

An imitation learning approach to train robots without the need for real human demonstrations

November 18 2019, by Ingrid Fadelli

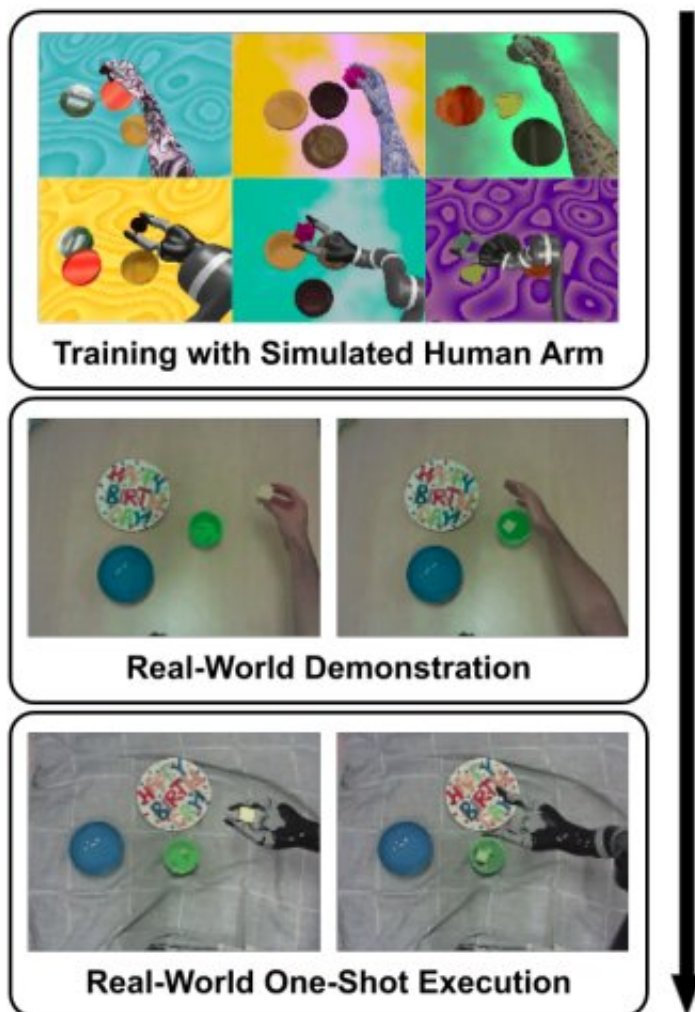


Figure explaining how the learning approach proposed by the researchers works.
Credit: Bonardi, James & Davison.

Most humans can learn how to complete a given task by observing another person perform it just once. Robots that are programmed to learn by imitating humans, however, typically need to be trained on a series of human demonstrations before they can effectively reproduce the desired behavior.

Researchers were recently able to teach robots to execute new tasks by having them observe a single human demonstration, using meta-learning approaches. However, these learning techniques typically require real-world data that can be expensive and difficult to collect.

To overcome this challenge, a team of researchers at Imperial College London has developed a new approach that enables one-shot imitation learning in robots without the need for real-world human demonstrations. Their approach, presented in [a paper pre-published on arXiv](#), uses algorithms known as task-embedded control networks (TecNets), which allow artificial agents to learn how to complete tasks from a single or multiple demonstrations, as well as artificially generated [training data](#).

"We show that with task-embedded control networks, we can infer control policies by embedding human demonstrations that can condition a control policy and achieve one-shot imitation learning," the researchers write in their paper.

The approach presented by the researchers does not require any interaction with real humans during the [robot's training](#). The method uses TecNets to infer control policies, embedding human demonstrations that can condition a given control policy and ultimately enable one-shot imitation learning.

To remove the need for real-world human demonstrations during training, the researchers used a dataset of videos simulating human demonstrations, which they generated using [PyRep](#), a recently released toolkit for robot learning research. Using PyRep, the researchers modeled a human-like 3-D arm and broke it down into shapes in order to reproduce movements that resemble those observed in humans.

They then created a dataset composed of videos in which this simulated arm performed a number of tasks and used it to train a robotic system. Ultimately, the robot was able to learn how to complete a task just by analyzing these simulation videos and a single human demonstration in the real-world.

"Importantly, we do not use a real human arm to supply demonstrations during training, but instead leverage domain randomization in an application that has not been seen before: sim-to-real transfer on humans," the researchers explain in their paper.

The team evaluated the new one-shot learning approach both in simulations and in the real-world, using it to train a robot to complete tasks that involved placing and pushing objects. Remarkably, their learning method achieved results comparable to those achieved using a more conventional imitation learning-based approach, even though it entails training a robot on artificially generated videos rather than real human demonstrations.

The researchers write, "We were able to achieve similar performance to a state-of-the-art alternative method that relies on thousands of training demonstrations collected in the real world, whilst also remaining robust to visual domain shifts, such as substantially different backgrounds."

The approach developed by this team of researchers could enable one-shot imitation learning for a number of robots without the need to collect

large quantities of real-world human demonstrations. This could save a lot of effort, resources and time for those trying to train robots using [imitation](#) learning. The researchers are now planning to investigate other actions that robots could be trained on using their approach.

"We hope to further investigate the variety of human actions that can be transferred from simulation to reality," the researchers wrote in their paper. "For example, in this work, we have shown that a human arm can be transferred, but would the same method work from demonstrations including the entire torso of a human?"

More information: Learning one-shot imitation from humans without humans. arXiv:1911.01103 [cs.RO]. arxiv.org/abs/1911.01103

© 2019 Science X Network

Citation: An imitation learning approach to train robots without the need for real human demonstrations (2019, November 18) retrieved 20 March 2024 from <https://techxplore.com/news/2019-11-imitation-approach-robots-real-human.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>
--