CS391R: Robot Learning

Perception and Decision Making: Architectures, Algorithms, and Applications

Prof. Yuke Zhu

Fall 2021

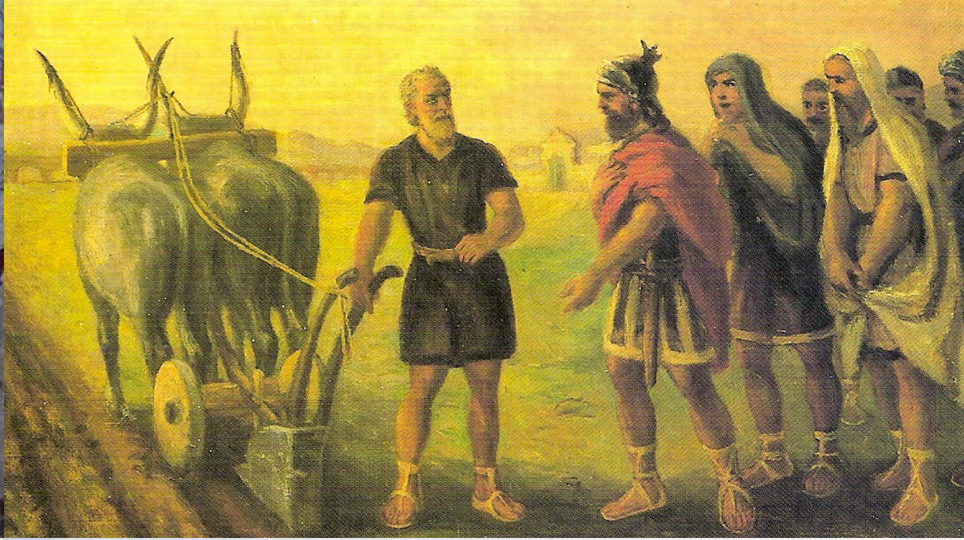


CLAAS


70
302331
302331
302159
DHL
microdrones.com
PACKSET

Traditional form of **automation**

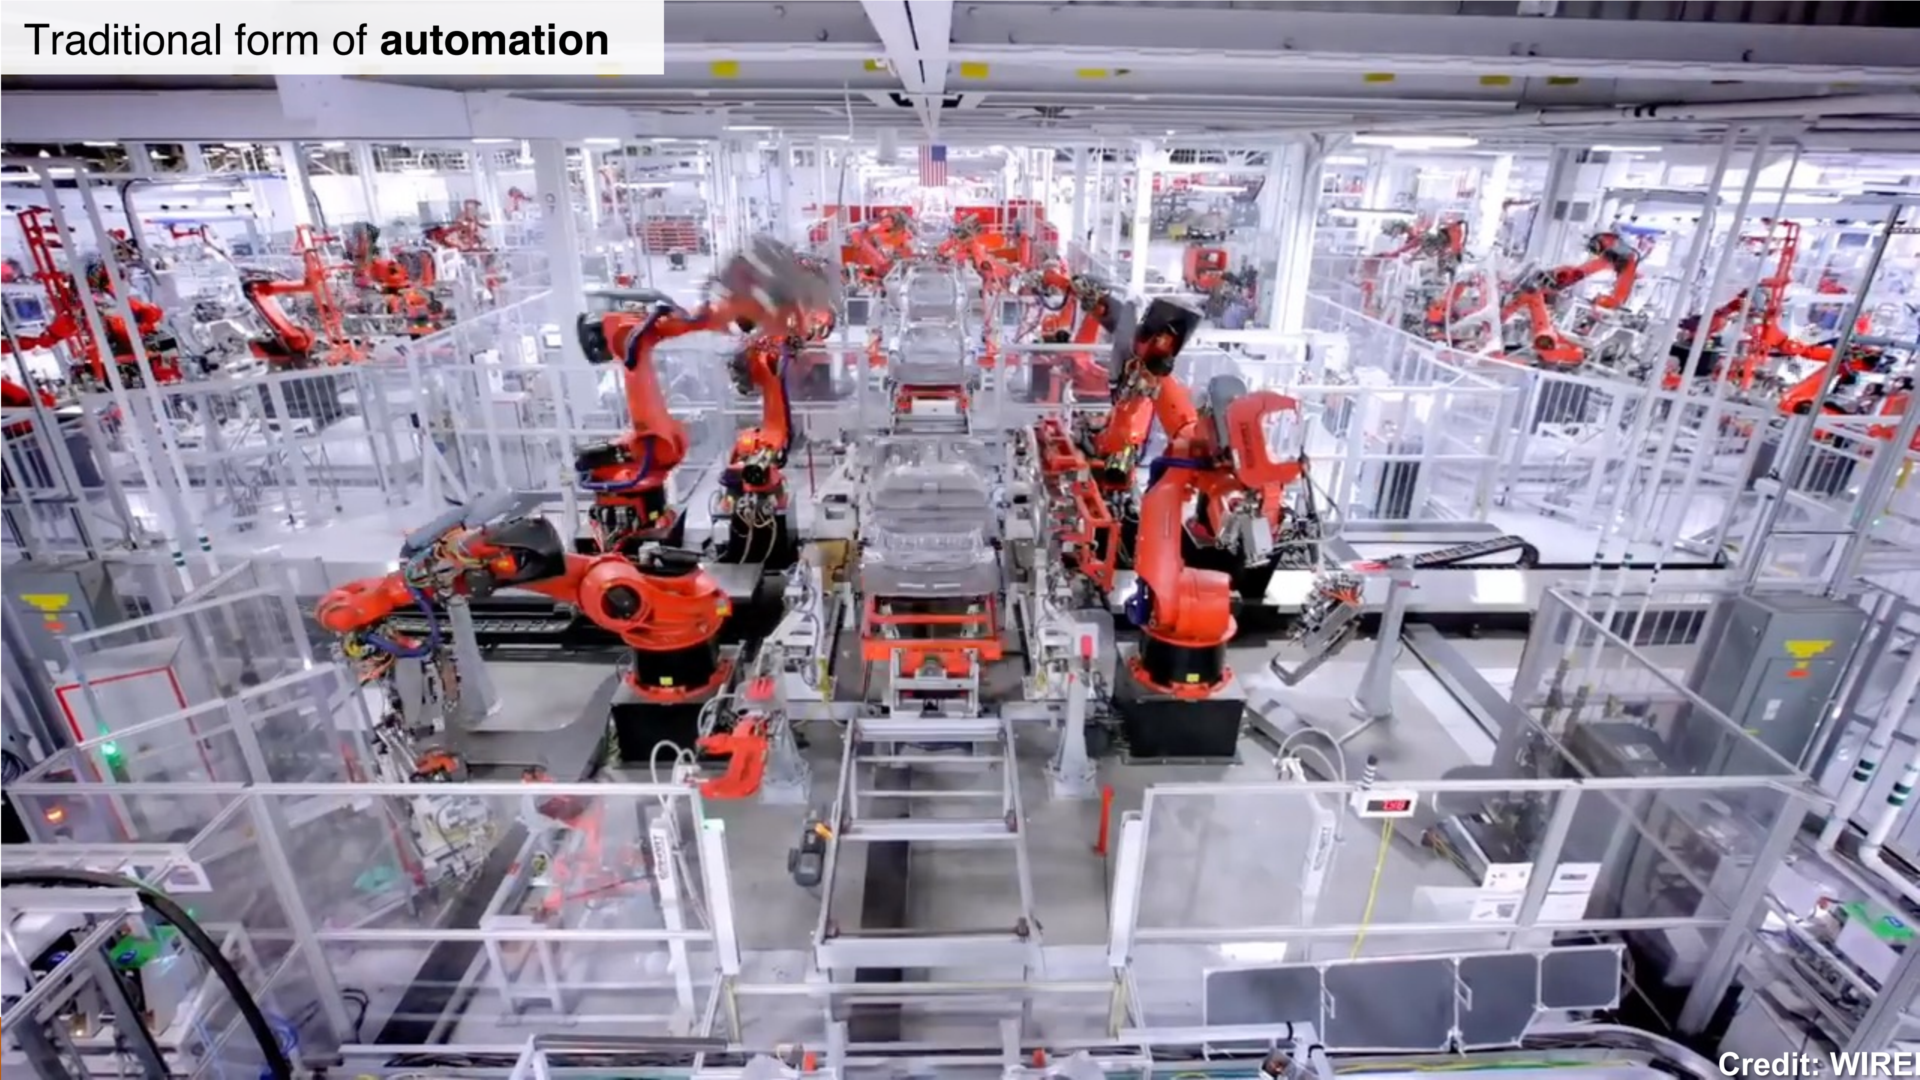

Credit: WIRE

General-purpose **robot autonomy**

Today's Agenda

- What is Robot Learning?
- Why studying Robot Learning now?
