

Learning Actions from Demonstration

Michael Tirtowidjojo, Matthew Frierson, Benjamin Singer, Palak Hirpara

October 2, 2016

Abstract

The goal of our project is twofold. First, we will design a controller layout that is user friendly, intuitive, and multi-functional. This new layout will allow the user to maneuver the arm, control teleoperation, and switch between the two modes. In addition, the buttons on the controller will allow for a future user to add extra functionality and customization. The second section of our project involves using the new controller to teach the robot, by demonstration, to be able to identify primitive actions such as poke, lift, and sweep.

1 Introduction

As robots become more prevalent in the everyday world, the importance of having robots that can be taught by humans and that can adapt to hostile situations has increased significantly. However, before a robot can tackle complex tasks, it first needs to master primitive actions. Then, these commands can be chained together to achieve more meaningful tasks. This problem can be simplified using learning from demonstration because this method of teaching allows a person with no prior experience in robotics to teach a robot.

2 Related Works

Our proposed project builds off several previous papers that utilize learning from demonstration techniques. “A Survey of Robot Learning From Demonstration” by Argall, Chernova, Veloso, and Browning provides a survey of research in robotic learning by demonstration and explains various methods to teach robots using this general method.¹ The authors note that by using this method, robots become more universally accessible to everyday users. The paper discusses two techniques for learning: generalizing after all actions have been completed, and updating a policy as training data becomes available.

A second paper, “Learning Task Sequences from Scratch: Application to the Control of Tools and Toys by a Humanoid Robot” by Arsenio discusses teaching complex action sequences

to a robot using visual observation of human teachers completing similar tasks.² He describes the learning process in three general steps. First, the robot recognizes an object using the objects color, luminance, and shape cues and generates object models. Second, it associates the object with a corresponding action. Lastly, the robot learns both the sequence of events that comprise a task as well as the objects being acted on. In the process, the robot uses it Markov chains and kinetics mapping to extract information about a particular task it is trying to learn.

Finally, “Learning Similar Tasks From Observation and Practice” by Bentivegna, Atkeson, and Cheng describes a case study in which the researchers taught a robot, using learning from demonstration and learning from practice, to solve a marble maze.³ This paper focuses on the use of “local representations” of a situation, descriptions which focus on features of the board local to the robot, and builds off of previous research in which the authors used “global representations,” descriptions that include the specific location of the ball on the board. They claim that using a global representation helps the robot develop skills specific to a particular situation, while a local representation allows a robot to generalize the data it collects. After completing their experiment, they noted that the robot achieved its goal much more slowly than its human teachers, indicating that practice would help the robot improve its skills. After the robot practiced the game sixty times, the researchers noticed both an increase in the average velocity of the robots movements and in the number of times it successfully completed the maze. To buttress their claims, the authors ran the same experiment using a software maze and demonstrated similar success.

3 Problem

3.1 Better Control System

The first problem our project will address is improving the control system for the arm and robots movements. The current setup is inconvenient to use: to teleoperate the robot, a human is required to bend over and use the keyboard keys, and to manipulate the arm, a human is required either to hold and move the arm or use an unintuitive joystick. Controlling the robot is tiring and uncomfortable. Although people who are familiar with the robot will not find the current setup too technically complicated, inexperienced users may find it difficult to control the robot. This is especially problematic if these inexperienced users are using a learning from demonstration technique to teach the robot. Moreover the current setup runs the risk of inexperienced humans overestimating the force the arm can take and inadvertently damaging it.

3.2 Generalizing Primitive Actions

Our second goal is to develop a method of generalizing primitive actions learned from demonstration. Currently, the robot can perform such primitive actions only if manually programmed

to do so. For example, the arm can be directed to move if it is passed parameters such as velocity and angle. However, it cannot classify the actions it performs; if the robot closes its fingers, it is not aware that it has completed a grasping action. We will solve this problem by developing a protocol for the robot to learn primitive actions. Using this protocol, the robot will be able to generalize these primitive actions to both combine actions to accomplish a larger goal and break down a task into discrete components.

4 Technical Approach

Our project consists of two elements: implementing a more intuitive joystick controller and teaching a robot actions through demonstration. These parts build on each other; we will use the joystick to demonstrate actions to the robot.

4.1 Joystick

We will adapt either a PS3 or XBox 360 controller to work with the robot. These controllers have both multiple small joysticks and buttons, providing for more functionality than the current setup. In order to design the button layout and connect the controllers to our robots, we plan to utilize joy and ps3joy, open-source ROS driver packages. Moreover, our new controller will allow other researchers to customize configurations for their own projects a project related to grasping can now configure the buttons to be grasp-specific.

4.2 Learning from Demonstration

Learning from demonstration requires the the robot be able to classify observed actions and events in order to generalize and plan. In our project, we intend to demonstrate variations on a set of actions and allow the robot to learn how to generalize those actions and classify different variations on them.

4.2.1 Method of Demonstration

As discussed in “A Survey of Robot Learning by Demonstration,” one can demonstrate either through direct control of the robot, via a joystick or physically moving the robot, or through sensors connected to a human teacher, a sensor suit. We intend to directly control the robot arm through our controller. As described by the survey paper, this method allows for direct mapping between demonstrated data and classifiers. Conversely, recording data using a sensor suit would require the data to be translated into a pattern, limited by the robots degrees of freedom, that the robot can use. Therefore, by demonstrating the actions on the robot itself, we bypass these onerous restrictions.

