Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Pick-and-place : Learning from virtual demonstration

by Matthew Ng Cher-Wai

Today’s Seminar

1. Introduction
   - What is VR?
   - What is Learning from Demonstration (LfD)?
2. Common limitations of LfD
3. VR Teleoperation (Proof)
4. Results of VR Teleoperation
5. Virtual to physical results
6. Conclusion

What is VR?

https://www.youtube.com/watch?v=1SlZvuhABGk&t=35s

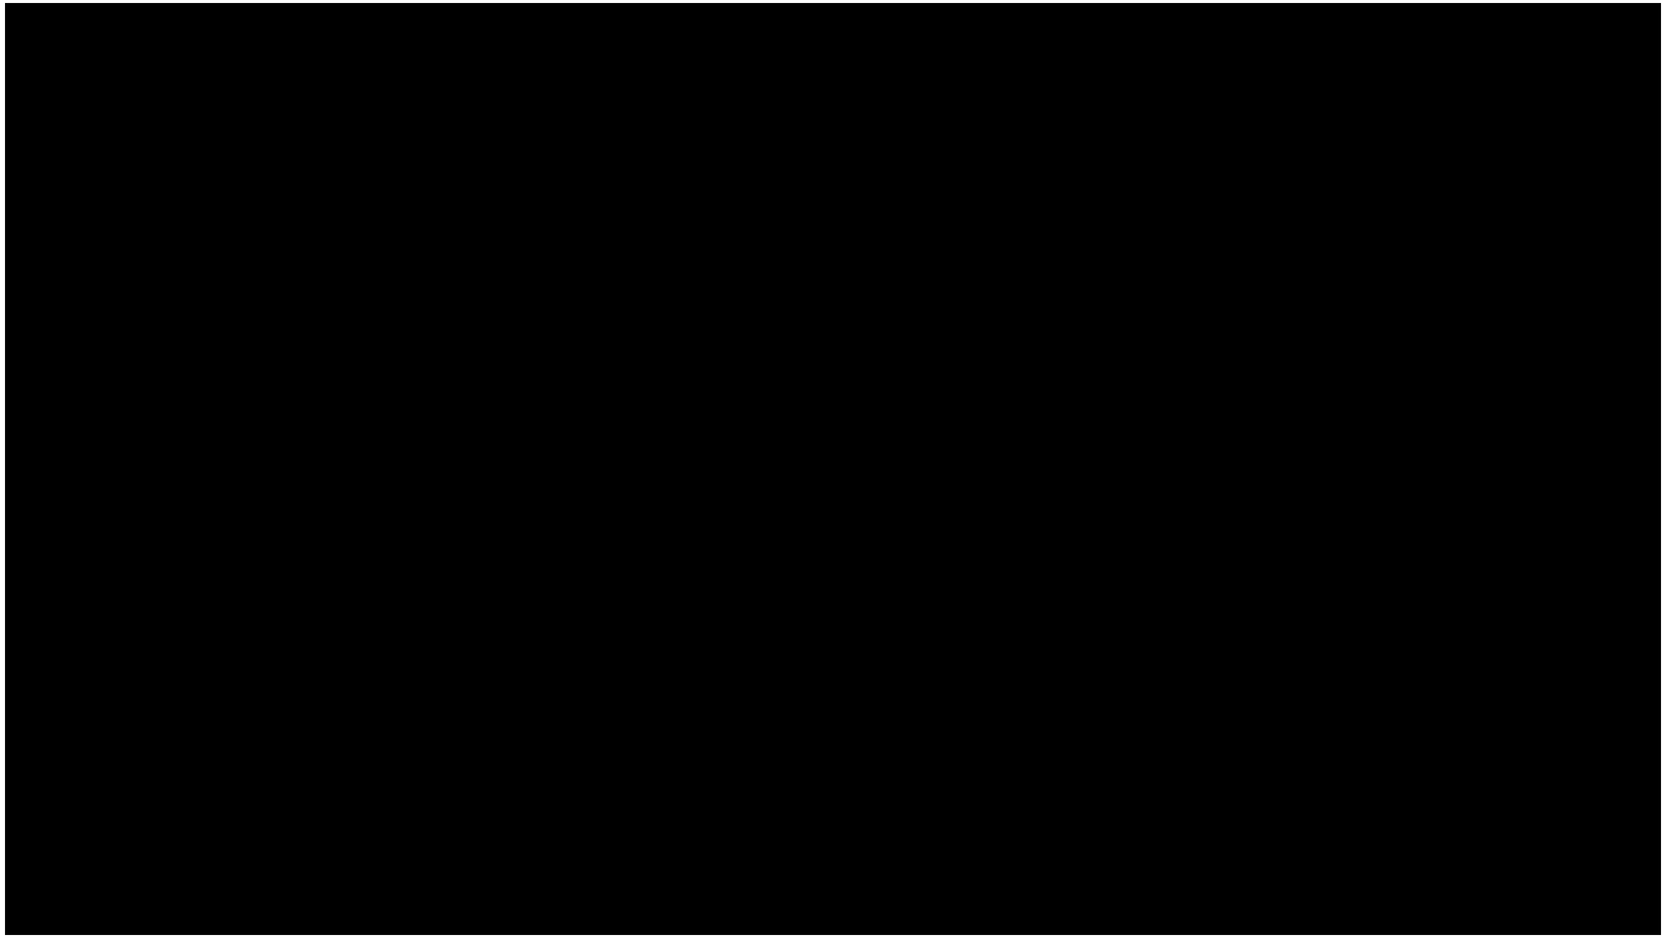

Matthew Ng (Pick-and-place : Learning from VR Demonstrations)

What is VR?

https://www.youtube.com/watch?v=bv7I8nMV914&t=19s

Matthew Ng (Pick-and-place : Learning from VR Demonstrations)

What is VR?

[7]

- Simulated environmental experience
- Headsets, sensors and controllers
- User is able to move, act and perform tasks within the virtual space
- Eg. Google Cardboard, HTC Vive, Virtuix Omni treadmill

What is Learning from Demonstration (LfD)?

- A method of teaching robots new tasks
- Does not utilize programming
