Learning from Crowdsourced Virtual Reality Demonstrations

David Whitney
Brown University
Providence, Rhode Island
david_whitney@brown.edu

Eric Rosen
Brown University
Providence, Rhode Island
eric_rosen@brown.edu

Stefanie Tellex
Brown University
Providence, Rhode Island
stefie10@cs.brown.edu

ABSTRACT

Learning from demonstration (LfD) has been a widely popular methodology for teaching robots how to perform manipulation tasks because it leverages human knowledge. However, collecting high quality demonstrations that can be used for learning robot policies can be time-consuming and difficult. Recently, some researchers have begun using consumer-grade virtual reality (VR) hardware as a more efficient means of teleoperating a robot for collecting demonstrations. Previous work in this space has focused on tasks and algorithms that require relatively little data due to the time-sink of demonstration generation. We propose a novel crowdsourcing framework that takes advantage of the large virtual reality gaming community. By treating these experienced VR users as citizen scientists, we will empower our robot with the demonstration data needed to complete complex manipulation tasks.

KEYWORDS

Human robot interaction, virtual reality, learning from demonstration, crowdsourcing

ACM Reference Format:
David Whitney, Eric Rosen, and Stefanie Tellex. 2018. Learning from Crowdsourced Virtual Reality Demonstrations. In *Proceedings of International Workshop on Virtual, Augmented and Mixed Reality for Human-Robot Interaction (VAM-HRI 2018).* ACM, New York, NY, USA, 3 pages. https://doi.org/10.475/123_4

1 INTRODUCTION

In Learning from Demonstration (LfD) a human teacher demonstrates a behavior that a robot then mimics. Sometimes, getting high quality demonstrations is easy to collect. For example, in LfD for self driving cars, the human and robot have the same action space (turning the wheel and pushing pedals), and recording the human behavior (wheel and pedal angle) is straightforward. In other tasks, however, like object manipulation with a high degree-of-freedom (DOF) robotic arm, collecting high quality demonstrations has proved more difficult. Human and robot arms do not have the same kinematics, making mappings from human motion to robot motion non-trivial. Additionally, recording human poses accurately requires tracking markers or other systems that are difficult to replicate outside of a lab setting. Compare this to the wheel angle measurement, which only requires an inexpensive encoder placed on the steering column. Finally, human demonstrators do not have the same sensors as the robot, and mapping these two observation spaces is also non-trivial.

Kinesthetic teaching avoids this problem by having a human physically move the robot in a desired motion. This solves the mapping issue, but requires demonstrators to be co-located with the robot they intent to teach. Teleoperation attempts to fix this issue by having human demonstrators control the robot through a remote interface, often a keyboard or joystick paired with a computer monitor. While this allows actual robot behaviors to be recorded, quickly and intuitively controlling a high DOF arm via a keyboard and monitor is difficult for untrained users.

The recent advent of consumer-grade virtual reality (VR) hardware has made practical VR teleoperation possible. VR immerses users in a 3D environment, making comprehension of 3D scenes easier than on a 2D monitor, and tracks user's hands, letting them specify locations and trajectories by simply moving their arms, as they would in the real world. Consumer-grade VR has become quite popular in recent years, with an estimated 400,000 HTC Vive systems sold in 2016 alone [15]. Leveraging these VR consumers as citizen scientists could alleviate the time cost of collecting demonstrations.

In this abstract, we propose that VR teleoperation will allow for crowdsourcing robotic demonstrations at scale, especially for tasks where high DOF arms are used. We will conduct an experiment by creating a virtual reality video game that supports the HTC Vive, a commercially available VR headset, and allow users to provide demonstrations for an object fetching task by having them control a virtual Baxter [14]. We will simulate the virtual environment in Unity [1] and use our publicly available ROS Reality [5] package to create our virtual robot. We will then train a deep neural network model on the demonstrations to teach the robot how to perform the task in real life.

2 RELATED WORKS

Our related works in broken into three sections: Learning from Demonstration (LfD), crowdsourcing, and virtual reality teleoperation.

2.1 Learning from Demonstration

Imitation learning, also known as Learning from Demonstration, has been a great tool for teaching robots. Leveraging human knowledge by recording demonstrations allows good initial points for quickly optimizing a robot policy [2, 3]. LfD has shown success in very varied tasks, from object grasping [6, 9] to autonomous driving [10, 11].

