CS391R: Robot Learning

Perception, Decision Making, and Research Frontiers

Prof. Yuke Zhu

Fall 2023


70
302331
302331
302159
DHL
microdrones.com
M
PACKSET

Traditional form of **automation**

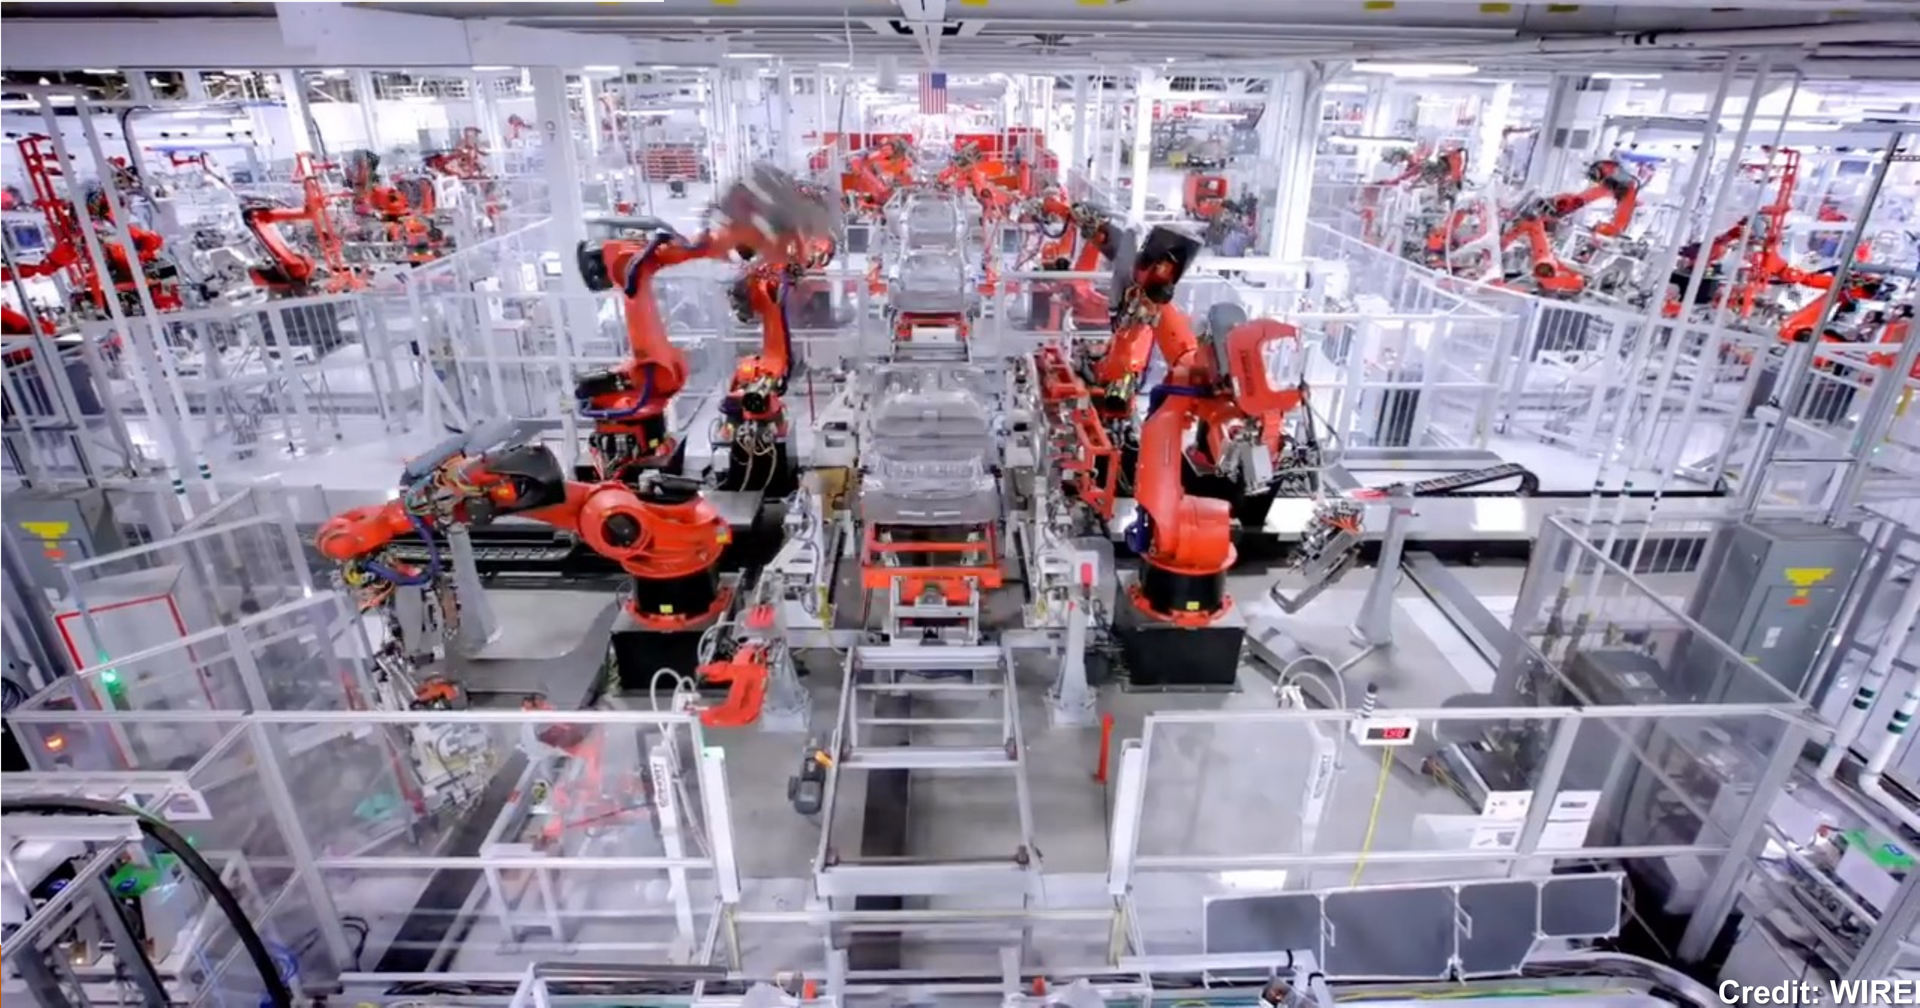

Credit: WIRED

General-purpose **robot autonomy**

Today’s Agenda

- What is Robot Learning?
- Why studying Robot Learning now?
- Course topic preview
- Logistics
- Student introduction

Special-Purpose Robot Automation

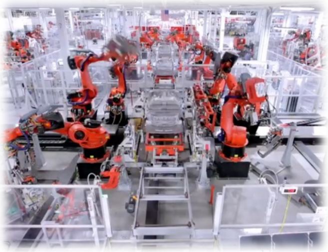

custom-built
robots

human expert
programming

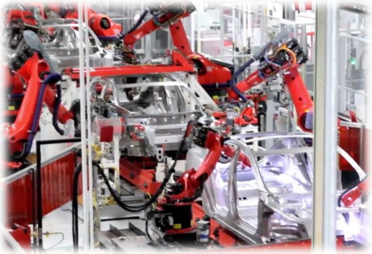

special-purpose
behaviors

General-Purpose Robot Autonomy

general-purpose
robots

?

general-purpose
behaviors

Special-Purpose Robot Automation

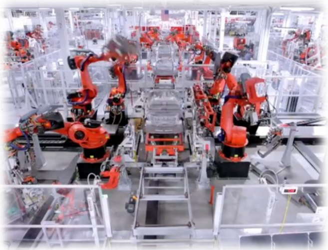

custom-built
robots

human expert
programming

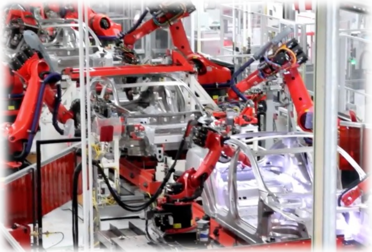

special-purpose
behaviors

General-Purpose Robot Autonomy

general-purpose
robots

Robot Learning

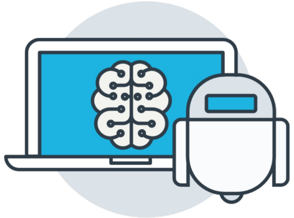

general-purpose
behaviors

General-Purpose Robot Autonomy: Imaginations

Unimate - The First Industrial Robot

British TV (1968)

