One Small Step for AlphaZero, One Giant Leap for Artificial Intelligence

By Don Sutherland

In 1997, an IBM supercomputer called "Deep Blue" defeated reigning world chess champion Garry Kasparov in a six-game match. Deep Blue's success was based on a combination of search techniques by which the supercomputer used its enormous power to evaluate past matches by Kasparov and almost countless arrangements of chess pieces. Deep Blue's mastery was a triumph of supercomputing power more than a function of its own "learning." Programmers gave it its "memory," but it did not learn on its own. But now that may have changed.

An article published in the December 7, 2018 edition of ***Science*** provides an example of a computer program called "AlphaGo Zero" or "AlphaZero" having engaged in self-learning of three complex games: Chess, Go, and Shogi. In these cases, AlphaZero was provided only with the rules of each of the games. It then "learned" through self-play. Afterward, it defeated some of the leading computer programs for each of these games. The outcomes were often decisive. The article reported:

Our results demonstrate that a general purpose reinforcement learning algorithm can learn tabula rasa—without domain-specific human knowledge or data, as evidenced by the same algorithm succeeding in multiple domains—superhuman performance across multiple challenging games.

Perhaps most remarkable of all was the system's development of a capacity to sacrifice chess pieces for short-term loss, but long-term gain. The article revealed:

In several games, AlphaZero sacrificed pieces for long-term strategic advantage, suggesting that it has a more fluid, context-dependent positional evaluation than the rule-based evaluations used by previous chess programs.

Finally, the system's ability to "learn" was replicated in other games. AlphaZero's algorithm was applied without any modification to the game of Shogi and only the rules of the game were given to it. Within hours, it outperformed the leading computer programs for that game.

The capacity to learn is a major step forward for the emerging field of Artificial Intelligence. A broadening of the capacity to fields beyond games, especially those that involve highly complex problems, could represent the next major area of progress.