CS391R: Robot Learning

Perception and Decision Making: Architectures, Algorithms, and Applications

Prof. Roberto Martín-Martín

Fall 2022

Today's Agenda

- A bit about me…
- What is Robot Learning?
- Why studying Robot Learning now?
- Course content overview
- Logistics
- Student introduction
- Final remarks

A bit about me…

- I'm from Madrid, Spain
- I studied in Berlin, Germany
- I moved to Stanford, US
- And now I'm here!

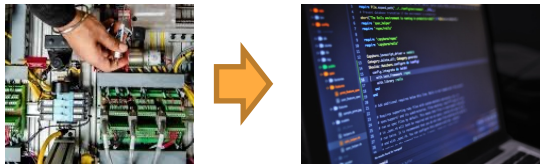

What about you…

- We will be using PollEv during the course
- PollEv.com/robertomartinmartin739

Where are you from?

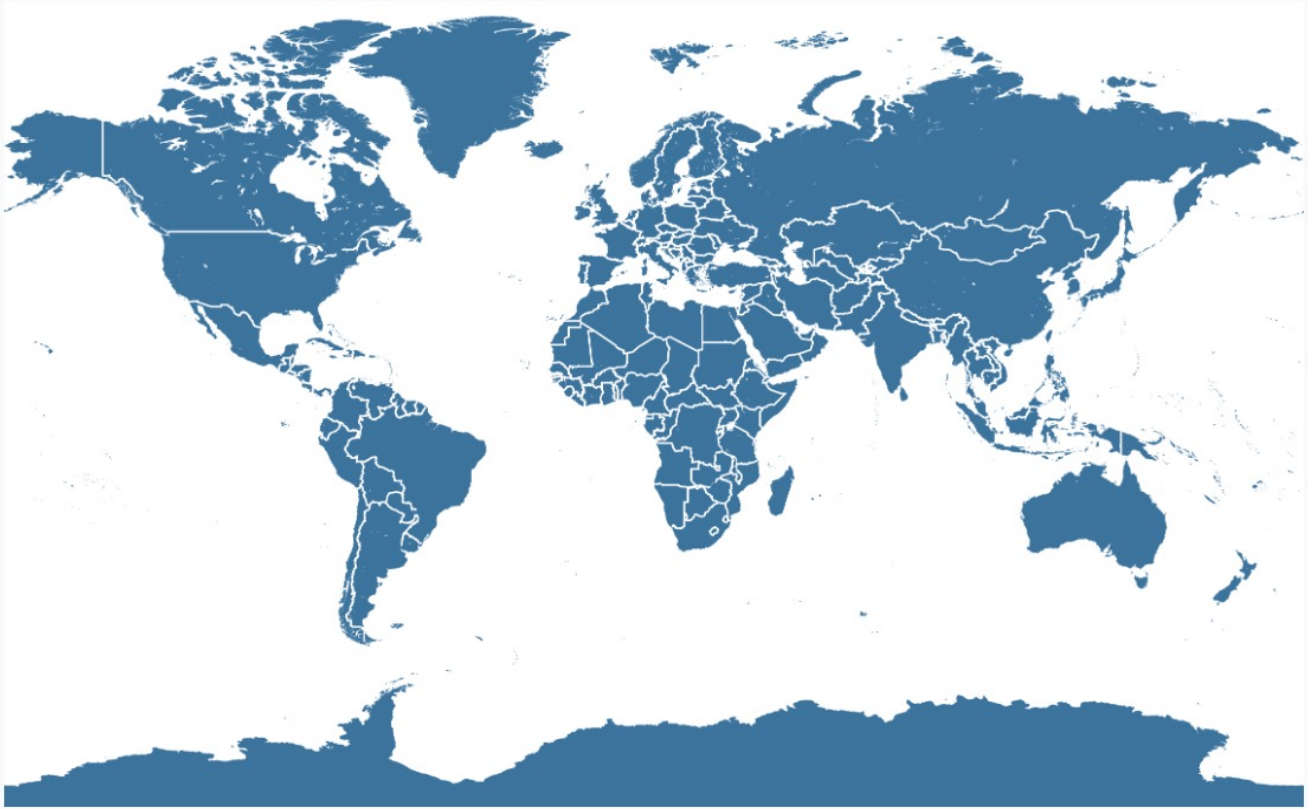

What are you studying?

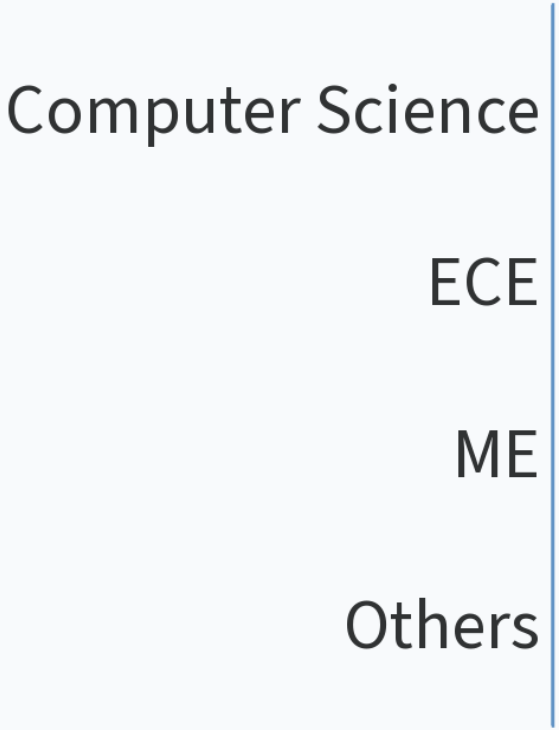

A bit more about me…

+

=

A bit more about me…

What is a robot?

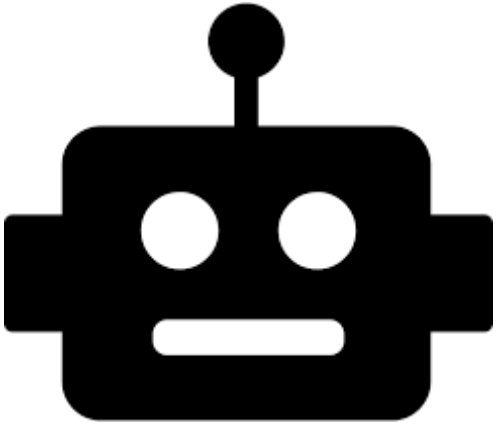

- A robot is an autonomous machine capable of **sensing** its environment, carrying out **computations** to make decisions, and **performing actions** in the real world.

