EC 500: Robot Learning and Vision for Navigation

Course Info:
Tue/Thu 3:30-5:15PM, EPC 206

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

Teaching Assistant:
Yang Hu, huuyang@bu.edu

Office Hours: Thu 5:15-6:15PM
Blackboard: https://learn.bu.edu

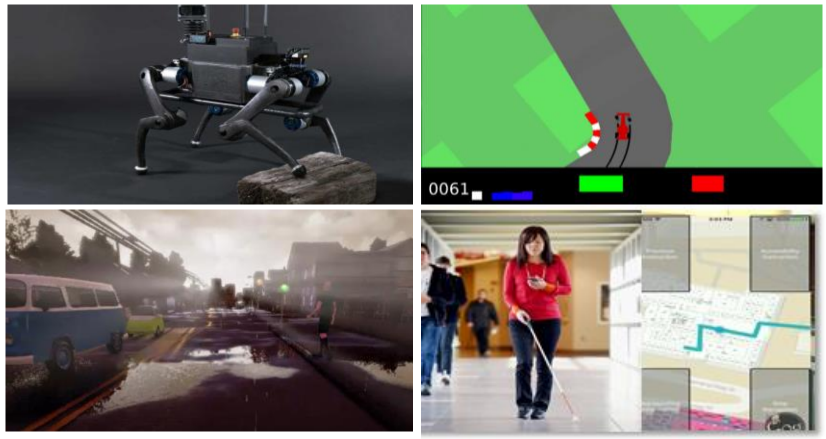

Course Description

In this class we will focus on one of society's most important problems: mobility. The task of autonomous navigation in complex and dynamic 3D environments is a fundamental (and urgent) engineering problem. There is a wide range of applications, from assistive systems to self-driving cars. Despite decades of research, existing perception, decision-making, and interaction algorithms for navigation perform poorly when compared to strategies employed by humans. **The goal of this class is to provide theoretical and experimental frameworks for understanding such limitations while exploring novel solutions.** We will use simulation environments and real-world datasets to analyze key questions, such as paradigms for the coupling of perception and control.

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 sensori-motor navigation
- Explore challenges in perception-based control pipelines
- Affordance-based and goal-directed representation of semantic scene, motion, and geometry
- Limitations of reinforcement, inverse reinforcement, and imitation learning methods
- Benefits of self-supervised techniques in computer vision
- Policy optimization techniques (model-based vs. model-free vs. evolutionary), sample efficiency
- Multi-modal algorithms for camera, LIDAR and radar perception
- Learning to reason in language-driven navigation
- Modeling human-in-the-loop systems, predicting behavior in social, dynamic, uncertain settings
- Project-based analysis, case-study implementation, and communication of research ideas
- Participation in a competition to put these ideas to the test (for awards and fame)

Learning from Anywhere

Class material and relevant readings to be posted prior to each class (a highly recommended book is Reinforcement Learning by Sutton and Barto, freely available at http://incompleteideas.net/book/RLbook2018.pdf). The class can be taken either as an in-person or remote format, however, please notify the instructor in the case of switching between remote and in-person class attendance. The instructor will make efforts to facilitate active participation by students who, due to time zone or illness, must take a class asynchronously. Classes will be recorded in accordance with LfA class recording policy, https://digital.bu.edu/lfa-classroom-recordings/, and shared via Blackboard. In particular, the class materials and recordings cannot be shared with anyone not registered in the course. Please contact the instructor if you require any additional accommodations.

Grading

Homework (30%): Homework problems can be worked on individually or in groups of two (1 report per group). Homework will require coding in Python and PyTorch. GPU-cluster access will be provided and are highly recommended. Make sure you are familiar with Python. Prior experience with PyTorch or Tensorflow is not required but a plus.

Presentations and participation (20%): Each student will get the opportunity to present. You will be graded based on your level of insight into the material, how well you relate the paper to other papers and lecture material, as well as how well you present the material to the class and lead a discussion.

Final Project (50%): Each student is required to work individually or in groups of two on a research project. The project requires a 1-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.

Tentative Schedule

Date	Topic	Notes	
9/3	Introduction: Why Navigation?	HW0: Due 9/8	
9/8	Deep Imitation Learning		
9/10	Deep Imitation Learning	HW1: Due in 2 weeks	
9/15	Affordances and Direct Perception		
9/17	Vehicle Dynamics and Localization		
9/22	Semantic Scene Understanding		
9/24	Semantic Scene Understanding	HW2: Due in 2 weeks	
9/29	Object Detection and Tracking		
10/1	Object Detection and Tracking	Project Proposal Reports Due	

10/6	Markov Decision Process	HW3: Due in 2 weeks	
10/8	Q-Learning and SARSA		
10/13	Double Q-Learning, Experience Replay Techniques	Competition Presentations	
10/15	Advanced RL Topics		
10/20	Human-in-the-Loop Reinforcement Learning		
10/22	Learning from Demonstrations (LfD)	Project Proposal Revision Due	
10/27	LfD/Inverse Reinforcement Learning		
10/29	LfD/Inverse Reinforcement Learning		
11/3	Model-based RL		
11/5	Midterm Project Updates		
11/10	Pedestrian Prediction and Social Robotics		
11/12	Perception-Action in Connected Systems		
11/17	Sim2Real		
11/19	Embodied Navigation		
11/24	Bio-Inspired and Evolutionary Techniques		
11/26	Thanksgiving		
12/1	Language and Navigation		
12/3	Ethics and Social Implications		
12/8	Research Presentations		
12/10	Research Presentations		
12/19	Submit Research Paper		