


DeepMind AI shows off winning cooperative team behavior

July 5 2018, by Nancy Owano



A schematic of the For The Win (FTW) agent architecture. The agent combines recurrent neural networks (RNNs) on fast and slow timescales, includes a shared memory module, and learns a conversion from game points to internal reward. Credit: DeepMind blog

The recent [Dota 2](#) feats of openAI, where their people spent months training a new AI system so bots could play together as a team, had emboldened the writing on the wall: Team play in multiplayer video games is a hot focal point in AI right now.

Researchers are taking on the challenge in how to make that happen on advanced levels. (Why is AI teamwork always described as such a challenge? Will Knight in *MIT Technology Review* said, the "Teamwork is extremely difficult to develop effectively in AI programs, because it involves dealing with a complex and ever-changing situation.")

This week you have proof of another star team-up, this time the work of DeepMind. They royally showed how AI can beat out human opponents in a multi-player game. DeepMind's scientists and engineers reminded us how AI agents can match and in some instances surpass humans in cooperative play.

Posting in The Download, *MIT Technology Review*'s Knight said AI agents were trained by DeepMind in a modified version of the Quake III Arena. Teamwork skills were on parade—the agents played Capture The Flag (CTF) on Quake III Arena.

For those not familiar with CTF and Quake III Arena, an overview from Rob LeFebvre can help. Writing in *Engadget*, he explained their challenge.

"The team focused on a capture the flag mode, one in which the map changes from match to match. Its AI agents had to learn general strategies to be able to adapt to each new map, something humans do easily. The agents also needed to both cooperate with team members as well as compete against the opposite team, and be able to adjust to different enemy play [styles](#)."

A test [tournament](#) involved humans playing CTF with and against trained agents.

The "For The Win" (FTW) agents learned to become stronger than the strong baseline methods, and exceeded the win-rate of the human players. In a survey among participants they were rated more collaborative than human participants.

As *MIT Technology Review* stated, "the AI agents can also collaborate with human players—and those players say the programs are better teammates than most [people](#)."

"Two teams of individual players compete on a given map with the goal of capturing the opponent team's flag while protecting their own. To gain tactical advantage they can tag the opponent team members to send them back to their spawn points. The team with the most flag captures after five minutes wins."

The Telegraph described the game as a "maze with the [goal](#) of capturing the opponent team's flag while protecting their own."

For Margi Murphy in *The Telegraph*, what is so interesting about this feat is that "DeepMind's AI appeared to pick up typically human psychological tactics in order to win games - rather than memorising the game map and possible outcomes from a move."

The For the Win (FTW) AI played nearly 450,000 games of Quake III Arena to gain its dominance over human players, said Khari Johnson in *The Verge*, and establish its understanding of how to effectively work with other machines and humans.

Johnson recapped the game outcomes.

Flags captured: "On average, human-machine teams captured 16 flags per game less than a team of two FTW agents."

Tagging performance: "Agents were found to be efficient than humans in tagging, achieving the tactic 80 percent of the time compared to humans 48 percent." (Tagging involved touching an opponent to send them back to their spawning point.)

Behavior: A survey of human participants "found FTW to be more collaborative than human teammates."

The DeepMind blog is interesting for anyone curious about the AI training process for such a scenario in multi-player video games.

"Rather than training a single agent, we train a population of agents, which learn by playing with each other, providing a diversity of teammates and opponents.

"Each agent in the population learns its own internal reward signal, which allows agents to generate their own internal goals, such as capturing a flag. A two-tier optimisation process optimises agents' internal rewards directly for winning, and uses reinforcement learning on the internal rewards to learn the agents' policies.

"Agents operate at two timescales, fast and slow, which improves their ability to use memory and generate consistent action sequences."

More information: Human-level performance in first-person multiplayer games with population-based deep reinforcement learning, arXiv:1807.01281 [cs.LG], arxiv.org/abs/1807.01281

Abstract

Recent progress in artificial intelligence through reinforcement learning

(RL) has shown great success on increasingly complex single-agent environments and two-player turn-based games. However, the real-world contains multiple agents, each learning and acting independently to cooperate and compete with other agents, and environments reflecting this degree of complexity remain an open challenge. In this work, we demonstrate for the first time that an agent can achieve human-level in a popular 3D multiplayer first-person video game, Quake III Arena Capture the Flag, using only pixels and game points as input. These results were achieved by a novel two-tier optimisation process in which a population of independent RL agents are trained concurrently from thousands of parallel matches with agents playing in teams together and against each other on randomly generated environments. Each agent in the population learns its own internal reward signal to complement the sparse delayed reward from winning, and selects actions using a novel temporally hierarchical representation that enables the agent to reason at multiple timescales. During game-play, these agents display human-like behaviours such as navigating, following, and defending based on a rich learned representation that is shown to encode high-level game knowledge. In an extensive tournament-style evaluation the trained agents exceeded the win-rate of strong human players both as teammates and opponents, and proved far stronger than existing state-of-the-art agents. These results demonstrate a significant jump in the capabilities of artificial agents, bringing us closer to the goal of human-level intelligence.

© 2018 Tech Xplore

Citation: DeepMind AI shows off winning cooperative team behavior (2018, July 5) retrieved 9 April 2024 from <https://techxplore.com/news/2018-07-deepmind-ai-cooperative-team-behavior.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.