


# Researchers introduce novel approach allowing robots to be trained in simulations of scanned home environments

August 1 2024, by Rachel Gordon

---



With the help of “digital twins,” RialTo, developed by MIT researchers Marcel Torne Villasevil (left) and Pulkit Agrawal, can guide a robot to practice in a simulated environment much faster than it would if it honed its skills in the real world. Credit: Mike Grimmett/MIT CSAIL

At the top of many automation wish lists is a particularly time-consuming task: chores.

The moonshot of many roboticists is cooking up the proper hardware and software combination so that a machine can learn "generalist" policies (the rules and strategies that guide robot behavior) that work everywhere, under all conditions.

Realistically, though, if you have a home robot, you probably don't care much about it working for your neighbors. MIT Computer Science and Artificial Intelligence Laboratory (CSAIL) researchers decided, with that in mind, to attempt to find a solution to easily train robust robot policies for very specific environments.

"We aim for robots to perform exceptionally well under disturbances, distractions, varying lighting conditions, and changes in object poses, all within a single environment," says Marcel Torne Villasevil, MIT CSAIL research assistant in the Improbable AI lab and lead author on a [paper](#) about the work, which appears on the preprint server *arXiv*.

"We propose a method to create digital twins on the fly using the latest advances in computer vision. With just their phones, anyone can capture a digital replica of the real world, and the robots can train in a simulated environment much faster than the real world, thanks to GPU parallelization. Our approach eliminates the need for extensive reward engineering by leveraging a few real-world demonstrations to jump-start the training process."

## **Taking your robot home**

RialTo, of course, is a little more complicated than just a simple wave of a phone and (boom!) home bot at your service. It begins by using your device to scan the target environment using tools like NeRFStudio,

ARCode, or Polycam. Once the scene is reconstructed, users can upload it to RialTo's interface to make detailed adjustments, add necessary joints to the robots, and more.

The refined scene is exported and brought into the simulator. Here, the aim is to develop a policy based on real-world actions and observations, such as one for grabbing a cup on a counter. These real-world demonstrations are replicated in the simulation, providing some valuable data for reinforcement learning.

"This helps in creating a strong policy that works well in both the simulation and the real world. An enhanced algorithm using reinforcement learning helps guide this process, to ensure the policy is effective when applied outside of the simulator," says Torne.

Testing showed that RialTo created strong policies for a variety of tasks, whether in controlled lab settings or more unpredictable real-world environments, improving 67% over imitation learning with the same number of demonstrations. The tasks involved opening a toaster, placing a book on a shelf, putting a plate on a rack, placing a mug on a shelf, opening a drawer, and opening a cabinet.

For each task, the researchers tested the system's performance under three increasing levels of difficulty: randomizing object poses, adding visual distractors, and applying physical disturbances during task executions. When paired with real-world data, the system outperformed traditional imitation-learning methods, especially in situations with lots of visual distractions or physical disruptions.

"These experiments show that if we care about being very robust to one particular environment, the best idea is to leverage digital twins instead of trying to obtain robustness with large-scale data collection in diverse environments," says Pulkit Agrawal, director of Improbable AI Lab,

MIT electrical engineering and computer science (EECS) associate professor, MIT CSAIL principal investigator, and senior author on the work.

As far as limitations, RialTo currently takes three days to be fully trained. To speed this up, the team mentions improving the underlying algorithms and using foundation models. Training in simulation also has its limitations, and currently it's difficult to do effortless sim-to-real transfer and simulate deformable objects or liquids.

## **The next level**

So what's next for RialTo's journey? Building on previous efforts, the scientists are working on preserving robustness against various disturbances while improving the model's adaptability to new environments.

"Our next endeavor is this approach to using pre-trained models, accelerating the learning process, minimizing human input, and achieving broader generalization capabilities," says Torne.

"We're incredibly enthusiastic about our 'on-the-fly' robot programming concept, where robots can autonomously scan their environment and learn how to solve specific tasks in simulation. While our current method has limitations—such as requiring a few initial demonstrations by a human and significant compute time for training these policies (up to three days)—we see it as a significant step towards achieving 'on-the-fly' robot learning and deployment," says Torne.

"This approach moves us closer to a future where robots won't need a preexisting policy that covers every scenario. Instead, they can rapidly learn new tasks without extensive real-world interaction. In my view, this advancement could expedite the practical application of robotics far

sooner than relying solely on a universal, all-encompassing policy."

"To deploy robots in the real world, researchers have traditionally relied on methods such as imitation learning from expert data, which can be expensive, or [reinforcement learning](#), which can be unsafe," says Zoey Chen, a computer science Ph.D. student at the University of Washington who wasn't involved in the paper.

"RialTo directly addresses both the safety constraints of real-world RL [robot learning], and efficient data constraints for data-driven learning methods, with its novel real-to-sim-to-real pipeline.

"This novel pipeline not only ensures safe and robust training in simulation before real-world deployment, but also significantly improves the efficiency of data collection. RialTo has the potential to significantly scale up [robot](#) learning and allows robots to adapt to complex real-world scenarios much more effectively."

"Simulation has shown impressive capabilities on real robots by providing inexpensive, possibly infinite data for policy learning," adds Marius Memmel, a computer science Ph.D. student at the University of Washington who wasn't involved in the work.

"However, these methods are limited to a few specific scenarios, and constructing the corresponding simulations is expensive and laborious. RialTo provides an easy-to-use tool to reconstruct real-world environments in minutes instead of hours.

"Furthermore, it makes extensive use of collected demonstrations during policy learning, minimizing the burden on the operator and reducing the sim2real gap. RialTo demonstrates robustness to object poses and disturbances, showing incredible real-world performance without requiring extensive simulator construction and data collection."

**More information:** Marcel Torne et al, Reconciling Reality through Simulation: A Real-to-Sim-to-Real Approach for Robust Manipulation, *arXiv* (2024). [DOI: 10.48550/arxiv.2403.03949](https://doi.org/10.48550/arxiv.2403.03949)

*This story is republished courtesy of MIT News ([web.mit.edu/newsoffice/](http://web.mit.edu/newsoffice/)), a popular site that covers news about MIT research, innovation and teaching.*

Provided by Massachusetts Institute of Technology

Citation: Researchers introduce novel approach allowing robots to be trained in simulations of scanned home environments (2024, August 1) retrieved 1 August 2024 from <https://techxplore.com/news/2024-08-approach-robots-simulations-scanned-home.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.