


Teaching AI agents navigation subroutines by feeding them videos

12 June 2019, by Ingrid Fadelli



Given an input image as shown at the top, the researchers propose a novel method to learn subroutines, short horizon policies that exhibit a coherent behavior (such as going left into a room), and affordances, what subroutines can be invoked where. Credit: Kumar, Gupta & Malik.

Researchers at UC Berkeley and Facebook AI Research have recently proposed a new approach that can enhance the navigation skills of machine learning models. Their method, presented in a paper pre-published on arXiv, allows models to acquire visuo-motor navigation subroutines by processing a series of videos.

"Every morning, when you decide to get a cup of coffee from the kitchen, you think of going down the hallway, turning left into the corridor and then entering the room on the right," the researchers wrote in their paper. "Instead of deciding the exact muscle torques, you plan at this higher level of abstraction by composing these reusable lower level visuo-motor subroutines to reach your goal."

These "visuo-motor subroutines" or "hierarchical abstractions" that humans create in their minds ultimately help them to effectively move within their surrounding environment. Reproducing a similar mechanism in computational agents could thus

significantly enhance their navigation and planning skills.

Approaches for training models on these hierarchical abstractions have so far fallen into two key categories: hand-design methods (i.e. classical planning) and reinforcement learning techniques. Both these types of approaches, however, have significant limitations. Classical planning strategies are often sub-optimal, while reinforcement learning methods can be unstable, as well as expensive to develop and train.

In their study, the researchers at UC Berkeley and Facebook introduced an alternative paradigm that allows models to acquire hierarchical abstractions by analyzing passive first-person observation data (i.e. videos). These videos are labeled with agent actions, which can ultimately help a robot to navigate its environment.

"We use an inverse [model](#) trained on small amounts of interaction data to pseudo-label the passive first person videos with agent actions," the researchers explained in their paper. "Visuo-motor subroutines are acquired from these pseudo-labeled videos by learning a latent intent-conditioned policy that predicts the inferred pseudo-actions from the corresponding image observations."

The researchers evaluated their approach and demonstrated that it can significantly enhance an agent's navigation capabilities. In their tests, their method successfully enabled the acquisition of a variety of visuo-motor subroutines from passive first-person videos.

"We demonstrate the utility of our acquired visuo-motor subroutines by using them as is for exploration and as sub-policies in a hierarchical RL framework for reaching point goals and semantic goals," the researchers wrote. "We also demonstrate behavior of our subroutines in the real

world, by deploying them on a real robotic platform."

The approach proposed by the researchers achieved remarkable performance on all of the metrics assessed by the researchers. In addition, it was found to outperform state-of-the-art learning based techniques that were trained on substantially bigger interaction samples, generating trajectories that covered the environment more thoroughly.

Moreover, while the new approach acquired hierarchical abstractions from a total of 45,000 interactions with the environment, the state-of-the-art techniques it was compared to achieved less satisfactory results after up to 10 million interactions. The researchers' method also outperformed hand-crafted baselines that were specifically designed to navigate the environment while avoiding obstacles.

"Successful learning from first-person videos allowed the agent to execute coherent trajectories, even though it had only ever executed random actions," the researchers wrote. "It also successfully learned the bias towards forward actions in navigation and the notion of obstacle avoidance, leading to a high maximum distance and a low collision rate."

The study carried out by this team of researchers introduces a viable and highly effective alternative to current methods for training AI agents on navigation subroutines. In the future, their approach could inform the development of robots with more advanced planning and [navigation](#) skills.

More information: Ashish Kumar et al. Learning navigation subroutines by watching videos. arXiv:1905.12612 [cs.RO]. arxiv.org/abs/1905.12612

ashishkumar1993.github.io/subroutines/

© 2019 Science X Network

APA citation: Teaching AI agents navigation subroutines by feeding them videos (2019, June 12) retrieved 26 September 2020 from <https://techxplore.com/news/2019-06-ai-agents-subroutines-videos.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.