to

input board positions, the “synaptic weights” in the network are updated based on the history of what worked well in millions of previous games. The Google DeepMind group noted early on that the organization of networks thus trained bore little or no resemblance to the organization of human brains (10). Indeed, the inscrutable connectivity of the trained AlphaZero networks may indicate why the detailed connectivity of brains has failed to shed as much light on neural operation as expected.

The take-home message may be that understanding neural circuitry in terms of steps specified by algorithms is not a path to success much beyond the challenge of tic-tac-toe. If one accepts that playing complex games empirically is a paradigm for playing the vastly more complex game of life, brain science could advance rapidly albeit in a framework different from the rule-based computation that has often driven the field. In any event, given what is at stake, the implications of gameplay for human brain function should be vigorously debated.

-
- 1 V. Mnih *et al.*, Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
 - 2 D. Silver *et al.*, Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
 - 3 D. Silver *et al.*, Mastering the game of Go without human knowledge. *Nature* **550**, 354–359 (2017).
 - 4 D. Silver *et al.*, A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* **362**, 1140–1144 (2018).
 - 5 F.-H. Hsu, *Behind Deep Blue: Building the Computer That Defeated the World Chess Champion* (Princeton University Press, 2002).
 - 6 L. V. Allis, “Searching for solutions in games and artificial intelligence,” PhD thesis, University of Limburg, Maastricht, The Netherlands (1994).
 - 7 D. Kahneman, *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge University Press, Cambridge, UK, 1982).
 - 8 W. S. McCulloch, W. Pitts, A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.* **5**, 115–133 (1943).
 - 9 J. A. Anderson, E. Rosenfeld, *Talking Nets: An Oral History of Neural Networks* (MIT Press, Cambridge, MA, 2000).
 - 10 E. Gibney, DeepMind algorithm beats people at classic video games. *Nature* **518**, 465–466 (2015).