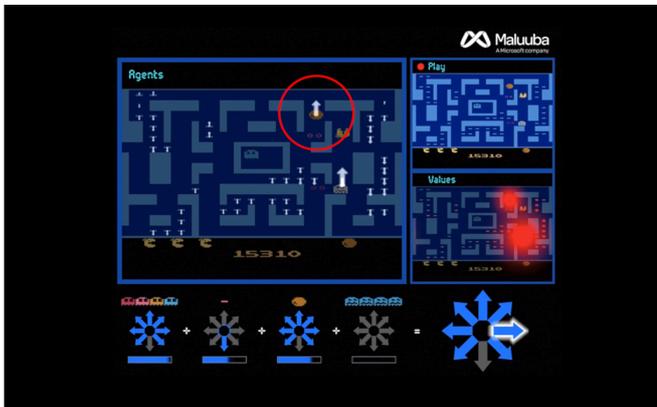


Microsoft AI seriously at play with Ms. Pac-Man

16 June 2017, by Nancy Owano



Credit: Microsoft

Microsoft AI has won the maximum score of 999,990 points playing Ms. Pac-Man, surpassing the best human high-score record by four times. No human or AI has ever achieved this score.

Yes, that is a big deal because those familiar with Ms. Pac-Man understand that it is a tough [game](#).

Katyanna Quach and Andrew Silver on Thursday in *The Register* noted Ms. Pac-Man was a "rather tricky game," tough work for an artificial brain.

How tough? "Computers can't play this game well since there are just too many possible game states to consider," they said.

Allison Linn in the Microsoft blog discussing this Ms. Pac-Man feat quoted an associate professor of computer science at McGill, Doina Precup, who said that was indeed special.

AI researchers use videogames to test their systems but, Precup said, have found Ms. Pac-Man among the most difficult to [crack](#).

Note this is the Atari 2600 version of Ms. Pac-Man that was played. Steve Golson, meanwhile, one of the co-creators of the arcade version, said in the blog the reason why Ms. Pac-Man had to be simple to grasp yet nearly impossible to conquer was that it was originally designed for arcade play, and they wanted people to keep dropping coins

The spotlight need at this point to turn to Maluuba, a Canadian deep learning startup acquired by Microsoft. They use something they call Hybrid Reward Architecture, a method parked under the umbrella of [reinforcement learning](#).

This architecture used more than 150 agents. Each worked in parallel with the other agents to master Ms. Pac-Man, said Linell. This is how the team got the high score—dividing the large problem of mastering Ms. Pac-Man into small pieces, distributed among AI agents. (Precup in the blog said that this idea of having them work on different pieces to achieve a common goal was interesting.)

Their paper discusses the strategy. "Hybrid Reward Architecture for Reinforcement Learning" is up on arXiv. The authors are from McGill and Microsoft Maluuba, in Montreal.

"One of the strengths of HRA is that it can exploit domain knowledge to a much greater extent than single-head methods," the authors wrote.

The Register had a helpful description of how the AI behaved in game play:

"Instead of a single bot trying to singlehandedly complete the game, the problem is shared between up to 163 sub-agents working in parallel for an oracle agent. This central oracle controls Ms. Pac-Man's [movements](#). When the oracle agent finds a new object – a pellet, ghost or fruit – it creates a sub-agent representing that object and assigns it a fixed weight. Pills and fruit get positive weights, whereas ghosts get negative weights."

The oracle aggregates expected rewards from sub-agents; the information is used to move Ms. Pac-Man in the directions that maximize the total reward.

All in all, Quach and Silver made the observation that the HRA is a proof of concept that did not have to learn the hard way. "It was born knowing everything it ever needed to know," said Quach and Silver, describing it as "a preprogrammed maze-searching algorithm." Hardcoded to solve Ms. Pac-Man, it may "be tough to adapt the design to other scenarios without starting all over again with another specialized model."

Then why should we care?

Allison Linn said their "method "could have broad implications for teaching AI agents to do [complex tasks](#) that augment human capabilities."

Precup in the blog said that was similar to some theories of how the brain works, "and it could have broad implications for teaching AIs to do complex tasks with limited [information](#)."

More information: arxiv.org/pdf/1706.04208.pdf

© 2017 Tech Xplore

APA citation: Microsoft AI seriously at play with Ms. Pac-Man (2017, June 16) retrieved 2 July 2022 from <https://techxplore.com/news/2017-06-microsoft-ai-ms-pac-man.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.