General-Purpose Robot Autonomy: Challenges

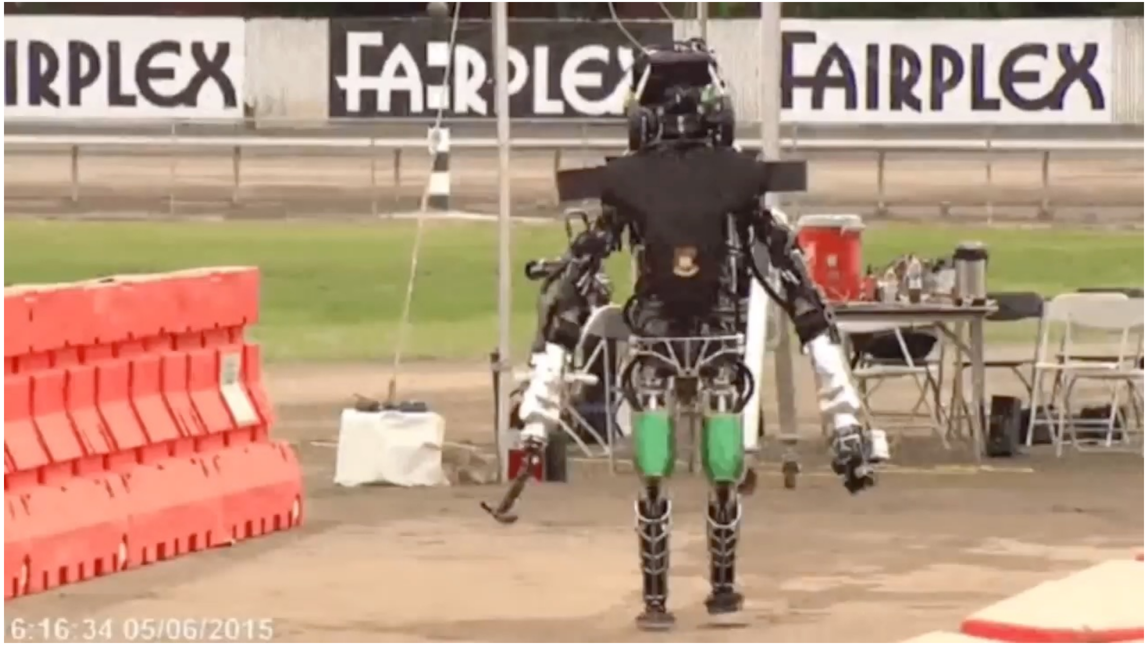

DARPA Robotics Challenge (2015)

“The Moravec's paradox”

General-Purpose Robot Autonomy: Progress

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

Manipulation
(OpenAI; 2019)

Locomotion
(ETHZ and ANYbotics; 2020)

Mobile Manipulation
(Google Robotics; 2023)

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)

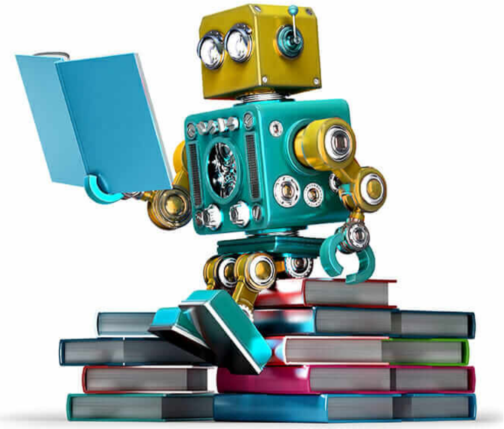

When **NOT** to Make Robots Learn?

