


Dog training methods help teach robots to learn new tricks

October 26 2020



"The question here was how do we get the robot to learn a skill?" said lead author Andrew Hundt, a PhD student working in Johns Hopkins' Computational Interaction and Robotics Laboratory. Credit: Will Kirk/Johns Hopkins University

With a training technique commonly used to teach dogs to sit and stay, Johns Hopkins University computer scientists showed a robot how to teach itself several new tricks, including stacking blocks. With the method, the robot, named Spot, was able to learn in days what typically takes a month.

By using [positive reinforcement](#), an approach familiar to anyone who's used treats to change a dog's behavior, the team dramatically improved the [robot](#)'s skills and did it quickly enough to make training robots for real-world work a more feasible enterprise. The findings are newly published in a paper called, "Good Robot!"

"The question here was how do we get the robot to learn a skill?" said lead author Andrew Hundt, a Ph.D. student working in Johns Hopkins' Computational Interaction and Robotics Laboratory. "I've had dogs so I know rewards work and that was the inspiration for how I designed the learning algorithm."

Unlike humans and animals that are born with highly intuitive brains, computers are blank slates and must learn everything from scratch. But true learning is often accomplished with trial and error, and roboticists are still figuring out how robots can learn efficiently from their mistakes.

The team accomplished that here by devising a reward system that works for a robot the way treats work for a dog. Where a dog might get a

cookie for a job well done, the robot earned numeric points.

Hundt recalled how he once taught his terrier mix puppy named Leah the command "leave it," so she could ignore squirrels on walks. He used two types of treats, ordinary trainer treats and something even better, like cheese. When Leah was excited and sniffing around the treats, she got nothing. But when she calmed down and looked away, she got the good stuff. "That's when I gave her the cheese and said, 'Leave it! Good Leah!'"

Similarly, to stack blocks, Spot the robot needed to learn how to focus on constructive actions. As the robot explored the blocks, it quickly learned that correct behaviors for stacking earned high points, but incorrect ones earned nothing. Reach out but don't grasp a block? No points. Knock over a stack? Definitely no points. Spot earned the most by placing the last block on top of a four-block stack.

The training tactic not only worked, it took just days to teach the robot what used to take weeks. The team was able to reduce the practice time by first training a simulated robot, which is a lot like a video game, then running tests with Spot.

"The robot wants the higher score," Hundt said. "It quickly learns the right behavior to get the best reward. In fact, it used to take a month of practice for the robot to achieve 100% accuracy. We were able to do it in two days."

Positive reinforcement not only worked to help the robot teach itself to stack blocks, with the point system the robot just as quickly learned several other tasks—even how to play a simulated navigation game. The ability to learn from mistakes in all types of situations is critical for designing a robot that could adapt to new environments.

"At the start the robot has no idea what it's doing but it will get better and better with each practice. It never gives up and keeps trying to stack and is able to finish the task 100% of the time," Hundt said.

The team imagines these findings could help train household robots to do laundry and wash dishes—tasks that could be popular on the open market and help seniors live independently. It could also help design improved self-driving cars.

"Our goal is to eventually develop robots that can do [complex tasks](#) in the real world—like product assembly, caring for the elderly and surgery," Hager said. "We don't currently know how to program tasks like that—the world is too complex. But work like this shows us that there is promise to the idea that robots can learn how to accomplish such [real-world](#) tasks in a safe and efficient way."

More information: Andrew Hundt et al, "Good Robot!": Efficient Reinforcement Learning for Multi-Step Visual Tasks with Sim to Real Transfer, *IEEE Robotics and Automation Letters* (2020). [DOI: 10.1109/LRA.2020.3015448](#)

Provided by Johns Hopkins University

Citation: Dog training methods help teach robots to learn new tricks (2020, October 26) retrieved 23 April 2024 from <https://techxplore.com/news/2020-10-dog-methods-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.
