CS391R: Robot Learning

Perception, Decision Making, and General-Purpose Autonomy

Prof. Yuke Zhu

Fall 2020

Robotics and COVID-19

Photos from the Internet

Today's Agenda

- Overview of general-purpose robot autonomy
- Why studying Robot Learning now?
- Research topics of Robot Learning
- Course content and logistics
- Student introduction

Special-Purpose Robot Automation

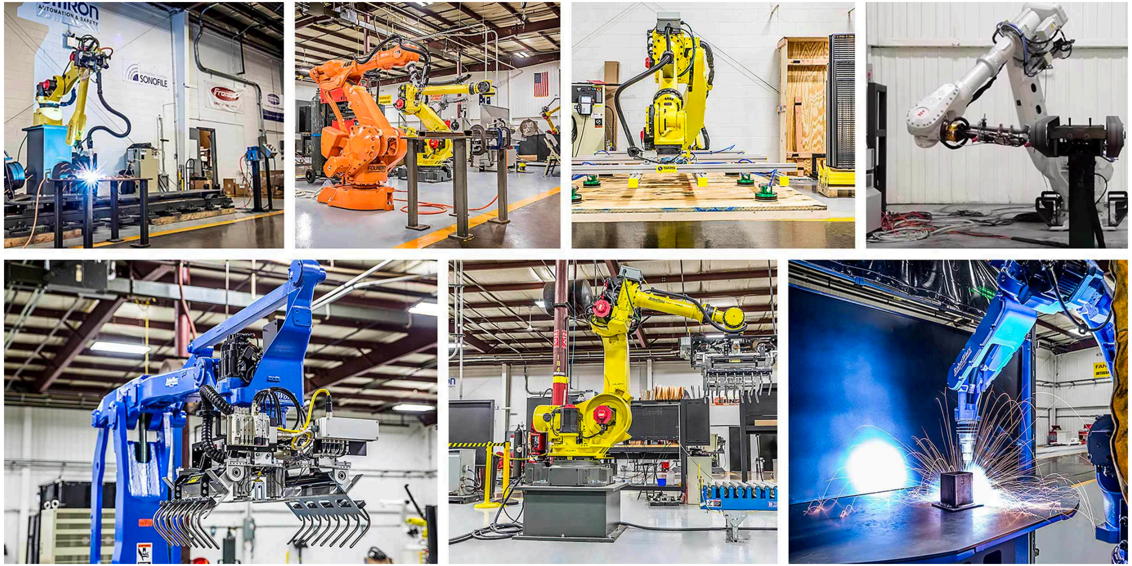

Structured Environments

Fixed Set of Tasks

Pre-programmed Procedures

General-Purpose Robot Autonomy … in the Wild

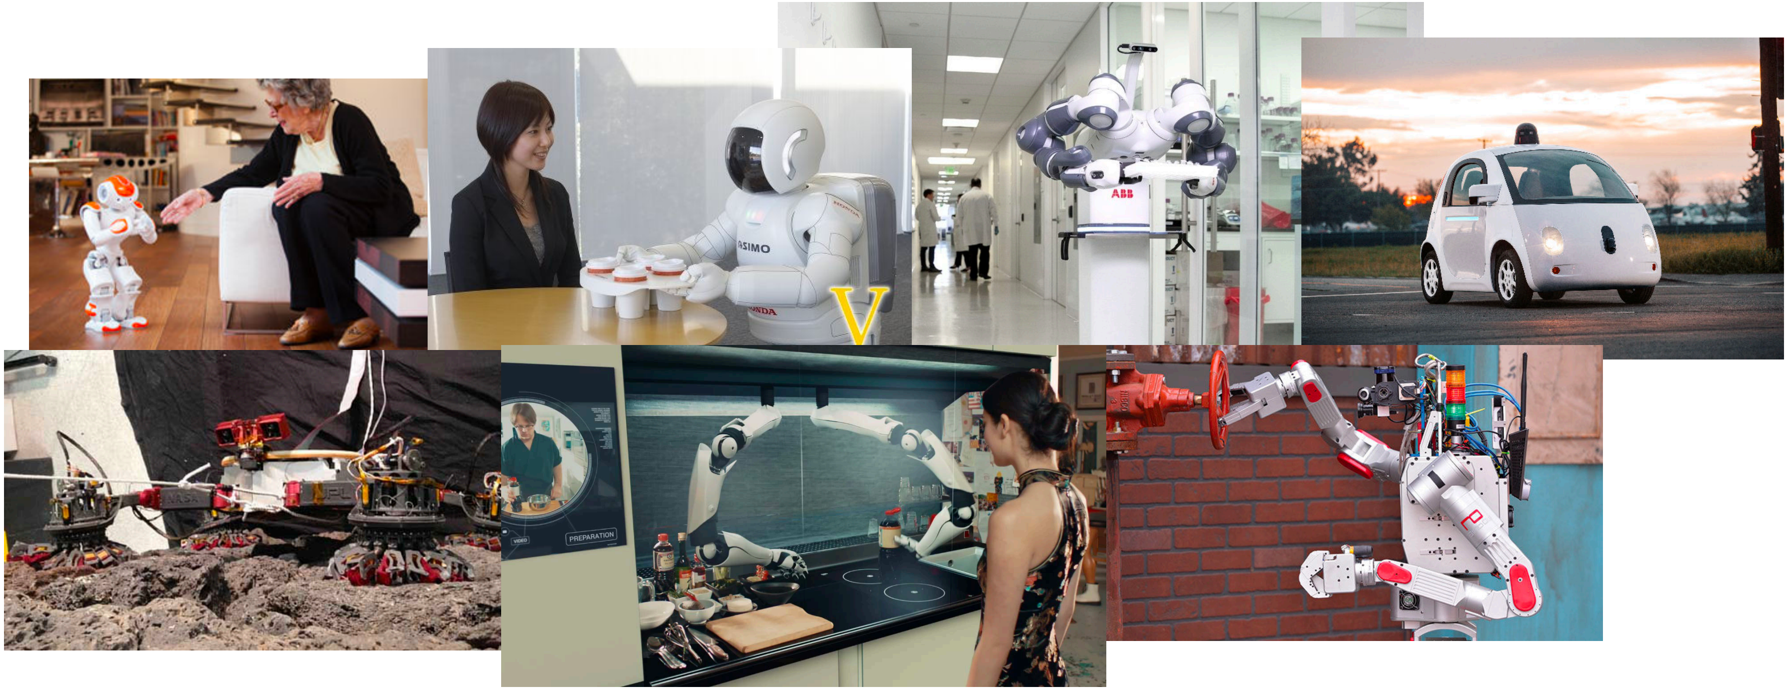

Unstructured Environments

Ever-changing Tasks

Human Involvement

Special-Purpose Robot Automation

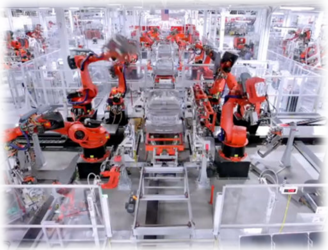

custom-built
robots

human expert
programming

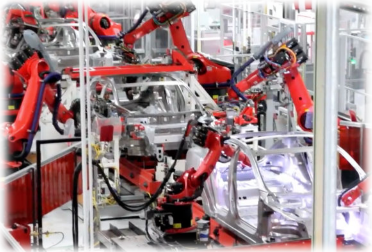

special-purpose
behaviors

General-Purpose Robot Autonomy

general-purpose
robots

?

general-purpose
behaviors

Special-Purpose Robot Automation

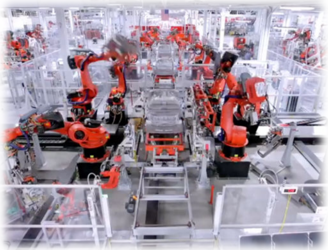