Learning is not the solution to every (generalization) 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.

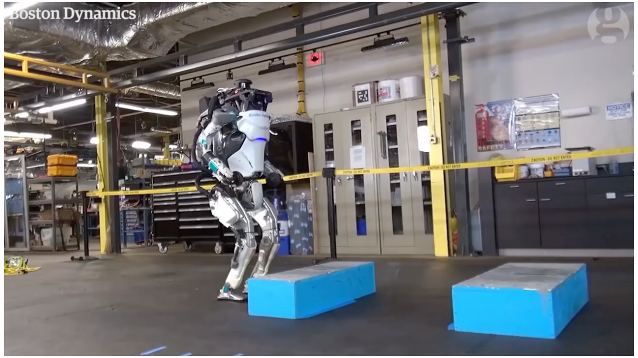

When to Make Robots Learn?

Learning is critical for deploying robots to the real world.

real-world variations

environmental uncertainty

needs for adaptation

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

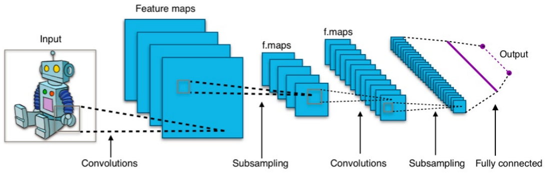

Artificial Intelligence

Recent breakthroughs in AI, such as deep learning (Turing awards 2018) in computer vision, natural language processing, etc.

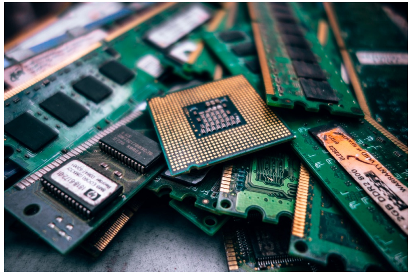

Compute & Data

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

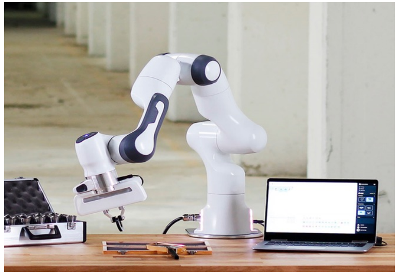

Robot Hardware

Reliable and affordable robot hardware that costs around annual salary of American workers

Robot Learning as a Growing Research Community

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

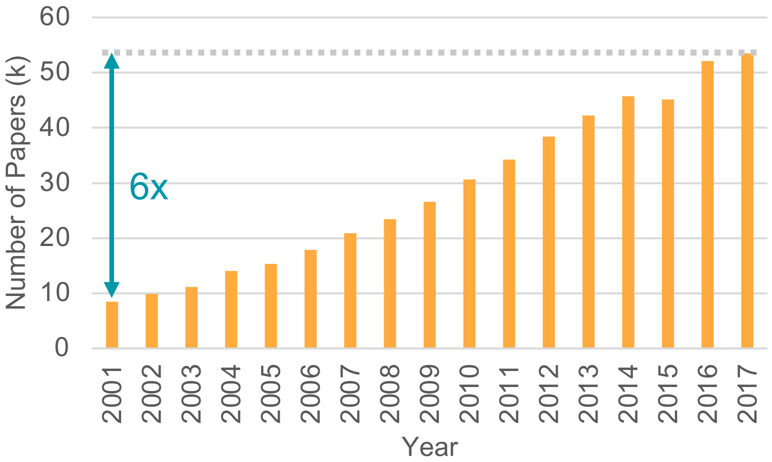

Growth of "**Robot Learning**" Publications

[Source: Google Scholar]

