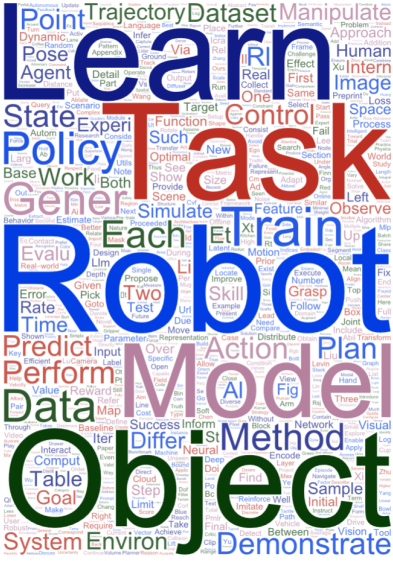

Robot Learning

Introduction

[word cloud from all CoRL 2023 papers]

Marc Toussaint
Technical University of Berlin
Summer 2024

What is this lecture about?

- Related Lectures:
  - Guanya Shi (CMU): Robot Learning `https://16-831-s24.github.io/lectures`
  - Erdem Biyik (USC): `https://liralab.usc.edu/csci699/`
  - Jan Peters (TU Darmstadt): `https://learn.ki-campus.org/courses/moocrobot-tud2021`
  - Yisong Yue & Hoang M. Le (CalTech): `https://sites.google.com/view/icml2018-imitation-learning/`

What is this lecture about?

- Shi's lecture (referenced below):

Let's Start!

What is "robot learning" and what is this class about?

❑ **Learning** to make **sequential** decisions in the **physical world**

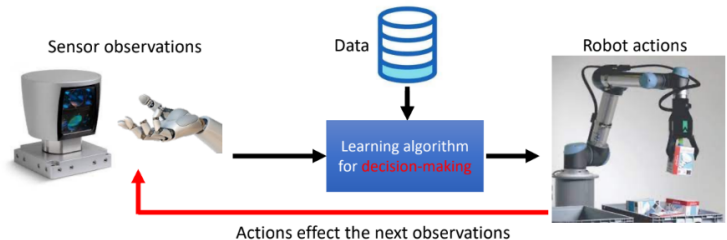

What is this lecture about?

- Shi's lecture (referenced below):

What is "Robot Learning"?

- **Learning** to make **sequential** decisions in the **physical world**
- **Learning**: Data-driven and improve from data
  - W/o learning & data: search & planning, classic control, optimal control, …
- **Sequential**: The current action/decision influences the next state therefore the next action/decision
  - W/o sequential: bandit, standard supervised learning, …
- **Physical world**: The robot needs to interact with the physical world in the closed-loop
  - A.k.a. "embodied intelligence"
  - W/o physical world: RL for games, LLMs…

What is this lecture about?

- In Shi's view:
  - Formalize the problem "making sequential decisions in a physical world" ($\rightarrow$ MDPs)
  - Focus on Learning in MDPs $\rightarrow$ Reinforcement Learning

What is this lecture about?

- However, the topic is much wider
- Robotics is a very wide field – **Learning can be applied almost anywhere**

What is this lecture about?

- Module description (Moses 41016) – Learning Outcomes
  - The students have a systematic understanding of the wide variety of contexts and problems settings in which machine learning methods can be applied within robotics.
  - They understand how the learning problems are mathematically formulated in these settings.
  - [They also learn about underlying ML methods to tackle these problems.]. . .
- Content
  - The term Robot Learning generally denotes the use of learning methods in the context of robotics, which is ubiquitous in modern robotics research. This course aims to provide a systematic introduction to the field, in particular to the various contexts and problem setting where machine learning can be applied and the specific learning methods themselves. This includes topics such as:
  - System identification, model learning, residual model learning
  - Imitation learning, behavior cloning, learning from demonstration
  - Reinforcement Learning (RL), skill learning, offline RL
  - Constraint learning, grasp learning, iterative learning control
  - Learning to predict plans, learning to warmstart MPC or optimization
  - Inverse RL
  - . . .

