

# DeepMind develops AI that demonstrates social learning capabilities

December 4 2023, by Bob Yirka

**Worlds**

Square worlds: 18m × 18m to 31m × 31m

Bumpy terrain uses a noisy height map.

Walls are invisible to the agents.

Obstacles have mesh colliders.

Density of all obstacle types, bumpiness and world size are *independent* variables.

**Games**

Goal colors are always unique.

Ex. 1: 4-cycle, 0 crossings, clockwise

Ex. 2: 5-cycle, 2 crossings

Ex. 3: Score changes on collision

Ex. 4: Two equally rewarding paths

Ex. 5: Paths can become quite complex

**Players**

Expert  
Scripted Bot or Human Trajectory

MEDAL-ADR  
Our Agents

**Actions**

- Move: forward, back, left or right
- Rotate: pitch and yaw
- Jump or crouch

**Observations**

LIDAR sensor (training & test)

AVATAR sensor (training only)

$\{rel\_pos\_x, rel\_pos\_y, rel\_pos\_z\}$

GoalCycle3D. A 3D physical simulated task space. Each task contains procedurally generated terrain, obstacles, and goal spheres, with parameters randomly sampled on task creation. Each agent is independently rewarded for visiting goals in a particular cyclic order, also randomly sampled on task creation. The correct order is not provided to the agent, so an agent must deduce the rewarding order either by experimentation or via cultural transmission from

an expert. Our task space presents navigational challenges of open-ended complexity, parameterized by world size, obstacle density, terrain bumpiness and a number of goals. Credit: *Nature Communications* (2023). DOI: 10.1038/s41467-023-42875-2

A team of AI researchers at Google's DeepMind project have developed a type of AI system that is able to demonstrate social learning capabilities. In their paper [published](#) in the journal *Nature Communications*, the group describes how they developed an AI application that showed it was capable of learning new skills in a virtual world by copying the actions of an implanted "expert."

Most AI systems, such as ChatGPT, gain their knowledge through exposure to huge amounts of data, such as from repositories on the Internet. But such an approach, those in the industry have noted, is not very efficient. Therefore many in the field continue to look for other ways to teach AI systems to learn.

One of the most popular approaches used by researchers is to attempt to mimic the process by which humans learn. Like traditional AI apps, humans learn by exposure to known elements in an environment and by following the examples of others who know what they are doing. But unlike AI apps, humans pick things up without the need for huge numbers of examples. A child can learn to play the game of Jacks, for example, after watching others play for just a few minutes—an example of cultural transmission. In this new effort, the research team has attempted to replicate this process using AI constrained to a virtual world.

The work by the team involved first building a virtual world (called GoalCycle3D) made up of uneven terrain upon which sat various

obstacles and multiple-colored spheres. They then added AI [agents](#), which were meant to travel through the virtual world by avoiding the obstacles and passing through the spheres. The agents were given learning modules but no other information about the world they would inhabit. They gained knowledge of how to proceed via reinforcement learning.

To get the agents to learn, they were given rewards and allowed to make their way through multiple similar virtual worlds, over and over. By doing this, the agents were able to make their way through the virtual world to a desired destination. The researchers then added another feature to the [virtual world](#), [expert](#) agents that already knew the best way to get from one place to another without running into obstacles. In the new scenario, the non-expert agents soon learned that the quickest way to get to a desired destination was to learn from an expert.

In watching the agents learn, the researchers found that they did so much more quickly with the expert and were able to better able to navigate other new similar virtual worlds by mimicking what they had learned from the expert in prior trials. They were also able to apply such skills (courtesy of memory modules) even in the absence of the expert—an example, the researchers claim, of social learning.

**More information:** Avishkar Bhoopchand et al, Learning few-shot imitation as cultural transmission, *Nature Communications* (2023). [DOI: 10.1038/s41467-023-42875-2](https://doi.org/10.1038/s41467-023-42875-2)

© 2023 Science X Network

Citation: DeepMind develops AI that demonstrates social learning capabilities (2023, December 4) retrieved 30 April 2024 from <https://techxplore.com/news/2023-12-deepmind-ai-social-capabilities.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.