custom-built
robots

human expert
programming

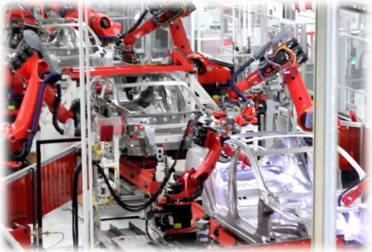

special-purpose
behaviors

General-Purpose Robot Autonomy

general-purpose
robots

Robot Learning

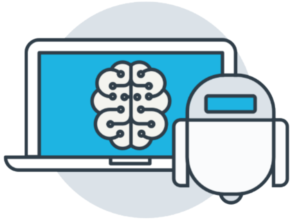

general-purpose
behaviors

General-Purpose Robot Autonomy: Imaginations

Unimate - The First Industrial Robot

British TV (1968)

General-Purpose Robot Autonomy: Challenges

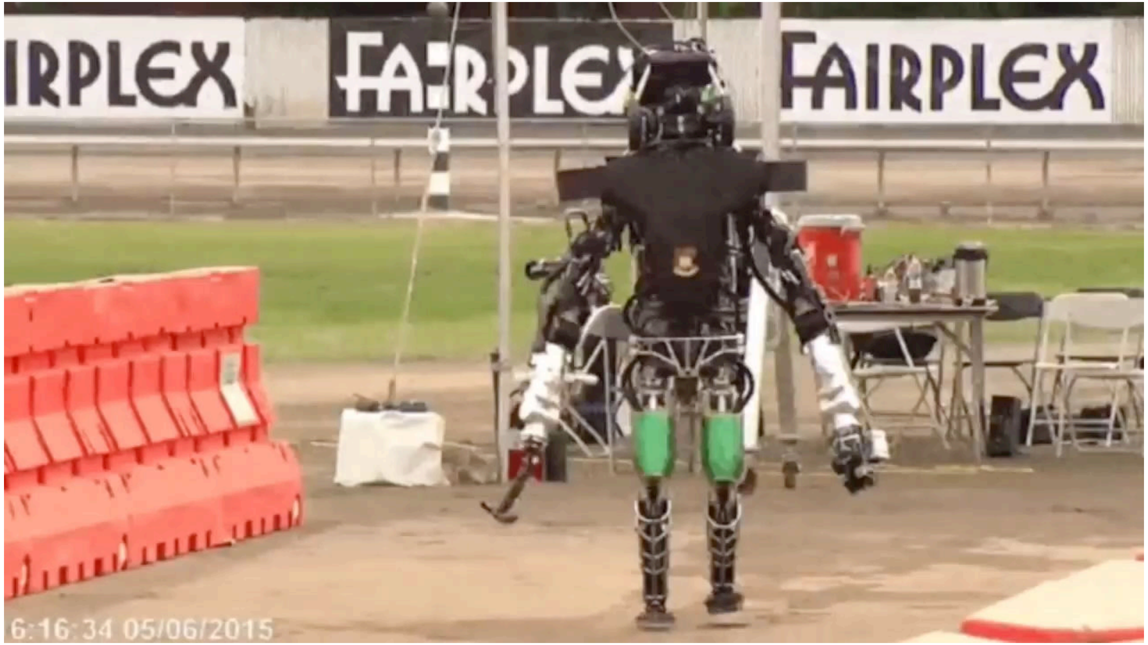

DARPA Robotics Challenge

(2015)

"The Moravec's paradox"

General-Purpose Robot Autonomy: Progress

We will learn the algorithms and techniques behind the latest progress.

Grasping (DexNet 4.0; 2019)

Locomotion (ANYmal; 2019)

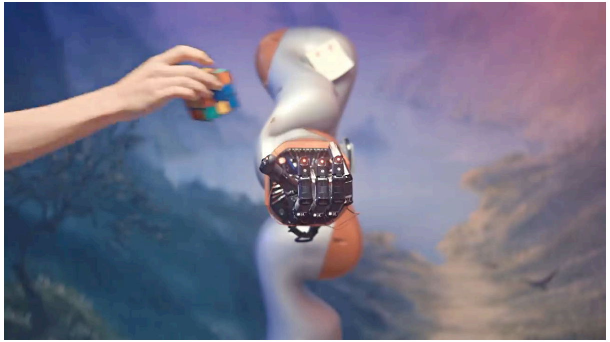

Manipulation (OpenAI; 2019)

What Makes Robot Learning Special?

Robots are **physically embodied** and **environmentally situated**.

[Sa et al. IROS 2014]

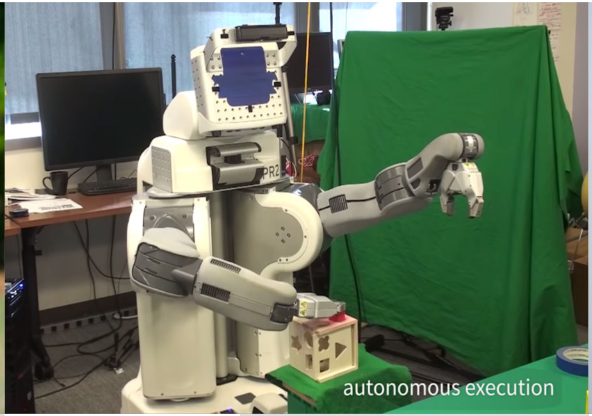

[Levine et al. JMLR 2016]

[Bohg et al. ICRA 2018]

What Makes Robot Learning Special?

