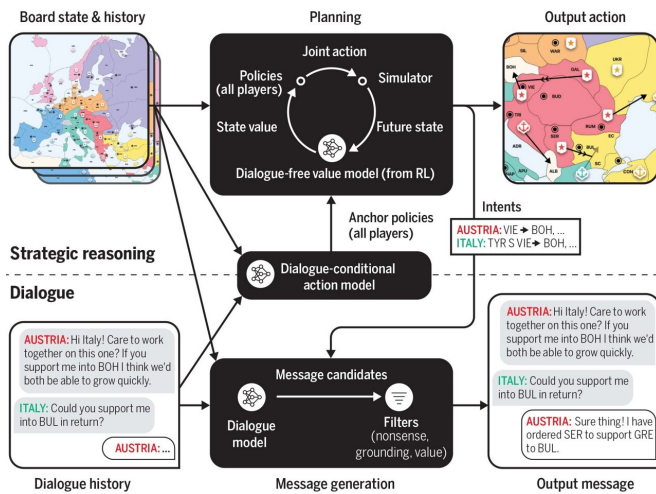


Meta team builds AI that plays 'Diplomacy' at very high level

23 November 2022, by Bob Yirka



Architecture of Cicero. Cicero predicts likely human actions for each player based on the board state and dialogue, using that as the starting point for a planning algorithm using RL-trained models. The output of planning is an action for the agent as well as beliefs about other players' actions, which are used to select intents for a dialogue model to condition on. Generated message candidates undergo several filtering steps before a final message is sent. Credit: *Science* (2022). DOI: 10.1126/science.ade9097

A group of researchers on the Meta Fundamental AI Research Diplomacy Team (FAIR), affiliated with multiple institutions in the U.S., has added a negotiation component to its Diplomacy-playing AI system. In a paper published in the journal *Science*, the team describes the work that went into building the system and its performance when pitted against anonymous human opponents.

The [game](#) of Diplomacy has been described as one of the most complex games ever devised—it recreates events surrounding World War I, when many nations were at war with many other nations, and not all of them were clearly aligned. In the game, [players](#) work to negotiate best possible

outcomes for the country they represent based on mutual agreements, battle outcomes, alignments and diplomacy.

As its name implies, one of the hallmarks of the game is diplomacy, which is a difficult-to-describe skill. In this new effort, the team at FAIR has improved on a previous version of an AI that played only the board part of the game by adding a module that engages in [real-time](#) English language [diplomacy](#) with human players.

The new system is called CICERO and was built around two main components: one that processes dialogue and another that strategizes. The strategizing engine was essentially the same as the one used in the prior version of the system. However, the team built a new engine that could be used to accept dialogue from other players, pass it along, and then receive strategic plans translated to dialogue sent out to other players. The [dialogue](#) engine was trained using data from 50,000 conversations in games played by humans.

Once the system was built and tested internally, the team logged on to an internet-based version of Diplomacy, where people were not known to one another. CICERO played the game anonymously, as would any of the human players. Thus, it not only had to play the game as a human would, it had to do so without other players noticing that they were playing against a computer system.

The researchers found that only one player became suspicious of the humanity of CICERO. Also, the AI system turned out to be a better player than most of its human opponents, scoring in the top 10% of players.

More information: Andrew Goff et al, Human-level play in the game of Diplomacy by combining language models with strategic reasoning, *Science* (2022). [DOI: 10.1126/science.ade9097](#)

Project page:

ai.facebook.com/research/cicero/diplomacy/

© 2022 Science X Network

APA citation: Meta team builds AI that plays 'Diplomacy' at very high level (2022, November 23)
retrieved 28 November 2022 from <https://techxplore.com/news/2022-11-meta-team-ai-plays-diplomacy.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.