- Course content overview
- Logistics
- Student introduction

Special-Purpose Robot Automation

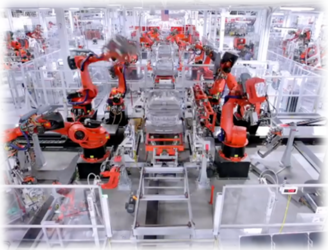

custom-built
robots

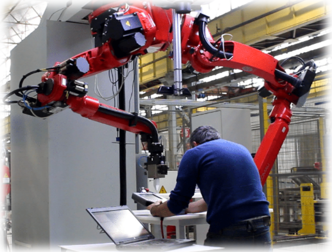

human expert
programming

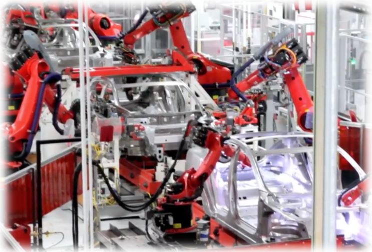

special-purpose
behaviors

General-Purpose Robot Autonomy

general-purpose
robots

?

general-purpose
behaviors

Special-Purpose Robot Automation

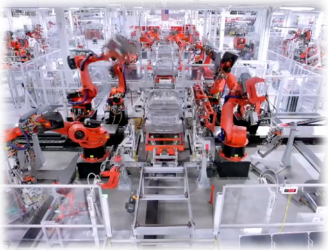

custom-built
robots

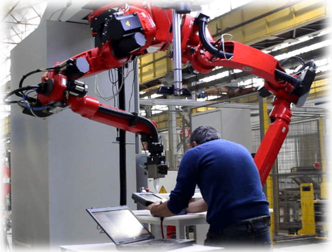

human expert
programming

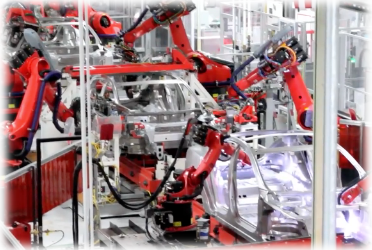

special-purpose
behaviors

General-Purpose Robot Autonomy

general-purpose
robots

Robot Learning

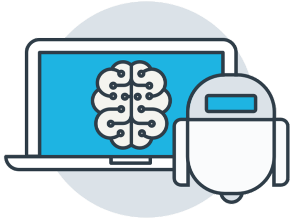

general-purpose
behaviors

General-Purpose Robot Autonomy: Imaginations

Unimate - The First Industrial Robot

British TV (1968)

