

An AI named Cicero can beat humans in Diplomacy, a complex alliance-building game. Here's why that's big news

November 28 2022, by Toby Walsh



Credit: AI-generated image ([disclaimer](#))

In a rare piece of good news from Meta, artificial intelligence researchers at the company have just announced [a scientific breakthrough](#). Their AI program named Cicero can now play the board game Diplomacy at a human level.

Now, before you get too excited, Cicero isn't playing at superhuman level. It was beaten by around 10% of the humans it played against. By comparison, in previous AI milestones, like AI beating humans in chess or Go, humans have long been completely surpassed.

DeepMind's Go-playing AI is, for example, a "Go god"—according to the Chinese grandmaster Ke Jie. Even the human Go world champion would now lose 100-0 to the computer.

Diplomacy is a simplified and abstract [game](#), involving rival armies and navies invading, or not invading, each others' territories. It's fair to say it lacks the complexity and subtlety of the sort of [diplomacy](#) undertaken in the corridors of the United Nations.

Nevertheless, the news of Cicero's performance was one in the eye for tech rivals such as Google, who owns DeepMind. The CEO and founder of DeepMind, Demis Hassabis, is a Diplomacy expert. He won the World Team Championship in 2004, and was 4th in the world in the 2006 World Championship.

I expect Hassabis would be able to beat Cicero easily because of some of the limitations I will point out shortly.

The game of Diplomacy

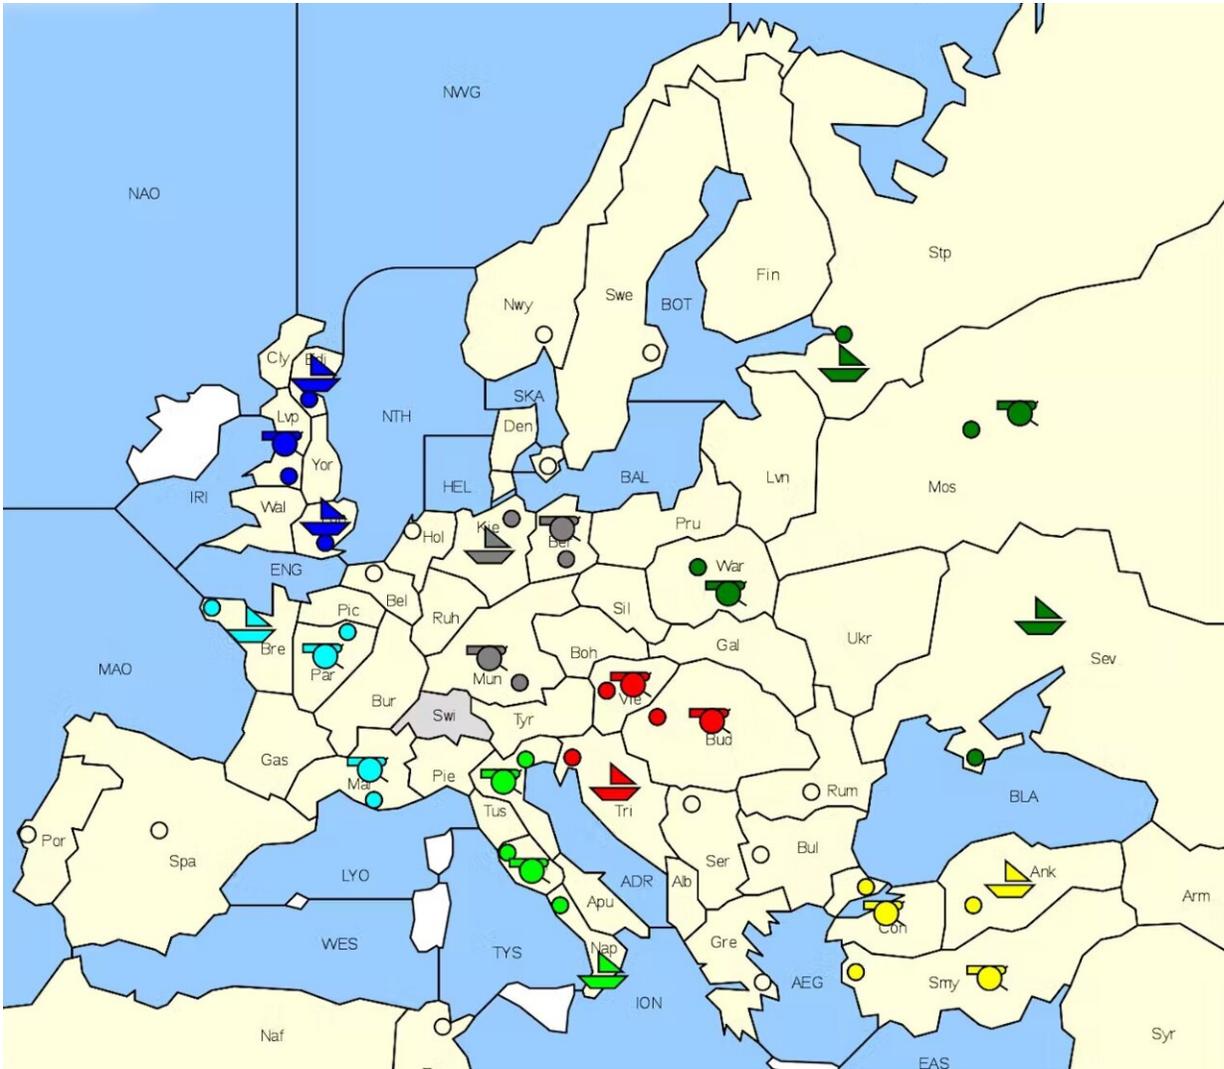
Diplomacy is what AI researchers call a "seven player, zero sum and deterministic game of imperfect information". A seven player game is much harder to solve than a two player game such as chess or Go. You must consider the many possible strategies of not one but six other players. This makes it much harder to write an AI to play the game.

