

A better way to control shape-shifting soft robots

May 9 2024



Reconfigurable soft robots. In this paper, we address challenges in controlling reconfigurable—or shape-shifting—robots, who can change their morphology to accomplish desired tasks. The figure illustrates a task where a circular robot needs to alter its body shape to fit within a confined chamber to manipulate the square cargo to the target point (right). We introduce a benchmark with 8 tasks for shape-shifting robots, which we dub "DittoGym" inspired by a shape-shifting; Pokémon (left). Credit: *arXiv* (2024). DOI: 10.48550/arxiv.2401.13231

Imagine a slime-like robot that can seamlessly change its shape to squeeze through narrow spaces, which could be deployed inside the human body to remove an unwanted item.

While such a [robot](#) does not yet exist outside a laboratory, researchers are working to develop reconfigurable soft robots for applications in health care, wearable devices, and industrial systems.

But how can one control a squishy robot that doesn't have joints, limbs, or fingers that can be manipulated, and instead can drastically alter its

entire shape at will? MIT researchers are working to answer that question.

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

They developed a [control algorithm](#) that can autonomously learn how to move, stretch, and shape a reconfigurable robot to complete a [specific task](#), even when that task requires the robot to change its morphology multiple times. The team also built a simulator to test control algorithms for deformable soft robots on a series of challenging, shape-changing tasks.

Their method completed each of the eight tasks they evaluated while outperforming other algorithms. The technique worked especially well on multifaceted tasks. For instance, in one test, the robot had to reduce its height while growing two tiny legs to squeeze through a narrow pipe, and then un-grow those legs and extend its torso to open the pipe's lid.

While reconfigurable soft robots are still in their infancy, such a technique could someday enable general-purpose robots that can adapt their shapes to accomplish diverse tasks.

"When people think about soft robots, they tend to think about robots that are elastic, but return to their original shape. Our robot is like slime and can actually change its morphology. It is very striking that our method worked so well because we are dealing with something very new," says Boyuan Chen, an electrical engineering and computer science (EECS) graduate student and co-author of a paper on this approach.

Chen's co-authors include lead author Suning Huang, an undergraduate student at Tsinghua University in China who completed this work while a visiting student at MIT; Huazhe Xu, an assistant professor at Tsinghua University; and senior author Vincent Sitzmann, an assistant professor of

EECS at MIT who leads the Scene Representation Group in the Computer Science and Artificial Intelligence Laboratory. The research will be presented at the International Conference on Learning Representations.

Controlling dynamic motion

Scientists often teach robots to complete tasks using a [machine-learning approach](#) known as reinforcement learning, which is a trial-and-error process in which the robot is rewarded for actions that move it closer to a goal.

This can be effective when the robot's moving parts are consistent and well-defined, like a gripper with three fingers. With a robotic gripper, a reinforcement learning algorithm might move one finger slightly, learning by trial and error whether that motion earns it a reward. Then it would move on to the next finger, and so on.

But shape-shifting robots, which are controlled by magnetic fields, can dynamically squish, bend, or elongate their entire bodies.

"Such a robot could have thousands of small pieces of muscle to control, so it is very hard to learn in a traditional way," says Chen.

To solve this problem, he and his collaborators had to think about it differently. Rather than moving each tiny muscle individually, their reinforcement learning algorithm begins by learning to control groups of adjacent muscles that work together.

Then, after the algorithm has explored the space of possible actions by focusing on groups of muscles, it drills down into finer detail to optimize the policy, or action plan, it has learned. In this way, the control algorithm follows a coarse-to-fine methodology.

"Coarse-to-fine means that when you take a random action, that random action is likely to make a difference. The change in the outcome is likely very significant because you coarsely control several muscles at the same time," Sitzmann says.

To enable this, the researchers treat a robot's action space, or how it can move in a certain area, like an image.

Their machine-learning model uses images of the robot's environment to generate a 2D action space, which includes the robot and the area around it. They simulate robot motion using what is known as the material-point-method, where the action space is covered by points, like image pixels, and overlaid with a grid.

The same way nearby pixels in an image are related (like the pixels that form a tree in a photo), they built their algorithm to understand that nearby action points have stronger correlations. Points around the robot's "shoulder" will move similarly when it changes shape, while points on the robot's "leg" will also move similarly, but in a different way than those on the "shoulder."

In addition, the researchers use the same machine-learning model to look at the environment and predict the actions the robot should take, which makes it more efficient.

Building a simulator

After developing this approach, the researchers needed a way to test it, so they created a simulation environment called DittoGym.

DittoGym features eight tasks that evaluate a reconfigurable robot's ability to dynamically change shape. In one, the robot must elongate and curve its body so it can weave around obstacles to reach a target point. In

another, it must change its shape to mimic letters of the alphabet.

"Our task selection in DittoGym follows both generic reinforcement learning benchmark design principles and the specific needs of reconfigurable robots. Each task is designed to represent certain properties that we deem important, such as the capability to navigate through long-horizon explorations, the ability to analyze the environment, and interact with external objects," Huang says.

"We believe they together can give users a comprehensive understanding of the flexibility of reconfigurable robots and the effectiveness of our [reinforcement learning](#) scheme."

Their algorithm outperformed baseline methods and was the only technique suitable for completing multistage tasks that required several shape changes.

"We have a stronger correlation between action points that are closer to each other, and I think that is key to making this work so well," says Chen.

While it may be many years before shape-shifting robots are deployed in the real world, Chen and his collaborators hope their work inspires other scientists not only to study reconfigurable soft robots but also to think about leveraging 2D action spaces for other complex control problems.

More information: Suning Huang et al, DittoGym: Learning to Control Soft Shape-Shifting Robots, *arXiv* (2024). [DOI: 10.48550/arxiv.2401.13231](https://doi.org/10.48550/arxiv.2401.13231)

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT

research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: A better way to control shape-shifting soft robots (2024, May 9) retrieved 17 July 2024 from <https://techxplore.com/news/2024-05-shifting-soft-robots.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.