- Allows for intuitive programming in more novel situations

Task at hand

- Two main papers in question today
- One to prove VR as a viable tool for learning from demonstration[2]
- Second to show how publicly sourced data can be used to train an intelligent robot[1]
- Pick-and-place, a general, all-purpose task with many applications

Common limitations of LfD

- Different action space[1]
- Must be on-site with demonstrators who are familiar with the robot
- Teleoperation – Done with keyboard and other input devices, requires robots to operate.
- Time consuming[1]

Solutions

- Teleoperation using Unity3D generated VR as the input

- Crowdsourcing for increased data sets and demonstrations

[2]

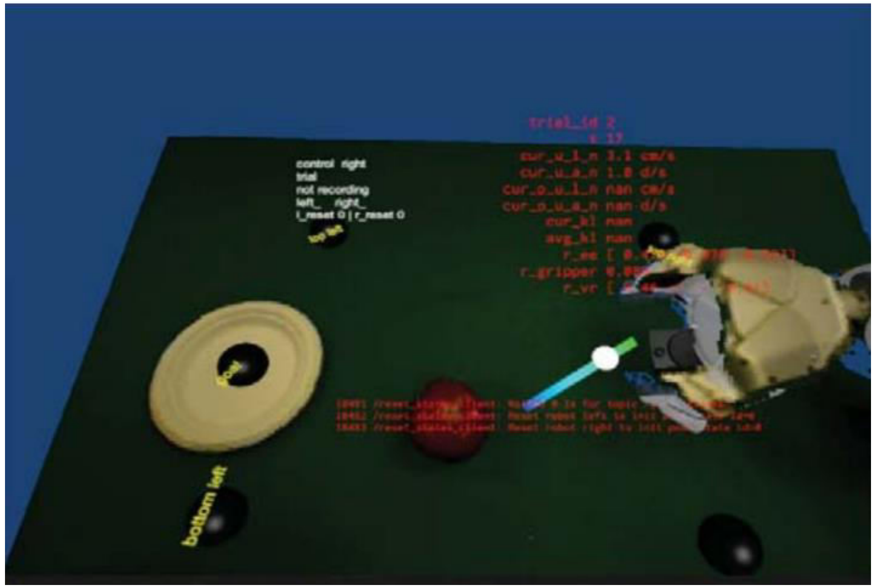

[2]

Matthew Ng (Pick-and-place : Learning from VR Demonstrations)

VR Teleoperation Test - Setup

- University of California, Berkeley [2]
- Proving that Learning by VR teleoperation could work.
- Using HTC Vive VR system with PR2 robot