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

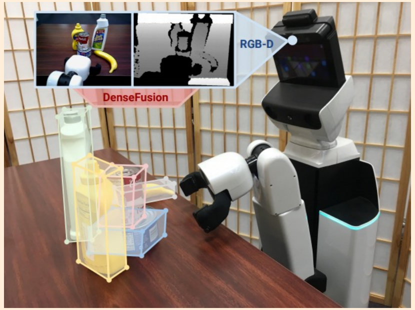

Topics #1-7

seeing and understanding the physical world

Part II: Robot Decision Making

Topics #8-14

planning and control for robot behaviors

Part III: Research Frontiers

Topics #15-18

latest research advances and open problems

Prerequisite: coursework / experience in AI and Machine Learning

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

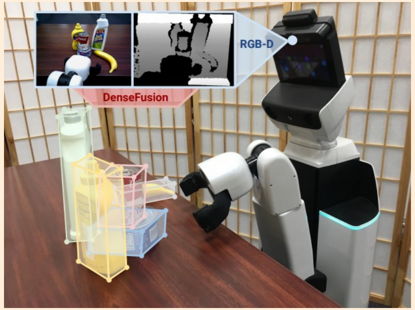

Topics #1-7

seeing and understanding the physical world

Part II: Robot Decision Making

Topics #8-14

planning and control for robot behaviors

Part III: Research Frontiers

Topics #15-18

latest research advances and open problems

Prerequisite: coursework / experience in AI and Machine Learning

Robot Perception

2D visual recognition

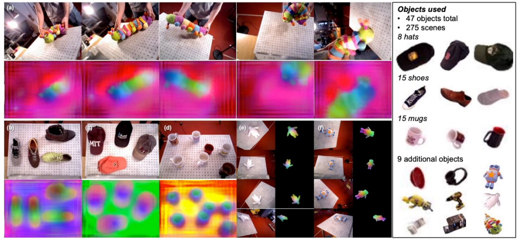

representation learning for robotics

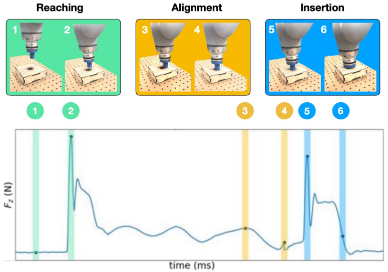

multimodal perception

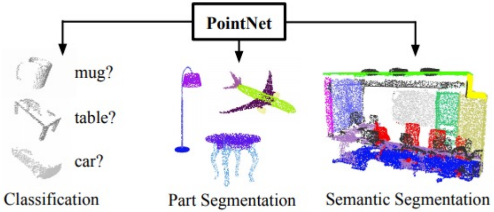

3D data processing

