


DeepMind researchers boost AI learning speed with UNREAL agent

21 November 2016, by Nancy Owano



"Our reinforcement learning agents have achieved breakthroughs in Atari 2600 games and the game of Go. Such systems, however, can require a lot of data and a long time to learn so we are always looking for ways to improve our generic learning algorithms."

So what did they do? They said they worked on "augmenting the standard deep reinforcement learning methods with two main additional tasks for our agents to perform during training."

The paper explained that "The first task involves the agent learning how to control the pixels on the screen, which emphasises learning how your actions affect what you will see rather than just prediction. This is similar to how a baby might learn to control their hands by moving them and observing the movements."

The authors said, "In the second task the agent is trained to predict the onset of immediate rewards from a short historical context."

Writing in Bloomberg, Jeremy Kahn said that they imbedded their technology "with attributes that function in a way similar to how animals are thought to dream."

How so? Kahn said, "One way the researchers achieved their results was by having Unreal [replay](#) its own past attempts at the game, focusing especially on situations in which it had scored points before. The researchers equated this in their paper to the way "animals dream about positively or negatively rewarding events more frequently."

Does the agent work well? The authors of the paper reported their results.

"Our agent significantly outperforms the previous state-of-the art on Atari, averaging 880% expert human performance, and a challenging suite of first-person, three-dimensional Labyrinth tasks leading

(Tech Xplore)—At Google's DeepMind, a team has made AI inroads in speed and performance.

The researchers behind the effort have put out a paper describing the maneuvers and the paper is on arXiv, titled "Reinforcement learning with unsupervised auxiliary tasks."

Authors are Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Leibo, David Silver and Koray Kavukcuoglu, of DeepMind, London.

They call their new agent UNREAL (Unsupervised Reinforcement and Auxiliary Learning). They said in the DeepMind blog that they tested this agent on a suite of 57 Atari games as well as a 3-D environment called Labyrinth, with 13 levels.

Liam Tung in *ZDNet* wrote about the new agent and pointed out that it is using the same learning methods used to master the game Go.

The [DeepMind blog](#), by Jaderberg, Mnih, and Czarnecki, said they were thinking how they could raise the bar on their own algorithms.

to a mean speedup in learning of 10x and averaging expert human performance, and a challenging suite of first-person, three-dimensional environments leading to a mean speedup in learning of 10x and averaging 87% expert human performance on Labyrinth."

Commenting more on those numbers, the blog entry said, "In Labyrinth, the result of using the auxiliary tasks - controlling the pixels on the screen and predicting when reward is going to occur - means that UNREAL is able to learn over 10x faster than our previous best A3C agent, and reaches far better performance. We can now achieve 87% of expert human performance averaged across the Labyrinth levels we considered, with super-human performance on a number of them."

of first-person, three-dimensional environments leading to a mean speedup in learning of 10x and averaging 87% expert human performance on Labyrinth

© 2016 Tech Xplore

Summing up, then, what is the significance of their research? The authors stated, "We have shown how augmenting a deep [reinforcement learning](#) agent with auxiliary control and reward prediction tasks can drastically improve both data efficiency and robustness to hyperparameter settings."

Tung said, "now they have a faster-learning agent and one that's also more [flexible](#)."

Looking into the future, they said they hoped their work will allow them "to scale up our agents to ever more complex environments," according to the blog.

More information: Reinforcement Learning with Unsupervised Auxiliary Tasks, arXiv:1611.05397 [cs.LG] arxiv.org/abs/1611.05397

Abstract

Deep reinforcement learning agents have achieved state-of-the-art results by directly maximising cumulative reward. However, environments contain a much wider variety of possible training signals. In this paper, we introduce an agent that also maximises many other pseudo-reward functions simultaneously by reinforcement learning. All of these tasks share a common representation that, like unsupervised learning, continues to develop in the absence of extrinsic rewards. We also introduce a novel mechanism for focusing this representation upon extrinsic rewards, so that learning can rapidly adapt to the most relevant aspects of the actual task. Our agent significantly outperforms the previous state-of-the-art on Atari, averaging 880%

APA citation: DeepMind researchers boost AI learning speed with UNREAL agent (2016, November 21) retrieved 27 September 2020 from <https://techxplore.com/news/2016-11-deepmind-boost-ai-unreal-agent.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.