VAM-HRI 2018, March 2018, Chicago, Illinois USA

ACM ISBN 123-4567-24-567/08/06.
https://doi.org/10.475/123_4

2.2 Crowdsourcing

Crowdsourcing demonstration for robot manipulation has proven successful in the past [4, 13]. Penaloza et al. [13] collected demonstrations via the web to learn robot decision policies. While learning from demonstrations suffers when the number of demonstrations is low, crowdsourcing unlocks the potential to access a much larger user base to draw from [12]. Careful consideration of the interface design also allows for non-experts to provide samples [8]. Furthermore, drawing from a wide range of users accounts for biases that any particular user may bring. Clair et al. [4] crowdsourced by collecting demonstrations using a 2D web interface that controlled a real robot. [Users moved a robot through a maze]. Controlling a real robot to collect demonstrations allows the human to experience the same real world conditions the robot encounters. However, the amount of people who can perform demonstrations at any given time is limited to how many real robots are available. Other works have used virtual environments to either simulate the robot or present digital datasets to collect labeling for robot learning [7].

2.3 Virtual Reality Teleoperation

Recent state-of-the-art work (including our own) has used virtual reality (VR) to effectively teleoperate a robot for object manipulation, generally by mapping the hand pose of the user to the end-effector pose of the robot [16, 17]. These VR systems are able to accurately track hand and head poses of the user without expensive laboratory motion capture system, as well as intuitively display 3D environments to the users. Despite only tracking head and hand pose, as opposed to full body pose (as most professional/industrial motion capture systems do), Zhang et al. [17] have shown that effective LfD can still be done with consumer-grade VR hardware. However, to our knowledge, all learning from demonstration works that have used VR as the interface have required individuals to come to the lab to teleoperate a real robot, and only collected around thirty minutes of data per task. We posit that these LfD algorithms would benefit from larger datasets, and that collecting demonstrations from owners of consumer VR headsets is an untapped data source.

3 PROPOSED WORK

In brief, our proposed work is to use demonstrations of untrained users operating a virtual robot to train complex autonomous behaviors.

Currently, our group has the capability to teleoperate a real robot over the Internet. The human user sees the current robot pose and environment (as captured by the robot's wrist cameras and a Kinect 2 mounted on the robot's head), and sends back goal coordinates for the robots end effectors (by moving their controller).

For our crowdsourced game, we will need to simulate the robot entirely in Unity. We have already written a URDF parser, so we can build the transform (TF) tree of the robot, but currently we rely on ROS for solving our inverse kinematics (IK) requests and updating the TF appropriately. There are two possible solutions to this problem.

Solution 1 is to maintain a connection to simulate a Baxter in Gazebo on a server, which will process IK requests from all users and move the simulated Baxter to the calculated joint angles, and return the new TF. This has the advantage of using the actual IK solver that Baxter uses, guaranteeing similar behavior. The downside is that we will need to maintain a server to handle the IK requests, which could have scaling issues.

Solution 2 is to implement our own IK solver in C# that will run in our game locally. We would lose the guarantee that our IK solutions will be close to ones we will get at test time, but will no longer need a connection to a ROS network. Currently, this is the solution we plan to pursue.

Once we have simulated our robot, we will need to create scenes for users to demonstrate tasks in. For this, we will use Unity to build a simulated lab environment with the relevant objects near our robot. The objects will be Unity physics objects, and we will use Unity's physics engine to simulate the contact dynamics of our scene.

We will record demonstrations by users. Users will start and end recording by hitting a button, as in [17]. Recordings will be the 6 DOF pose and velocity of the user's controllers, as well as the 6 DOF pose of the relevant game objects in the scene. This will allow us to reconstruct the demonstration after the fact. The demonstrations will be sent over the Internet to an anonymized database, hosted on AWS or Google Cloud.

For our initial round of learning, we will use a convolution neural network similar in architecture to [17]. Their network takes as input an RGBD image, which we will need to simulate from our Unity demonstrations. Fortunately, this can be done by placing a virtual camera in the scene and recording its color image and depth mask.

