


DeepMind thinkers test architectures on puzzle game and spaceship navigation game

July 26 2017, by Nancy Owano



An agent playing the spaceship task. The red lines indicate trajectories that are

executed in the environment while blue and green depict imagined trajectories.
Credit: Deepmind blog

(Tech Xplore)—Agents that imagine and plan: that is the title of a DeepMind discussion earlier this month from six DeepMind team members.

They make a case in that direction for DeepMind research work that was done on two papers, which are now on arXiv.

Imagining the consequences of your actions before you take them is a powerful tool of human cognition, they said. (We do it every day. Like that pen teetering on the edge of your tabletop right now. You can "imagine" it falling so you move it to a safer position.)

"If our algorithms are to develop equally sophisticated behaviours, they too must have the capability to 'imagine' and reason about the future."

As Thomas Claburn pointed out in *The Register*, DeepMind "has found that instilling its software agents with something like imagination helps them [learn](#) better."

Something *like* imagination.

If you study news sites reporting on the research, you notice they carefully word what it is and what it is not but using phrases such as imagination-"like" and imagination-"based."

They fundamentally are looking at novel techniques for improving deep reinforcement learning.

Claburn translated what they are doing in very clear language:

"Reinforcement learning is a form of machine learning. It involves a software agent that learns by interacting with a specific environment, usually through trial and error. Deep learning is a form of [machine](#) that involves algorithms inspired by the human brain, called neural networks. And the two techniques can be used together." DeepMind's work, he said, tries to offer the best of both worlds.

In the real world, complexity rules. And speaking of rules, in real life they are not so clearly defined. Stuff happens you cannot easily predict. As for planning strategies, we know too well that a one-size-fits-all answer does not always work in real environments.

How well can agents take on complexities? The team members have two papers that deal with all this. They described a family of approaches for imagination-based planning. Architectures were introduced for new ways for agents to learn and construct plans to maximize the efficiency of a task.

One of the more interesting feature descriptions about these agents is that "they can learn different strategies to construct plans. They do this by choosing between continuing a current imagined trajectory or restarting from scratch."

They also can use different imagination models, "with different accuracies and computational costs."

The researchers tested the architectures on tasks, including the puzzle game Sokoban and a spaceship navigation game.

Sokoban video notes: An agent played Sokoban from a pixel representation, not knowing the rules of the game.

Notes on the spaceship task: The DeepMind posting carried a visual of "An agent playing the spaceship task. The red lines indicate trajectories that are executed in the environment while blue and green depict imagined trajectories."

Results?

"For both tasks, the imagination-augmented agents outperform the imagination-less baselines considerably: they learn with less experience and are able to deal with the imperfections in modelling the environment."

As Alejandro Tauber in *TNW* said, "the type of imagination described in these papers is nowhere near what humans are capable of, but it does show that AIs can and benefit from being able to efficiently imagine different scenarios before [acting](#)."

About the papers: "Imagination-Augmented Agents for Deep Reinforcement Learning" was submitted this month on *arXiv*.

These agents use approximate environment models by 'learning to interpret' their imperfect predictions, they said, and their algorithm can be trained directly on low-level observations with little domain knowledge.

"Without making any assumptions about the structure of the environment model and its possible imperfections, our approach learns in an end-to-end way to extract useful knowledge gathered from model simulations – in particular not relying exclusively on simulated returns."

The other paper is "Learning model-based planning from scratch," also submitted this month and up on *arXiv*. "We show that our architecture can learn to solve a challenging continuous control problem, and also

learn elaborate planning strategies in a discrete maze-solving task."

More information: Imagination-Augmented Agents for Deep Reinforcement Learning, arXiv:1707.06203 [cs.LG]
arxiv.org/abs/1707.06203

Learning model-based planning from scratch, arXiv:1707.06170 [cs.AI]
arxiv.org/abs/1707.06170

Deepmind blog: deepmind.com/blog/agents-imagine-and-plan/

© 2017 Tech Xplore

Citation: DeepMind thinkers test architectures on puzzle game and spaceship navigation game (2017, July 26) retrieved 9 April 2024 from <https://techxplore.com/news/2017-07-deepmind-thinkers-architectures-puzzle-game.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>
--