

Chess: How to spot a potential cheat

September 29 2022, by Michael K. Cohen



Credit: Robbe Jaspers from Pexels

A few years ago, the chess website [Chess.com](https://www.chess.com) temporarily [banned](#) U.S. grandmaster [Hans Niemann](#) for playing chess moves online that the site suspected had been suggested to him by a computer program. It had reportedly previously [banned](#) his mentor Maxim Dlugy.

And at the [Sinquefield Cup](#) earlier this month, world champion [Magnus Carlsen](#) resigned without comment after playing a poor game against 19-year-old Niemann. He has since said this was because he believes [Niemann has continued to cheat](#) recently.

Another participant, the Russian Grandmaster Ian Nepomniachtchi, [called](#) Niemann's performance "more than impressive." While Nieman has admitted to sometimes having cheated in previous online games, he [has strongly denied](#) ever cheating at a live chess tournament.

But how does Chess.com, the world's biggest chess website, decide that a player has probably cheated? It can't show the world the code it uses, or else would-be cheaters would know exactly how to avoid detection. The website states: "Though legal and practical considerations prevent Chess.com from revealing the full set of data, metrics and tracking used to evaluate games in our fair-play tool, we can say that at the core of Chess.com's system is a [statistical model](#) that evaluates the probability of a human player matching an engine's top choices, and surpassing the confirmed clean play of some of the greatest chess players in history."

Luckily, research can shed light on which approach the website may be using.

Humans v AI

When AI company [DeepMind](#) developed the program AlphaGo, which could play the strategy game Go, [it was taught to predict which moves a human would make](#) from any given position.

Predicting human moves is a supervised learning problem, the bread and butter of machine learning. Given lots of examples of positions from human games (the dataset) and an example of a human move from each such position (the label), machine learning algorithms can be trained to

predict labels at new data points. So DeepMind taught its AI to estimate the probability that a human would make any given move from any given position.

AlphaGo famously beat human rival Lee Sedol in 2017. One of the AI's famous moves in the game was Move 37. As lead researcher David Silver noted in [the documentary AlphaGo](#), "AlphaGo said there was a 1/10,000 probability that Move 37 would have been played by a human player."

So according to that machine learning model of human Go players, if you saw a person play Move 37, it would be evidence that they didn't come up with the idea themselves. But of course, it wouldn't be proof. Any human *could* make that move.

To become very confident that someone cheats at a game, you have to look at lots of moves. For example, researchers have investigated how lots of moves from a player can be [analyzed collectively to detect anomalies](#).

Chess.com openly uses machine learning to predict which moves might be made by a human in any given position. In fact, it has different models of individual famous chess players, and you [can actually play against them](#). Presumably, similar models are used to detect cheating.

A recent study suggested that, in addition to predicting how likely a human would be to make a certain move, it's also important [to account for how good that move is](#). This matches Chess.com's statement that it evaluates whether moves "surpass ... confirmed clean play" from the greats.

But how do you measure which moves are better than others? In theory, a chess position is either "winning" (you can guarantee a win), "losing"

(the other player can) or "drawing" (neither can), and a good move would be any move that doesn't make your position worse. But realistically, although computers are much better at calculating and picking future moves than humans, for many positions not even they can tell for sure whether a position is winning, losing or drawing. And they certainly could never prove it—a proof would generally require too many calculations, examining every leaf of an exponential game tree.

So what people and computers do is use "[heuristics](#)" (gut guesses) to assess the "value" of different positions—estimating which player they think will win. This can also be cast as a machine learning problem where the dataset is lots of board positions and the labels are who won—which trains the algorithm to predict who will win from a given position.

Typically, machine learning models used for this purpose do some thinking about the next few likely moves, consider what positions are accessible to both players, and then use "gut feeling" about those future positions to inform their evaluation of the current position.

But who wins from a given position depends on how good the players are. So the model's evaluation of a particular game will depend on who was playing the games that made it into the training dataset. Typically, when chess commentators talk about the "objective value" of different positions, they mean who is likely to win from a given position when both sides are being played by the very best chess AIs available. But this measure of value isn't always the most useful when considering a position that human players will have to carry out in the end. So it's not clear exactly what Chess.com (or we) should consider to be a "good move."

If I were cheating at chess and made a few moves suggested by a chess engine, it might not even help me win. Those moves might be setting up

a brilliant attack that would never occur to me, so I would squander it unless I asked the chess engine to play the rest of the [game](#) for me. (Lichess.org tells me I've played 3,049 [Blitz games](#) at the time of writing, and my not-very-good [ELO rating](#) of 1632 means you can expect me to miss good tactics left and right.)

Detecting cheating is hard. If you're playing online and you're wondering if your opponent is cheating, you really aren't going to be able to tell with any measure of certainty—because you haven't seen millions of human games played with radically varying styles. It's a problem where machine learning models trained with huge amounts of data have a big advantage. Ultimately, they may be critical for the ongoing integrity of [chess](#).

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: Chess: How to spot a potential cheat (2022, September 29) retrieved 13 March 2024 from <https://techxplore.com/news/2022-09-chess-potential.html>

<p>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.</p>
--