CSC2621 – Imitation Learning for Robotics

Winter 2019

Instructor:	Florian Shkurti	**Lectures:** Fri, 1-3pm, AB107	
Email:	florian@cs.toronto.edu	**Office Hours:** Tue, 4-6pm, PT283E	
TA:	TBA	**TA Office Hours:** TBA	

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

Overview: This graduate-level seminar 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. Students who are interested in using real robot hardware and have shown sufficient progress in their final project, are encouraged to talk to the instructor about how to best arrange this.

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*

Course Communications:

- The official discussion board for the course is Quercus https://q.utoronto.ca. Announcements will be posted there, too.
- Email the instructor or the TA with CSC2621 in the subject line
- You are welcome to provide anonymous feedback / suggestions for improvement any time during the semester: https://www.surveymonkey.com/r/LJJV5LY

Grading Policy: 1x assignment (20%), 1x paper presentation (20%), and 1x course project (60%). The grade of the course project consists of a proposal (10%), midterm progress report (10%), project presentation (10%), 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. Query-Efficient Policy Imitation via Novel State Detection
4. Imitation Learning Combined with Reinforcement Learning, Control, and Planning #1
5. Imitation Learning Combined with Reinforcement Learning, Control, and Planning #2
6. Imitation as Program Induction and Modular Decomposition of Demonstrations
7. Reading Week
8. Inverse Reinforcement Learning #1
9. Inverse Reinforcement Learning #2
10. Shared Autonomy for Robot Control with Human in-the-Loop
11. Adversarial Imitation Learning
12. Project Presentations #1
13. Project Presentations #2

Important Due Dates (tentative):

Assignment ..... Feb 1, 2019, by 6pm EST
Project Proposal ..... Feb 6, 2019, by 6pm EST
Reading Week ..... Feb 18-22, 2019
Midterm Project Report ..... Mar 4, 2019, by 6pm EST
Project Presentations ..... Mar 29 and Apr 5, 2019
Project Report and Code ... Apr 10, 2019, by 6pm EST

Class Attendance Policy: Regular attendance and class participation with questions is essential. I expect you to be engaged and actively contribute to in-class discussions.