- Are dishwashers robots?
- Are flushing toilets robots?

- Demonstrate intelligent behavior
  - Adaptability
  - Complex reasoning
  - Learning!

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)

Definition #4

Methods to extract autonomously patterns of information from sensory data in order to make optimal decisions on how to act in the world to change it and achieve goals

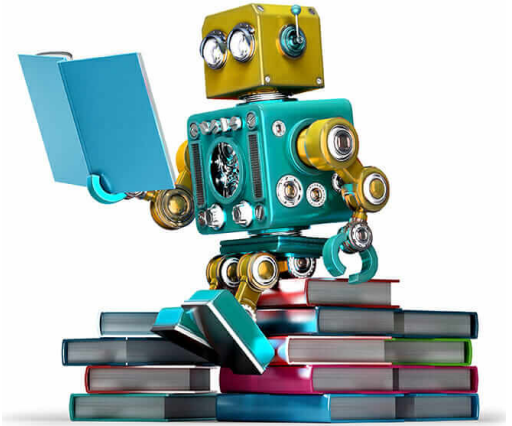

When **NOT** to Make Robots Learn?

Learning is not a solution to every problem in robotics.

Repetitive, non-changing processes

Do not “reinvent the wheel”!

We have very good models of physical processes

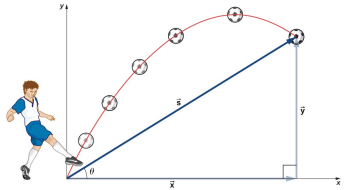

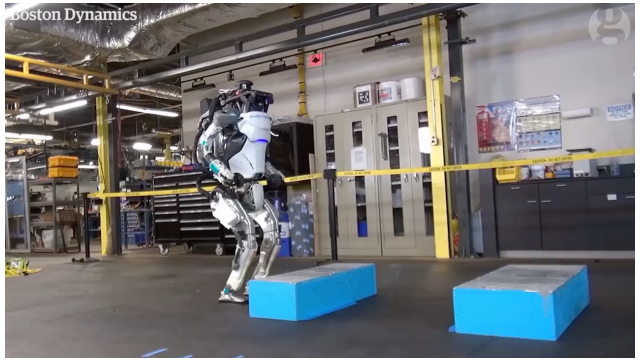

When to Make Robots Learn?

Learning is critical to bring robots to unstructured environments

structured

unstructured

When to Make Robots Learn?

Learning is critical to bring robots to unstructured environments

unstructured

Our secret weapon?

learning!

variable

uncertain

dynamic

Today's Agenda

- A bit about me…
- What is Robot Learning?
- Why studying Robot Learning now?
- Course content overview
- Logistics
- Student introduction
- Final remarks

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

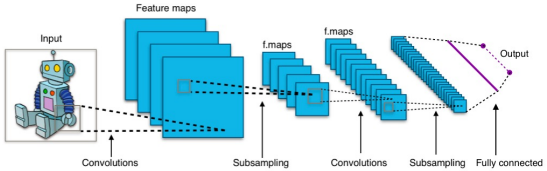

Artificial Intelligence

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

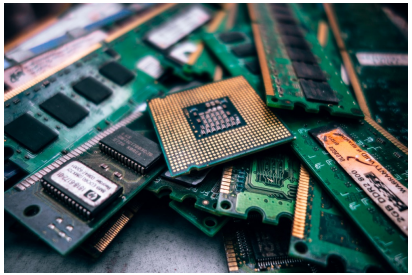

Computing Power

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

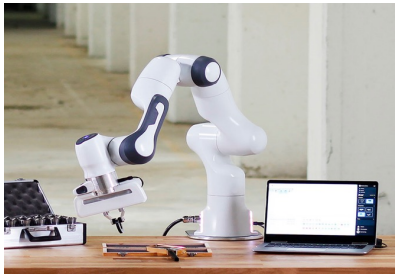

Robot Hardware

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.

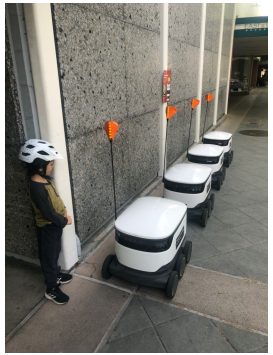

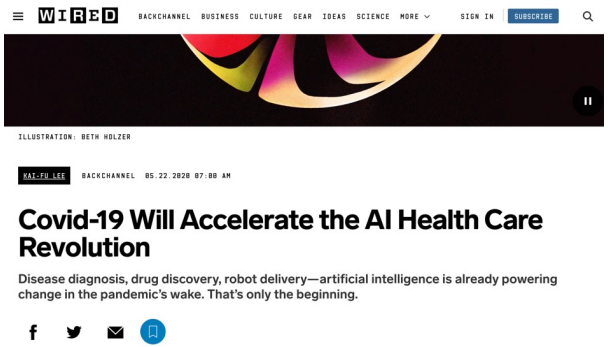

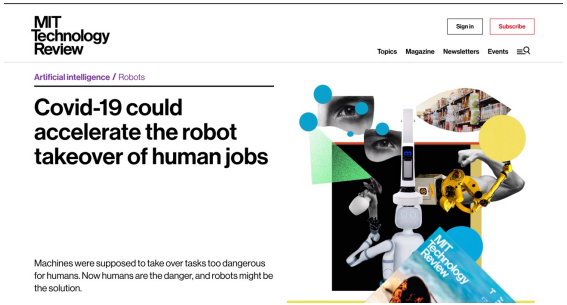

https://www.therobotreport.com/tag/coronavirus/

Robot Learning as a Growing Research Community

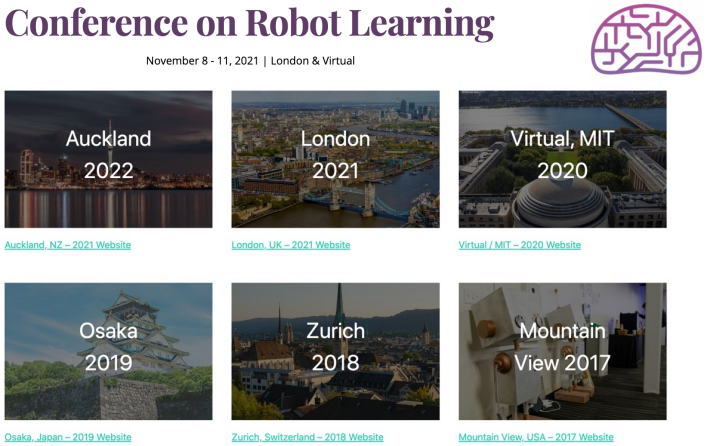

Conference on Robot Learning is **6** years old.