Robots are **physically embodied** and **environmentally situated**.

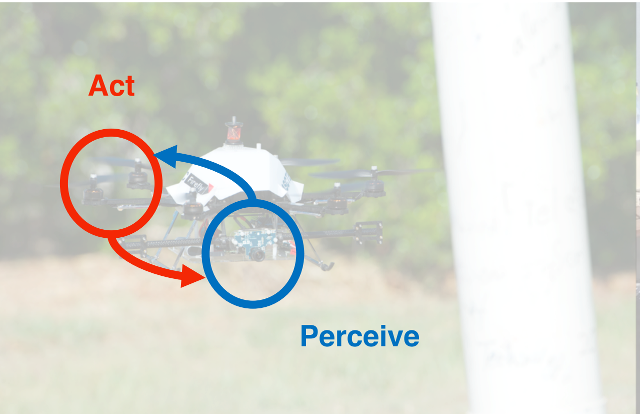

[Sa et al. IROS 2014]

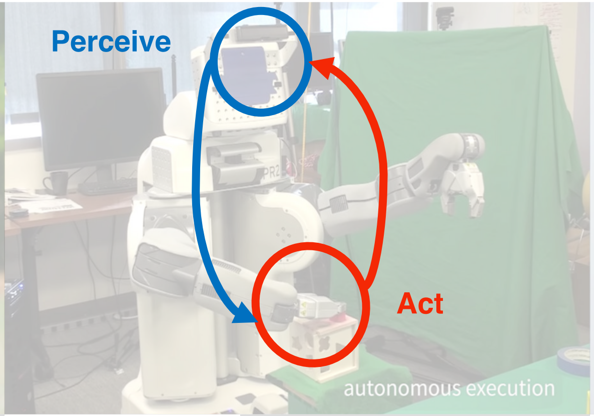

[Levine et al. JMLR 2016]

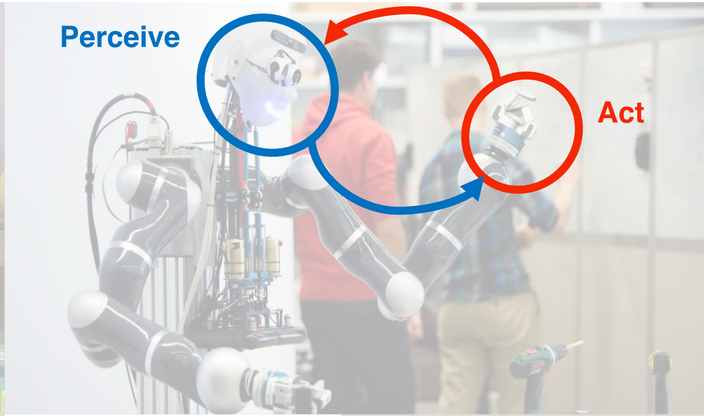

[Bohg et al. ICRA 2018]

A key challenge in Robot Learning is to close the perception-action loop.

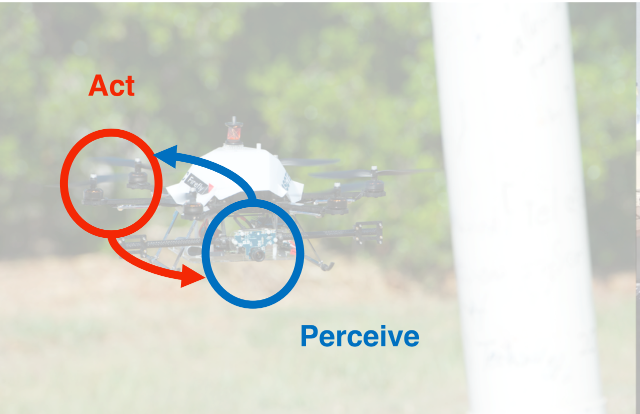

[Sa et al. IROS 2014]

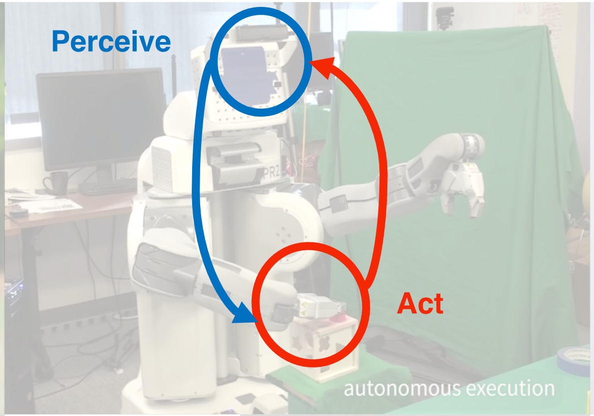

[Levine et al. JMLR 2016]

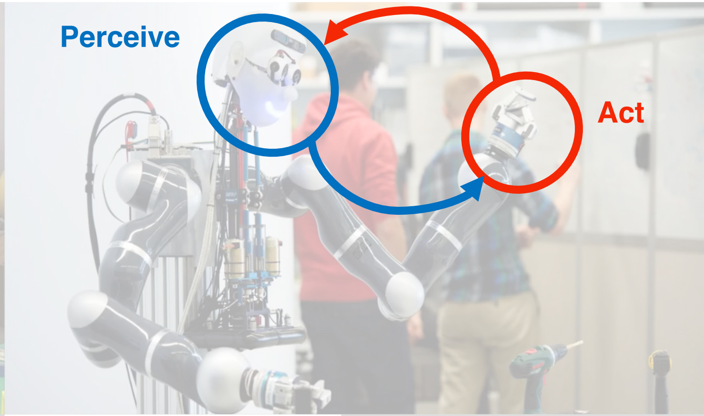

[Bohg et al. ICRA 2018]

Now is the best time to study and work on Robot Learning.

