Dynamic and Distributed Decision Making (MIE567S)

Course Description:
This course is to provide fundamental concepts and mathematical frameworks for sequential decision making of a team of decision makers in the presence of uncertainty. Topics include Markov decision processes, reinforcement learning, theory of games, multi-agent reinforcement learning and decentralized Markov decision processes. The course is technical by nature and for advanced students with strong mathematical background and programming skills.

Textbook: No textbook

References: Reinforcement Learning: An Introduction, Second Ed., R. Sutton and A. Barto, MIT Press, 2018
Markov Decision Processes in Artificial Intelligence, edited by Sigaud and Buffet, Wiley, 2008
Games of Strategy, Dixit and Skeath, Reiley, W.W. North & Company, 2014

Instructor: Chi-Guhn Lee
Phone: (416) 946-7867
E-mail: cglee@mie.utoronto.ca

TA: TBA

List of topics:

1. Markov decision processes
   a. Finite MDP (Ch 3 of Sutton and Barto)
   b. Algorithms (Ch 4 of Sutton and Barto)
2. Reinforcement Learning
   a. Sampling-based approach to MDP (Ch 5 of Sutton and Barto)
   b. Various RL algorithms (Ch 6, 7, 8 of Sutton and Barto)
3. Game theory and Multi-agent Reinforcement Learning
   a. Basic concepts and theories (Dixit et al.)
   b. Stochastic games (Ch 8 of Sigaud and Buffet)
   c. Multi-agent reinforcement learning

Learning Objectives:

By successfully completing the course, each student will be able to

- Formulate and solve dynamic decision problem as a Markov decision process
- Solve complex MDP problems using reinforcement learning methods
- Understand fundamental concepts and theories in multi-agent reinforcement learning
- Implement the standard and refined algorithms for multi-agent learning and decision problems

Course Web: Quercus (MIE567) at q.utoronto.ca

Bulletin Board: Please post your course related questions in the "discussion board (Q&A)" of our course web.
Email questions won't be answered.

Office Hours: I will stay longer after lecture on Thursday

Tutorial Hours: First tutorial Second week of the semester
A TA will be available to answer questions related to course materials and any other course related issues.

Evaluation:

35 %	Final Examination
45 %	4 Projects (11.25% each)
20 %	Mid-term exam

Important Dates:

Exam	Week prior to the reading week (2-hour in class)
Final Exam:	TBA

Policies

1. Complaints regarding marks will have to be submitted to the grading TA via email using the following subject line within a week from the distribution of graded works.

   `Subject line: MIE567 Grading Issues from (your student #)`

2. Academic integrity

   Students are expected to conduct themselves in accordance with the highest ethical standards of the Profession of Engineering and evince academic integrity in all their pursuits and activities at the university. As such, in accordance with the General Academic Regulations on Academic Integrity, students are reminded that plagiarism or any other form of cheating in examinations, term tests, assignments, projects, or laboratory reports is subject to serious academic penalty (e.g. suspension or expulsion from the faculty or university). A student found guilty of contributing to cheating by another student is also subject to serious academic penalty.