

Machine learning guarantees robots' performance in unknown territory

November 17 2020, by Molly Sharlach



Princeton researchers tested a new machine learning approach for guaranteeing robots' safety and success in unfamiliar settings. Experiments included programming a small drone called a Parrot Swing to avoid obstacles while flying down a 60-foot-long corridor. Credit: Video by the Intelligent Robot Motion Lab; GIF by Josh Cartagena

A small drone takes a test flight through a space filled with randomly



placed cardboard cylinders acting as stand-ins for trees, people or structures. The algorithm controlling the drone has been trained on a thousand simulated obstacle-laden courses, but it's never seen one like this. Still, nine times out of 10, the pint-sized plane dodges all the obstacles in its path.

This experiment is a proving ground for a pivotal challenge in modern robotics: the ability to guarantee the safety and success of automated robots operating in novel environments. As engineers increasingly turn to machine learning methods to develop adaptable robots, new work by Princeton University researchers makes progress on such guarantees for robots in contexts with diverse types of obstacles and constraints.

"Over the last decade or so, there's been a tremendous amount of excitement and progress around machine learning in the context of robotics, primarily because it allows you to handle rich sensory inputs," like those from a robot's camera, and map these complex inputs to actions, said Anirudha Majumdar, an assistant professor of mechanical and aerospace engineering at Princeton.

However, robot control algorithms based on machine learning run the risk of overfitting to their training data, which can make algorithms less effective when they encounter inputs that differ from those they were trained on. Majumdar's Intelligent Robot Motion Lab addressed this challenge by expanding the suite of available tools for training robot control policies, and quantifying the likely success and safety of robots performing in novel environments.

In three new papers, the researchers adapted machine learning frameworks from other arenas to the field of robot locomotion and manipulation. They turned to generalization theory, which is typically used in contexts that map a single input onto a single output, such as automated image tagging. The new methods are among the first to apply



generalization theory to the more complex task of making guarantees on robots' performance in unfamiliar settings. While other approaches have provided such guarantees under more restrictive assumptions, the team's methods offer more broadly applicable guarantees on performance in novel environments, said Majumdar.

In the first paper, a proof of principle for applying the machine learning frameworks, the team tested their approach in simulations that included a wheeled vehicle driving through a space filled with obstacles, and a robotic arm grasping objects on a table. They also validated the technique by assessing the obstacle avoidance of a small drone called a Parrot Swing (a combination quadcopter and fixed-wing airplane) as it flew down a 60-foot-long corridor dotted with cardboard cylinders. The guaranteed success rate of the drone's control policy was 88.4%, and it avoided obstacles in 18 of 20 trials (90%).

The work, published Oct. 3 in the *International Journal of Robotics Research*, was coauthored by Majumdar; Alec Farid, a graduate student in mechanical and aerospace engineering; and Anoopkumar Sonar, a computer science concentrator from Princeton's Class of 2021.

When applying machine learning techniques from other areas to robotics, said Farid, "there are a lot of special assumptions you need to satisfy, and one of them is saying how similar the environments you're expecting to see are to the environments your policy was trained on. In addition to showing that we can do this in the robotic setting, we also focused on trying to expand the types of environments that we could provide a guarantee for."

"The kinds of guarantees we're able to give range from about 80% to 95% success rates on new environments, depending on the specific task, but if you're deploying [an unmanned aerial vehicle] in a real environment, then 95% probably isn't good enough," said Majumdar. "I



see that as one of the biggest challenges, and one that we are actively working on."

Still, the team's approaches represent much-needed progress on generalization guarantees for robots operating in unseen environments, said Hongkai Dai, a senior research scientist at the Toyota Research Institute in Los Altos, California.

"These guarantees are paramount to many safety-critical applications, such as self-driving cars and autonomous drones, where the training set cannot cover every possible scenario," said Dai, who was not involved in the research. "The guarantee tells us how likely it is that a policy can still perform reasonably well on unseen cases, and hence establishes confidence on the policy, where the stake of failure is too high."

In two other papers, to be presented Nov. 18 at the virtual Conference on Robot Learning, the researchers examined additional refinements to bring robot control policies closer to the guarantees that would be needed for real-world deployment. One paper used imitation learning, in which a human "expert" provides training data by manually guiding a simulated robot to pick up various objects or move through different spaces with obstacles. This approach can improve the success of machine learningbased control policies.

To provide the <u>training data</u>, lead author Allen Ren, a graduate student in mechanical and aerospace engineering, used a 3-D computer mouse to control a simulated robotic arm tasked with grasping and lifting drinking mugs of various sizes, shapes and materials. Other imitation learning experiments involved the arm pushing a box across a table, and a simulation of a wheeled robot navigating around furniture in a home-like environment.

The researchers deployed the policies learned from the mug-grasping



and box-pushing tasks on a robotic arm in the laboratory, which was able to pick up 25 different mugs by grasping their rims between its two finger-like grippers—not holding the handle as a human would. In the box-pushing example, the policy achieved 93% success on easier tasks and 80% on harder tasks.

"We have a camera on top of the table that sees the <u>environment</u> and takes a picture five times per second," said Ren. "Our policy training simulation takes this image and outputs what kind of action the robot should take, and then we have a controller that moves the arm to the desired locations based on the output of the model."

A third paper demonstrated the development of vision-based planners that provide guarantees for flying or walking robots to carry out planned sequences of movements through diverse environments. Generating control policies for planned movements brought a new problem of scale—a need to optimize vision-based policies with thousands, rather than hundreds, of dimensions.

"That required coming up with some new algorithmic tools for being able to tackle that dimensionality and still be able to give strong generalization guarantees," said lead author Sushant Veer, a postdoctoral research associate in mechanical and aerospace engineering.

A key aspect of Veer's strategy was the use of motion primitives, in which a policy directs a robot to go straight or turn, for example, rather than specifying a torque or velocity for each movement. Narrowing the space of possible actions makes the planning process more computationally tractable, said Majumdar.

Veer and Majumdar evaluated the vision-based planners on simulations of a drone navigating around obstacles and a four-legged robot traversing rough terrain with slopes as high as 35 degrees—"a very challenging



problem that a lot of people in robotics are still trying to solve," said Veer.

In the study, the legged <u>robot</u> achieved an 80% success rate on unseen test environments. The researchers are working to further improve their policies' guarantees, as well as assessing the policies' performance on real robots in the laboratory.

More information: Anirudha Majumdar et al, PAC-Bayes control: learning policies that provably generalize to novel environments, *The International Journal of Robotics Research* (2020). DOI: <u>10.1177/0278364920959444</u>

Provided by Princeton University

Citation: Machine learning guarantees robots' performance in unknown territory (2020, November 17) retrieved 2 May 2024 from <u>https://techxplore.com/news/2020-11-machine-robots-unknown-territory.html</u>

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.