CSC2626 – Imitation Learning for Robotics

Fall 2022

Instructor:	Florian Shkurti
Email:	csc2626-instructor@cs.toronto.edu
Lectures:	Wed 11am-1pm ET, GB244 + Zoom
Office Hours:	Mon 12-1pm ET, SF3328 + Zoom
TAs:	Jonathan Lorraine, Skylar Hao, Mohamed Khodeir
Email:	csc2626-tas@cs.toronto.edu
Office Hours:	TBA, Zoom

Course Page: http://www.cs.toronto.edu/~florian/courses/csc2626w22

Zoom Link: See the course's Quercus home page.

Hybrid Course Delivery: Lectures will be held in person in GB244, livestreamed on Zoom, recorded, and uploaded on Quercus after class. Office hours from the instructor will be held in person in SF3328. Students have the option to attend them virtually on Zoom, but they will not be recorded. TA-led office hours will be offered on Zoom only. We will use Piazza for questions and discussions, and Quercus for course announcements. Class participation and questions during lectures and office hours are very much encouraged.

Communication for the course:

- The official discussion board for the course is Quercus https://q.utoronto.ca. All course announcements will be posted there.
- Email the instructor or the TAs with "CSC2626" in the subject line, otherwise your email might get mislabeled and potentially not seen.
- You are welcome to provide anonymous feedback / suggestions for improvement any time during the semester: https://www.surveymonkey.com/r/LJJV5LY

Overview: This graduate-level course will examine some of the most important papers in imitation learning for robot control, placing more emphasis on developments in the last 10 years. Its purpose is to familiarize students with the frontiers of this research area, to help them identify open problems, and to enable them to make a novel contribution. The majority of lectures, particularly after the first two weeks of introductory material, will consist of in-class student presentations. This course will broadly cover the following areas:

- Imitating the policies of demonstrators (people, expensive algorithms, optimal controllers)
- Connections between imitation learning, optimal control, and reinforcement learning
- Learning the cost functions that best explain a set of demonstrations
- Shared autonomy between humans and robots for real-time control

The course involves a significant final project component, which will likely involve the use of robot simulators (see the course webpage for suggestions on simulators).

Prerequisites: You need to be comfortable with introductory machine learning concepts (such as from CSC411/ECE521 or equivalent), linear algebra, basic multivariate calculus, intro to probability. You also need to have strong programming skills in Python. Note: if you don't meet all the prerequisites above please contact the instructor by email. Optional, but recommended: experience with neural networks, such as from CSC321 or equivalent, and introductory-level familiarity with reinforcement learning and control.

Main References: There is no required textbook for this course. In-class discussions will be based on research papers. The following are optional, but recommended textbooks:

- Aude Billard, Sylvain Calinon, Rudiger Dillmann, Stefan Schaal, *Robot programming by demonstration*.
- Sonia Chernova, Andrea Thomaz, *Robot learning from human teachers*.
- Takayuki Osa, Joni Pajarinen, Gerhard Neumann, Andrew Bagnell, Pieter Abbeel, Jan Peters, *An algorithmic perspective on imitation learning*

Grading Policy: 2x assignments (50%) and 1x course project (50%). The grade of the course project consists of a proposal (10%), midterm progress report (5%), project presentation (5%), and a final report with code at the end of the term (30%).

Tentative Course Outline By Week:

1. Imitation vs. Robust Behavioral Cloning
2. Intro to Optimal Control and Model-Based Reinforcement Learning
3. Batch / Offline Reinforcement Learning
4. Imitation Learning Guided by Optimal Control Experts
5. Imitation as Program Induction and Modular Decomposition of Demonstrations
6. Inverse Reinforcement Learning #1
7. Shared Autonomy for Robot Control with Human in-the-Loop
8. Adversarial Imitation Learning
9. Reading Week
10. Imitation Learning Combined with Reinforcement Learning and Planning
11. Inverse Reinforcement Learning #2
12. Rewards and Value Alignment
13. Project Presentations

Important Due Dates (tentative):

Assignment 1	Sep 27, 2022, by 6pm ET
Assignment 2	Oct 18, 2022, by 6pm ET
Project Proposal	Oct 25, 2022, by 6pm ET
Midterm Report	Nov 20, 2022, by 6pm ET
Project Presentations	Dec 7, 2022, in class
Final Report and Code	Dec 12, 2022, by 6pm ET