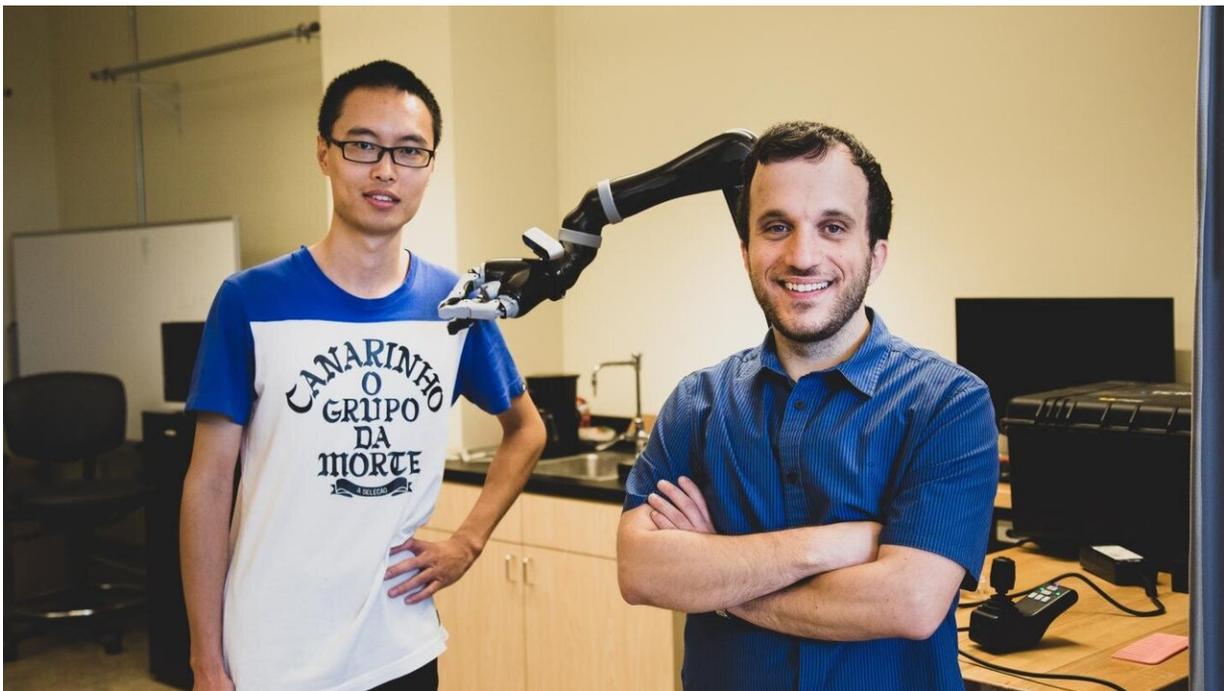


Showing robots 'tough love' helps them succeed, finds new study

November 6 2019



USC PhD student Jiali Duan (left) and Stefanos Nikolaidis, an assistant professor in computer science, use reinforcement learning, a technique in which artificial intelligence programs 'learn' from repeated experimentation. Credit: Haotian Mai.

According to a new study by USC computer scientists, to help a robot succeed, you might need to show it some tough love.

In a computer-simulated manipulation task, the researchers found that training a [robot](#) with a human adversary significantly improved its grasp of objects.

"This is the first [robot learning](#) effort using adversarial human users," said study co-author Stefanos Nikolaidis, an assistant professor of computer science.

"Picture it like playing a sport: if you're playing tennis with someone who always lets you win, you won't get better. Same with robots. If we want them to learn a manipulation task, such as grasping, so they can help people, we need to challenge them."

The study, "Robot Learning via Human Adversarial Games," was presented Nov. 4 at the International Conference on Intelligent Robots and Systems. USC Ph.D. students Jiali Duan and Qian Wang are lead authors, advised by Professor C. C. Jay Kuo, with additional co-author Lerrel Pinto from Carnegie Mellon University.

Learning from practice

Nikolaidis, who joined the USC Viterbi School of Engineering in 2018, and his team use reinforcement learning, a technique in which artificial intelligence programs "learn" from repeated experimentation.

Instead of being limited to completing a small range of repetitive tasks, such as industrial robots, the robotic system "learns" based on previous examples, in theory increasing the range of tasks it can perform.