Motivation

- OpenAI / Figure robot: `https://www.youtube.com/watch?v=Sq1QZB5baNw`
- Boston Dynamics: `https://www.youtube.com/watch?v=tF4DML7FIWk`

Motivation

- OpenAI / Figure robot: https://www.youtube.com/watch?v=Sq1QZB5baNw
- Boston Dynamics: https://www.youtube.com/watch?v=tF4DML7FIWk

- CoRL 2023 award/finalist papers:
  - https://hshi74.github.io/robocook/
  - https://mimic-play.github.io/
  - https://robot-parkour.github.io/

The State-of-the-Art in Robot Learning

- Conference on Robot Learning `https://www.corl.org/`
- Robotics: Science and Systems Conference `https://roboticsconference.org/`
- ICRA, IROS, L4C conferences
- NeurIPS, ICML conferences

- The meta-goal of this lecture:
  Enable you to read & understand papers at these conferences

- The meta-goal of this lecture:
  Enable you to read & understand papers at these conferences

- Some of the lectures will directly discuss essential research papers

Planned Lectures

- Taxonomy (today)
- Robotics Primer & Machine Learning Primer
- Dynamics Learning / System Identification
- Imitation Learning
- *Method Lecture:* Diffusion & other policy representations
- Reinforcement Learning & variants (several lectures)

Planned Lectures

- Taxonomy (today)
- Robotics Primer & Machine Learning Primer
- Dynamics Learning / System Identification
- Imitation Learning
- *Method Lecture:* Diffusion & other policy representations
- Reinforcement Learning & variants (several lectures)
- Safe Learning, Multi-Robot Learning
- Constraint Learning, Grasping/Manipulation Learning, Affordance Learning
- *Method Lecture:* Robotics/3D ML: Rotation encodings, PointNet, SE(3)-Equivariant
- *Method Lecture:* Black-Box Optimization, CMA, CEM
- Plan Prediction Learning (from MPC to Language Models)
- Online adaptation
- *Method Lecture:* Generative models (PCA, auto encoder, VAE, GANs, diffusion, stochastic outputs in transformers)

Organization

Organization

- 6 LPs (180h, 12h/w, 15 weeks)
- Lectures, weekly, in person
- Tutorials, weekly:
  - Weekly exercise sheets, mix of analytic/coding, to be discussed in the tutorials

- ISIS as central webpage
- Contact:
  - Office (grades/etc): Ilaria Cicchetti-Nilsson <office@lis.ut-berlin.de>

Ilaria Cicchetti-Nilsson

Assignments & Exam

- Tutorial exercises are a mix of analytic and coding problems. **Voting System:**
  - When attending a tutorial, students mark in an ISIS questionnaire which exercises they have worked on
  - Students are randomly selected to present their solutions (no need for correct solutions – just something to present and discuss)
  - When not attending: upload pdf notes/solutions on ISIS

- **Exam prerequisite:**
  - at least 50% votes in the exercises

- The written exam will be about analytical problems, determines final grade (no portfolio)

Prerequisites

- Module description:
  - Knowledge in Machine Learning
  - Fundamentals in AI (esp. Markov Decision Processes)
  - Foundations of robotics
  - Basic programming skills

Prerequisites

- Module description:
  - Knowledge in Machine Learning
  - Fundamentals in AI (esp. Markov Decision Processes)
  - Foundations of robotics
  - Basic programming skills
- Self-Checks:
  - Maths, AI, ML & Robotics lectures:
    `https://www.user.tu-berlin.de/mtoussai/teaching/Lecture-Maths.pdf`
    `https://www.user.tu-berlin.de/mtoussai/teaching/Lecture-AI.pdf`
    `https://www.user.tu-berlin.de/mtoussai/teaching/Lecture-MachineLearning.pdf`
    `https://www.user.tu-berlin.de/mtoussai/teaching/Lecture-Robotics.pdf`
  - ML: not only pyTorch.. but also *Hastie et al: The Elements of Statistical Learning*?
    `https://hastie.su.domains/Papers/ESLII.pdf`
  - For reference:
    `https://www.user.tu-berlin.de/mtoussai/teaching/#reference-material`

Prerequisites

- Module description:
  - Knowledge in Machine Learning
  - Fundamentals in AI (esp. Markov Decision Processes)
  - Foundations of robotics
  - Basic programming skills
- Self-Checks:
  - Maths, AI, ML & Robotics lectures:
    `https://www.user.tu-berlin.de/mtoussai/teaching/Lecture-Maths.pdf`
    `https://www.user.tu-berlin.de/mtoussai/teaching/Lecture-AI.pdf`
    `https://www.user.tu-berlin.de/mtoussai/teaching/Lecture-MachineLearning.pdf`
    `https://www.user.tu-berlin.de/mtoussai/teaching/Lecture-Robotics.pdf`
  - ML: not only pyTorch.. but also *Hastie et al: The Elements of Statistical Learning*?
    `https://hastie.su.domains/Papers/ESLII.pdf`
  - For reference:
    `https://www.user.tu-berlin.de/mtoussai/teaching/#reference-material`
- Numeric coding in Python (numpy)

Module description (Moses 41016)

- Grading
  - graded, written exam, English (90min)
- This module is used in the following module lists:
  - Automotive Systems (M. Sc.)
  - Computer Engineering (M. Sc.)
  - Computer Science (Informatik) (M. Sc.)
  - Elektrotechnik (M. Sc.)