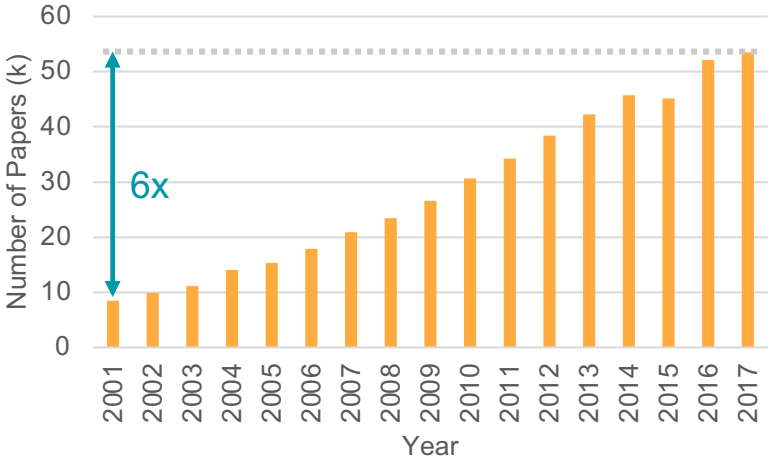

Growth of "**Robot Learning**" Publications

[Source: Google Scholar]

Today's Agenda

- A bit about me…
- What is Robot Learning?
- Why studying Robot Learning now?
- Course content overview
- Logistics
- Student introduction
- Final remarks

Course Content

Robot Learning:
The Good, The Bad and The Ugly

The Good: Exciting Recent Progress

Grasping (DexNet 4.0; 2019)

Locomotion (ANYmal; 2020)

Manipulation (OpenAI; 2019)

We will learn the algorithms and techniques that enabled the latest progress.

The Bad: Unmet Expectations

Unimate - The First Industrial Robot

British TV (1968)

The Bad: Unmet Expectations

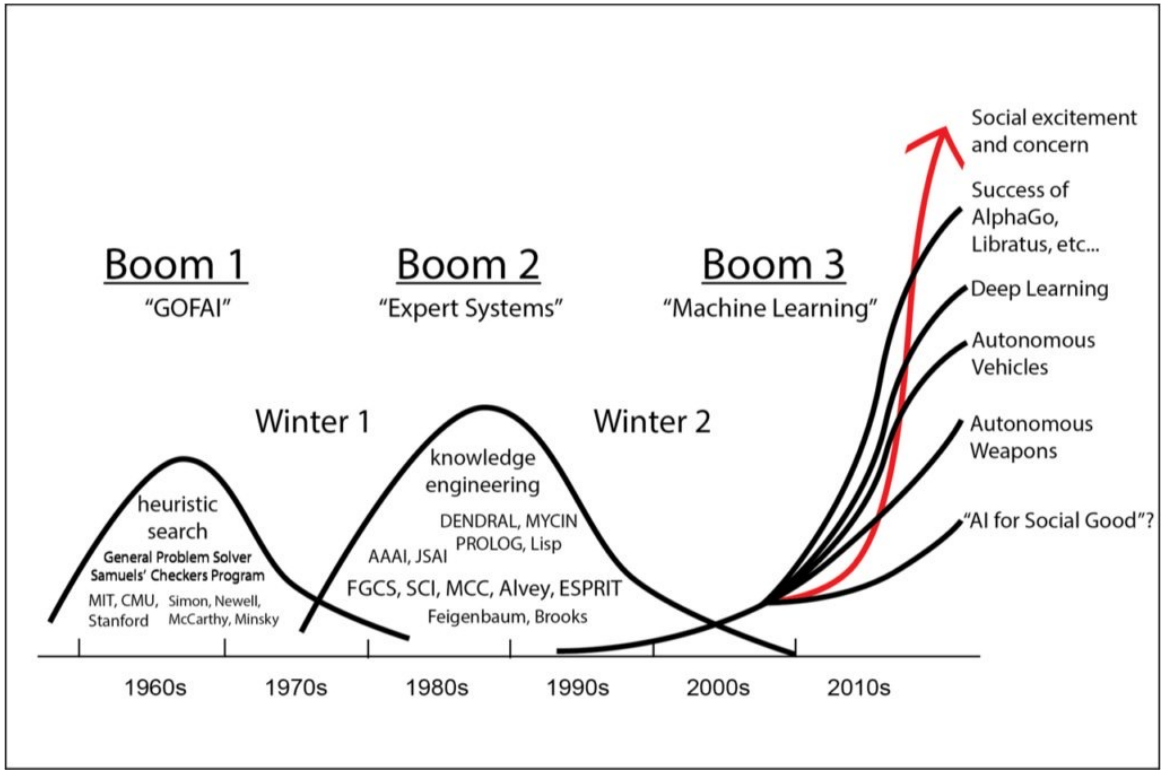

(AI) winter is coming?

(I really don't think so…)

The Ugly: Robot Learning is Challenging!

