


A new algorithm to help robots practice skills independently to adapt to unfamiliar environments

August 8 2024, by Alex Shipp



A new algorithm developed by researchers at MIT CSAIL helps robots practice skills on their own. In experiments, it guided a quadruped with sweeping and placing various items. Credits: Alex Shipp/MIT CSAIL

The phrase "practice makes perfect" is usually reserved for humans, but it's also a great maxim for robots newly deployed in unfamiliar environments.

Picture a [robot](#) arriving in a warehouse. It comes packaged with the skills it was trained on, like placing an object, and now it needs to pick items from a shelf it's not familiar with. At first, the machine struggles with this, since it needs to get acquainted with its new surroundings. To improve, the robot will need to understand which skills within an overall task it needs improvement on, then specialize (or parameterize) that action.

A human onsite could program the robot to optimize its performance, but researchers from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) and The AI Institute have developed a more effective alternative. Presented at the [Robotics: Science and Systems Conference](#) last month, their "Estimate, Extrapolate, and Situate" (EES) algorithm enables these machines to practice on their own, potentially helping them improve at useful tasks in factories, households, and hospitals.

The study is [published](#) on the *arXiv* preprint server.

Sizing up the situation

To help robots get better at activities like sweeping floors, EES works with a vision system that locates and tracks the machine's surroundings. Then, the algorithm estimates how reliably the robot executes an action (like sweeping) and whether it would be worthwhile to practice more. EES forecasts how well the robot could perform the overall task if it refines that particular skill, and finally, it practices. The vision system subsequently checks whether that skill was done correctly after each attempt.

EES could come in handy in places like a hospital, factory, house, or coffee shop. For example, if you wanted a robot to clean up your living room, it would need help practicing skills like sweeping. According to Nishanth Kumar SM '24 and his colleagues, though, EES could help that robot improve without [human intervention](#), using only a few practice trials.

"Going into this project, we wondered if this specialization would be possible in a reasonable amount of samples on a real robot," says Kumar, co-lead author of a paper describing the work, Ph.D. student in electrical engineering and computer science, and a CSAIL affiliate.

"Now, we have an algorithm that enables robots to get meaningfully better at specific skills in a reasonable amount of time with tens or hundreds of [data points](#), an upgrade from the thousands or millions of samples that a standard reinforcement learning algorithm requires."

See Spot sweep

EES's knack for efficient learning was evident when implemented on Boston Dynamics' Spot quadruped during research trials at The AI Institute. The robot, which has an arm attached to its back, completed manipulation tasks after practicing for a few hours. In one demonstration, the robot learned how to securely place a ball and ring on a slanted table in roughly three hours.

In another, the algorithm guided the machine to improve at sweeping toys into a bin within about two hours. Both results appear to be an upgrade from previous frameworks, which would have likely taken more than 10 hours per task.

"We aimed to have the robot collect its own experience so it can better choose which strategies will work well in its deployment," says co-lead

author Tom Silver SM '20, Ph.D. '24, an [electrical engineering](#) and computer science (EECS) alumnus and CSAIL affiliate who is now an assistant professor at Princeton University.

"By focusing on what the robot knows, we sought to answer a key question: In the library of skills that the robot has, which is the one that would be most useful to practice right now?"

EES could eventually help streamline autonomous practice for robots in new deployment environments, but for now, it comes with a few limitations. For starters, they used tables that were low to the ground, which made it easier for the robot to see its objects.

Kumar and Silver also 3D printed an attachable handle that made the brush easier for Spot to grab. The robot didn't detect some items and identified objects in the wrong places, so the researchers counted those errors as failures.

Giving robots homework

The researchers note that the practice speeds from the physical experiments could be accelerated further with the help of a simulator. Instead of physically working at each skill autonomously, the robot could eventually combine real and virtual practice. They hope to make their system faster with less latency, engineering EES to overcome the imaging delays the researchers experienced. In the future, they may investigate an algorithm that reasons over sequences of practice attempts instead of planning which skills to refine.

"Enabling robots to learn on their own is both incredibly useful and extremely challenging," says Danfei Xu, an assistant professor in the School of Interactive Computing at Georgia Tech and a research scientist at NVIDIA AI, who was not involved with this work.

"In the future, home robots will be sold to all sorts of households and expected to perform a wide range of tasks. We can't possibly program everything they need to know beforehand, so it's essential that they can learn on the job. However, letting robots loose to explore and learn without guidance can be very slow and might lead to unintended consequences.

"The research by Silver and his colleagues introduces an algorithm that allows robots to practice their skills autonomously in a structured way. This is a big step towards creating home robots that can continuously evolve and improve on their own."

Silver and Kumar's co-authors are The AI Institute researchers Stephen Proulx and Jennifer Barry, plus four CSAIL members: Northeastern University Ph.D. student and visiting researcher Linfeng Zhao, MIT EECS Ph.D. student Willie McClinton, and MIT EECS professors Leslie Pack Kaelbling and Tomás Lozano-Pérez. Their work was supported, in part, by The AI Institute, the U.S. National Science Foundation, the U.S. Air Force Office of Scientific Research, the U.S. Office of Naval Research, the U.S. Army Research Office, and MIT Quest for Intelligence, with high-performance computing resources from the MIT SuperCloud and Lincoln Laboratory Supercomputing Center.

More information: Nishanth Kumar et al, Practice Makes Perfect: Planning to Learn Skill Parameter Policies, *arXiv* (2024). [DOI: 10.48550/arxiv.2402.15025](https://doi.org/10.48550/arxiv.2402.15025)

Provided by Massachusetts Institute of Technology

Citation: A new algorithm to help robots practice skills independently to adapt to unfamiliar environments (2024, August 8) retrieved 8 August 2024 from

<https://techxplore.com/news/2024-08-algorithm-robots-skills-independently-unfamiliar.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.