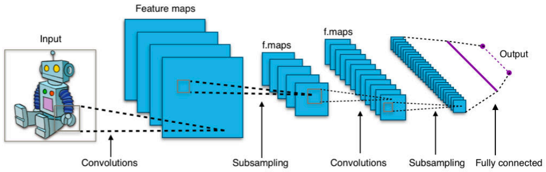

Artificial Intelligence

Recent breakthroughs in machine learning and computer vision, e.g., deep learning (Turing awards 2018)

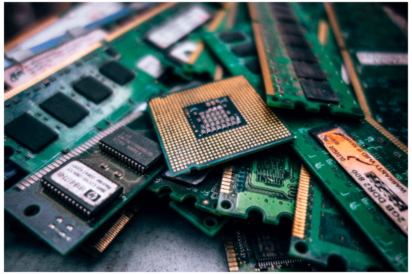

Computing Power

Your smartphone is millions of times more powerful than all of NASA's combined computing in 1969.

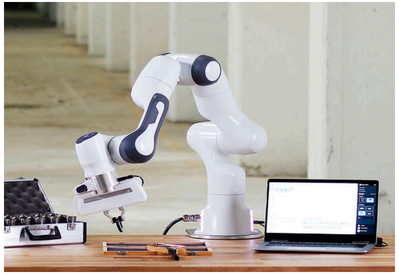

Robot Hardware

More reliable and affordable cobot hardware that costs around annual salary of American workers

Now is the best time to study and work on Robot Learning.

Positive and negative **societal impacts** of robot learning research is an important part of our in-class discussions.

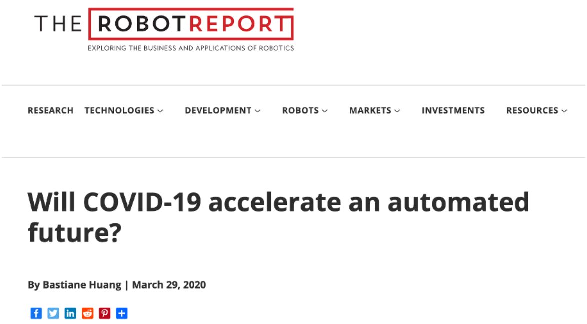

THE ROBOTREPORT
EXPLORING THE BUSINESS AND APPLICATIONS OF ROBOTICS

RESEARCH TECHNOLOGIES DEVELOPMENT ROBOTS MARKETS INVESTMENTS RESOURCES

Will COVID-19 accelerate an automated future?

By Bastiane Huang | March 29, 2020

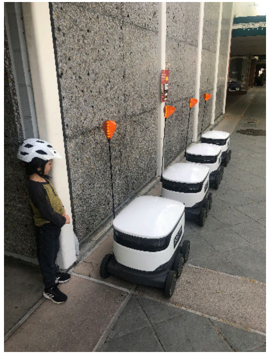

BBC Sign in News Sport Reel Worklife Travel Future Mor

NEWS

Home US Election Coronavirus Video World US & Canada UK Business Tech

Technology

Coronavirus: Will Covid-19 speed up the use of robots to replace human workers?

By Zoe Thomas
Technology reporter

19 April 2020

Share

Coronavirus pandemic

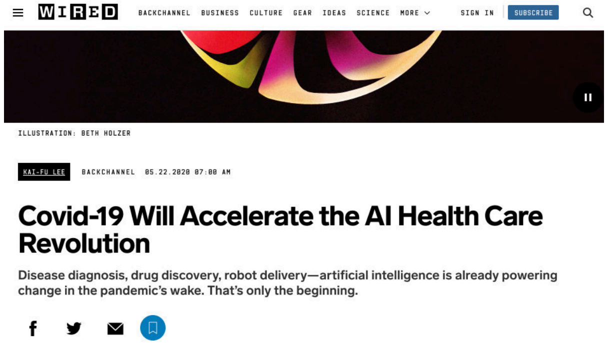

WIRED BACKCHANNEL BUSINESS CULTURE GEAR IDEAS SCIENCE MORE SIGN IN SUBSCRIBE

ILLUSTRATION: BETH HOLZER

KAI-FU LEE BACKCHANNEL 05.22.2020 07:00 AM

Covid-19 Will Accelerate the AI Health Care Revolution

Disease diagnosis, drug discovery, robot delivery—artificial intelligence is already powering change in the pandemic's wake. That's only the beginning.

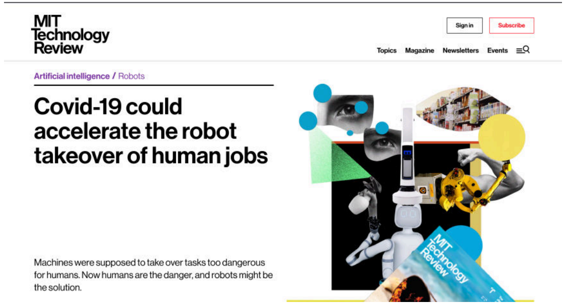

MIT Technology Review

Sign in Subscribe

Topics Magazine Newsletters Events

Artificial intelligence / Robots

Covid-19 could accelerate the robot takeover of human jobs

Machines were supposed to take over tasks too dangerous for humans. Now humans are the danger, and robots might be the solution.

https://www.therobotreport.com/tag/coronavirus/

Robot Learning as a Growing Research Community

Conference on Robot Learning is **3** years old.

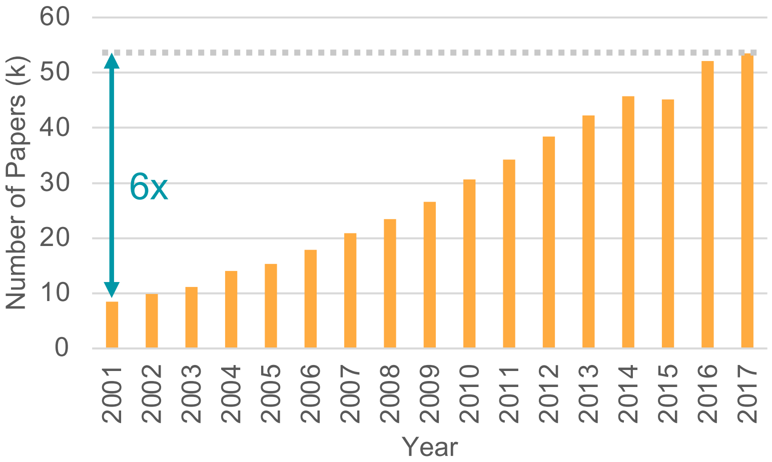

Growth of "**Robot Learning**" Publications

[Source: Google Scholar]

When **NOT** to Make Robots Learn?

Learning is not for every problem in robotics.

Harnessing the priors and structures of a problem goes a long way…

Learning is most effective when used in conjunction with modeling.

