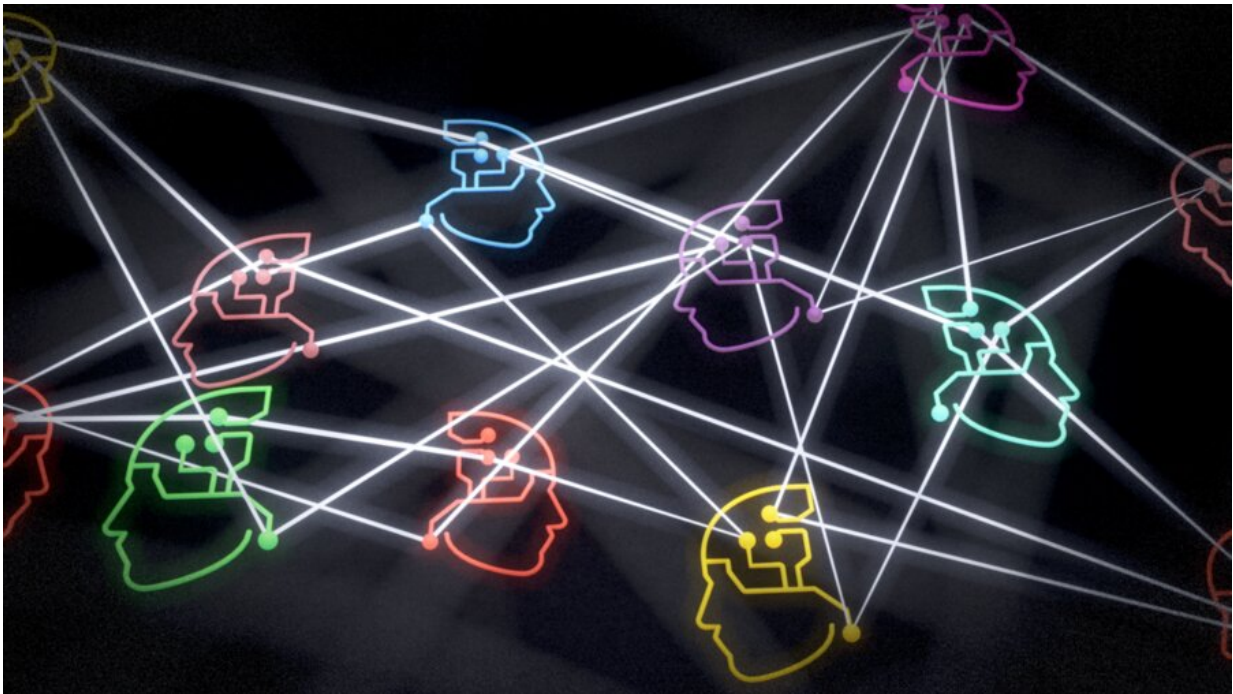# Is diversity the key to collaboration? New AI research suggests so

May 26 2022, by Kylie Foy

New research suggests that training an artificial intelligence model with mathematically "diverse" teammates improves its ability to collaborate with other AI it has never worked with before. Credit: Bryan Mastergeorge

As artificial intelligence gets better at performing tasks once solely in the hands of humans, like driving cars, many see teaming intelligence as a next frontier. In this future, humans and AI are true partners in high-stakes jobs, such as performing complex surgery or defending from

missiles. But before teaming intelligence can take off, researchers must overcome a problem that corrodes cooperation: humans often do not like or trust their AI partners.

Now, new research points to diversity as being a key parameter for making AI a better team player.

MIT Lincoln Laboratory researchers have found that training an AI model with mathematically "diverse" teammates improves its ability to collaborate with other AI it has never worked with before, in the card game Hanabi. Moreover, both Facebook and Google's DeepMind concurrently published independent work that also infused diversity into training to improve outcomes in human-AI collaborative games.

Altogether, the results may point researchers down a promising path to making AI that can both perform well and be seen as good collaborators by human teammates.

"The fact that we all converged on the same idea—that if you want to cooperate, you need to train in a diverse setting—is exciting, and I believe it really sets the stage for the future work in cooperative AI," says Ross Allen, a researcher in Lincoln Laboratory's Artificial Intelligence Technology Group and co-author of a paper detailing this work, which was recently presented at the International Conference on Autonomous Agents and Multi-Agent Systems.

