


Giving robots a sense of touch

June 5 2017, by Larry Hardesty



A GelSight sensor attached to a robot's gripper enables the robot to determine precisely where it has grasped a small screwdriver, removing it from and inserting it back into a slot, even when the gripper screens the screwdriver from the robot's camera. Credit: Robot Locomotion Group at MIT

Eight years ago, Ted Adelson's research group at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) unveiled a new

sensor technology, called GelSight, that uses physical contact with an object to provide a remarkably detailed 3-D map of its surface.

Now, by mounting GelSight [sensors](#) on the grippers of robotic arms, two MIT teams have given robots greater sensitivity and dexterity. The researchers presented their work in two papers at the International Conference on Robotics and Automation last week.

In one paper, Adelson's group uses the data from the GelSight sensor to enable a [robot](#) to judge the hardness of surfaces it touches—a crucial ability if household robots are to handle everyday objects.

In the other, Russ Tedrake's Robot Locomotion Group at CSAIL uses GelSight sensors to enable a robot to manipulate smaller objects than was previously possible.

The GelSight sensor is, in some ways, a low-tech solution to a difficult problem. It consists of a block of transparent rubber—the "gel" of its name—one face of which is coated with metallic paint. When the paint-coated face is pressed against an [object](#), it conforms to the object's shape.

The metallic paint makes the object's surface reflective, so its geometry becomes much easier for computer vision algorithms to infer. Mounted on the sensor opposite the paint-coated face of the rubber block are three colored lights and a single camera.

"[The system] has colored lights at different angles, and then it has this reflective material, and by looking at the colors, the computer ... can figure out the 3-D shape of what that thing is," explains Adelson, the John and Dorothy Wilson Professor of Vision Science in the Department of Brain and Cognitive Sciences.

In both sets of experiments, a GelSight sensor was mounted on one side of a robotic gripper, a device somewhat like the head of a pincer, but with flat gripping surfaces rather than pointed tips.

Contact points

For an autonomous robot, gauging objects' softness or hardness is essential to deciding not only where and how hard to grasp them but how they will behave when moved, stacked, or laid on different surfaces. Tactile sensing could also aid robots in distinguishing objects that look similar.

In previous work, robots have attempted to assess objects' hardness by laying them on a flat surface and gently poking them to see how much they give. But this is not the chief way in which humans gauge hardness. Rather, our judgments seem to be based on the degree to which the contact area between the object and our fingers changes as we press on it. Softer objects tend to flatten more, increasing the contact area.

The MIT researchers adopted the same approach. Wenzhen Yuan, a [graduate student](#) in [mechanical engineering](#) and first author on the paper from Adelson's group, used confectionary molds to create 400 groups of silicone objects, with 16 objects per group. In each group, the objects had the same shapes but different degrees of hardness, which Yuan measured using a standard industrial scale.

Then she pressed a GelSight sensor against each object manually and recorded how the contact pattern changed over time, essentially producing a short movie for each object. To both standardize the data format and keep the size of the data manageable, she extracted five frames from each movie, evenly spaced in time, which described the deformation of the object that was pressed.

Finally, she fed the data to a neural network, which automatically looked for correlations between changes in contact patterns and hardness measurements. The resulting system takes frames of video as inputs and produces hardness scores with very high accuracy. Yuan also conducted a series of informal experiments in which human subjects palpated fruits and vegetables and ranked them according to hardness. In every instance, the GelSight-equipped robot arrived at the same rankings.

Yuan is joined on the paper by her two thesis advisors, Adelson and Mandayam Srinivasan, a senior research scientist in the Department of Mechanical Engineering; Chenzhuo Zhu, an undergraduate from Tsinghua University who visited Adelson's group last summer; and Andrew Owens, who did his PhD in [electrical engineering](#) and computer science at MIT and is now a postdoc at the University of California at Berkeley.

Obstructed views

The paper from the Robot Locomotion Group was born of the group's experience with the Defense Advanced Research Projects Agency's Robotics Challenge (DRC), in which academic and industry teams competed to develop control systems that would guide a humanoid robot through a series of tasks related to a hypothetical emergency.

Typically, an [autonomous robot](#) will use some kind of computer vision system to guide its manipulation of objects in its environment. Such systems can provide very reliable information about an object's location—until the robot picks the object up. Especially if the object is small, much of it will be occluded by the robot's gripper, making location estimation much harder. Thus, at exactly the point at which the robot needs to know the object's location precisely, its estimate becomes unreliable. This was the problem the MIT team faced during the DRC, when their robot had to pick up and turn on a power drill.

"You can see in our video for the DRC that we spend two or three minutes turning on the drill," says Greg Izatt, a graduate student in electrical engineering and computer science and first author on the new paper. "It would be so much nicer if we had a live-updating, accurate estimate of where that drill was and where our hands were relative to it."

That's why the Robot Locomotion Group turned to GelSight. Izatt and his co-authors—Tedrake, the Toyota Professor of Electrical Engineering and Computer Science, Aeronautics and Astronautics, and Mechanical Engineering; Adelson; and Geronimo Mirano, another graduate student in Tedrake's group—designed control algorithms that use a computer vision system to guide the robot's gripper toward a tool and then turn location estimation over to a GelSight sensor once the robot has the tool in hand.

In general, the challenge with such an approach is reconciling the data produced by a vision system with data produced by a tactile sensor. But GelSight is itself camera-based, so its data output is much easier to integrate with visual data than the data from other tactile sensors.

In Izatt's experiments, a robot with a GelSight-equipped gripper had to grasp a small screwdriver, remove it from a holster, and return it. Of course, the data from the GelSight sensor don't describe the whole screwdriver, just a small patch of it. But Izatt found that, as long as the vision system's estimate of the screwdriver's initial position was accurate to within a few centimeters, his algorithms could deduce which part of the screwdriver the GelSight sensor was touching and thus determine the screwdriver's position in the robot's hand.

"I think that the GelSight technology, as well as other high-bandwidth tactile sensors, will make a big impact in robotics," says Sergey Levine, an assistant professor of electrical engineering and computer science at the University of California at Berkeley. "For humans, our sense of

touch is one of the key enabling factors for our amazing manual dexterity. Current robots lack this type of dexterity and are limited in their ability to react to surface features when manipulating objects. If you imagine fumbling for a light switch in the dark, extracting an object from your pocket, or any of the other numerous things that you can do without even thinking—these all rely on touch sensing."

"Software is finally catching up with the capabilities of our sensors," Levine adds. "Machine learning algorithms inspired by innovations in deep learning and computer vision can process the rich sensory data from sensors such as the GelSight to deduce object properties. In the future, we will see these kinds of learning methods incorporated into end-to-end trained manipulation skills, which will make our robots more dexterous and capable, and maybe help us understand something about our own sense of touch and motor control."

More information: Tracking Objects with Point Clouds From Vision and Touch. [groups.csail.mit.edu/robotics- ... c_papers/Izatt16.pdf](https://groups.csail.mit.edu/robotics-...c_papers/Izatt16.pdf)

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: Giving robots a sense of touch (2017, June 5) retrieved 9 April 2024 from <https://techxplore.com/news/2017-06-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.
