
Reinforcement Learning (MIE1630)

Course Description:

This course is to provide fundamental concepts and mathematical frameworks for reinforcement learning. Specific topics include Markov decision processes, tabular reinforcement learning, policy gradient methods, and function approximation such as deep reinforcement learning. Optional topics are distributional reinforcement learning, model-based methods, off-line learning, inverse reinforcement learning and multi-agent reinforcement learning. The course is designed to allow research students to experience research on reinforcement learning from the perspective of methodological development or application of reinforcement learning to applications.

Textbook: No textbook

References: Reinforcement Learning: An Introduction, Second Ed., R. Sutton and A. Barto, MIT Press, 2018
Markov Decision Processes - Discrete Stochastic Dynamic Programming, Martin L. Puterman, Wiley, 1994
A list of papers is in Appendix

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

TA: To be announced

List of topics:

1. Basic concepts and simple algorithms of Reinforcement Learning (3~4 weeks)
 - a. Markov decision processes
 - b. Monte Carlo-based learning
 - c. Tabular methods: Q-learning, SARSA
 - d. Model-based learning: Dyna-Q
 - e. *Entropy, KL-divergence, Mutual Information and Variational Inference*
2. Policy gradient methods and deep reinforcement learning (2~3 weeks)
 - a. Policy gradient theorem
 - b. *Deterministic policy gradient methods*
 - c. Deep learning and reinforcement learning
 - d. *Natural policy gradient method and trust region policy optimization*
3. *Advanced topics: some of the following topics to be covered. (2~3 weeks)*
 - a. *Distributional RL*
 - b. *Soft RL*
 - c. *Model-based learning*
 - d. *Off-line learning*

Lecture Hours: Monday 13:10 – 16:00

Course Web: Quercus (MIE1630) 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.

Project	Team-based
	3 members per team
	4 deliverables: proposal, initial paper submission, reviews (individual), final paper submission
Evaluation:	15 % Project proposal and presentation
	30 % Review report
	30% Review scores received
	25 % Final paper and presentation
Important Dates:	Proposal submission 11:59 pm on Sep/27 (Fri)
	Proposal PPT submission 11:59 pm on Sep/29 (Sun)
	Proposal presentation in class on Sep/30 (Mon)
	Paper submission 11:59 pm on Nov/24 (Sun)
	Review begins 9:00 am on Nov/25 (Mon)
	Review submission 11:59 pm on Dec/2 (Mon)
	Final paper submission: 11:59 pm on Dec/6 (Fri)
	Final presentation in class on Dec/9 (Mon)
	No lectures on Oct/28 (reading week), Dec/2 (replaced by Dec/9)