General-Purpose Robot Autonomy: Challenges

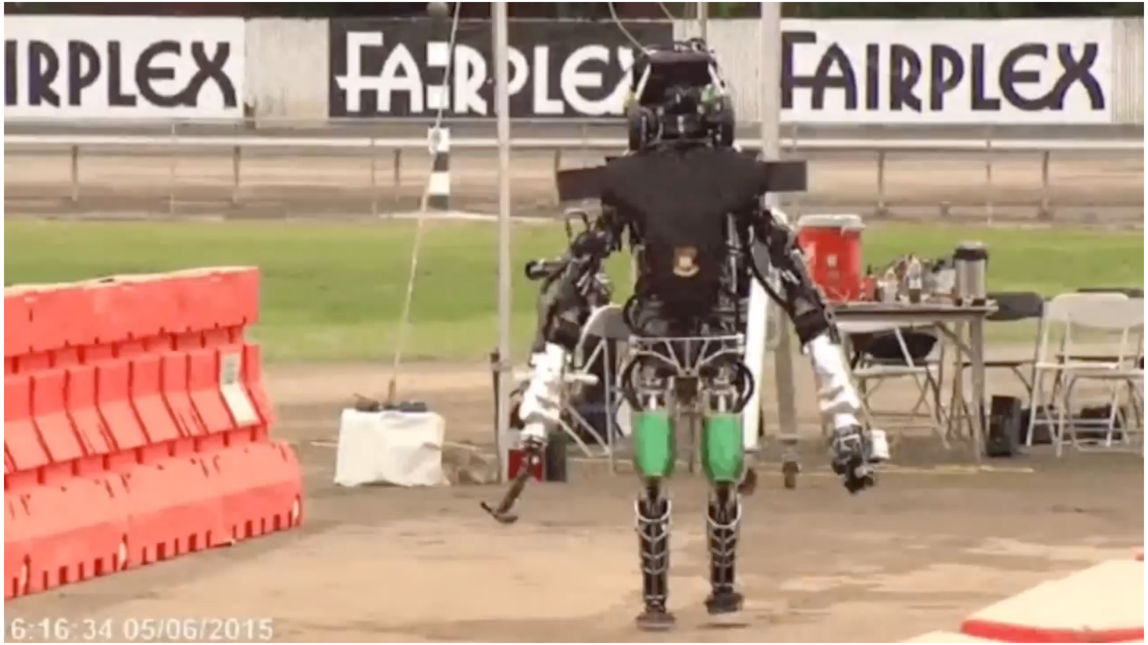

DARPA Robotics Challenge

(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; 2020)

Manipulation (OpenAI; 2019)

What is Robot Learning?

Definition #1

The study of machine learning algorithms and principles with their applications to robotics problems

Definition #2

The study of methods and principles that make robots learn from data

Definition #3

The research field at the intersection of machine learning and robotics (copied from Wikipedia)

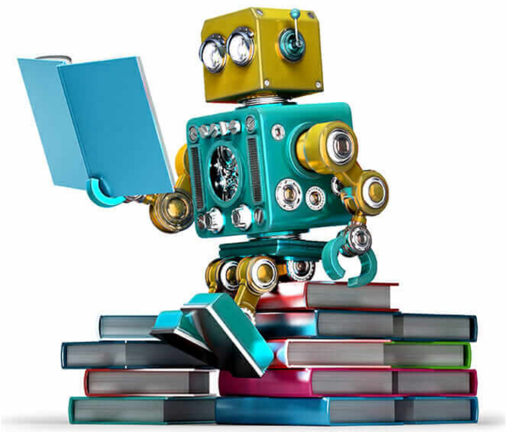

When **NOT** to Make Robots Learn?

Learning is not a solution to 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.