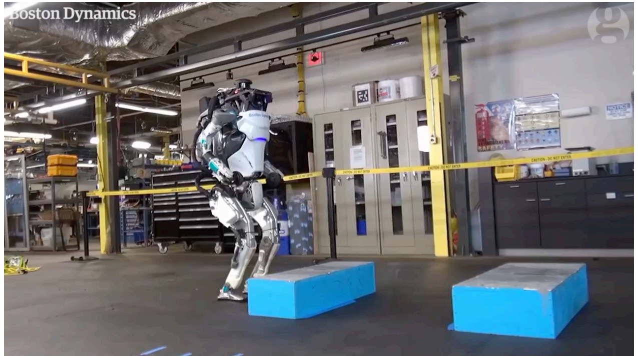

When to Make Robots Learn?

Learning is critical for getting robots to work in the real world.

object variation

environment uncertainty

adaptation

You learn **CS391R: Robot Learning** so that

Robots learn **faster** and **better**.

Key Ingredients of General-Purpose Robot Autonomy

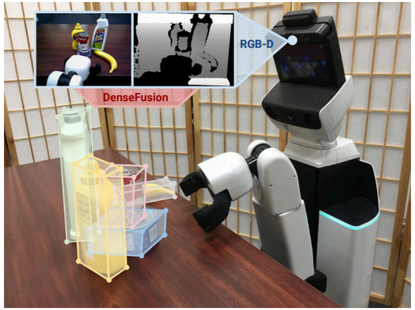

Perception

seeing and understanding 3D environments

Decision Making

planning and control for long-term interactions

Intelligence

cognitive reasoning & fast adaptation to new situations

Real-World Systems

physical embodiment & hardware constraints

Course Content

We review the Robot Learning literature in these topics.