4.2.2 Demonstrating Actions

We intend to teach the robot a set of actions: {Lift, Poke, Circle, Grasp, Sweep, Push, Punch, Rotate, Wave} by demonstrating several variations of each action. These variations will allow the robot to generalize over various examples: a “lift” action encompasses lifting to different heights or starting from different points. We intend to collect the velocity and force readings from each of the arm joints and from the end effector and pass these data points to the classifier. The aim of these demonstrations is to help the robot recognize actions after observing a new action, the robot should be able to correctly label the movement. Thus, testing the classifier and learning from demonstration demands providing new variations on the set of actions to confirm that the robot can correctly classify the new variations.

5 Evaluation

Our project evaluation too has multiple aspects: evaluating the controller and evaluating the learning from the actions demonstrated to the robot. It is important to note, however, that the controller goal is secondary to our larger learning by demonstration project. Thus, our evaluation of the joystick will not be as rigorous as our learning from demonstration tests.

5.1 Joystick Evaluation

Our primary metric in measuring the efficacy of our new controller is its ease of use compared to the current setup. In addition, the new controller needs to be simple enough for a beginner user to understand. Thus, our measurement revolves around the new controllers simplicity in comparison to the previous controller. We will capture this measurement by timing new users completing certain tasks. In our experiment, the user will be timed using each controller while performing five different actions and tasks. Our new controller will be judged a success if, on average, users are able to complete at least 75% of the assigned tasks more quickly.

5.1.1 Joystick Tasks

Each participant will complete the five tasks described in Fig. 1. The participant will be randomly given a joystick setup to use first. He or she will then complete each of the five tasks in a random order, and then redo the sequence using the other controller. We can then compare how long the users take to complete each task with each setup.

5.2 Learning from Demonstration Evaluation

After the robot has been trained through demonstrations to be able to classify a given basic set of actions, we will teach the robot how to interpret and classify a variation on each action in the set. To test how well the robot has learned and generalized the demonstrated actions, the

Move forward	Move the arm straight forward to a point without moving up or down
Move forward and up	Move the arm to a point that is in front of and above the robot
Move diagonally	Move the arm to a point that is in front, above, and to the side of the robot
Pick up an object	Direct the arm to an object and grasp the object with the fingers
Bring object to self	Starting from having an object grasped, lift the object and move it back to the robot

Figure 1: Five tasks that will be used to evaluate the effectiveness of our new controller.

robot will be tasked with classifying variations on each action. An action will be considered successfully learned if, given the set of test variations, the robot is able to correctly classify 80% correctly. Once we determine the success metric for each action, we will measure our overall success using the average success rate of all trained actions.

Lift a different distance	Lift starting to the side of the robot	Lift starting above the robot	Lift while holding an object
			

Figure 2: Example variations on a trained “Lift” action. A lift action is defined as a robot vertically raising its arm while not moving in any other direction. The diagram represents the trained action as the red arrow and the variation of the action as the green arrow. From left to right: lift a different distance lifts further up than in training, lift starting to the side starts the arm to the side of the robot and lifts, lift starting above has starts above the trained example, and lift while holding an object holds an object in the robot arm while lifting.

6 Expected Contribution and Future Research

We anticipate that future research in this area will be able to build off of our work by using the notion of learning variations of universal primitives to teach robots to complete complex tasks. We would like to provide the groundwork for a method of learning from demonstration that focuses on building abstractions of more basic actions and uses those abstractions to complete more complicated goals.

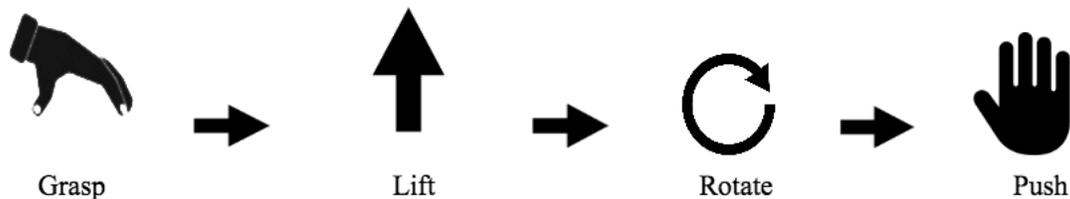


Figure 3: Simple movements can be chained together to perform complex tasks. In this drawing, we have the commands chained together to have a robot put an object on a shelf. The arm robot grasps an object, lifts it, rotates it toward a shelf, and then pushes the object onto the shelf, completing the task.

7 Conclusion

There are two elements to our project: building a better controller for the arm so that it the robot can more easily be used for research and teaching the robot, using learning from demonstration, how to identify and classify simple motor movements. We anticipate that this first step toward robots learning more complex tasks will both make the robots easier to use for people who have no prior knowledge of or experience with robotics and will allow researchers to be able to more easily teach robots complex sequences of actions.

References and Notes

1. Argall, B., Chernova, S, Veloso M, and Browning B. (2009). Survey of Robot Learning from Demonstration. *Robotics and Autonomous Systems*, 2009, 469-483.
2. Arsenio, A. (2004). Learning task sequences from scratch: Applications to the control of tools and toys by a humanoid robot. *Proceedings of the 2004 IEEE International Conference on Control Applications*, 2004, 1-6.

3. Bentivegna, S., Atkeson, C, and Cheng, G. (2004). Learning tasks from observation and practice. *Robotics and Autonomous Systems*, 2009, 163-169.