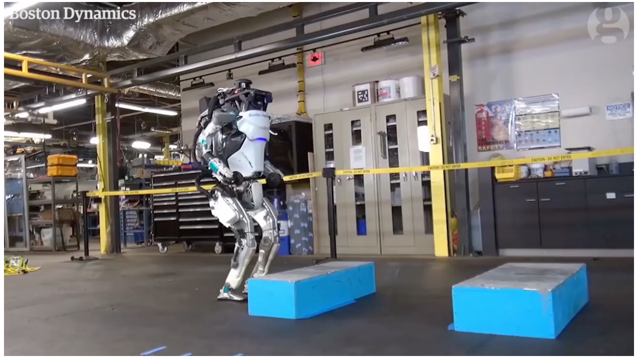

When to Make Robots Learn?

Learning is critical for taking robots to the real world.

object variation

environment uncertainty

adaptation

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

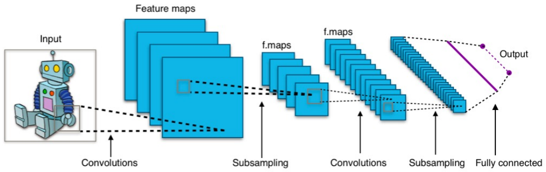

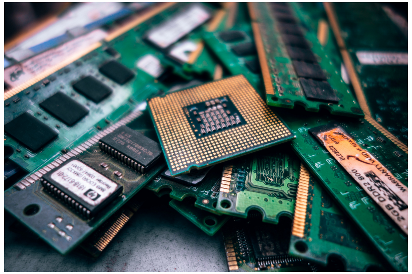

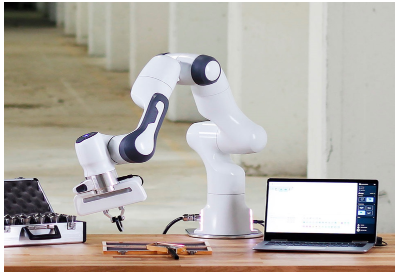

Artificial Intelligence

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 **4** years old.

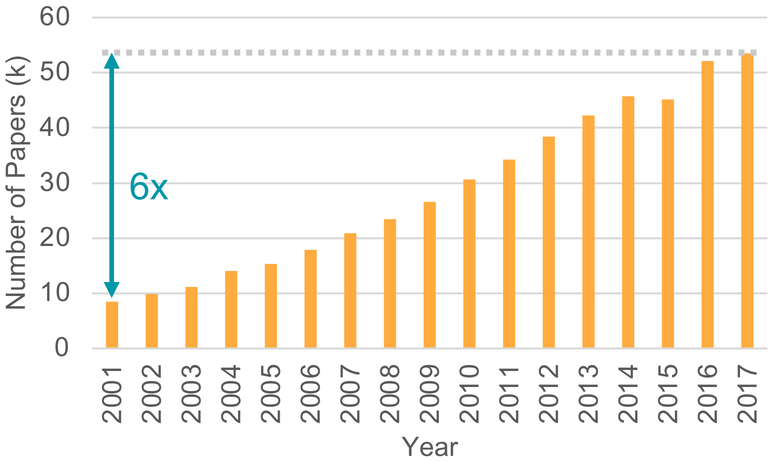

Growth of "**Robot Learning**" Publications

[Source: Google Scholar]

