

Pushy robots learn the fundamentals of object manipulation

22 October 2019, by Rob Matheson



A key to compiling the novel Omnipush dataset was building modular objects (pictured) that enabled the robotic system to capture a vast diversity of pushing behavior. The central pieces contain markers on their centers and points so a motion-detection system can detect their position within a millimeter. Credit: Massachusetts Institute of Technology

MIT researchers have compiled a dataset that captures the detailed behavior of a robotic system physically pushing hundreds of different objects. Using the dataset—the largest and most diverse of its kind—researchers can train robots to "learn" pushing dynamics that are fundamental to many complex object-manipulation tasks, including reorienting and inspecting objects, and uncluttering scenes.

To capture the data, the researchers designed an automated system consisting of an industrial robotic arm with precise control, a 3-D motion-tracking system, depth and traditional cameras, and software that stitches everything together. The arm pushes around modular objects that can be adjusted for weight, shape, and mass distribution. For each push, the system captures how those characteristics affect the [robot's](#) push.

The dataset, called "Omnipush," contains 250 different pushes of 250 objects, totaling roughly 62,500 unique pushes. It's already being used by researchers to, for instance, build models that help robots predict where objects will land when they're pushed.

"We need a lot of rich data to make sure our robots can learn," says Maria Bauza, a [graduate student](#) in the Department of Mechanical Engineering (MechE) and first author of a paper describing Omnipush that's being presented at the upcoming International Conference on Intelligent Robots and Systems. "Here, we're collecting data from a real robotic system, [and] the objects are varied enough to capture the richness of the pushing phenomena. This is important to help robots understand how pushing works, and to translate that information to other similar objects in the real world."

Joining Bauza on the paper are: Ferran Alet and Yen-Chen Lin, graduate students in the Computer Science and Artificial Intelligence Laboratory and the Department of Electrical Engineering and Computer Science (EECS); Tomas Lozano-Perez, the School of Engineering Professor of Teaching Excellence; Leslie P. Kaelbling, the Panasonic Professor of Computer Science and Engineering; Phillip Isola, an assistant professor in EECS; and Alberto Rodriguez, an associate professor in MechE.

Diversifying data

Why focus on pushing behavior? Modeling pushing dynamics that involve friction between objects and surfaces, Rodriguez explains, is critical in higher-level robotic tasks. Consider the visually and technically impressive robot that can play Jenga, which Rodriguez recently co-designed. "The robot is performing a complex task, but the core of the mechanics driving that task is still that of pushing an object affected by, for instance, the friction between blocks," Rodriguez says.

Omnipush builds on a similar dataset built in the Manipulation and Mechanisms Laboratory (MCube) by Rodriguez, Bauza, and other researchers that captured pushing data on only 10 objects. After making the dataset public in 2016, they gathered feedback from researchers. One complaint was lack of object diversity: Robots trained on the dataset struggled to generalize information to new objects. There was also no video, which is important for computer vision, video prediction, and other tasks.

For their new dataset, the researchers leverage an industrial robotic arm with precision control of the velocity and position of a pusher, basically a vertical steel rod. As the arm pushes the objects, a "Vicon" motion-tracking system—which has been used in films, virtual reality, and for research—follows the objects. There's also an RGB-D camera, which adds depth information to captured video.

The key was building modular objects. The uniform central pieces, made from aluminum, look like four-pointed stars and weigh about 100 grams. Each central piece contains markers on its center and points, so the Vicon system can detect its pose within a millimeter.

Smaller pieces in four shapes—concave, triangular, rectangular, and circular—can be magnetically attached to any side of the central piece. Each piece weighs between 31 to 94 grams, but extra weights, ranging from 60 to 150 grams, can be dropped into little holes in the pieces. All pieces of the puzzle-like objects align both horizontally and vertically, which helps emulate the friction a single object with the same shape and mass distribution would have. All combinations of different sides, weights, and mass distributions added up to 250 unique objects.

For each push, the arm automatically moves to a random position several centimeters from the object. Then, it selects a random direction and pushes the object for one second. Starting from where it stopped, it then chooses another random direction and repeats the process 250 times. Each push records the pose of the object and RGB-D video, which can be used for various video-

prediction purposes. Collecting the data took 12 hours a day, for two weeks, totaling more than 150 hours. Human intervention was only needed when manually reconfiguring the objects.

The objects don't specifically mimic any real-life items. Instead, they're designed to capture the diversity of "kinematics" and "mass asymmetries" expected of real-world objects, which model the physics of the motion of real-world objects. Robots can then extrapolate, say, the physics model of an Omnipush object with uneven mass distribution to any real-world object with similar uneven weight distributions.

"Imagine pushing a table with four legs, where most weight is over one of the legs. When you push the table, you see that it rotates on the heavy leg and have to readjust. Understanding that [mass distribution](#), and its effect on the outcome of a push, is something robots can learn with this set of objects," Rodriguez says.

Powering new research

In one experiment, the researchers used Omnipush to train a model to predict the final pose of pushed objects, given only the initial pose and description of the push. They trained the model on 150 Omnipush objects, and tested it on a held-out portion of objects. Results showed that the Omnipush-trained model was twice as accurate as models trained on a few similar datasets. In their paper, the researchers also recorded benchmarks in accuracy that other researchers can use for comparison.

Because Omnipush captures video of the pushes, one potential application is video prediction. A collaborator, for instance, is now using the dataset to train a robot to essentially "imagine" pushing objects between two points. After training on Omnipush, the robot is given as input two video frames, showing an object in its starting position and ending position. Using the starting position, the robot predicts all future video frames that ensure the object reaches its ending position. Then, it pushes the object in a way that matches each predicted video frame, until it gets to the frame with the ending position.

"The robot is asking, "If I do this action, where will the object be in this frame?" Then, it selects the action that maximizes the likelihood of getting the object in the position it wants," Bauza says. "It decides how to move objects by first imagining how the pixels in the image will change after a push."

"Omnipush includes precise measurements of object motion, as well as visual data, for an important class of interactions between robot and objects in the world," says Matthew T. Mason, a professor of [computer science](#) and robotics at Carnegie Mellon University. "Robotics researchers can use this data to develop and test new robot learning approaches ... that will fuel continuing advances in robotic manipulation."

More information: Omnipush: accurate, diverse, real-world dataset of pushing dynamics with RGB-D video: arxiv.org/abs/1910.00618

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: Pushy robots learn the fundamentals of object manipulation (2019, October 22) retrieved 18 January 2021 from <https://techxplore.com/news/2019-10-pushy-robots-fundamentals.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.