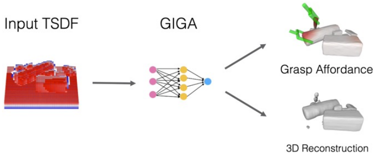

neural fields

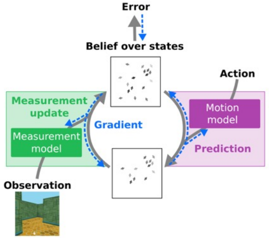

state estimation

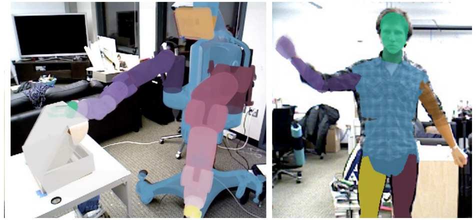

visual tracking

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

Topics #1-7

seeing and understanding the physical world

Part II: Robot Decision Making

Topics #8-14

planning and control for robot behaviors

Part III: Research Frontiers

Topics #15-18

latest research advances and open problems

Prerequisite: coursework / experience in AI and Machine Learning

Robot Decision Making

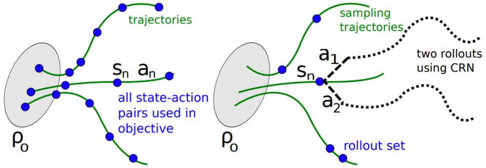

model-free RL

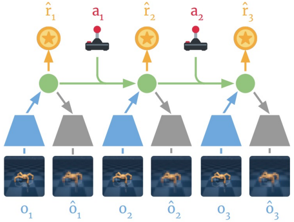

model-based RL

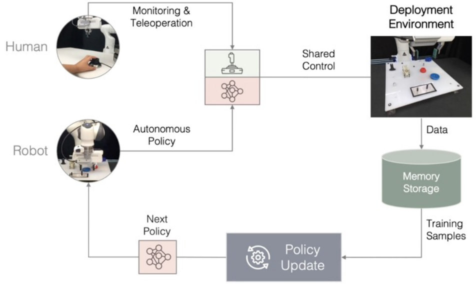

human-in-the-loop learning

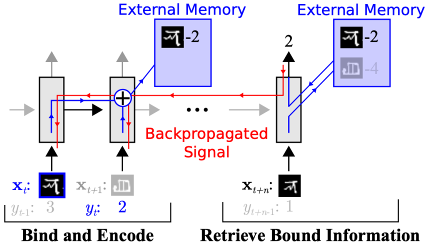

learning to learn

imitation as supervised learning

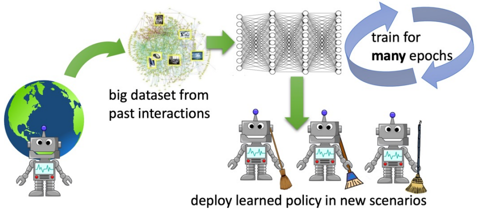

offline RL

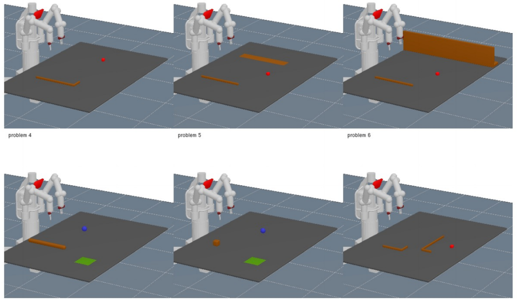

task and motion planning

Course Content

We review the Robot Learning literature in these topics.

