


# Solving the challenges of robotic pizza-making

March 31 2022, by Adam Zewe



Researchers from MIT and elsewhere have created a framework that could enable a robot to effectively complete complex manipulation tasks with deformable objects, like dough or cloth, that require many tools and take a long time to complete. Credit: Massachusetts Institute of Technology

Imagine a pizza maker working with a ball of dough. She might use a

spatula to lift the dough onto a cutting board then use a rolling pin to flatten it into a circle. Easy, right? Not if this pizza maker is a robot.

For a robot, working with a deformable object like dough is tricky because the shape of dough can change in many ways, which are difficult to represent with an equation. Plus, creating a new shape out of that dough requires multiple steps and the use of different tools. It is especially difficult for a robot to learn a manipulation task with a long sequence of steps—where there are many possible choices—since learning often occurs through trial and error.

Researchers at MIT, Carnegie Mellon University, and the University of California at San Diego, have come up with a better way. They created a framework for a robotic manipulation system that uses a two-stage learning process, which could enable a robot to perform complex dough-manipulation tasks over a long timeframe. A "teacher" algorithm solves each step the robot must take to complete the task. Then, it trains a "student" [machine-learning](#) model that learns abstract ideas about when and how to execute each skill it needs during the task, like using a rolling pin. With this knowledge, the system reasons about how to execute the skills to complete the entire task.

The researchers show that this method, which they call DiffSkill, can perform complex manipulation tasks in simulations, like cutting and spreading dough, or gathering pieces of dough from around a cutting board, while outperforming other machine-learning methods.

Beyond pizza-making, this method could be applied in other settings where a robot needs to manipulate deformable objects, such as a caregiving robot that feeds, bathes, or dresses someone elderly or with motor impairments.

"This method is closer to how we as humans plan our actions. When a

human does a long-horizon task, we are not writing down all the details. We have a higher-level planner that roughly tells us what the stages are and some of the intermediate goals we need to achieve along the way, and then we execute them," says Yunzhu Li, a graduate student in the Computer Science and Artificial Intelligence Laboratory (CSAIL), and author of a paper presenting DiffSkill.

Li's co-authors include lead author Xingyu Lin, a graduate student at Carnegie Mellon University (CMU); Zhiao Huang, a graduate student at the University of California at San Diego; Joshua B. Tenenbaum, the Paul E. Newton Career Development Professor of Cognitive Science and Computation in the Department of Brain and Cognitive Sciences at MIT and a member of CSAIL; David Held, an assistant professor at CMU; and senior author Chuang Gan, a research scientist at the MIT-IBM Watson AI Lab. The research will be presented at the International Conference on Learning Representations.

## **Student and teacher**

The "teacher" in the DiffSkill framework is a trajectory [optimization algorithm](#) that can solve short-horizon tasks, where an object's initial state and target location are close together. The trajectory optimizer works in a simulator that models the physics of the real world (known as a differentiable physics simulator, which puts the "Diff" in "DiffSkill"). The "teacher" algorithm uses the information in the simulator to learn how the dough must move at each stage, one at a time, and then outputs those trajectories.

Then the "student" neural network learns to imitate the actions of the teacher. As inputs, it uses two camera images, one showing the dough in its current state and another showing the dough at the end of the task. The neural network generates a high-level plan to determine how to link different skills to reach the goal. It then generates specific, short-horizon

trajectories for each skill and sends commands directly to the tools.

The researchers used this technique to experiment with three different simulated dough-manipulation tasks. In one task, the robot uses a spatula to lift dough onto a cutting board then uses a rolling pin to flatten it. In another, the robot uses a gripper to gather dough from all over the counter, places it on a spatula, and transfers it to a cutting board. In the third task, the robot cuts a pile of dough in half using a knife and then uses a gripper to transport each piece to different locations.

## **A cut above the rest**

DiffSkill was able to outperform popular techniques that rely on reinforcement learning, where a robot learns a task through trial and error. In fact, DiffSkill was the only method that was able to successfully complete all three dough manipulation tasks. Interestingly, the researchers found that the "student" neural network was even able to outperform the "teacher" algorithm, Lin says.

"Our framework provides a novel way for robots to acquire new skills. These skills can then be chained to solve more complex tasks which are beyond the capability of previous robot systems," says Lin.

Because their method focuses on controlling the tools (spatula, knife, rolling pin, etc.) it could be applied to different robots, but only if they use the specific tools the researchers defined. In the future, they plan to integrate the shape of a tool into the reasoning of the "student" network so it could be applied to other equipment.

The researchers intend to improve the performance of DiffSkill by using 3D data as inputs, instead of images that can be difficult to transfer from simulation to the real world. They also want to make the [neural network](#) planning process more efficient and collect more diverse training data to

enhance DiffSkill's ability to generalize to new situations. In the long run, they hope to apply DiffSkill to more diverse tasks, including cloth manipulation.

**More information:** DiffSkill: Skill Abstraction from Differentiable Physics for Deformable Object Manipulations with Tools.  
[openreview.net/pdf?id=Kef8cKdHWpP](https://openreview.net/pdf?id=Kef8cKdHWpP)

*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: Solving the challenges of robotic pizza-making (2022, March 31) retrieved 17 June 2024 from <https://techxplore.com/news/2022-03-robotic-pizza-making.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.