

AVID: a framework to enhance imitation learning in robots

January 3 2020, by Ingrid Fadelli



Sample stills of model-generated translations (top row are real human images, bottom row are fake robot images). Credit: Smith et al.


In recent years, research teams worldwide have been using reinforcement learning (RL) to teach robots how to complete a variety of tasks. Training these algorithms, however, can be very challenging, as it also requires substantial human efforts in properly defining the tasks that the robot is to complete.

One way to teach robots how to complete specific tasks is through humans demonstrations. While this may seem straightforward, it can be very difficult to implement, mainly because robots and humans have very different bodies, thus they are capable of different movements.

Researchers at the University of California Berkeley have recently developed a new framework that could help to overcome some of the challenges encountered when training robots via imitation learning (i.e., using human demonstrations). Their framework, called AVID, is based on two deep-learning models developed in previous research.

"When developing AVID, we built largely on two recent works, [CycleGAN](#) and [SOLAR](#), which introduced approaches to address fundamental limitations that have precluded learning from human videos in domain-shift and training on a physical robot from visual input, respectively," Laura Smith, one of the researchers who carried out the study, told TechXplore.

Instead of using techniques that do not take into account the differences between a robot and a human user's body, Smith and her colleagues used Cycle-GAN, a technique that can transform images on a pixel level. Using Cycle-GAN, their method converts human demonstrations of how to complete a given task into videos of a robot completing the same task. They then used these videos to develop a reward function for a RL algorithm.



Sample stills of model-generated translations (top row are real human images, bottom row are fake robot images). Credit: Smith et al.

"AVID works by having the robot observe a human perform some task then imagine what it'd look like for itself to perform the same thing," Smith explained. "To learn how to actually achieve this imagined success, we let the robot learn by trial and error."

Using the framework developed by Smith and her colleagues, a robot learns tasks one stage at a time, resetting each stage and trying it again without requiring a human user's intervention. The [learning process](#) thus becomes largely automated, with the robot learning new skills with minimal human intervention.

"A key advantage of our approach is that the human teacher can interact with the robot student while it's learning," Smith explained.

"Furthermore, we design our training framework to be amenable to learning long-term behavior with minimal effort."

The researchers evaluated their approach in a series of trials and found that it can effectively teach robots how to complete complex tasks, such

as operating a coffee machine, simply by processing 20 minutes of raw human demonstration videos and practicing the new skill for 180 minutes. In addition, AVID outperformed all other techniques to which it was, including imitation ablation, pixel-space ablation, and behavioral cloning approaches.

"What we found is that we can leverage CycleGAN to effectively make videos of human demonstrations comprehensible to the robot without requiring a tedious data collection process," Smith said. "We also show that exploiting the multi-stage nature of temporally extended tasks lets us learn robust behavior while making training easy. We see our work as a meaningful step toward bringing real-world deployment of autonomous robots within reach as it gives us a very natural, intuitive way for us to teach them."

The new learning framework introduced by Smith and her colleagues enables a different type of imitation learning, where a robot learns to complete one higher-level goal at a time, focusing on what it finds most challenging in every step. Moreover, instead of requiring human users to reset the scene after each practice trial, it allows robots to reset the scene automatically and continue practicing. In the future, AVID could enhance imitation learning processes, allowing developers to train robots faster and more effectively.

"One of the main limitations of our work so far is that we require data collection and training of the CycleGAN for each new scene the robot might encounter. We hope to be able to treat the CycleGAN training as a one-time, upfront cost such that [training](#) once on a large corpus of data can allow the robot to very quickly pick up a very wide variety of skills with a few demonstrations and a little practice."

More information: AVID: Learning multi-stage tasks via pixel-level translation of human videos.

arXiv:1912.04443 [cs.RO]. arxiv.org/abs/1912.04443

SOLAR: Deep structured representations for model-based reinforcement learning.

arXiv: 1808.09105 [cs.LG]. arxiv.org/abs/1808.09105

Unpaired image-to-image translation using Cycle- consistent adversarial networks. arXiv:1703.10593 [cs.CV]. arxiv.org/abs/1703.10593

© 2020 Science X Network

Citation: AVID: a framework to enhance imitation learning in robots (2020, January 3) retrieved 2 May 2024 from <https://techxplore.com/news/2020-01-avid-framework-imitation-robots.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.