At first, we will use their network and attempt to replicate their results using the demonstrations we record ourselves, as an efficacy test of our system. Second, we will release our game on the Internet, and attempt to learn manipulation behaviors from crowdsourced data. Zhang et al. [17] stated they could collect roughly 200 demonstrations in half an hour. When we posted a video of our previous work to the Vive subreddit, we received 101 upvotes and 32 comments, many asking to try our system. If even half of those users used our system, we would collect demonstrations at the rate of 20,000 per hour (50 users $\times$ 400 demonstrations per hour).

One issue with crowdsourced data is quality control. We will need mechanisms to filter our bad (or even malicious) demonstrations. Currently, we plan to filter our demonstrations that start or end too far from the relevant game objects, or that are too short or long compared to the mean or median demonstration length.

We will evaluate the effectiveness of our algorithm on a real Baxter robot. We will attempt each task a fixed number of times, and record our completion rate.

4 CONCLUSION

Consumer-grade VR hardware is now a reality. It makes interacting with and controlling 3D virtual robots intuitive enough for non-experts. We believe using these VR interfaces for robot learning holds promise for solving many complex tasks.

REFERENCES

[1] [n. d.]. Unity. ([n. d.]). https://unity3d.com/
[2] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and autonomous systems* 57, 5 (2009), 469–483.
[3] Aude Billard, Sylvain Calinon, Ruediger Dillmann, and Stefan Schaal. 2008. Robot programming by demonstration. In *Springer handbook of robotics*. Springer,

1371–1394.
[4] Aaron St Clair, Carl Saldanha, Adrian Boteanu, and Sonia Chernova. 2016. Interactive hierarchical task learning via crowdsourcing for robot adaptability. In *Refereed workshop Planning for Human-Robot Interaction: Shared Autonomy and Collaborative Robotics at Robotics: Science and Systems, Ann Arbor, Michigan. RSS.*
[5] Brown University H2R Lab. 2017. ROS Reality. (2017). https://github.com/h2r/ros_reality
[6] Micha Hersch, Florent Guenter, Sylvain Calinon, and Aude Billard. 2008. Dynamical system modulation for robot learning via kinesthetic demonstrations. *IEEE Transactions on Robotics* 24, 6 (2008), 1463–1467.
[7] Kaijen Hsiao and Tomas Lozano-Perez. 2006. Imitation learning of whole-body grasps. In *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on.* IEEE, 5657–5662.
[8] Ashesh Jain, Debarghya Das, Jayesh K Gupta, and Ashutosh Saxena. 2015. Planit: A crowdsourcing approach for learning to plan paths from large scale preference feedback. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on.* IEEE, 877–884.
[9] David Kent, Morteza Behrooz, and Sonia Chernova. 2014. Crowdsourcing the construction of a 3d object recognition database for robotic grasping. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on.* IEEE, 4526–4531.
[10] Markus Kuderer, Shilpa Gulati, and Wolfram Burgard. 2015. Learning driving styles for autonomous vehicles from demonstration. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on.* IEEE, 2641–2646.
[11] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436.
[12] Sarah Osentoski, Christopher Crick, Grayin Jay, and Odest Chadwicke Jenkins. 2010. Crowdsourcing for closed loop control. In *Proc. of the NIPS Workshop on Computational Social Science and the Wisdom of Crowds, NIPS.*
[13] Christian I Penaloza, Sonia Chernova, Yasushi Mae, and Tatsuo Arai. 2013. Robot reinforcement learning using crowdsourced rewards. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).*
[14] Rethink Robotics. [n. d.]. Baxter Wiki. ([n. d.]). http://sdk.rethinkrobotics.com/wiki/Main_Page
[15] SuperData Research. 2017. *Virtual Reality Consumers.* Technical Report. SuperData Research. https://superdata-research.myshopify.com/products/report-virtual-reality-consumers?variant=36950222286
[16] David Whitney, Eric Rosen, Elizabeth Phillips, George Konidaris, and Stefanie Tellex. [n. d.]. Comparing Robot Grasping Teleoperation across Desktop and Virtual Reality with ROS Reality. ([n. d.]).
[17] Tianhao Zhang, Zoe McCarthy, Owen Jow, Dennis Lee, Ken Goldberg, and Pieter Abbeel. 2017. Deep Imitation Learning for Complex Manipulation Tasks from Virtual Reality Teleoperation. *arXiv preprint arXiv:1710.04615* (2017).