EC 518: Robot Learning

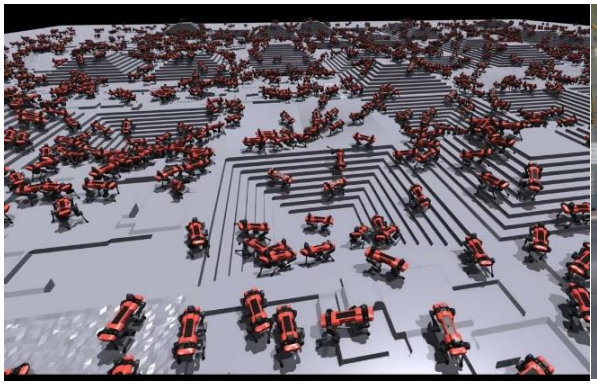

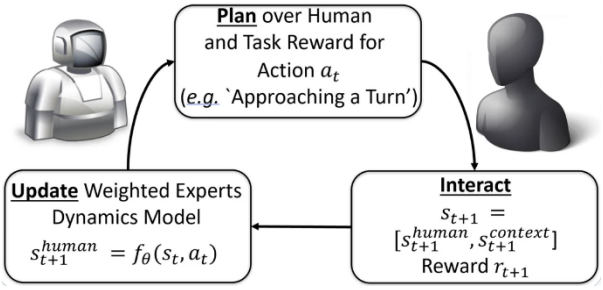

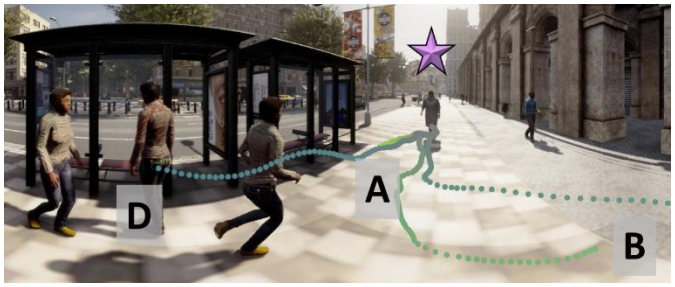

Course Info:
Mon/Wed 4:30-6:15PM, PHO 205

Instructor Name:
Eshed Ohn-Bar, eohnbar@bu.edu

Office Hours: Wed 6:15-7:15PM
Blackboard: https://learn.bu.edu

Course Description

How can we design robotic systems that robustly operate in complex, dynamic, and uncertain real-world environments? Towards this overarching goal, this class will discuss recent developments in machine learning and perception for robotics. Specifically, we will study advanced concepts in perception, decision-making, and interaction algorithms in order to provide theoretical and experimental frameworks for understanding current limitations in state-of-the-art approaches for robot learning. Topics will include sensorimotor paradigms for perception and action, robot reinforcement learning, imitation learning, inverse reinforcement learning, exploration, hierarchical learning, model-based approaches, and human-machine interaction. We will particularly focus on one of society's most important and urgent problems, mobility, to analyze and address key challenges in state-of-the-art robot learning techniques.

Prerequisites

We will extensively leverage knowledge in Machine Learning, Linear Algebra, Python, and PyTorch. It is highly recommended to have have prior experience with these topics.

Learning Objective

We will cover advanced techniques in computer vision and robot learning for real-world systems. At the end of this course, you will gain an understanding of fundamental paradigms and challenges in learning-based, state-of-the-art autonomous systems:

- How to formulate learning algorithms for general robotics and sensorimotor navigation
- Explore challenges in perception-based control pipelines
- Affordance-based and goal-directed representation of semantic scene, motion, and geometry
- Limitations (and complementarity) of reinforcement, inverse reinforcement, and imitation learning approaches
- Policy optimization (e.g., model-based vs. model-free) and sample efficiency analysis
- Multi-modal algorithms for multi-sensor (camera, LIDAR, and radar perception) and multi-task learning
- Human-in-the-loop learning
- Reading and presenting recent research papers
- Participation in a competition to put these ideas to the test (for awards and fame)

ChatGPT Policy

You are welcomed to use ChatGPT as you would like. However, we ask that you clearly describe and attribute which components were generated by ChatGPT and how it was used.

Grading

Class Midterms (30%): Two midterm exams will be given and will draw on concepts from presented material and papers.

Homework (20%): Homework will require coding in Python and PyTorch. GPU-cluster access will be provided and are highly recommended. Familiarity with Python will be assumed. Prior experience with PyTorch or Tensorflow is not required but is a plus.

Presentations and Participation (15%): Each student will get the opportunity to present. You will be graded based on your level of insight into the material, clarity and depth of presentation, how well you relate the paper to other papers and lecture material, as well as a lead discussion.

Final Project (35%): Each student is required to on a final research project. The project requires a 2-page proposal including the relevant literature survey, a proposal presentation, a milestone review, a 6-10 page-long final report, and a final presentation/demo. Projects are expected to be research level, uncovering new knowledge, and could be done either in simulation or on a real platform.

Tentative Schedule

Day	Topic	Notes	
1	Introduction	HW0: Due 9/11	
2	Deep Imitation Learning		
3	Deep Imitation Learning	HW1: Due in 2 weeks	
4	Affordances and Direct Perception		
5	Dynamics and Localization		
6	Semantic Scene Understanding		
7	Semantic Scene Understanding	HW2: Due in 2 weeks	
8	Object Detection and Tracking		
9	Object Detection and Tracking	Project Proposal Reports Due	
10	Markov Decision Process	HW3: Due in 2 weeks	
11	Q-Learning and SARSA		
12	Double Q-Learning and Deep RL for Robotics	Competition Presentations	
13	TRPO, DDPG, PPO	Project Proposal Revision Due	
14	LfD/Inverse Reinforcement Learning		
15	LfD/Inverse Reinforcement Learning		
16	Human-in-the-Loop Learning		
17	Human-in-the-Loop Learning		
18	Midterm Project Updates		
19	Model-based RL		
20	Model-based RL		
21	Sim2Real		
22	Bio-Inspired and Evolutionary Techniques		
24	Incorporating Language		
25	Partial Observability		
26	Human-Robot Interaction		
27	Ethics and Social Implications		
28	Research Presentations		
	Research Paper Submission		