VR Teleoperation Test - Parameters

- Object localization
- High-precision control
- Handling contact
- Multi stage tasks (e.g. Place a toy into a bowl then push the bowl)

VR Teleoperation Test - Results

Task	Reaching	Grasping	Pushing	Plane	Cube	Nail	Grasp & Place	Grasp – Drop - Push	Cloth
Test	91.6%	97.2%	98.9%	87.5%	85.7%	87.5%	96.0%	**83.3%**	97.4%
Demo time (min)	13.7	11.1	16.9	**25.0**	12.7	13.6	12.3	14.5	10.1
Avg Length	41	37	58	47	37	38	68	87	60
#demo	200	180	175	319	206	215	109	100	100

Matthew Ng (Pick-and-place : Learning from VR Demonstrations)

VR Teleoperation Test - Evaluation

- Obtained good success rates (83.3 – 98.9%) with <30 minutes of demo time
- Achieves tractable sample efficiency
- The simple imitation learning algorithm can train successful control policies for a range of real-world manipulation tasks

Proposed VR Solution

- Researchers from Brown University, Rhode island
- Improvement upon teleoperation in California paper
- Using VR simulation as the data collection method[1]
- Crowdsourcing VR application for faster data collection[1]

Solution Overview

- Perspective of robot taken from wrist cameras and Kinect 2 on head
- Virtual representation of Baxter Robot to be the “Avatar”
- Public user data will be recorded and stored as data on a AWS or Google cloud.

Solution Overview

- Recordings will consist of 6 DOF poses and velocity of VR controllers.
- Recordings will be used to train a convolution neural network.[2]

Solution Option 1 – (ROS)

- Hosting the simulation of Baxter on Gazebo[1]

