

TRI Taking on the Hard Problems in Manipulation Research Toward Making Human-Assist Robots Reliable and Robust

June 27, 2019

Image not found or type unknown



Contributed Story by Dr. Russ Tedrake, TRI Vice President of Robotics Research:

“Wouldn’t it be amazing to have a robot in your home that could work with you to put away the groceries, fold the laundry, cook your dinner, do the dishes, and tidy up before the guests come over? For some of us, a robot assistant – a teammate – might only be a convenience. But for others, including our growing population of older people, applications like this could be the difference between living at home or in an assisted care facility. Done right, we believe these robots will amplify and augment human capabilities, allowing us to enjoy longer, healthier lives.

Decades of prognostications about the future – largely driven by science fiction novels and popular entertainment – have encouraged public expectations that someday home robots will happen. Companies have been trying for years to deliver on such forecasts and figure out how to safely introduce ever more capable robots into the unstructured home environment.

Despite this age of tremendous technological progress, the robots we see in homes to date are primarily vacuum cleaners and toys. Most people don’t realize how far today’s best robots are from being able to do basic household tasks. When they see heavy use of robot arms in factories or impressive videos on YouTube showing what a robot can do, they might reasonably expect these robots could be used in the home now.

Why haven’t home robots materialized as quickly as some have come to expect?

One big challenge is reliability. Consider:

- If you had a robot that could load dishes into the dishwasher for you, what if it broke a dish once a week?
- Or, what if your child brings home a “#1 DAD!” mug that she painted at the local art studio, and after dinner, the robot discards that mug into the trash because it didn’t recognize it as an actual mug?

A major barrier for bringing robots into the home are core unsolved problems in manipulation that prevent reliability. As I presented this week at the “Robotics: Science and Systems” (RSS) conference, the Toyota Research Institute (TRI) is working on fundamental issues in robot manipulation to tackle these unsolved reliability challenges. We have been pursuing a unique combination of robotics capabilities focused on dexterous tasks in an unstructured environment.

Unlike the sterile, controlled and programmable environment of the factory, the home is a “wild west” – unstructured and diverse. We cannot expect lab tests to account for every different object that a robot will see in your home. This challenge is sometimes referred to as “open-world manipulation,” as a callout to “open-world” computer games. Despite recent strides in artificial intelligence (AI) and machine learning (ML), it is still very hard to engineer a system that can deal with the complexity of a home environment and guarantee that it will (almost) always work correctly.

Here is a [demonstration video](#) showing how we are exploring the challenge of robustness that addresses the reliability gap. We are using a robot loading dishes in a dishwasher as an example task. Our pursuit is not to design a robot that loads the dishwasher, but rather we use this task as a means to develop the tools and algorithms that can in turn be applied in many different applications. Our focus is not on hardware, which is why we are using a factory robot arm in this demonstration rather than designing one that would be more appropriate for the home kitchen.

The robot in our demonstration uses stereo cameras mounted around the sink and deep learning algorithms to perceive objects in the sink. There are many robots out there today that can pick up almost any object— *random object clutter clearing* has become a standard benchmark robotics challenge. In clutter clearing, the robot doesn’t require much understanding about an object — perceiving the basic geometry is enough. For example, the

algorithm doesn't need to recognize if the object is a plush toy, a toothbrush, or a coffee mug. Given this, these systems are also relatively limited with what they can do with those objects; for the most part, they can only pick up the objects and drop them in another location only. In the robotics world, we sometimes refer to these robots as "pick and drop."

Loading the dishwasher is actually significantly harder than what most roboticists are currently demonstrating, and it requires considerably more understanding about the objects. Not only does the robot have to recognize a mug or a plate or "clutter," but it has to also understand the shape, position, and orientation of each object in order to place it accurately in the dishwasher. TRI's work-in-progress shows not only that this is possible, but that it can be done with robustness that allows the robot to continuously operate for hours without disruption.