Part I: Robot Perception

Topics #1-7

seeing and understanding

the physical world

Part II: Robot Decision Making

Topics #8-14

planning and control for

robot behaviors

Part III: Research Frontiers

Topics #15-18

latest research advances

and open problems

Prerequisite: coursework / experience in AI and Machine Learning

Research Frontiers

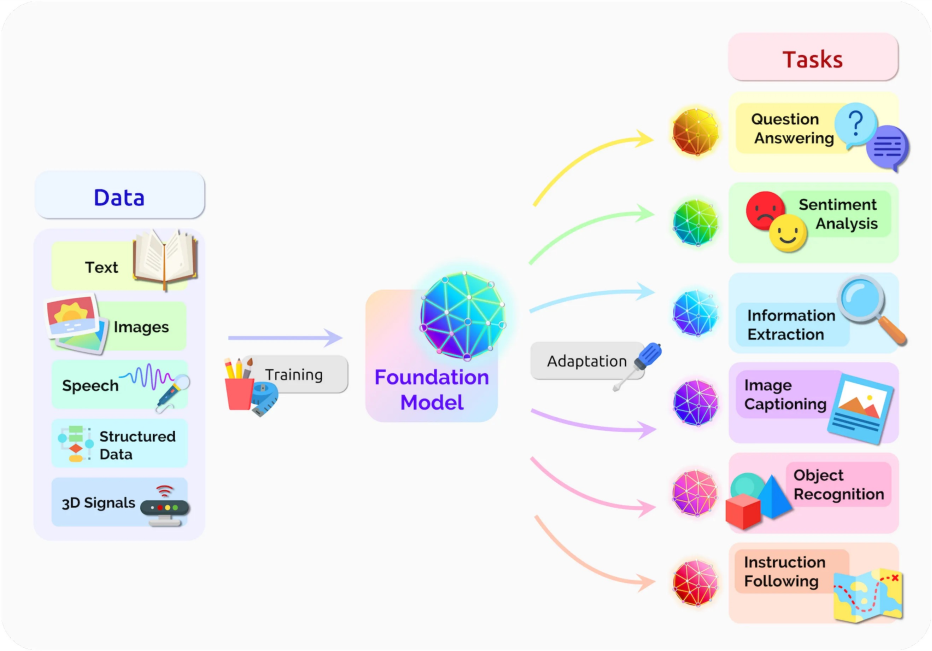

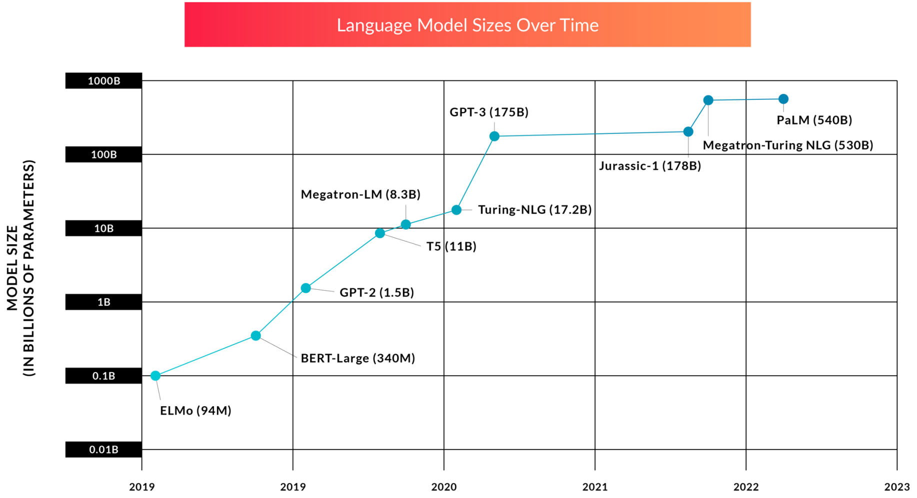

"the era of big models"

Research Frontiers

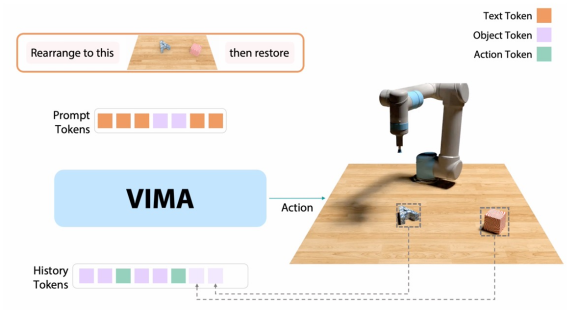

attention architectures

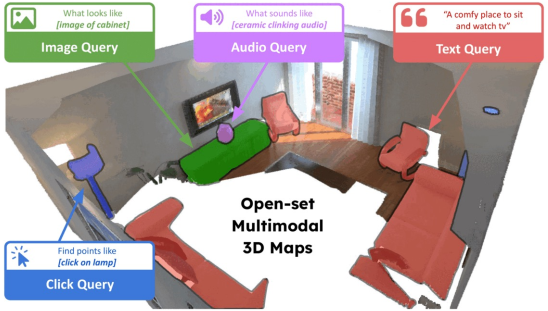

language in robotics

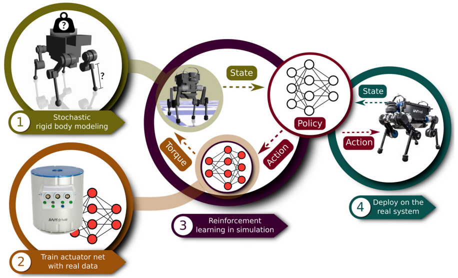

sim-to-real gap

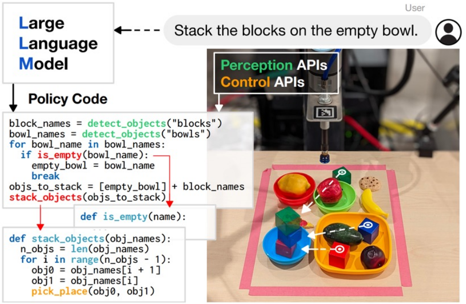

program synthesis for embodied agents

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: GDC 1.406 (in person)

