DeepMind's new AI app plays Stratego at expert level
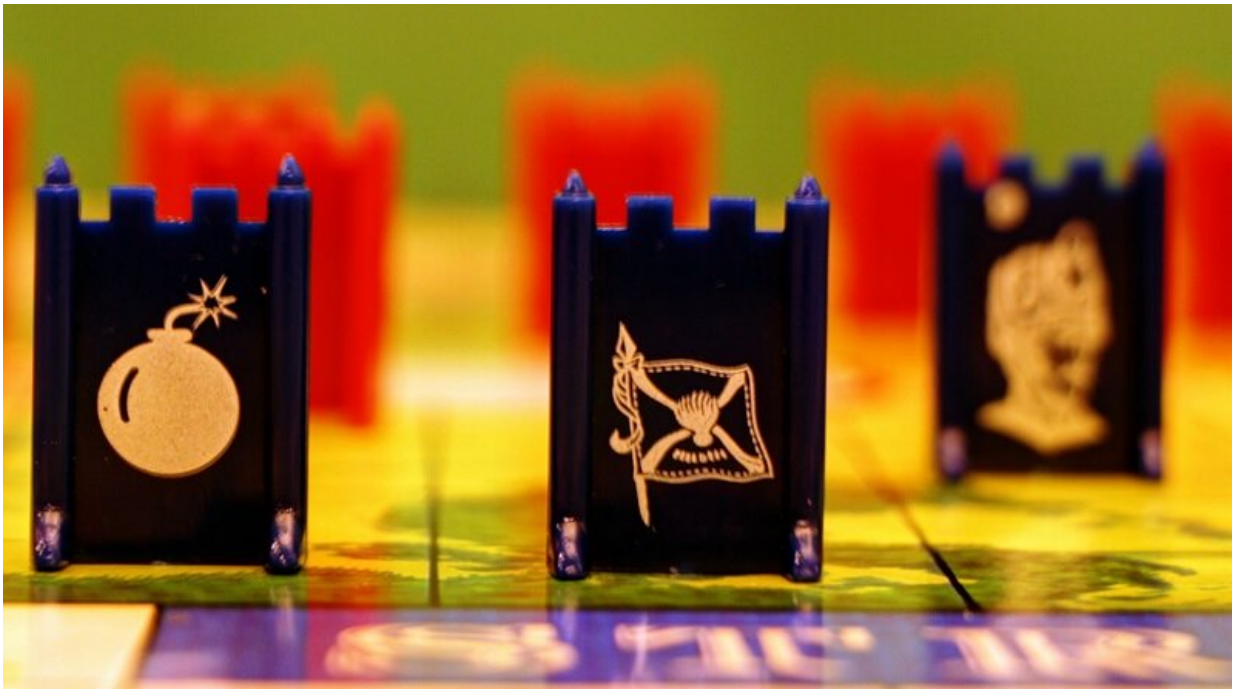
December 2 2022, by Bob Yirka

Stratego pieces. Credit: Wikipedia, <u>CC BY-SA</u>

A team of researchers at DeepMind Technologies Ltd., has created an AI application called "DeepNash" that is able to play the game Stratego at an expert level. In their paper published in the journal *Science*, the group describes the unique approach they took to improve the app's level of play.

Stratego is a two-player board game and is considered to be difficult to master. The goal for each player is to capture their opponent's flag, which is hidden among their initial 40 game pieces. Each of the game pieces is marked with a power ranking—higher-ranked players defeat lower-ranked players in face-offs. Making the game more difficult is that neither player can see the markings on the opponent's game pieces until they meet face-to-face.

Prior research has shown that the complexity of the game is higher than that of chess or go, with 10^{535} possible game scenarios. This level of complexity makes it extremely challenging for computer experts attempting to create Stratego-playing AI systems. In this new effort, the researchers took a different approach, creating an app capable of beating most human and other AI systems.

As with other AI systems designs, DeepNash first learned to play Stratego by playing itself many times—in this case, 5.5 billion times—equivalent to hundreds of years of playing time for a human. After it learned how to play, the researchers did not have it attempt to learn strategies from master human players, or even to play against other opponents in general.

Instead, the researchers devised an algorithm that worked toward an optimal strategy for each move rather than perfection. The algorithm was based on game theory: An optimal strategy would give DeepNash a 50/50 chance of success at a minimum on any given move—far better than humans could hope to achieve.

Testing showed that the team had found a way to improve the odds of an AI app playing Stratego—it achieved an 84% winning record while playing 50 times on an online gaming platform, and in so doing, became one of the top three players on the site. And the human opponents were never told they were playing against a computer.

More information: Julien Perolat et al, Mastering the game of Stratego with model-free multiagent reinforcement learning, *Science* (2022). DOI: 10.1126/science.add4679