Chess engine sacrifices mastery to mimic human play

January 26 2021, by Melanie Lefkowitz

Credit: Pixabay/CC0 Public Domain

When it comes to chess, computers seem to have nothing left to prove.

Since IBM's Deep Blue defeated world chess champion Garry Kasparov in 1997, advances in artificial intelligence have made chess-playing computers more and more formidable. No human has beaten a computer

in a chess tournament in 15 years.

In new research, a team including Jon Kleinberg, the Tisch University Professor of Computer Science, developed an artificially intelligent chess engine that doesn't necessarily seek to beat humans—it's trained to play like a human. This not only creates a more enjoyable chess-playing experience, it also sheds light on how computers make decisions differently from people, and how that could help humans learn to do better.

"Chess sits alongside virtuosic musical instrument playing and mathematical achievement as something humans study their whole lives and get really good at. And yet in chess, computers are in every possible sense better than we are at this point," Kleinberg said. "So chess becomes a place where we can try understanding human skill through the lens of super-intelligent AI."

Kleinberg is a co-author of "Aligning Superhuman AI With Human Behavior: Chess as a Model System," presented at the Association for Computing Machinery SIGKDD Conference on Knowledge Discovery and Data Mining, held virtually in August. In December, the Maia chess engine, which grew out of the research, was released on the free online chess server lichess.org, where it was played more than 40,000 times in its first week. Agadmator, the most-subscribed chess channel on YouTube, talked about the project and played two live games against Maia.

"Current chess AIs don't have any conception of what mistakes people typically make at a particular ability level. They will tell you all the mistakes you made—all the situations in which you failed to play with machine-like precision—but they can't separate out what you should work on," said co-author Ashton Anderson, assistant professor at the University of Toronto. "Maia has algorithmically characterized which

mistakes are typical of which levels, and therefore which mistakes people should work on and which mistakes they probably shouldn't, because they are still too difficult."

The paper's other co-authors are Reid McIlroy-Young, doctoral student at the University of Toronto, and Siddhartha Sen of Microsoft Research.

As artificial intelligence approaches or surpasses human abilities in a range of areas, researchers are exploring how to design AI systems with human collaboration in mind. In many fields, AI can inform or improve human work—for example, in interpreting the results of medical imaging—but algorithms approach problems very differently from humans, which makes learning from them difficult or, potentially, even dangerous.

In this project, the researchers sought to develop AI that reduced the disparities between human and algorithmic behavior by training the computer on the traces of individual human steps, rather than having it teach itself to successfully complete an entire task. Chess—with hundreds of millions of recorded moves by online players at every skill level—offered an ideal opportunity to train AI models to do just that.

"Chess been described as the `fruit fly' of AI research," Kleinberg said. "Just as geneticists often care less about the fruit fly itself than its role as a model organism, AI researchers love chess, because it's one of their model organisms. It's a self-contained world you can explore, and it illustrates many of the phenomena that we see in AI more broadly."

Training the AI model on individual human chess moves, rather than on the larger problem of winning a game, taught the computer to mimic human behavior. It also created a system that is more adjustable to different skill levels—a challenge for traditional AI.

Within each skill level, Maia matched human moves more than 50% of the time, with its accuracy growing as skill increases—a higher rate of accuracy than two popular chess engines, Stockfish and Leela. Maia was also able to capture what kinds of mistakes players at specific skill levels make, and when people reach a level of skill where they stop making them.

To develop Maia, the researchers customized Leela, an open-source system based on Deep Mind's AlphaZero program, that makes chess decisions with the same kinds of neural networks used to classify images or language. They trained different versions of Maia on games at different skill levels in order to create nine bots designed to play humans with ratings between 1100 and 1900 (ranging from the skill of more novice players to strong amateur players).

"Our model didn't train itself on the best move—it trained itself on what a human would do," Kleinberg said. "But we had to be very careful—you have to make sure it doesn't search the tree of possible moves too thoroughly, because that would make it too good. It has to just be laser-focused on predicting what a person would do next."

More information: Reid McIlroy-Young et al. Aligning Superhuman AI with Human Behavior, *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2020). DOI: 10.1145/3394486.3403219

Provided by Cornell University

Citation: Chess engine sacrifices mastery to mimic human play (2021, January 26) retrieved 5 May 2024 from https://techxplore.com/news/2021-01-chess-sacrifices-mastery-mimic-human.html

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.