Office Hours

Instructor: By appointment via email (GDC 3.422)

TA: 9:30-10:30am Wednesdays (GDC 3.416)

TA: Soroush Nasiriany

Logistics

Date	Topic
Week 1 Tue, Aug 22	Lecture Course Overview ▪ Building Machines That Learn and Think Like People. Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, Samuel J. Gershman (2016) ◦ Intelligence without Reason. Rodney Brooks (1991)
	Part I: Robot Perception
Week 1 Thu, Aug 24	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 Tue, Aug 29	2D Visual Recognition • You Only Look Once: Unified, Real-Time Object Detection. Joseph Redmon, Santosh Divvala, Ross Girshick, Ali Farhadi (2015) • Segment Anything. Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, Ross Girshick (2023) ◦ End-to-End Object Detection with Transformers. Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, Sergey Zagoruyko (2020) ◦ Mask R-CNN. Kaiming He, Georgia Gkioxari, Piotr Dollar, Ross Girshick (2017)
Week 15 Tue, Nov 28	Spotlight Final Project Spotlights I
Week 15 Thu, Nov 30	Spotlight Final Project Spotlights II
Week 16 Fri, Dec 8	No Class

■ **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)
- **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 2023) students to submit the paper reviews. The paper reviews must be submitted by 9:59pm the previous night for each class of student presentations in order to receive a grade.

- **No late date** - more than 15 presentation classes (can skip some)
- Check our policy on using **AI writing tools** like ChatGPT
- Have energy to do more? **Top-scored 15** for grading
- **Class attendance and participation** is required for review grades

Logistics

ed discussion

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 2023) students to submit the paper reviews. The paper reviews must be submitted by 9:59pm the previous night for each class of student presentations in order to receive a grade.

- **No late date** - more than 15 presentation classes (can skip some)
- Check our policy on using **AI writing tools** like ChatGPT
- 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 14.
- Project Milestone (5%). Due Thu Oct 19.
- Final Report (25%). Due Fri Dec 8.
- Spotlight Talk (5%). Week 15.

Hands-on experience of robot learning research

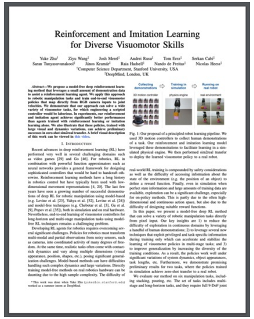
Reinforcement and Imitation Learning for Diverse Visuomotor Skills

+

Logistics

Grading Policy

Student presentation (20%)

Paper reviews (30%)

Course project (40%)

In-class participation (10%)

Default project: **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

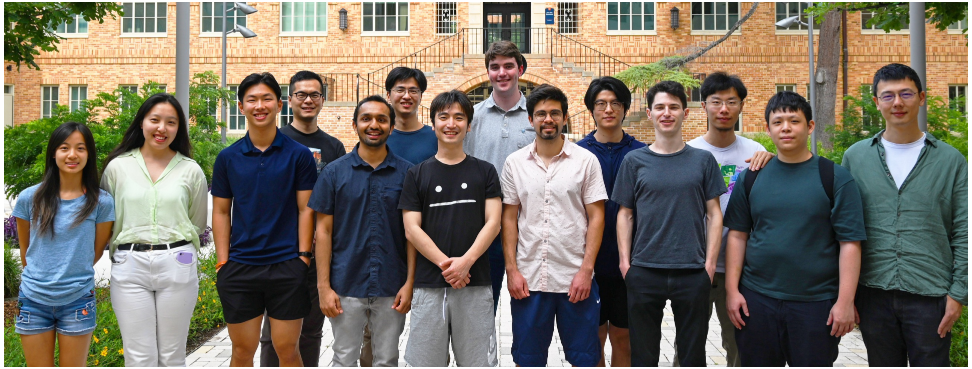

Mission: Building General-Purpose Robot Autonomy in the Wild

TEXAS Robotics

https://robotics.utexas.edu/