

Atari master: New AI smashes Google DeepMind in video game challenge

January 31 2019



Credit: CC0 Public Domain

A new breed of algorithms has mastered Atari video games 10 times faster than state-of-the-art AI, with a breakthrough approach to problem solving.

Designing AI that can negotiate planning problems, especially those

where rewards are not immediately obvious, is one of the most important research challenges in advancing the field.

A famous 2015 study showed Google DeepMind AI learnt to play Atari video games like Video Pinball to human level, but notoriously failed to learn a path to the first key in 1980s video [game](#) Montezuma's Revenge due to the game's complexity.

In the new method developed at RMIT University in Melbourne, Australia, computers set up to autonomously play Montezuma's Revenge learnt from mistakes and identified sub-goals 10 times faster than Google DeepMind to finish the game.

Associate Professor Fabio Zambetta from RMIT University unveils the new approach this Friday at the 33rd AAAI Conference on Artificial Intelligence in the United States.

The method, developed in collaboration with RMIT's Professor John Thangarajah and Michael Dann, combines "carrot-and-stick" reinforcement learning with an intrinsic motivation approach that rewards the AI for being curious and exploring its environment.

"Truly intelligent AI needs to be able to learn to complete tasks autonomously in ambiguous environments," Zambetta says.

"We've shown that the right kind of algorithms can improve results using a smarter approach rather than purely brute forcing a problem end-to-end on very powerful computers.

"Our results show how much closer we're getting to autonomous AI and could be a key line of inquiry if we want to keep making substantial progress in this field."

Zambetta's method rewards the system for autonomously exploring useful sub-goals such as 'climb that ladder' or 'jump over that pit', which may not be obvious to a computer, within the context of completing a larger mission.

Other state-of-the-art systems have required human input to identify these sub-goals or else decided what to do next randomly.

"Not only did our algorithms autonomously identify relevant tasks roughly 10 times faster than Google DeepMind while playing Montezuma's Revenge, they also exhibited relatively human-like behaviour while doing so," Zambetta says.

"For example, before you can get to the second screen of the game you need to identify sub-tasks such as climbing ladders, jumping over an enemy and then finally picking up a key, roughly in that order.

"This would eventually happen randomly after a huge amount of time but to happen so naturally in our testing shows some sort of intent.

"This makes ours the first fully autonomous sub-goal-oriented agent to be truly competitive with state-of-the-art agents on these games."

Zambetta said the system would work outside of video games in a wide range of tasks, when supplied with raw visual inputs.

"Creating an algorithm that can complete [video](#) games may sound trivial, but the fact we've designed one that can cope with ambiguity while choosing from an arbitrary number of possible actions is a critical advance.

"It means that, with time, this technology will be valuable to achieve goals in the real world, whether in self-driving cars or as useful robotic

assistants with natural language recognition," he says.

Deriving Subgoals Autonomously to Accelerate Learning in Sparse Reward Domains (attached) will be presented at the 33rd AAAI Conference on Artificial Intelligence in Honolulu, Hawaii on 1 February 2019.

Provided by RMIT University

Citation: Atari master: New AI smashes Google DeepMind in video game challenge (2019, January 31) retrieved 10 April 2024 from <https://techxplore.com/news/2019-01-atari-master-ai-google-deepmind.html>

<p>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.</p>
--