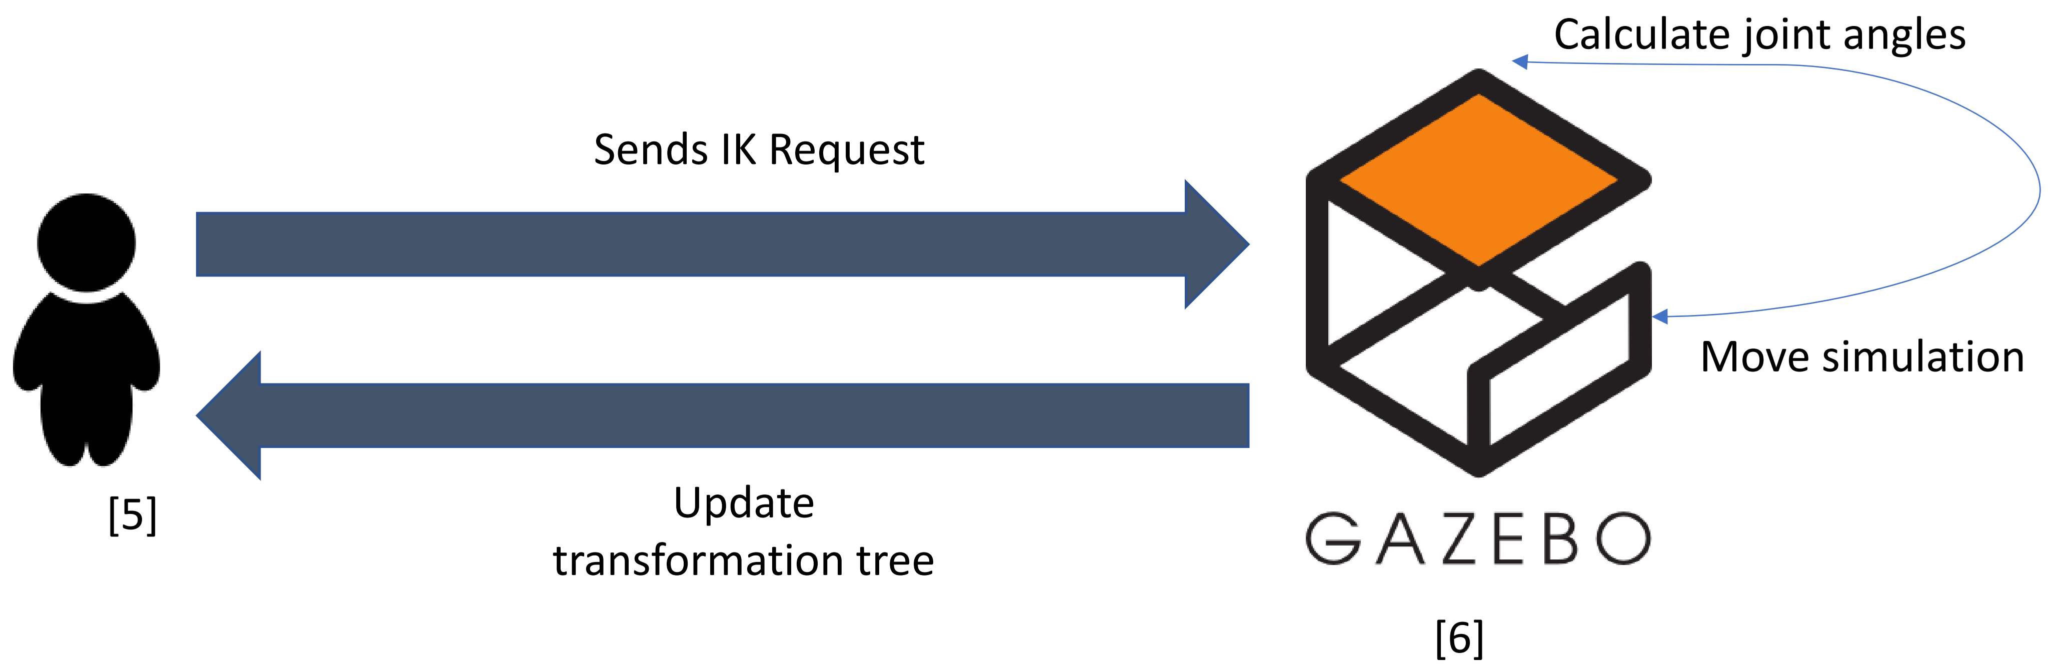

Solution Option 1 – (ROS)

- Guarantees high accuracy due to usage of the Inverse Kinematics (IK) solver which Baxter has. No mismatch in compatibility.

- Requires an ROS Server to be active to handle IK requests.

Solution Option 2 – (Homebrew)

- Homebrewed IK solver in C#[1]
- No longer requires constant server connection as IK solver is in game.
- Actual IK solver within Baxter will be slightly different, thus losing some degrees of accuracy.

Conversion of results from virtual to physical

- Using the recordings, extracting information for each demonstration will be possible.
- Input for CNN in [2] is a RGB-D image. We can obtain such from the Unity3D simulation.
- Using a virtual camera to record color image and depth mask.

Problems in Solution (?)

- People on the internet aren’t very pleasant
- Users might make malicious demonstrations
- Public might not be interested

Problems in Solution (?)

- People aside, proof of VR input as learning method is on a real robot.
- Reading in the RGB-D images from a simulation is still unproven.
- Might not have the correct accuracy as the real-life sensors.

Problems in Solution (?)

- When the video of the researchers previous work to the Vive subreddit[1]

- The post received 101 upvotes and 32 comments[1]

- Many of aforementioned comments were to try out their system

- Assume half have the time to participate (still have ~ 50 testers)

Conclusion

- Teleoperation is proven to be sufficiently effective.
- Solid outline of how to convert crowdsourced data into workable information for a CNN to learn
- Public response decent enough to collect a sizeable sample for demonstrations
- Untested, but promising

Questions?

References

- [1] ‘’Learning from crowdsourced virtual reality demonstrations’’, published in the Proceedings of the 1st International Workshop on Virtual, Augmented, and Mixed Reality for HRI At: Chicago, IL, USA in March 2018 by Eric Rosen, David Whitney and Stefanie Tellex
- [2] “Deep Imitation Learning for Complex Manipulation Tasks from Virtual Reality Teleoperation’’ published in the 2018 IEEE International Conference on Robotics and Automation (ICRA) May 21-25, 2018, Brisbane, Australia by Tianhao Zhang, Zoe McCarthy, Owen Jow, Dennis Lee, Xi Chen, Ken Goldberg and Pieter Abbeel.
- [3]Vectorstock.com
- [4]ROS.org
- [5]3dwarehouse.com
- [6]Gazebosim.org
- [7] https://www.mysn.de/vr-headsets/vive-pro-full-kit