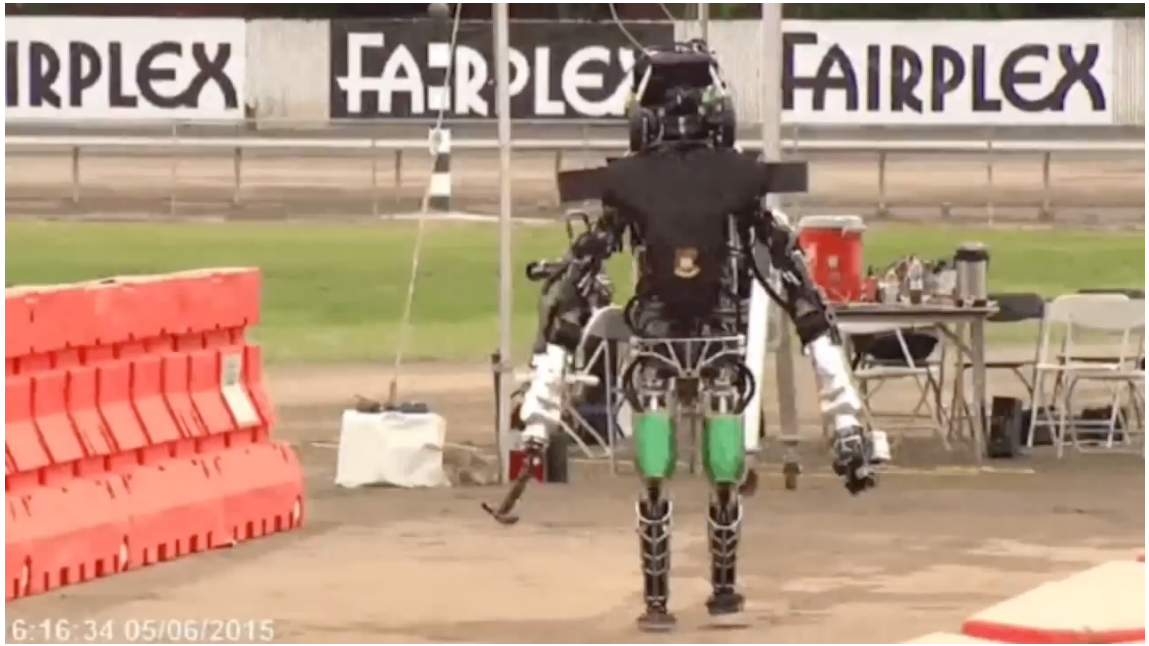

DARPA Robotics Challenge

(2015)

“The Moravec's paradox”

Moravec's paradox (1988):

"it is comparatively easy to make computers exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year-old when it comes to perception and mobility"

[Hans Moravec (1988), *Mind Children*, Harvard University Press]

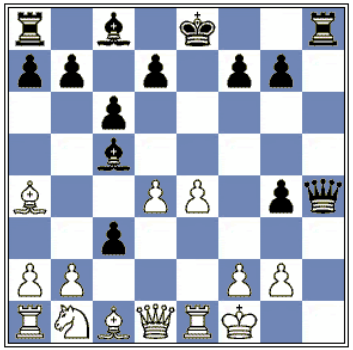

[source: getty]

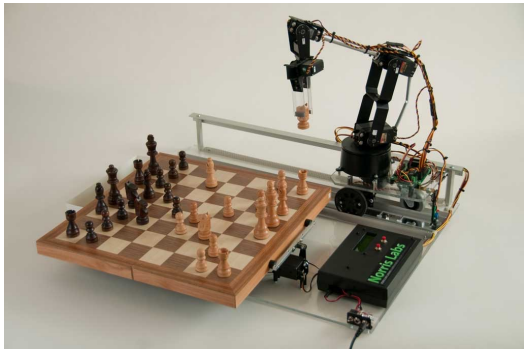

[source: Norris Labs]