Part I

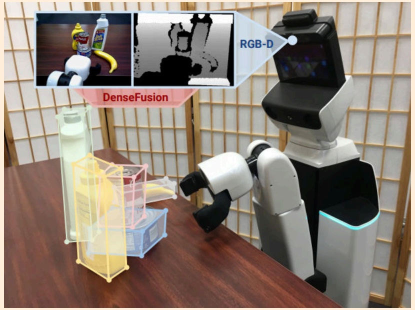

Perception

seeing and understanding 3D environments

Part II

Decision Making

planning and control for long-term interactions

Part III

Intelligence

cognitive reasoning & fast adaptation to new situations

Part IV

Real-World Systems

physical embodiment & hardware constraints

Prerequisite: coursework / experience in AI and Machine Learning

Course Content

Part I

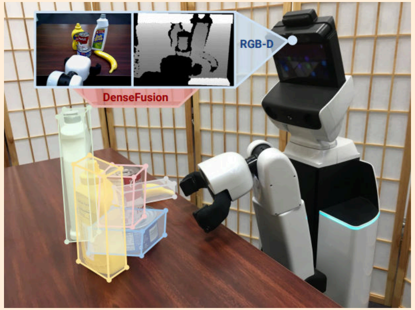

Perception

seeing and understanding 3D environments

Part II

Decision Making

planning and control for long-term interactions

Part III

Intelligence

cognitive reasoning & fast adaptation to new situations

Part IV

Real-World Systems

physical embodiment & hardware constraints

Course Content: Perception

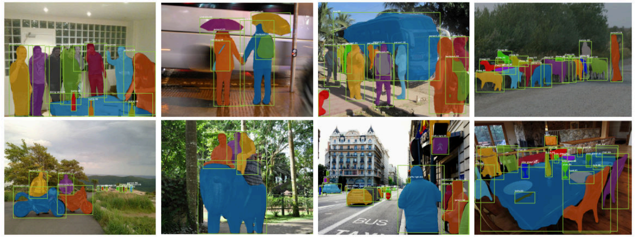

object detection

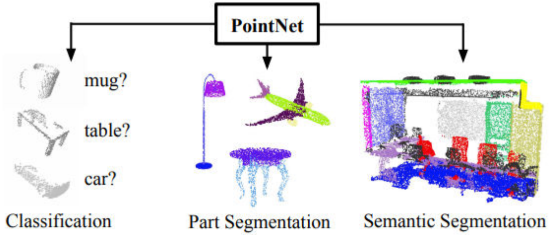

3d point cloud

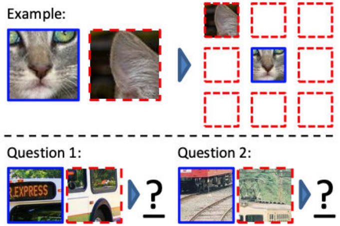

unsupervised visual learning

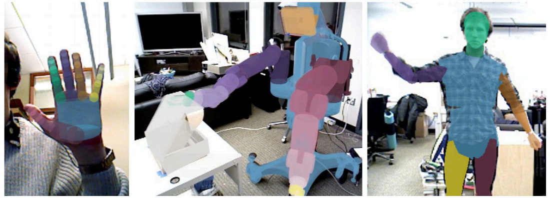

visual tracking

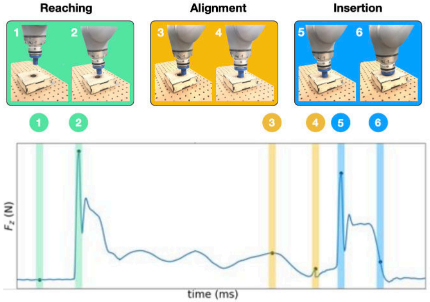

multimodal understanding

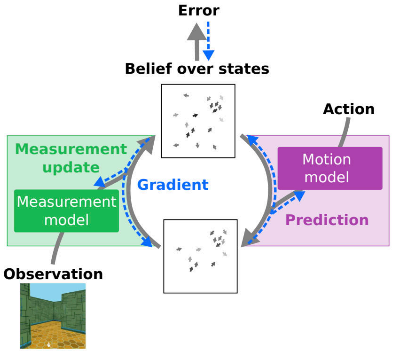

recursive state estimation

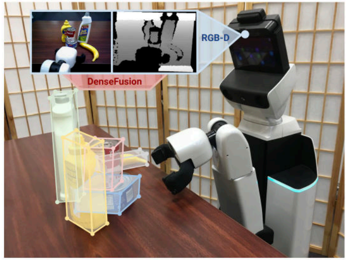

pose estimation

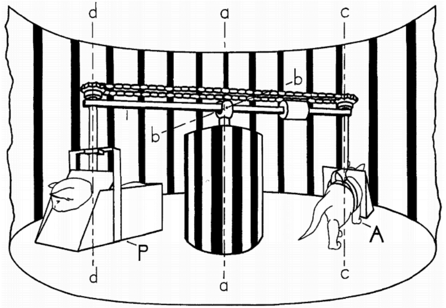

interactive perception

Course Content

Part I

Perception

seeing and understanding 3D environments

Part II

Decision Making

planning and control for long-term interactions

Part III

Intelligence

cognitive reasoning & fast adaptation to new situations

Part IV

Real-World Systems

physical embodiment & hardware constraints

Course Content: Decision Making

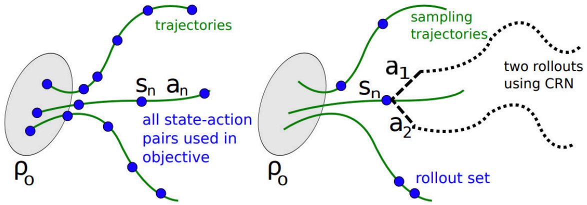

model-free reinforcement learning

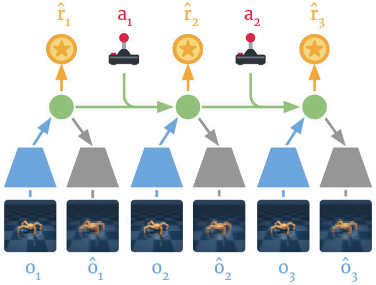

model-based reinforcement learning

imitation as supervised learning

inverse reinforcement learning

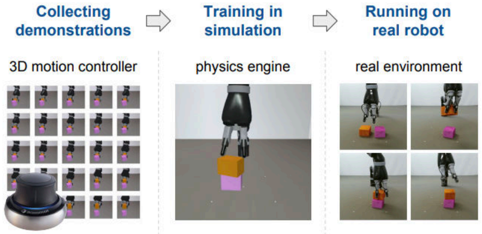

adversarial imitation learning

Course Content

Part I

Perception

seeing and understanding 3D environments

Part II

Decision Making

planning and control for long-term interactions

Part III

Intelligence

cognitive reasoning & fast adaptation to new situations

Part IV

Real-World System

physical embodiment & hardware constraints

Course Content: Intelligence

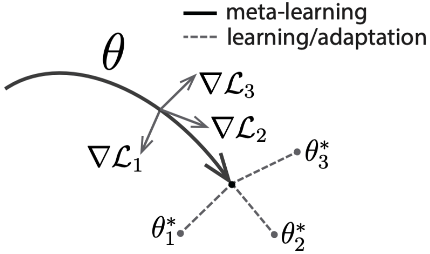

learning to learn:
meta-learning

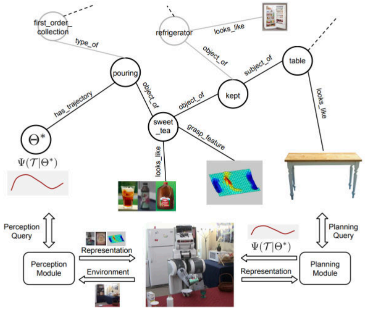

learning to learn:
lifelong learning

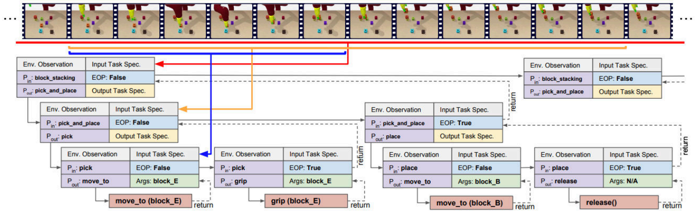

compositionality: hierarchy

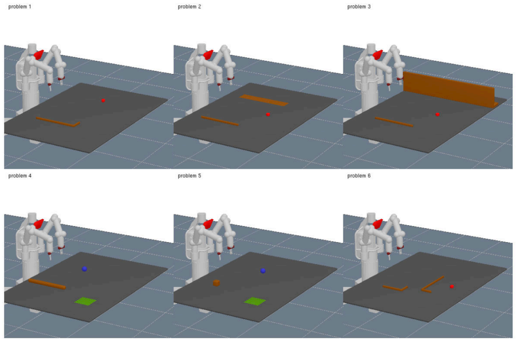

compositionality: task and motion

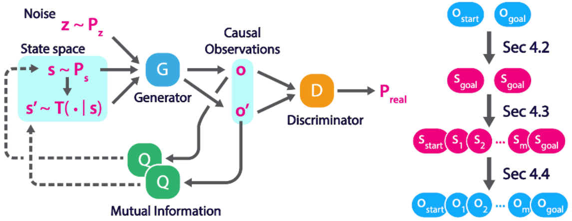

(a) Causal InfoGAN model

(b) Planning paradigm

causal reasoning

Course Content

Part I

Perception

seeing and understanding 3D environments

Part II

Decision Making

planning and control for long-term interactions

Part III

Intelligence

cognitive reasoning & fast adapting to new situations

Part IV

Real-World Systems

physical embodiment & hardware constraints

Course Content: Systems

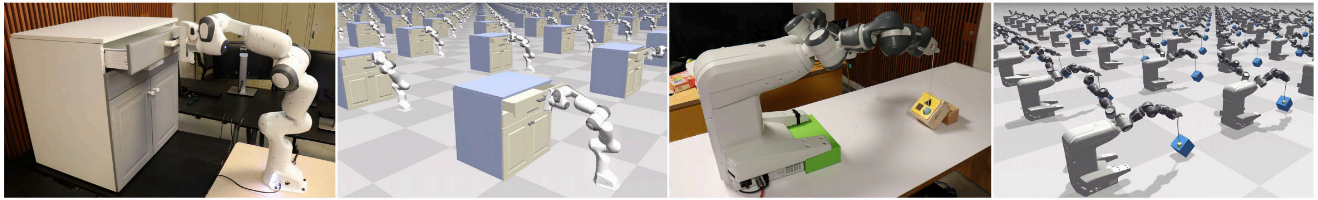

simulation-reality gap

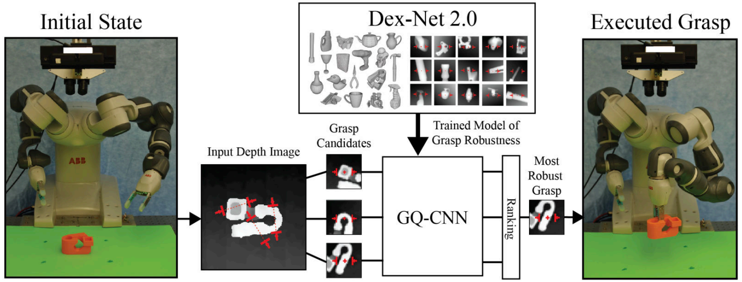

data-driven robotic grasping

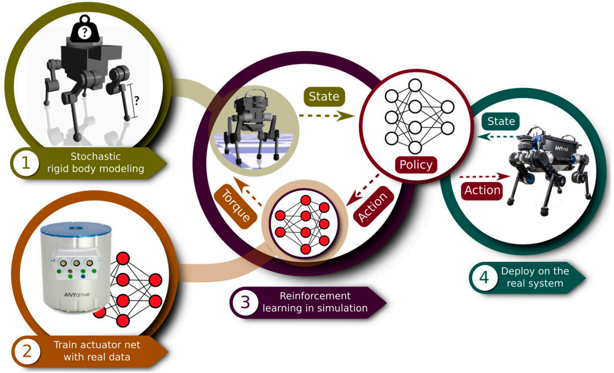

building robotic systems

Learning Objectives

- understand the potential and societal impact of general-purpose robot autonomy in the real world, the technical challenges arising from building it, and the role of machine learning and AI in addressing these challenges;
- get familiar with a variety of model-driven and data-driven principles and algorithms on robot perception and decision making;
- be able to evaluate, communicate, and apply advanced AI-based techniques to robotics problems.

… through literature reviews, research presentations, and course projects

Learning Objectives

Get a taste of Robot Learning research in the full circle

Logistics

Lectures

Time: 2:00-3:30pm CT, Tuesdays and Thursdays

Location: Online (Zoom links on Canvas)

Office Hours

Instructor: TBA next week

TA: TBA next week

Fill in the time zone survey on Canvas!

https://www.cs.utexas.edu/~yukez/cs391r_fall2020/

Logistics

	Part I: Robot Perception	
Week 2 Tue, Sept 1	Lecture Overview of Robot Perception ▪ The Limits and Potentials of Deep Learning for Robotics. Niko Sünderhauf, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, Peter Corke (2018) ◦ A Sensorimotor Account of Vision and Visual Consciousness. Kevin O'Regan and Alva Noë (2001)	
