

# **Once is enough: Helping robots learn quickly in new environments**

December 13 2023, by Greg Hardesty



Resultant Language Conditioned Policy generated using

"Robot pushing red button"



Resultant Video Conditioned Policy generated using expert video demonstration

Using only one video or textual demonstration of a task, roboclip performed two to three times better than other imitation learning (il) methods. Credit: Sontakke et al

Alone at home, your bones creaky due to old age, you crave a cool beverage. You turn to your robot and say, "Please get me a tall glass of



water from the refrigerator." Your AI-trained companion obliges. Soon, your thirst is quenched.

While this scenario still is a decade or more away in terms of a seamless <u>real-world</u> application, a new research paper led by USC computer science student Sumedh A. Sontakke, with his advisors Assistant Professor Erdem B1y1k and Professor Laurent Itti, opens the door wider to this potential reality with a new online algorithm they created called RoboCLIP.

Aging populations and caregivers stand to benefit the most from future work based on RoboCLIP, which dramatically reduces how much data is needed to train robots by allowing anyone to interact with them through language or videos—at least, for now, in <u>computer simulations</u>.

"To me, the most impressive thing about RoboCLIP is being able to make our robots do something based on only one video demonstration or one language description," says Biyik, a roboticist who joined USC Viterbi's Thomas Lord Department of Computer Science in August 2023 and leads the Learning and Interactive Robot Autonomy Lab (Lira Lab).

#### Learning quickly with few demonstrations

The <u>paper</u>, titled "RoboCLIP: One Demonstration is Enough to Learn Robot Policies," is published on the *arXiv* preprint server and will be presented by Sontakke at the <u>37th Conference on Neural Information</u> <u>Processing Systems (NeurIPS)</u>, Dec. 10-16 in New Orleans.

"The large amount of data currently required to get a robot to successfully do the task you want it to do is not feasible in the real world, where you want robots that can learn quickly with few demonstrations," Sontakke explains.



To get around this notoriously <u>difficult problem</u> in <u>reinforcement</u> <u>learning</u>—a subset of AI in which a machine learns by trial and error how to behave to get the best reward—the researchers tested RoboCLIP.

The result?

Using only one video or textual demonstration of a task, RoboCLIP performed two to three times better than other imitation learning (IL) methods.

Future research is needed before this study translates into a world where robots can learn quickly with few demonstrations or instructions—such as fetching you a tall glass of chilled water—but RoboCLIP represents a significant step forward in IL research, Sontakke and Biyik said.

Right now, IL methods require many demonstrations, massive datasets, and substantial human supervision for a robot to master a task in computer simulations.

Now it can learn from just one, the RoboCLIP research shows.

## Performing well 'out of the box'

RoboCLIP was inspired by advances in the field of generative AI and video-language models (VLMs), which are pretrained on large amounts of video and textual demonstrations, Sontakke and Biyik explained. The new algorithm harnesses the power of these VLM embeddings to train robots.

A handful of experimental videos on the <u>RoboCLIP website</u> show the method's effectiveness.

In the videos, a robot—in computer simulations—pushes a red button,



closes a black box, and closes a green drawer after being instructed with a single video demonstration or a textual description (for example, "Robot pushing red button").

"Out of the box," Biyik says, "RoboCLIP has performed well."

#### Two years in the making

Sontakke said the genesis of the research paper dates back two years ago.

"I started thinking about household tasks like opening doors and cabinets," he said. "I didn't like how much data I needed to collect before I could get the robot to successfully do the task I cared about. I wanted to avoid that, and that's where this project came from."

Collaborating with Sontakke, Biyik and Itti on the RoboCLIP paper were two USC Viterbi graduates, Sebastien M.R. Arnold, now at Google Research, and Karl Pertsch, now at UC Berkeley and Stanford University. Jesse Zhang, a fourth-year Ph.D. candidate in computer sciences at USC Viterbi, also worked on the RoboCLIP project.

## 'Key innovation'

"The key innovation here is using the VLM to critically 'observe' simulations of the virtual robot babbling around while trying to perform the task, until at some point it starts getting it right—at that point, the VLM will recognize that progress and reward the virtual robot to keep trying in this direction," Itti explained.

"The VLM can recognize that the virtual robot is getting closer to success when the textual description produced by the VLM observing the



robot motions becomes closer to what the user wants," Itti added. "This new kind of closed-loop interaction is very exciting to me and will likely have many more future applications in other domains."

Besides the <u>aging population</u> who will rely on robots to improve their daily lives, RoboCLIP could lead to applications that could help anyone.

Think of those DIY videos you look up on YouTube to figure out how to fix a busted garbage disposal or malfunctioning microwave.

Could you simply, in the future, ask your <u>robot</u> assistant to perform such tasks while you slumber on the couch?

The possibilities are intriguing, Biyik and Sontakke said.

**More information:** A Sontakke et al, RoboCLIP: One Demonstration is Enough to Learn Robot Policies, *arXiv* (2023). <u>DOI:</u> <u>10.48550/arxiv.2310.07899</u>

#### Provided by University of Southern California

Citation: Once is enough: Helping robots learn quickly in new environments (2023, December 13) retrieved 8 May 2024 from <u>https://techxplore.com/news/2023-12-robots-quickly-environments.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.