But creating general-purpose robots is notoriously challenging, due in part to the amount of training required. Robotic systems need to see a huge number of examples to learn how to manipulate an [object](#) in a human-like manner.

For instance, OpenAI's impressive robotic system learned to solve a Rubik's cube with a humanoid hand, but required the equivalent of 10,000 years of simulated training to learn to manipulate the cube.

More importantly, the robot's dexterity is very specific. Without extensive training, it can't pick up an object, manipulate it with another grip, or grasp and handle a different object.

"As a human, even if I know the object's location, I don't know exactly how much it weighs or how it will move or behave when I pick it up, yet we do this successfully almost all of the time," said Nikolaidis.

"That's because people are very intuitive about how the world behaves, but the robot is like a newborn baby."

In other words, [robotic systems](#) find it hard to generalize, a skill which humans take for granted. This may seem trivial, but it can have serious consequences. If assistive robotic devices, such as grasping robots, are to fulfill their promise of helping people with disabilities, robotic systems must be able to operate reliably in real-world environments.

Human in the loop

One line of research that's been quite successful in overcoming this issue is having a "human in the loop." In other words, the human provides feedback to the robotic system by demonstrating the ability to complete the task.

But, until now, these algorithms have made a strong assumption of a cooperating human supervisor assisting the robot.

"I've always worked on human-robot collaboration, but in reality, people won't always be collaborators with robots in the wild," said Nikolaidis.

As an example, he points to a study by Japanese researchers, who set a robot loose in a public shopping complex and observed children "acting violently" towards it on several occasions.

So, thought Nikolaidis, what if we leveraged our human inclination to make things harder for the robot instead? Rather than showing it how to better grasp an object, what if we tried to pull it away? By adding challenge, the thinking goes, the system would learn to be more robust to real world complexity.

Element of challenge

The experiment went something like this: in a computer simulation, the robot attempts to grasp an object. The human, at the computer, observes the simulated robot's grasp. If the grasp is successful, the human tries to snatch the object from the robot's grasp, using the keyboard to signal direction.

Adding this element of challenge helps the robot learn the difference between a weak grasp (say, holding a bottle at the top), versus a firm grasp (holding it in the middle), which makes it much harder for the human adversary to snatch away.

It was a bit of a crazy idea, admits Nikolaidis, but it worked.

The researchers found the system trained with the human adversary rejected unstable grasps, and quickly learned robust grasps for these objects. In an experiment, the model achieved a 52 percent grasping success rate with a human adversary versus a 26.5 percent grasping success rate with a human collaborator.

"The robot learned not only how to grasp objects more robustly, but also to succeed more often in with new objects in a different orientation,

because it has learned a more stable grasp," said Nikolaidis.

They also found that the model trained with a human adversary performed better than a simulated adversary, which had 28 percent grasping success rate. So, robotic systems learn best from flesh-and-blood adversaries.

"That's because humans can understand stability and robustness better than learned adversaries," explained Nikolaidis.

"The robot tries to pick up stuff and, if the human tries to disrupt, it leads to more stable grasps. And because it has learned a more stable grasp, it will succeed more often, even if the object is in a different position. In other words, it's learned to generalize. That's a big deal."

Finding a balance

Nikolaidis hopes to have the system working on a real robot arm within a year. This will present a new challenge—in the real world, the slightest bit of friction or noise in a robot's joints can throw things off. But Nikolaidis is hopeful about the future of adversarial learning for robotics.

"I think we've just scratched the surface of potential applications of learning via adversarial human games," said Nikolaidis.

"We are excited to explore human-in-the-loop adversarial learning in other tasks as well, such as obstacle avoidance for robotic arms and mobile robots, such as self-driving cars."

This begs the question: how far are we willing to take adversarial learning? Would we be willing to kick and beat robots into submission? The answer, said Nikolaidis, lies in finding a balance of tough love and

encouragement with our robotics counterparts.

"I feel that tough love—in the context of the algorithm that we propose—is again like a sport: it falls within specific rules and constraints," said Nikolaidis.

"If the human just breaks the robot's gripper, the robot will continuously fail and never learn. In other words, the robot needs to be challenged but still be allowed to succeed in order to learn."

Provided by University of Southern California

Citation: Showing robots 'tough love' helps them succeed, finds new study (2019, November 6) retrieved 20 April 2024 from <https://techxplore.com/news/2019-11-robots-tough.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.