Week 2 Thu, Sept 3	Object Detection • Mask R-CNN. Kaiming He, Georgia Gkioxari, Piotr Dollar, Ross Girshick (2017) • You Only Look Once: Unified, Real-Time Object Detection. Joseph Redmon, Santosh Divvala, Ross Girshick, Ali Farhadi (2015) ◦ Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun (2015) ◦ CornerNet: Detecting Objects as Paired Keypoints. Hei Law, Jia Deng (2018)	
Week 3 Tue, Sept 8	3D Point Cloud Processing • PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. Charles R. Qi, Hao Su, Kaichun Mo, Leonidas J. Guibas (2016) • Dynamic Graph CNN for Learning on Point Clouds. Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, Justin M. Solomon (2018) ◦ PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. Charles R. Qi, Li Yi, Hao Su, Leonidas J. Guibas (2017) ◦ 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. Christopher Choy, JunYoung Gwak, Silvio Savarese (2019)	
Week 15 Tue, Dec 1	Spotlight Final Project Spotlights I	
Week 15 Thu, Dec 3	Spotlight Final Project Spotlights II	
Week 16 Fri, Dec 11	No Class	Final Report Due

Instructor Lectures

overview of research topics

Student Presentations

presentation of research papers

Final Project Spotlights

spotlight talks of course projects

Logistics

Part I: Robot Perception

Week 2 Tue, Sept 1	Lecture Overview of Robot Perception ▪ The Limits and Potentials of Deep Learning for Robotics. Niko Sünderhauf, Oliver Brock, Walter Scheirer, Raia Hadsell, Dieter Fox, Jürgen Leitner, Ben Upcroft, Pieter Abbeel, Wolfram Burgard, Michael Milford, Peter Corke (2018) ○ A Sensorimotor Account of Vision and Visual Consciousness. Kevin O'Regan and Alva Noë (2001)
Week 2 Thu, Sept 3	Object Detection • Mask R-CNN. Kaiming He, Georgia Gkioxari, Piotr Dollar, Ross Girshick (2017) • You Only Look Once: Unified, Real-Time Object Detection. Joseph Redmon, Santosh Divvala, Ross Girshick, Ali Farhadi (2015) ○ Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun (2015) ○ CornerNet: Detecting Objects as Paired Keypoints. Hei Law, Jia Deng (2018)
Week 3 Tue, Sept 8	3D Point Cloud Processing • PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. Charles R. Qi, Hao Su, Kaichun Mo, Leonidas J. Guibas (2016) • Dynamic Graph CNN for Learning on Point Clouds. Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, Justin M. Solomon (2018) ○ PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. Charles R. Qi, Li Yi, Hao Su, Leonidas J. Guibas (2017) ○ 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. Christopher Choy, JunYoung Gwak, Silvio Savarese (2019)
Week 15 Tue, Dec 1	Spotlight Final Project Spotlights I
Week 15 Thu, Dec 3	Spotlight Final Project Spotlights II
Week 16 Fri, Dec 11	No Class — Final Report Due

▪ **Required Readings (No Review)**
overview or survey papers with lectures

● **Required Readings**
key papers that will be discussed in class

○ **Optional Readings**
recommended papers for in-depth reviews

Logistics

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

20% each

- At least **one presentation** for each student (chances to do more)
- Length: **25min (± 1min) + 5min Q&A**
- Format: problem formulation, technical approach, results, … (see **slide template** for more details)
- Followed by **5-10min in-class discussions**
- Email the slides to the TA and the instructor **seven days** (EOD) prior to the presentation date
- **Presentation recordings** posted in Canvas (protected under FERPA)
- **Breakout rooms** and **in-class discussions** will NOT be recorded.

Logistics

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

1.5% each x 20 reviews

- Due by **9:59pm** the previous night of each student presentation
- Write a review for **one paper** from the required readings (2 choices for each class)
- Online review form in R:SS format

ROBOTICS SCIENCE AND SYSTEMS

CS391R: Paper Review Form

This form is used for CS391R (Fall 2020) students to submit the paper reviews. The paper reviews must be submitted by 11:59pm the previous night for each class of student presentations in order to receive a grade.

- **No late date** - but more than 20 presentation classes (feel free to skip some)
- Have energy to do more? **Top-scored 20** for grading
- **Class attendance and participation** is required for review grades

Logistics

Online Learning Considerations

For students studying online in a different time zone, they may submit their written responses to the discussion questions on Canvas within 48hrs after the class as an alternative to in-class attendance.