Diplomacy is also a game of imperfect information, because players make moves simultaneously. Unlike games such as chess or Go, where

you know everything about your opponent's moves, players in Diplomacy make moves not knowing what their opponents are about to do. They must therefore predict their opponents' next actions. This also adds to the challenge of writing an AI to play it.

Finally, Diplomacy is a zero sum game in which if you win, I lose. And the outcome is deterministic and not dependent on chance. Nonetheless, before victory or defeat, it still pays for players to form alliances and team up on each other. Indeed, one of the real challenges in playing the game is managing the informal negotiations with other players before making simultaneous moves.

The main reason Cicero's performance is a [scientific breakthrough](#) is that it can both play the game well, and also perform these informal negotiations. This combination of natural language processing and strategic reasoning is a first for any game-playing AI.



This digital Diplomacy board shows the different land and sea territories players must traverse. Credit: Wikimedia

Beating Cicero

A close reading of the [paper Meta published](#) about Cicero in the journal *Science* offers a couple of clues about how you can beat it.

First, Cicero is almost entirely honest (unlike the [famous Roman](#) it's

named after). On the other hand, Diplomacy is a game of betrayal and dishonesty. Players offer to form alliances but often instantly renege on these deals. Cicero does not. It always plays straight.

Honesty is a surprisingly effective strategy in Diplomacy—but not if your opponents know you will never betray them. This is the catch. Cicero played anonymously, so its human opponents probably wouldn't have worked this out. But if you know this fact, it will be easy to take advantage.

Second, Cicero (this time like his namesake) is very talkative. Expert players of Diplomacy exchange twice the number of messages with other players than non-experts. The trick is to form alliances, and reassure your opponents of your intent. Cicero also exchanges twice the number messages of the human players it tends to beat.

Of course, being a bot, it is much easier for Cicero to handle six simultaneous conversations. And this, I would say, is an unfair advantage of being a computer in this scenario.

What next?

It's not clear how Meta intends to build on this research. A computer that can reason about the beliefs, goals, and intentions of others, as well as persuade and build relationships through dialogue, is a powerful tool. It's one that could be easily misused. Let's not forget how several years ago Facebook (which is owned by Meta) came in for a lot of justified criticism for an experiment to [manipulate users' emotions](#).

Yet it's hard to say exactly what the real-world applications of Cicero might be. After all, diplomacy in the real world is neither zero sum nor deterministic. Two countries can both agree not to go to war, and both will win.

Then there are multitudes of random factors that can change an outcome. The Spanish Armada, for example, lost more ships to unexpected summer storms than to enemy fire.

Whatever Meta's intent, the breakthrough is another example of how large tech companies are taking over the AI race with billion dollar investments that can't be matched by the public sector. Cicero was produced by a team of more than 25 researchers. Nobody working in a university has these sorts of resources to throw at solving a board game.

As an AI researcher at one of those universities, I am conflicted. I'm reminded of a famous graffito at Pompeii [which said](#)

Suti Ciciiro vapla bis

"You will like Cicero, or you will be whipped".

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

Citation: An AI named Cicero can beat humans in Diplomacy, a complex alliance-building game. Here's why that's big news (2022, November 28) retrieved 14 April 2024 from <https://techxplore.com/news/2022-11-ai-cicero-humans-diplomacy-complex.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.