


Giving robots a better feel for object manipulation

17 April 2019, by Rob Matheson



A new “particle simulator” developed by MIT researchers improves robots’ abilities to mold materials into simulated target shapes and interact with solid objects and liquids. This could give robots a refined touch for industrial applications or for personal robotics—such as shaping clay or rolling sticky sushi rice. Credit: Massachusetts Institute of Technology

A new learning system developed by MIT researchers improves robots' abilities to mold materials into target shapes and make predictions about interacting with solid objects and liquids. The system, known as a learning-based particle simulator, could give industrial robots a more refined touch—and it may have fun applications in personal robotics, such as modelling clay shapes or rolling sticky rice for sushi.

In robotic planning, physical simulators are models that capture how different materials respond to force. Robots are “trained” using the models, to predict the outcomes of their interactions with objects, such as pushing a solid box or poking deformable clay. But traditional learning-based simulators mainly focus on rigid objects and are unable to handle fluids or softer objects. Some more accurate physics-based simulators can

handle diverse materials, but rely heavily on approximation techniques that introduce errors when robots interact with objects in the real world.

In a paper being presented at the International Conference on Learning Representations in May, the researchers describe a [new model](#) that learns to capture how small portions of different materials—“particles”—interact when they’re poked and prodded. The [model](#) directly learns from data in cases where the underlying physics of the movements are uncertain or unknown. Robots can then use the model as a guide to predict how liquids, as well as rigid and deformable materials, will react to the force of its touch. As the robot handles the objects, the model also helps to further refine the robot's control.

In experiments, a robotic hand with two fingers, called “RiceGrip,” accurately shaped a deformable foam to a desired configuration—such as a “T” shape—that serves as a proxy for sushi rice. In short, the researchers' model serves as a type of “intuitive physics” brain that robots can leverage to reconstruct three-dimensional objects somewhat similarly to how humans do.

“Humans have an intuitive physics model in our heads, where we can imagine how an object will behave if we push or squeeze it. Based on this intuitive model, humans can accomplish amazing manipulation tasks that are far beyond the reach of current robots,” says first author Yunzhu Li, a graduate student in the Computer Science and Artificial Intelligence Laboratory (CSAIL). “We want to build this type of intuitive model for robots to enable them to do what humans can do.”

“When children are 5 months old, they already have different expectations for solids and liquids,” adds co-author Jiajun Wu, a CSAIL graduate student. “That’s something we know at an early age, so maybe that’s something we should try to model for robots.”

Joining Li and Wu on the paper are: Russ Tedrake, a CSAIL researcher and a professor in the Department of Electrical Engineering and Computer Science (EECS); Joshua Tenenbaum, a professor in the Department of Brain and Cognitive Sciences; and Antonio Torralba, a professor in EECS and director of the MIT-IBM Watson AI Lab.

Dynamic graphs

A key innovation behind the model, called the "particle interaction network" (DPI-Nets), was creating dynamic interaction graphs, which consist of thousands of nodes and edges that can capture complex behaviors of so-called particles. In the graphs, each node represents a particle. Neighboring nodes are connected with each other using directed edges, which represent the interaction passing from one particle to the other. In the simulator, particles are hundreds of small spheres combined to make up some liquid or a deformable object.

The graphs are constructed as the basis for a machine-learning system called a graph neural network. In training, the model over time learns how particles in different materials react and reshape. It does so by implicitly calculating various properties for each particle—such as its mass and elasticity—to predict if and where the particle will move in the graph when perturbed.

The model then leverages a "propagation" technique, which instantaneously spreads a signal throughout the graph. The researchers customized the technique for each type of material—rigid, deformable, and liquid—to shoot a signal that predicts particles positions at certain incremental time steps. At each step, it moves and reconnects particles, if needed.

For example, if a solid box is pushed, perturbed particles will be moved forward. Because all particles inside the box are rigidly connected with each other, every other particle in the object moves the same calculated distance, rotation, and any other dimension. Particle connections remain intact and the box moves as a single unit. But if an area of deformable foam is indented, the effect will be different. Perturbed particles move forward a lot,

surrounding particles move forward only slightly, and particles farther away won't move at all. With liquids being sloshed around in a cup, particles may completely jump from one end of the graph to the other. The graph must learn to predict where and how much all affected particles move, which is computationally complex.

Shaping and adapting

In their paper, the researchers demonstrate the model by tasking the two-fingered RiceGrip robot with clamping target shapes out of deformable foam. The robot first uses a depth-sensing camera and object-recognition techniques to identify the foam. The researchers randomly select particles inside the perceived shape to initialize the position of the particles. Then, the model adds edges between particles and reconstructs the foam into a dynamic graph customized for deformable materials.

Because of the learned simulations, the robot already has a good idea of how each touch, given a certain amount of force, will affect each of the particles in the graph. As the [robot](#) starts indenting the foam, it iteratively matches the real-world position of the particles to the targeted position of the particles. Whenever the particles don't align, it sends an error signal to the model. That signal tweaks the model to better match the real-world physics of the material.

Next, the researchers aim to improve the model to help robots better predict interactions with partially observable scenarios, such as knowing how a pile of boxes will move when pushed, even if only the boxes at the surface are visible and most of the other boxes are hidden.

The researchers are also exploring ways to combine the model with an end-to-end perception module by operating directly on images. This will be a joint project with Dan Yamins's group; Yamin recently completed his postdoc at MIT and is now an assistant professor at Stanford University. "You're dealing with these cases all the time where there's only partial information," Wu says. "We're extending our model to learn the dynamics of all [particles](#), while only seeing a small portion."

More information: Yunzhu Li et al. Learning Particle Dynamics for Manipulating Rigid Bodies, Deformable Objects, and Fluids, arXiv:1810.01566 [cs.LG]. arxiv.org/abs/1810.01566

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

APA citation: Giving robots a better feel for object manipulation (2019, April 17) retrieved 27 May 2022 from https://techxplore.com/news/2019-04-robots_1.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.