

# THESIS DEFENSE

## Data Centric Robot Learning

Friday, August 16, 2019  
4305 Newell Simon Hall  
4:00 p.m.

### Thesis Committee:



**Abhinav  
Gupta**  
Chair



**Christopher  
Atkeson**



**Martial  
Hebert**



**Siddhartha  
Srinivasa**  
University of  
Washington



**Pieter  
Abbeel**  
University of  
California, Berkeley



## Lerrel Joseph Pinto

### Abstract

While robotics has made tremendous progress over the last few decades, most success stories are still limited to carefully engineered and precisely modeled environments. Getting these robots to work in the complex and diverse world that we live in has proven to be a difficult challenge. Interestingly, one of the most significant

successes in the last decade of AI has been the use of Machine Learning (ML) to generalize about and robustly handle diverse situations. From detecting objects in cluttered visual scenes to translating sentences between languages, ML has become an integral part of modern technology. So why don't we just apply current learning algorithms to robots? One of the biggest reason is the lack of large-scale data. A key ingredient that sparked progress in Computer Vision and Natural Language Processing was internet-scale data: 1-10 million labelled images for computer vision; 100 million word-pairs for language translation. Unfortunately, this scale of data is not available for robots. Hence, effectively integrating learning to robots requires us to rethink our use of robotic data.

In this thesis we draw inspiration from other data driven fields of AI (Computer Vision, Natural Language Processing, etc.) to develop data-centric robot learning techniques. However, unlike the aforementioned fields, robotics involves interactions with real hardware systems, which presents three unique challenges for large-scale learning. The first challenge is the physical nature of robotics where every piece of data needs to be executed on a real system. To scalably collect large amounts of data with minimal human supervision, we present 'self-supervised' robot learning techniques where the robot both collects and labels real-world data. From manipulating unseen objects in new homes to avoiding previously unseen obstacles while flying, we demonstrate that large-scale, self-supervised data, when combined with off-the-shelf ML tools, can produce generalizable robotic skills.

The second challenge is that real robots are slow, which limits the amount of data we can collect even with 'self-supervised' techniques. This limitation necessitates the development of algorithms that can efficiently utilize available robot data. We do this by instilling notions of robustness via adversaries, sharing representations across multiple tasks, and using curricula to efficiently search high-dimensional action spaces.

Finally, in several applications, real robot data may not be easy to collect. In such scenarios, we can turn to simulated models of the real world. Towards this, we present algorithms to learn transferable and generalizable skills in simulations.