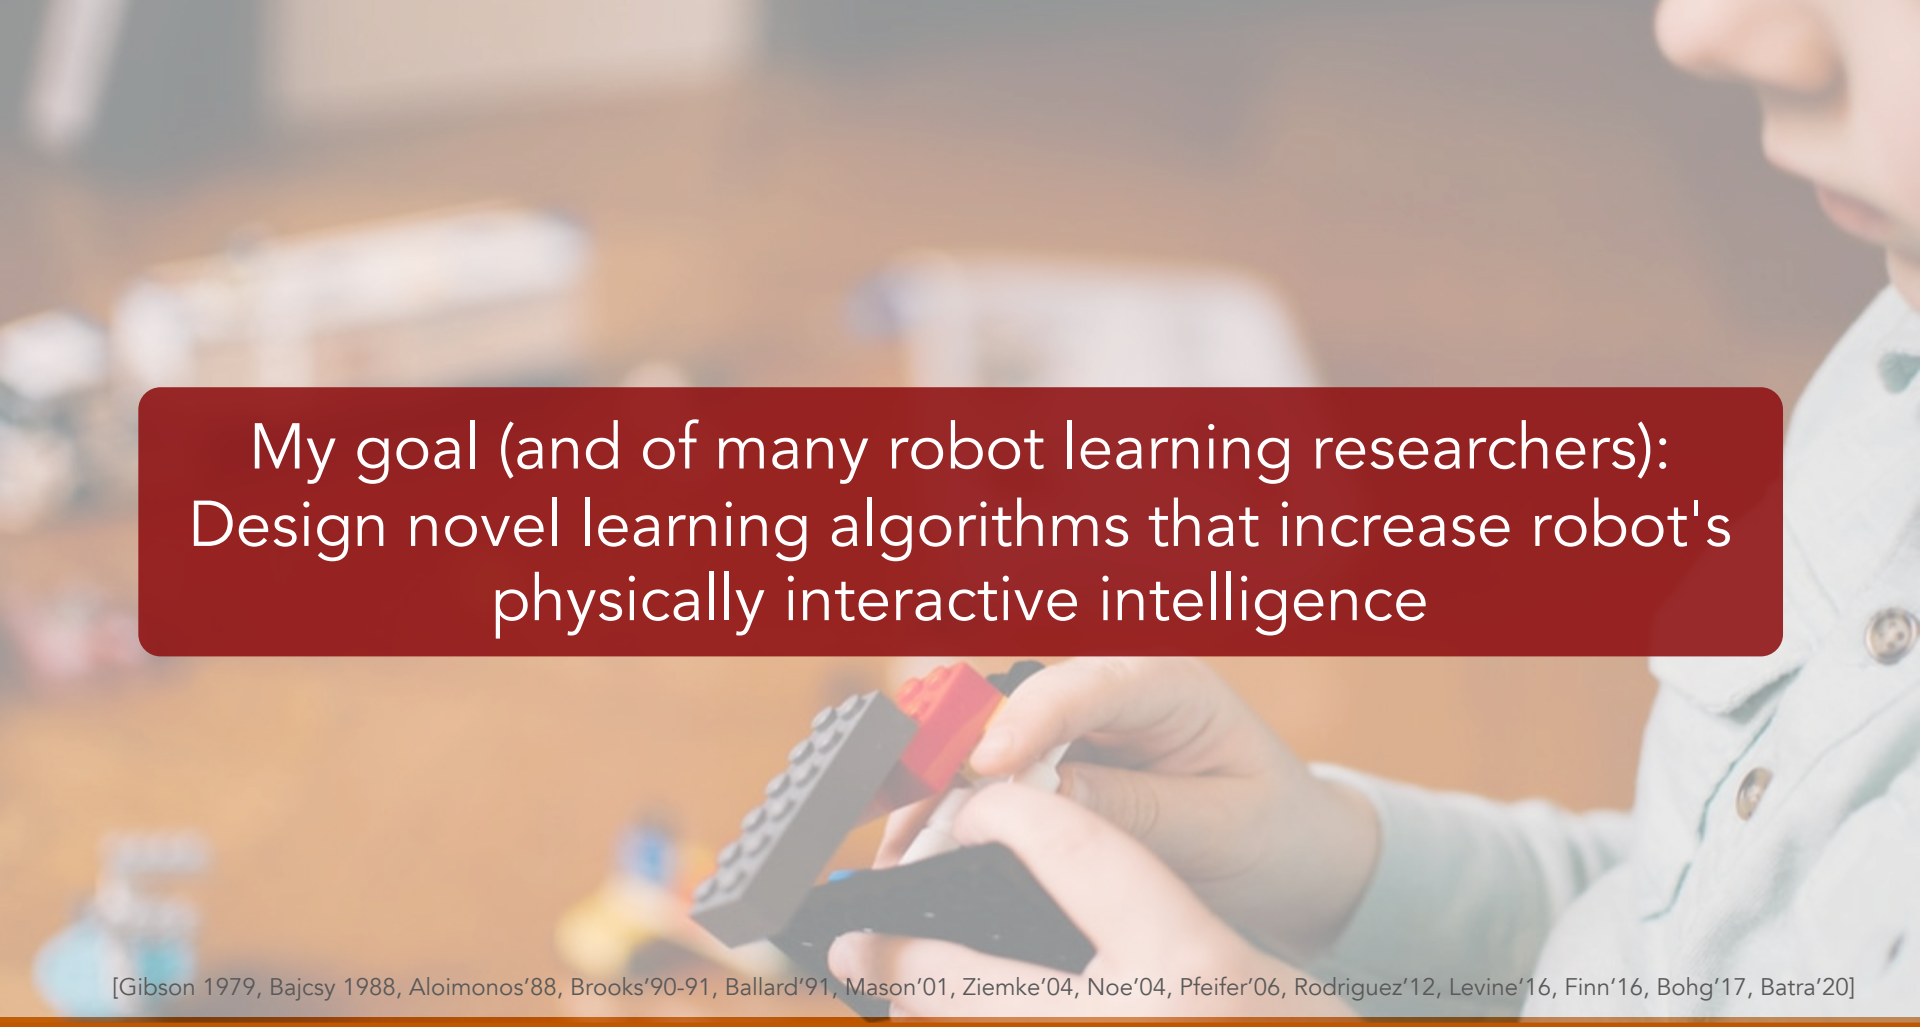

My goal (and of many robot learning researchers): Design novel learning algorithms that increase robot's physically interactive intelligence

[Gibson 1979, Bajcsy 1988, Aloimonos'88, Brooks'90-91, Ballard'91, Mason'01, Ziemke'04, Noe'04, Pfeifer'06, Rodriguez'12, Levine'16, Finn'16, Bohg'17, Batra'20]