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

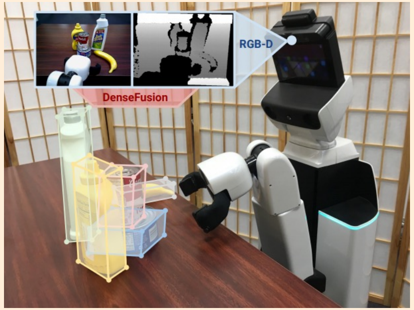

Topic 1-10

seeing and understanding

the physical world

Part II: Robot Decision Making

Topic 11-20

planning and control of

robot behaviors

Prerequisite: coursework / experience in AI and Machine Learning

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

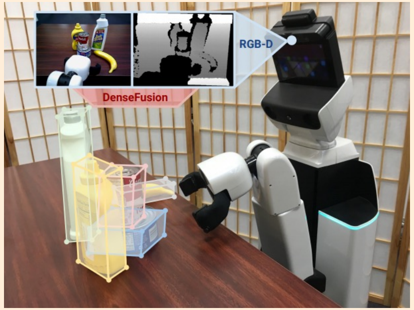

Topic 1-10

seeing and understanding

the physical world

Part II: Robot Decision Making

Topic 11-20

planning and control of

robot behaviors

Prerequisite: coursework / experience in AI and Machine Learning

Robot Perception

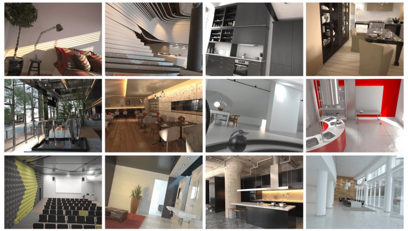

synthetic data for robot perception

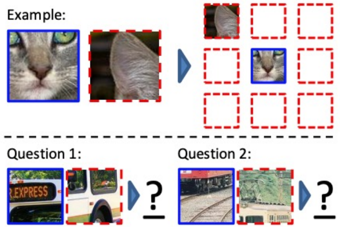

self-supervised visual learning

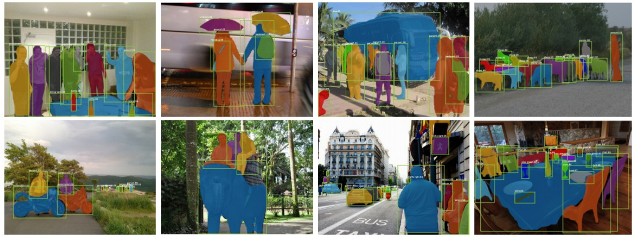

2D object detection

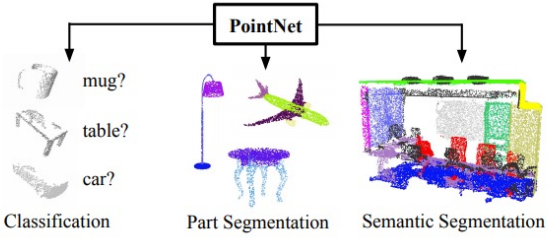

3D data processing

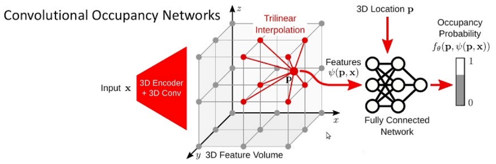

implicit neural representations

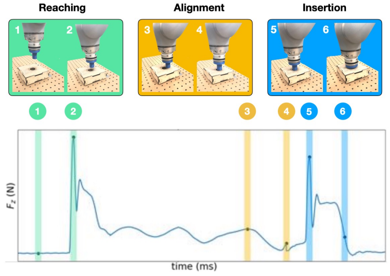

multimodal understanding

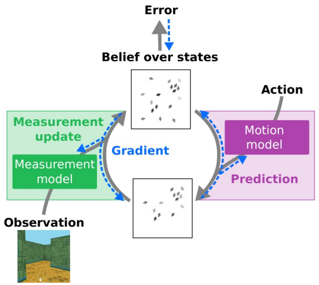

recursive state estimation

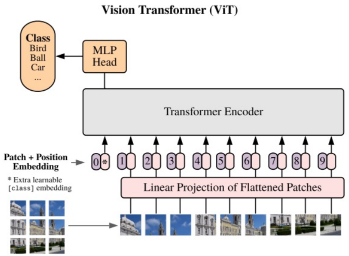

attention architectures

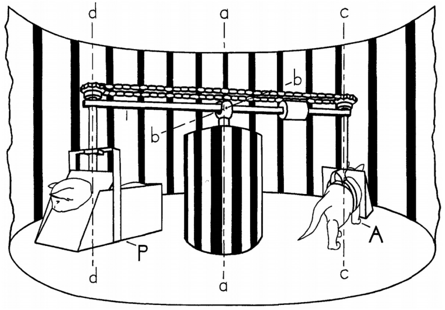

interactive perception

Course Content

We review the Robot Learning literature in these topics.