Adapting to different behaviors

To develop cooperative AI, many researchers are using Hanabi as a testing ground. Hanabi challenges players to work together to stack cards in order, but players can only see their teammates' cards and can only give sparse clues to each other about which cards they hold.

In [a previous experiment](#), Lincoln Laboratory researchers tested one of the world's best-performing Hanabi AI models with humans. They were surprised to find that humans strongly disliked playing with this AI model, calling it a confusing and unpredictable teammate. "The conclusion was that we're missing something about human preference, and we're not yet good at making models that might work in the real world," Allen says.

The team wondered if cooperative AI needs to be trained differently. The type of AI being used, called [reinforcement learning](#), traditionally learns how to succeed at complex tasks by discovering which actions yield the highest reward. It is often trained and evaluated against models similar to itself. This process has created unmatched AI players in competitive games like Go and StarCraft.

But for AI to be a successful collaborator, perhaps it has to not only care about maximizing reward when collaborating with other AI agents, but also something more intrinsic: understanding and adapting to others' strengths and preferences. In other words, it needs to learn from and adapt to diversity.

How do you train such a diversity-minded AI? The researchers came up with "Any-Play." Any-Play augments the process of training an AI Hanabi agent by adding another objective, besides maximizing the game score: the AI must correctly identify the play-style of its training partner.

This play-style is encoded within the training partner as a latent, or hidden, variable that the agent must estimate. It does this by observing differences in the behavior of its partner. This objective also requires its partner to learn distinct, recognizable behaviors in order to convey these differences to the receiving AI agent.

Though this method of inducing diversity is not new to the field of AI,

the team extended the concept to collaborative games by leveraging these distinct behaviors as diverse play-styles of the game.

"The AI agent has to observe its partners' behavior in order to identify that secret input they received and has to accommodate these various ways of playing to perform well in the game. The idea is that this would result in an AI agent that is good at playing with different play styles," says first author and Carnegie Mellon University Ph.D. candidate Keane Lucas, who led the experiments as a former intern at the laboratory.

Playing with others unlike itself

The team augmented that earlier Hanabi model (the one they had tested with humans in their prior experiment) with the Any-Play training process. To evaluate if the approach improved collaboration, the researchers teamed up the model with "strangers"—more than 100 other Hanabi models that it had never encountered before and that were trained by separate algorithms—in millions of two-player matches.

The Any-Play pairings outperformed all other teams, when those teams were also made up of partners who were algorithmically dissimilar to each other. It also scored better when partnering with the original version of itself not trained with Any-Play.

The researchers view this type of evaluation, called inter-algorithm cross-play, as the best predictor of how cooperative AI would perform in the real world with humans. Inter-algorithm cross-play contrasts with more commonly used evaluations that test a model against copies of itself or against models trained by the same algorithm.

"We argue that those other metrics can be misleading and artificially boost the apparent performance of some algorithms. Instead, we want to know, 'if you just drop in a partner out of the blue, with no prior

knowledge of how they'll play, how well can you collaborate?' We think this type of evaluation is most realistic when evaluating cooperative AI with other AI, when you can't test with humans," Allen says.

Indeed, this work did not test Any-Play with humans. However, research published by DeepMind, simultaneous to the lab's work, used a similar diversity-training approach to develop an AI agent to play the collaborative game Overcooked with humans. "The AI agent and humans showed remarkably good cooperation, and this result leads us to believe our approach, which we find to be even more generalized, would also work well with humans," Allen says. Facebook similarly used diversity in training to improve collaboration among Hanabi AI agents, but used a more complicated algorithm that required modifications of the Hanabi game rules to be tractable.

Whether inter-algorithm cross-play scores are actually good indicators of human preference is still a hypothesis. To bring human perspective back into the process, the researchers want to try to correlate a person's feelings about an AI, such as distrust or confusion, to specific objectives used to train the AI. Uncovering these connections could help accelerate advances in the field.

"The challenge with developing AI to work better with humans is that we can't have humans in the loop during training telling the AI what they like and dislike. It would take millions of hours and personalities. But if we could find some kind of quantifiable proxy for human preference—and perhaps diversity in training is one such proxy —then maybe we've found a way through this challenge," Allen says.

 More information: Keane Lucas, Ross E. Allen, Any-Play: An Intrinsic Augmentation for Zero-Shot Coordination. arXiv:2201.12436v1 [cs.AI], arxiv.org/abs/2201.12436

Provided by Massachusetts Institute of Technology