Key Ingredients in Robot Learning

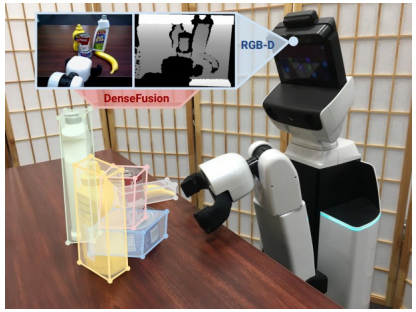

Perception

seeing and understanding 3D environments

Decision Making

planning and control for long-term interactions

Real-World Systems

physical embodiment opportunities & constraints

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

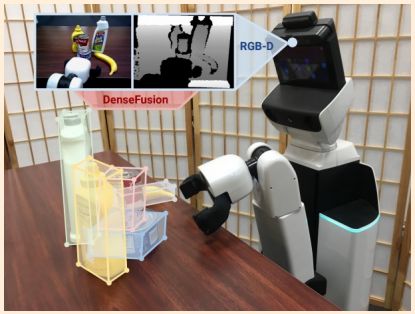

Topic 1-7

seeing and understanding the physical world

Part II: Robot Decision Making

Topic 8-17

planning and control of robot behaviors

Part III: Real-World Systems

Topic 18-20

physical embodiment challenges, sim/real

Prerequisite: coursework / experience in AI and Machine Learning

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

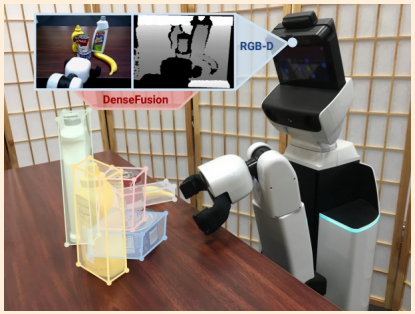

Topic 1-7

seeing and understanding the physical world

Part II: Robot Decision Making

Topic 8-17

planning and control of robot behaviors

Part III: Real-World Systems

Topic 18-20

physical embodiment challenges, sim/real

Prerequisite: coursework / experience in AI and Machine Learning

Robot Perception

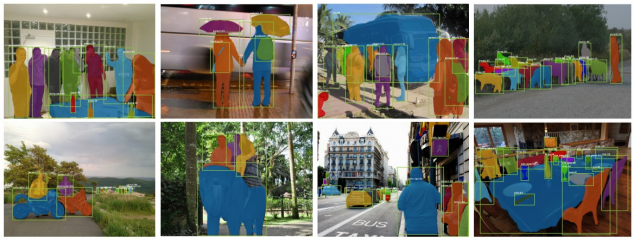

2D object detection

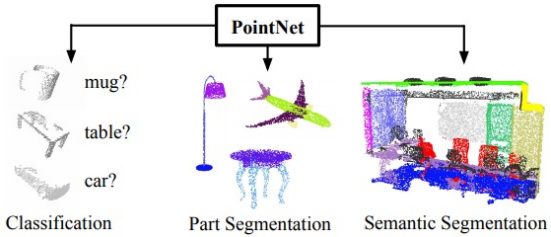

3D data processing

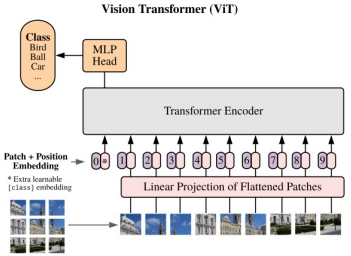

attention architectures

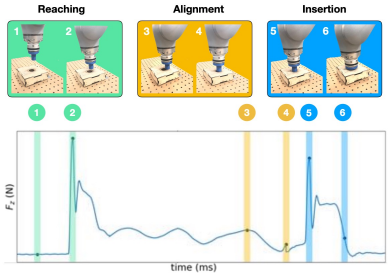

multimodal understanding

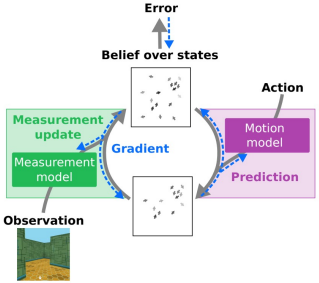

recursive state estimation

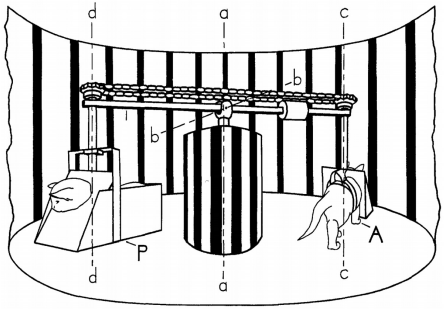

interactive perception

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

Topic 1-7

seeing and understanding the physical world

Part II: Robot Decision Making

Topic 8-17

planning and control of robot behaviors

Part III: Real-World Systems

Topic 18-20

physical embodiment challenges, sim/real

Prerequisite: coursework / experience in AI and Machine Learning

Robot Decision Making

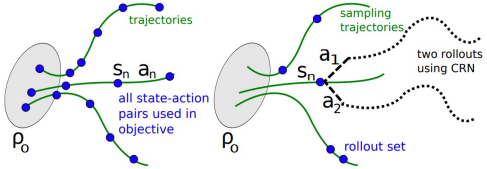

model-free RL

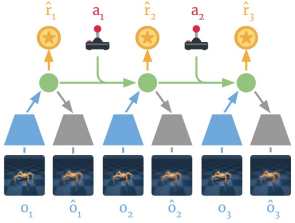

model-based RL

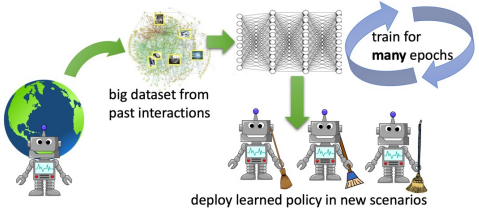

offline RL

inverse RL

imitation as supervised learning

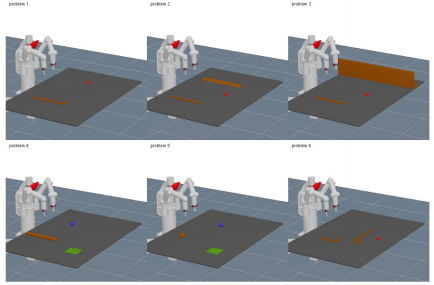

task and motion planning

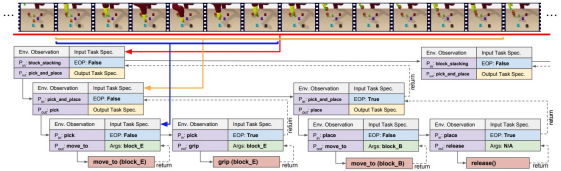

hierarchical policy learning

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

Topic 1-7

seeing and understanding the physical world

Part II: Robot Decision Making

Topic 8-17

planning and control of robot behaviors

Part III: Real-World Systems

Topic 18-20

physical embodiment challenges, sim/real

Prerequisite: coursework / experience in AI and Machine Learning

Real-World Robot Learning Systems

building robotic systems

robot learning challenges
in the real-world

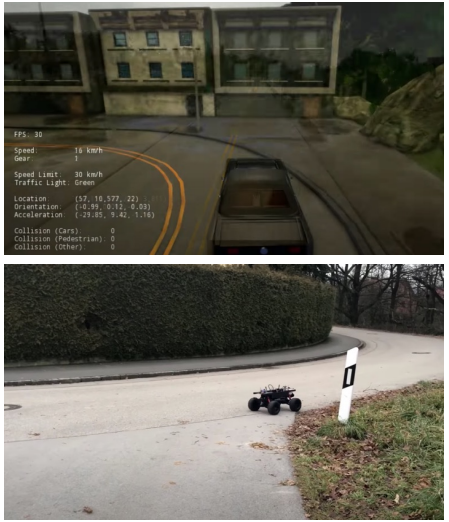

sim2real

Learning Objectives

- understand the potential and societal impact of robotics and robot learning in the real world, the technical challenges in 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 summaries/quizzes, research presentations, and course projects

Learning Objectives

Get a taste of Robot Learning research in the full circle

Today's Agenda

- A bit about me…
- What is Robot Learning?
- Why studying Robot Learning now?
- Course content overview
- Logistics
- Student introduction
- Final remarks

Logistics

Lectures

Time: 3:30-5:00pm CT, Tuesdays and Thursdays

Location: In-person (**preferred**!) or on zoom (links on Canvas)

Office Hours

Roberto: 5-6pm Tuesdays (GDC 3.404) or by appointment

TAs:

Yifeng: Thursdays 10:30-11:30 am

Jeff: Thursdays 5-6 pm

https://www.cs.utexas.edu/~robertom/cs391r_fall2022/

Logistics

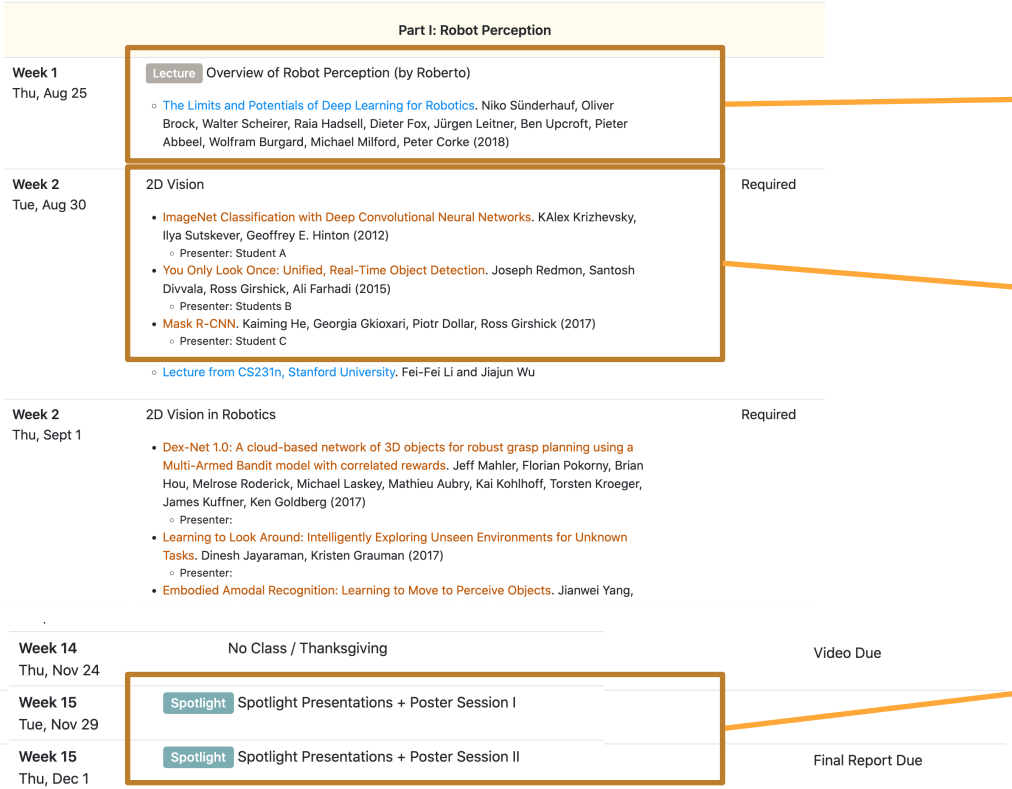

Part I: Robot Perception

Week 1
Thu, Aug 25

Lecture Overview of Robot Perception (by Roberto)

- 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)

Week 2
Tue, Aug 30

2D Vision — Required

- ImageNet Classification with Deep Convolutional Neural Networks. KAlex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton (2012)
  - Presenter: Student A
- You Only Look Once: Unified, Real-Time Object Detection. Joseph Redmon, Santosh Divvala, Ross Girshick, Ali Farhadi (2015)
  - Presenter: Students B
- Mask R-CNN. Kaiming He, Georgia Gkioxari, Piotr Dollar, Ross Girshick (2017)
  - Presenter: Student C
- Lecture from CS231n, Stanford University. Fei-Fei Li and Jiajun Wu

Week 2
Thu, Sept 1

2D Vision in Robotics — Required

- Dex-Net 1.0: A cloud-based network of 3D objects for robust grasp planning using a Multi-Armed Bandit model with correlated rewards. Jeff Mahler, Florian Pokorny, Brian Hou, Melrose Roderick, Michael Laskey, Mathieu Aubry, Kai Kohlhoff, Torsten Kroeger, James Kuffner, Ken Goldberg (2017)
  - Presenter:
- Learning to Look Around: Intelligently Exploring Unseen Environments for Unknown Tasks. Dinesh Jayaraman, Kristen Grauman (2017)
  - Presenter:
- Embodied Amodal Recognition: Learning to Move to Perceive Objects. Jianwei Yang,

Week 14
Thu, Nov 24

No Class / Thanksgiving — Video Due

Week 15
Tue, Nov 29

Spotlight Spotlight Presentations + Poster Session I

Week 15
Thu, Dec 1

Spotlight Spotlight Presentations + Poster Session II — Final Report Due

Instructor Lectures

overview of research topics

Student Presentations

presentation of research papers

Final Project Spotlights + Posters

spotlight talks and posters of course projects

https://www.cs.utexas.edu/~robertom/cs391r_fall2022/

Logistics

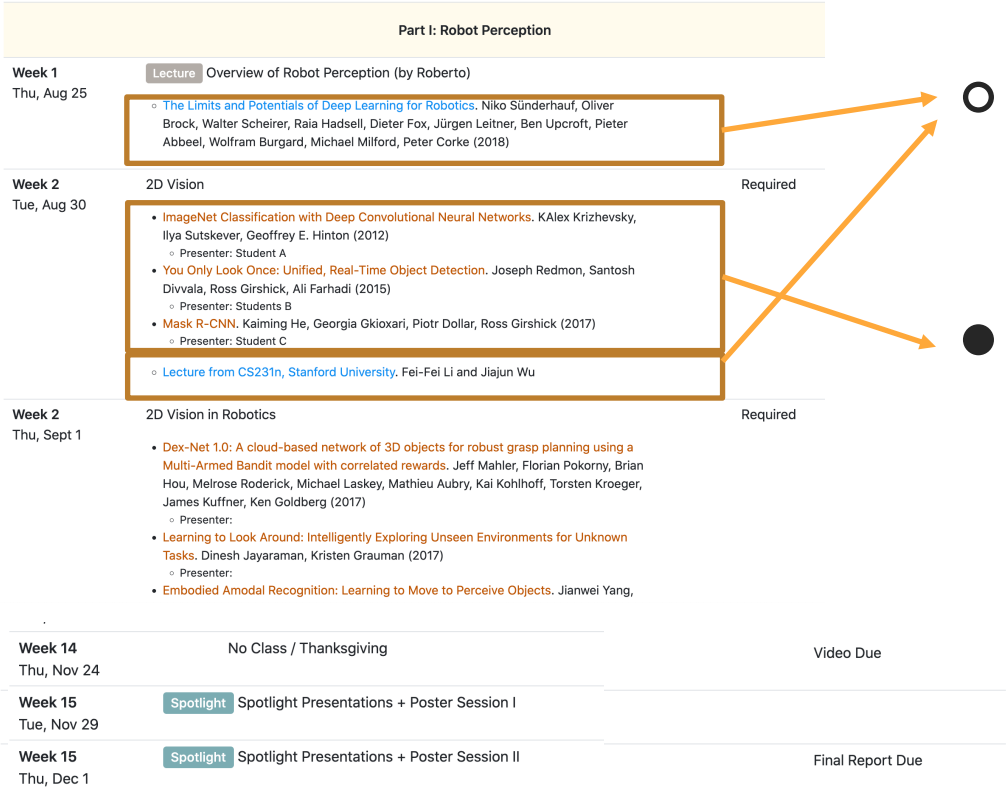

Part I: Robot Perception

Week 1 Thu, Aug 25
Lecture Overview of Robot Perception (by Roberto)
- 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)