Part I: Perception

Topic 1-10

seeing and understanding
the physical world

Part II: Decision Making

Topic 11-20

planning and control of
robot behaviors

Prerequisite: coursework / experience in AI and Machine Learning

Robot Decision Making

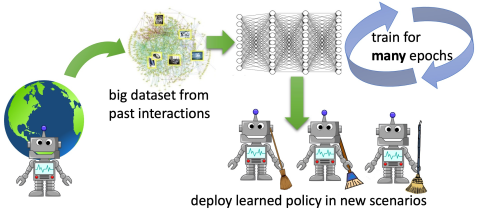

offline RL

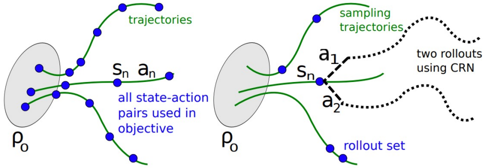

model-free RL

inverse RL

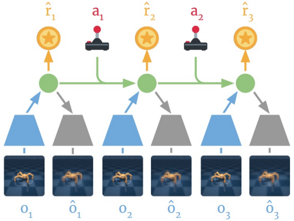

model-based RL

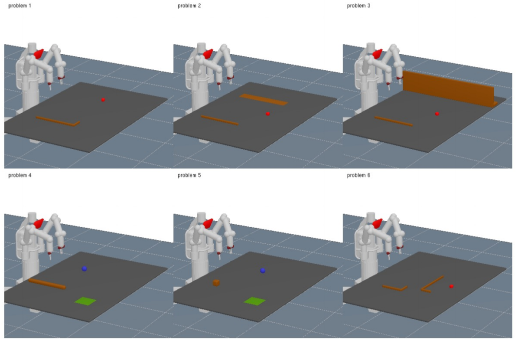

task and motion planning

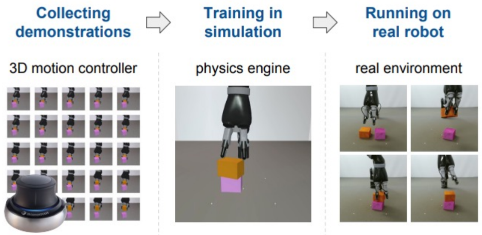

adversarial IL

imitation as supervised learning

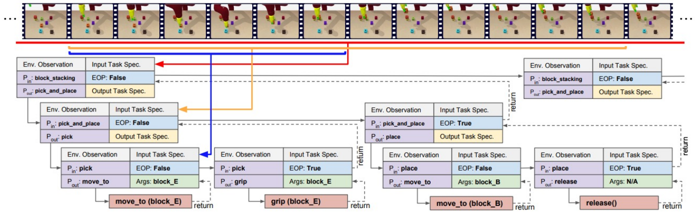

hierarchical policy & neural programming

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: 9:30-11:00am CT, Tuesdays and Thursdays

Location: Online or in-person (Zoom links on Canvas)

Office Hours

Instructor: 3-4pm Mondays (GDC 3.422) or by appointment

TA: 4-5pm Wednesdays (GDC 3.516)

Instruction Modality

Now to September 17

Online lectures + in-person office hours

After September 17

Adjusting plans based on university policy

In-Person Experiences

Office hours, instructor/TA meetings by appointment, GDC 4.302

Logistics

Part I: Robot Perception

Week 2
Tue, Aug 31

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 2

2D 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)

- CornerNet: Detecting Objects as Paired Keypoints. Hei Law, Jia Deng (2018)
- Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun (2015)