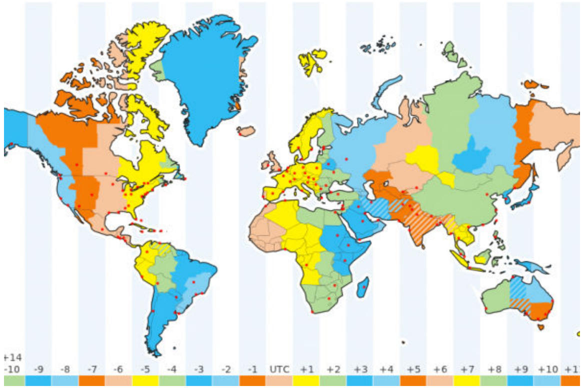

1.5% each x 20 reviews

- Due by **9:59pm** the previous night of each student presentation
- Write a review for **one paper** from the required readings (2 choices for each class)
- Online review form in R:SS format

ROBOTICS SCIENCE AND SYSTEMS

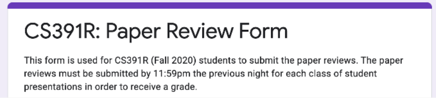

CS391R: Paper Review Form

This form is used for CS391R (Fall 2020) students to submit the paper reviews. The paper reviews must be submitted by 11:59pm the previous night for each class of student presentations in order to receive a grade.

- **No late date** - but more than 20 presentation classes (feel free to skip some)
- Have energy to do more? **Top-scored 20** for grading
- **Class attendance and participation** is required for review grades

Logistics

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

40%

- Project Proposal (5%). Due Thu Sept 17.
- Project Milestone (5%). Due Thu Oct 22.
- Final Report (25%). Due Fri Dec 11.
- Spotlight Talk (5%). Week 15.

Hands-on experience of robot learning research

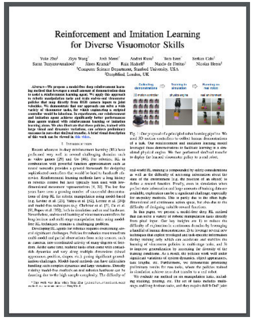
Reinforcement and Imitation Learning for Diverse Visuomotor Skills

+

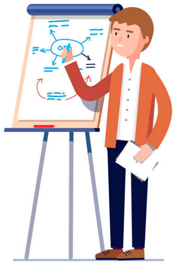

Logistics

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

Tutorials, computing resources, potential project ideas, …

CS391R - Fall 2020 Logistics Syllabus Course Project

Course Project

The primary objective of the course project is to give you in-depth, hands-on experiences applying AI-based techniques to practical robot learning problems. A successful project topic should involve at least one, ideally both, of the two critical components: a **perception** component, i.e., processing raw sensory data, and a **decision making** component, i.e., controlling robot actions, for example,

- Learning vision-based robot manipulation with deep reinforcement methods;
- Self-supervised representation learning of visual and tactile data;
- Model-based object pose estimation for 6-DoF grasping from RGB-D images.

Potential projects can have the following flavors:

- **Improve an existing approach.** You can select a paper you are interested in, reimplement it, and improve it with what you learned in the course.
- **Apply an algorithm to a new problem.** You will need to understand the strengths and weaknesses of an existing algorithm from research work, reimplement it, and apply it to a new problem.
- **Stress test existing approaches.** This kind of project involves a thorough comparison of several existing approaches to a robot learning problem.
- **Design your own approach.** In these kinds of projects, you come up with an entirely new approach to a specific problem. Even the problem may be something that has not been considered before.
- **Mix and Match approaches.** For these projects, you typically combine approaches that have been developed separately to address a larger and more complex problem.
- **Join a research project.** You can join an existing Robot Learning project with UT faculty and researchers. You are expected to articulate your own contributions in your project reports (more detail below).

You may work individually or pair up with one teammate on the project, and grades will be calibrated by team size. Projects of a larger scope are expected for teams of two. Your project may be related to research in another class project as long as consent is granted by instructors of both classes; however, you must clearly indicate in the project proposal, milestone, and final reports the exact portion of the project that is being counted for this course. In this case, you must prepare separate reports for each course, and submit your final report for the other course as well.

Grading Policy

The course project is worth 40% of the total grade. The following shows the breakdown:

- Project Proposal (5%). Due Thu Sept 17.
- Project Milestone (5%). Due Thu Oct 15.
- Final Report (25%). Due Fri Dec 11.

More details on the website!

Logistics

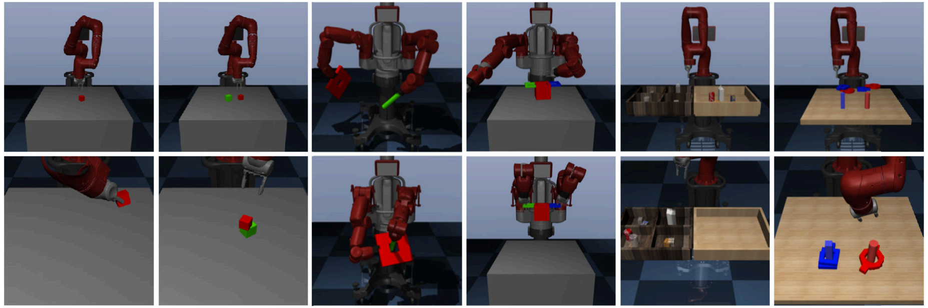

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

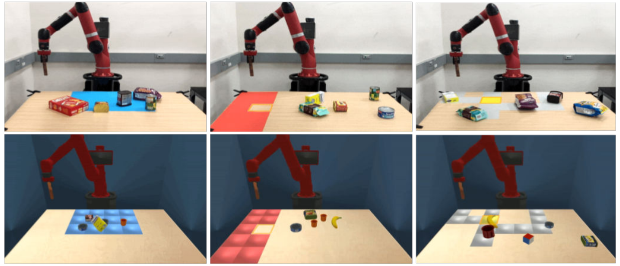

Do Robot Learning with Simulated Environments

Tutorials, computing resources, potential project ideas, …

Logistics

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

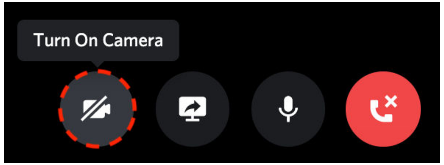

Tell Us About Yourself