Week 2 Tue, Aug 30
2D Vision — Required
- ImageNet Classification with Deep Convolutional Neural Networks. KAlex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton (2012)
  - Presenter: Student A
- You Only Look Once: Unified, Real-Time Object Detection. Joseph Redmon, Santosh Divvala, Ross Girshick, Ali Farhadi (2015)
  - Presenter: Students B
- Mask R-CNN. Kaiming He, Georgia Gkioxari, Piotr Dollar, Ross Girshick (2017)
  - Presenter: Student C
- Lecture from CS231n, Stanford University. Fei-Fei Li and Jiajun Wu

Week 2 Thu, Sept 1
2D Vision in Robotics — Required
- Dex-Net 1.0: A cloud-based network of 3D objects for robust grasp planning using a Multi-Armed Bandit model with correlated rewards. Jeff Mahler, Florian Pokorny, Brian Hou, Melrose Roderick, Michael Laskey, Mathieu Aubry, Kai Kohlhoff, Torsten Kroeger, James Kuffner, Ken Goldberg (2017)
  - Presenter:
- Learning to Look Around: Intelligently Exploring Unseen Environments for Unknown Tasks. Dinesh Jayaraman, Kristen Grauman (2017)
  - Presenter:
- Embodied Amodal Recognition: Learning to Move to Perceive Objects. Jianwei Yang,

Week 14 Thu, Nov 24
No Class / Thanksgiving — Video Due

Week 15 Tue, Nov 29
Spotlight Spotlight Presentations + Poster Session I

Week 15 Thu, Dec 1
Spotlight Spotlight Presentations + Poster Session II — Final Report Due

- **Optional Readings**

  recommended papers and tutorials for better understanding of materials

- **Required Readings**

  key papers that will be discussed in class

Logistics

Grading Policy

Student presentation (20%)

Paper abstracts/summaries and Quizzes (30%)

Course project (40%)

In-class participation (10%)

Enroll in the sheet we will send later today (firs come, first serve)

20% each

- **One presentation** per student
- Length: **15min (strict!) + 5min Q&A**
- Format: problem formulation, technical approach, results, … (see **slide template** for more details)
- Email the slides to the TA and the instructor **seven days** (EOD) prior to the presentation date
- Followed by final debate per session
- Email open-ended questions related to the paper for debate
- **Presentation recordings** posted in Canvas (protected under FERPA)
- **In-class discussions** will NOT be recorded.

Logistics

Grading Policy

Student presentation (20%)

Paper abstracts/summaries and Quizzes (30%)

Course project (40%)

In-class participation (10%)

20% each

- **Provide feedback to your peers!**

Mind your biases!

Be respectful and constructive

Logistics

Grading Policy

Student presentation (20%)

Paper abstracts/summaries and Quizzes (30%)

Course project (40%)

In-class participation (10%)

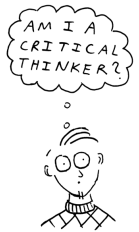

0.75% each x 40 abstracts/summaries = 30%

- Due by **9:59pm** the previous night of each student presentation
- Write an abstract for **each paper** from the required readings (three per class)
- **No late date** - but more than 40 papers (60, feel free to skip some)
- Have energy to do more? **Top-scored 40** for grading
- **Quizzes** about the 3 papers to be presented: multiple choice, very easy (if you have read the papers)
- **50% correct answers in the paper Quiz, class attendance and participation** is required for abstract grades

Logistics

Grading Policy

Student presentation (20%)

Paper abstracts/summaries and Quizzes (30%)

Course project (40%)

In-class participation (10%)

40%

- Project Proposal (5%). Due Thu Sept 15.
- Project Milestone (5%). Due Thu Oct 20.
- Final Report (20%). Due Fri Dec 9.
- Spotlight Talk and Interactive Poster session (10%). Week 15.

Hands-on experience of robot learning research

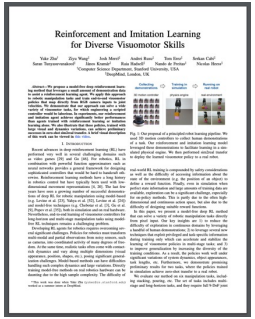

+

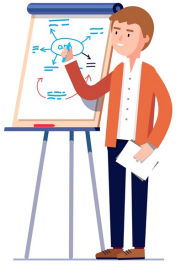

Logistics

Grading Policy

Student presentation (20%)

Paper abstracts/summaries and Quizzes (30%)

Course project (40%)

In-class participation (10%)

Tutorials, computing resources, project instructions, …

list of project topics online on canvas

Discussions on Piazza to form groups
1 group per project topic

Enroll in the sheet we will send later today (firs come, first serve)

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

Alternative projects require instructor approval.

Logistics

Grading Policy

Student presentation (20%)

Paper abstracts/summaries and Quizzes (30%)

Course project (40%)

In-class participation (10%)

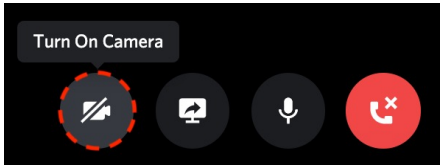

Tell Us About Yourself

- Name (with pronouns)
- Background (academic)
- Why are you interested in Robot Learning?
- What would you like to get out the course?

Robotics beyond CS391F

Be part of the Robotics + AI revolution.

RobIn
ROBOTIC INTERACTIVE INTELLIGENCE LAB

Study and develop novel intelligent algorithms

Combine machine learning, computer vision, robotics

TEXAS Robotics

https://robotics.utexas.edu/

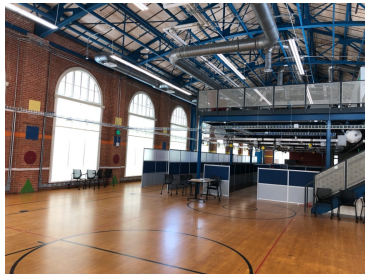

Summary

- We are going to learn the most influential works in robot learning, perception and decision making
- We will discuss the most important challenges in robotics and the opportunities machine learning provides to overcome them
- You will practice your presentation and debate skills
- You will apply what you learn on a research project

- Find a group+project before Friday!
- Sign-up for research presentation before Friday!