Week 3
Tue, Sept 7

3D Data Processing

- PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. Charles R. Qi, Li Yi, Hao Su, Leonidas J. Guibas (2017)
- Dynamic Graph CNN for Learning on Point Clouds. Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, Justin M. Solomon (2018)

- PointCNN: Convolution On X-Transformed Points. Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, Baoquan Chen (2018)
- 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. Christopher Choy, JunYoung Gwak, Silvio Savarese (2019)

Week 15 Tue, Nov 30	Spotlight Final Project Spotlights I	Video Due Nov 29
Week 15 Thu, Dec 2	Spotlight Final Project Spotlights II	
Week 16 Fri, Dec 10	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, Aug 31

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 2

2D 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)
- CornerNet: Detecting Objects as Paired Keypoints. Hei Law, Jia Deng (2018)
- Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun (2015)

Week 3
Tue, Sept 7

3D Data Processing

- PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. Charles R. Qi, Li Yi, Hao Su, Leonidas J. Guibas (2017)
- Dynamic Graph CNN for Learning on Point Clouds. Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, Justin M. Solomon (2018)
- PointCNN: Convolution On X-Transformed Points. Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, Baoquan Chen (2018)
- 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. Christopher Choy, JunYoung Gwak, Silvio Savarese (2019)

Week 15 Tue, Nov 30	Spotlight Final Project Spotlights I	Video Due Nov 29
Week 15 Thu, Dec 2	Spotlight Final Project Spotlights II	
Week 16 Fri, Dec 10	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: **20min (± 2min) + 3min 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%)

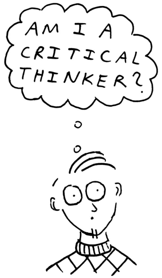

2% each x 15 reviews

- 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 15 presentation classes (feel free to skip some)
- Have energy to do more? **Top-scored 15** for grading
- **Class attendance and participation** is required for review grades

Logistics

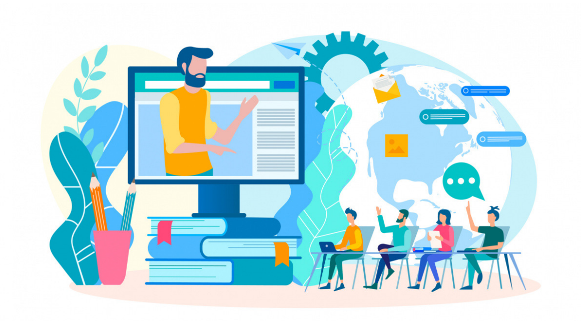

piazza

2% each x 15 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 15 presentation classes (feel free to skip some)
- Have energy to do more? **Top-scored 15** 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 16.
- Project Milestone (5%). Due Thu Oct 21.
- Final Report (25%). Due Fri Dec 10.
- Spotlight Talk (5%). Week 15.

Hands-on experience of

robot learning research

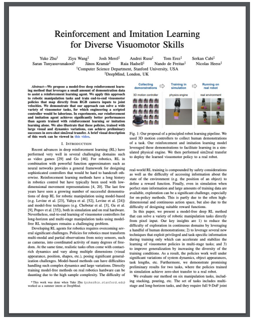
Reinforcement and Imitation Learning
for Diverse Visuomotor Skills

+

Logistics

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

project platform: **robosuite** (robosuite.ai)

Tutorials, computing resources, project instructions, …

Alternative projects require instructor approval.

Logistics

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

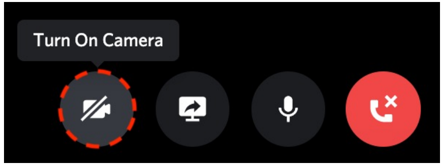

Tell Us About Yourself

Robotics beyond CS391R

Be part of the Robotics + AI revolution.

UT Robot Perception & Learning Lab

Mission: Building General-Purpose Robot Autonomy in the Wild

TEXAS Robotics

https://robotics.utexas.edu/