A Geometry-Based Approach for Learning from Demonstrations for Manipulation

John Schulman, Jonathan Ho, Cameron Lee, Pieter Abbeel *

One of the barriers to widespread deployment of autonomous robots—whether in factories, hospitals, or homes—is the substantial amount of effort and skill it takes to program them to perform new tasks. A large part of this effort must be repeated whenever the task is changed.

When we teach children to tie their shoelaces, we don't need to provide an extremely precise description of the task, nor do we need to enumerate all possible states and transitions—we teach them by example. Wouldn't it be nice if we could teach robots the same way? (But without the tears and frustration.)

This paper is concerned with teaching robots to perform manipulation tasks by demonstration. In other words, a human performs the task one or more times (by teleoperation or directly guiding the robot's end-effector), and a learning algorithm extracts the essence of these demonstrations so the robot can perform the task autonomously in a new situation.

We introduce a new approach for teaching robots from demonstrations, which emphasizes the geometry of the manipulated objects. Our method is motivated by problems where there is high variability in the manipulated objects, i.e., in their shapes, sizes, and poses. Such variability is often unavoidable when manipulating deformable objects, because of their high-dimensional configuration spaces. Our main running example is tying knots in rope, which is a proxy for several motivating tasks of practical importance, such as tying fasteners around wire bundles (common in aerospace applications) and surgical suturing. We also consider a variety of manipulation tasks involving household objects.

Our approach starts out by finding a non-rigid registration between the geometry (points, curves, surfaces) of the training scene and the testing scene, i.e., the new configuration of objects that the robot must act on. While registration is only concerned with the objects and their environment, we show that it is possible to meaningfully extrapolate to the entire space for our particular choice of non-rigid registration. This in turn enables using the extrapolated registration, which informally could be called a "space warping," to transform (=generalize) the robot tools' pose trajectories from the demonstration scene to the new scene. While non-rigid registration has been extensively studied for applications such as 3D modeling and medical image analysis, to ensure meaningful gripper pose trajectories we present a new objective criterion for the registration procedure that is more directly suitable for robotic applications.

Our method for generalizing trajectories can be used as part of a closed-loop system for completing an extended task such as knot tying. In brief, the robot repeats the following steps: (1) looks up the nearest situation in the demonstrations by comparing the point clouds, (2) adapts the demonstrated trajectory to the current situation, and (3) executes it, with the help of inverse kinematics and possibly motion planning. We experimentally validate this method in the setting of knot tying, enabling a two-armed robot to reliably tie several kinds of knots. This vision-based knot-tying procedure can handle different starting conditions and rope parameters, and it recovers from errors.

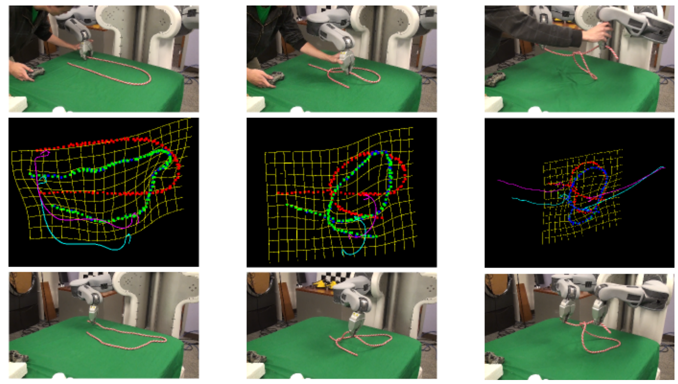

Figure 1: Robot learning and executing overhand knot, the simplest of five knots taught using our algorithm. Top row: human demonstration by kinesthetic teaching. Middle row: red: point cloud from the demonstration, blue: point cloud from the test, green: warping function applied to demonstration points (which for an accurate registration mostly fall right on top of blue), yellow lines: warping function applied to a uniform grid in demonstration. Bottom row: the robot applies the learned trajectory in the test situation. Note that this procedure is split into three steps, and registration is performed three times.

Using the same method for generalizing trajectories, we enable a robot to perform a variety of tasks in the domain of household chores, in which the robot must adapt to variations in the shapes and sizes of the manipulated objects. Most of these tasks involve a tool or held object, respectively; our framework naturally handles these tasks by treating the held object as the end-effector of the robot. Similarly to how we adapt to variation in the target objects, we can adapt to variation in the tool by using non-rigid registration to identify the end-effector frame of the new tool.

Our approach to trajectory generalization has the following desirable attributes, compared to some other learning methods:

- It operates directly on point clouds–the outputs of our sensor hardware, rather than requiring the user or programmer to define the relevant frames or some featurization of the input.
- It requires a very small number of demonstrations.

*UC Berkeley, EECS Department