

AI improves control of robot arms

October 26 2020, by Katharine Miller



Advances in robotic automation allow for more precise, intuitive movement.
Credit: Stanford University

More than one million American adults use wheelchairs fitted with robot arms to help them perform everyday tasks such as dressing, brushing their teeth, and eating. But the robotic devices now on the market can be hard to control. Removing a food container from a refrigerator or

opening a cabinet door can take a long time. And using a robot to feed yourself is even harder because the task requires fine manipulation.

A team of Stanford researchers has now developed a novel way to control assistive robotic arms that is both more intuitive and faster than existing approaches. In experiments, their [robot](#) controller allowed subjects to more efficiently cut tofu and shovel it onto a plate, or stab a marshmallow, scoop it in icing, and dip it in sprinkles.

"Feeding is one of my favorite problems to work on because it is difficult from a robotics standpoint; it requires precise manipulation and it is such a fundamental task—you need to feed yourself every day," says Dorsa Sadigh, assistant professor of computer science and of electrical engineering at Stanford and an affiliated faculty member of the Stanford Institute for Human-Centered Artificial Intelligence (HAI). It's also exciting: "You can see the benefits of it right in front of your eyes," she says.

Sadigh's team, which included engineering graduate student Hong Jun Jeon and computer science postdoctoral scholar Dylan P. Losey, developed a controller that blends two [artificial intelligence algorithms](#). The first, which was developed by Sadigh's group, enables control in two dimensions on a joystick without the need to switch between modes. It uses contextual cues to determine whether a user is reaching for a doorknob or a drinking cup, for example. Then, as the [robot arm](#) nears its destination, the second algorithm kicks in to allow more precise movements, with control shared between the human and the robot.

Reducing Dimensionality

Typical assistive robots now on the market have 6-7 joints. To control each of them a user switches between various modes on the joystick, which is unintuitive, mentally tiring, and takes a lot of time, Sadigh says.

"These robots are out there in the wild, but it's still really challenging to use them."

The team wondered: Can a joystick that gives commands in only two directions (up/down; left/right) nevertheless control a multi-jointed robot smoothly and quickly? For an answer, they turned to a process called dimensionality reduction. In any given context, a robot arm doesn't actually have to move every joint in every possible direction to accomplish a particular task. A smaller set of movements typically suffices. "The key insight is that, conditioned on certain limitations, such as context, the robot will know that pushing right on the joystick means a specific thing, such as picking up a cup," Sadigh explains. "Without me telling it, the robot will figure out the most important thing to pay attention to, given the context."

The process of dimensionality reduction begins with a human (i.e., a graduate student) moving the robot arm through various task-specific motions, essentially training it how to move in a more fluid and useful way in a given context. This high-dimensional dataset is then fed through a neural network (an autoencoder) that first compresses the data into two dimensions and then decodes that compressed representation to try to recreate the initial expert data. "That is how you make sure the compression works—because it is able to reproduce the expert data," Sadigh says.

The next step is where the magic happens: A person gives two-dimensional instructions on a joystick and the robot is able to recreate the more complex, context-dependent actions that the expert trained it to do. In experiments, when users controlled the robot with this "latent action" algorithm alone, they could pick up an egg, an apple, and a cup of flour and drop them in a bowl (making an "apple pie," so to speak) faster than an existing approach that required mode shifting on a joystick. Despite the increased speed, users found the interface

unpredictable. "It does the right thing, but users are not sure why," Sadigh says.

Adding Shared Autonomy

The latent action controller also wasn't very precise. To address that problem, the team blended the latent action algorithm with one called shared autonomy. Here, the novelty lay in the way the team integrated the two algorithms. "It's not an add-on," Sadigh says. "The system is trained all together."

In shared autonomy, the robot begins with a set of "beliefs" about what the controller is telling it to do and gains confidence about the goal as additional instructions are given. Since robots aren't actually sentient, these beliefs are really just probabilities. For example, faced with two cups of water, a robot might begin with a belief that there's an even chance it should pick up either one. But as the joystick directs it toward one cup and away from the other, the robot gains confidence about the goal and can begin to take over—sharing autonomy with the user to more precisely control the robot arm. The amount of control the robot takes on is probabilistic as well: If the robot has 80 percent confidence that it's going to cup A rather than cup B, it will take 80 percent of the control while the human still has 20 percent, Sadigh explains.

To test the integrated algorithms, the team conducted experiments in which users controlled a robot arm fitted with a fork. Their tasks (shown in this video): Cut and scoop tofu, or stab a marshmallow, scoop it in icing, and dip it in sprinkles. The result: The controller that used the combined algorithm (latent action with shared autonomy) was both faster and easier for users to control than the latent action algorithm alone or the standard controller either alone or with shared autonomy.

There's still a lot of work to do before the team's algorithms will impact

the lives of people living with disabilities, Jeon says. The system will need to be trained to use computer vision and to function in numerous contexts. And eventually there should be a study in which a large sample of people with disabilities have an opportunity to test the controller. In the long run, the hope is that AI-based assistive robotics will make the lives of disabled people easier. "It's empowering," Jeon says. "It gives people more agency over what they can do."

More information: Shared Autonomy with Learned Latent Actions. *Robotics: Science and Systems*.

www.roboticsproceedings.org/rss16/p011.pdf

Provided by Stanford University

Citation: AI improves control of robot arms (2020, October 26) retrieved 20 March 2024 from <https://techxplore.com/news/2020-10-ai-robot-arms.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.