Our manipulation robot has a relatively simple hand — a two-fingered gripper. The hand can make relatively simple grasps on a mug, but its ability to pick up a plate is more subtle. Plates are large and may be stacked, so we have to execute a complex "contact-rich" maneuver that slides one gripper finger under and between plates in order to get a firm hold. This is a simple example of the type of dexterity that humans achieve easily, but that we rarely see in robust robotics applications.

Silverware can also be tricky—it is small and shiny, which makes it hard to see with a machine learning camera. Plus, given that the robot hand is relatively large compared to the smaller sink, the robot occasionally needs to stop and nudge the silverware to the center of the sink in order to do the pick. Our system can also detect if an object is not a mug, plate or silverware and, labeling it as "clutter," and move it to a "discard" bin.

Connecting all of these pieces is a sophisticated *task planner*, which is constantly deciding what task the robot should execute next. This task planner decides if it should pull out the bottom drawer of the dishwasher to load some plates, pull out the middle drawer for mugs, or pull out the top drawer for silverware. Like the other components, we have made it resilient — if the drawer gets suddenly closed when it was needed to be open, the robot will stop, put down the object on the counter top, and pull the drawer back out to try again. This response shows how different this capability is than a typical precision, repetitive factory robot, which are typically isolated from human contact and environmental randomness.

The cornerstone of TRI's approach is the use of simulation. Simulation gives us a principled way to engineer and test systems of this complexity with incredible task diversity and machine learning and artificial intelligence components. It allows us to understand what level of performance the robot will have in your home with your mugs, even though we haven't been able to test in your kitchen during our development. An exciting achievement is that we have made great strides in making simulation robust enough to handle the visual and mechanical complexity of this dishwasher loading task and on closing the "sim to real" gap. We are now able to design and test in simulation and have confidence that the results will transfer to the real robot. At long last, we have reached a point where we do nearly all of our development in simulation, which has traditionally not been the case for robotic manipulation research.

We can run many more tests in simulation and more diverse tests. We are constantly generating random scenarios that will test the individual components of the dish loading plus the end-to-end performance.

Let me give you a simple example of how this works. Consider the task of extracting a single mug from the sink. We generate scenarios where we place the mug in all sorts of random configurations, testing to find "corner cases" — rare situations where our perception algorithms or grasping algorithms might fail. We can vary material properties and lighting conditions. We even have algorithms for generating random, but reasonable, shapes of the mug, generating everything from a small espresso cup to a portly cylindrical coffee mug.

We conduct simulation testing through the night, and every morning we receive a report that gives us new failure cases that we need to address. Early on, those failures were relatively easy to find, and easy to fix. Sometimes

they are failures of the simulator — something happened in the simulator that could never have happened in the real world – and sometimes they are problems in our perception or grasping algorithms. We have to fix all of these failures.

As we continue down this road to robustness, the failures are getting more rare and more subtle. The algorithms that we use to find those failures also need to get more advanced. The search space is so huge, and the performance of the system so nuanced, that finding the corner cases efficiently becomes our core research challenge. Although we are exploring this problem in the kitchen sink, the core ideas and algorithms are motivated by, and are applicable to, related problems such as verifying automated driving technologies.

The next piece of our work focuses on the development of algorithms to automatically “repair” the perception algorithm or controller whenever we find a new failure case. Because we are using simulation, we can test our changes against not only this newly discovered scenario, but also make sure that our changes also work for all of the other scenarios that we’ve discovered in the preceding tests. Of course, it’s not enough to fix this one test. We have to make sure we also do not break all of the other tests that passed before. It’s possible to imagine a not-so-distant future where this repair can happen directly in your kitchen, whereby if one robot fails to handle your mug correctly, then all robots around the world learn from that mistake.

We are committed to achieving dexterity and reliability in open-world manipulation. Loading a dishwasher is just one example in a series of experiments we will be using at TRI to focus on this problem. It’s a long journey, but ultimately it will produce capabilities that will bring more advanced robots into the home. When this happens, we hope that older adults will have the help they need to age in place with dignity, working with a robotic helper that will amplify their capabilities, while allowing more independence, longer.”