ECE 1508: Reinforcement Learning

Course Logistics

Ali Bereyhi

ali.bereyhi@utoronto.ca

Department of Electrical and Computer Engineering

University of Toronto

Fall 2025

Welcome to ECE 1508!

Great pleasure to see you in ECE 1508

Reinforcement Learning

Instructor: Ali Bereyhi

- *Office:* BA 7208 at Bahen Centre for Information Technology
- *Email:* `ali.bereyhi@utoronto.ca`

Teaching Assistants: TBA

- *There will be tutorial lectures*

Where and When?

- *Tuesdays at 5:00 PM till 7:00 PM* at **BA-1170**
- *Fridays at 5:00 PM till 7:00 PM* at **BA-1180**

Quercus and Piazza

We got a Quercus page

- You have been automatically enrolled
- Also you got registered at the Piazza page
  - ↳ You can login through the Quercus page
- We though use the Course Page to share the course materials

Please! Feel free to ask questions on Piazza!

What Do We Learn?

In nutshell: we learn Reinforcement Learning!

You may wonder how do we learn it? Well! in 3 steps

- *Step 1:* **Fundamentals of Reinforcement Learning**
  - We try to get understand the underlying framework
  - We understand what the main problem is
    - ↳ *We see that it's a sort of optimization problem*
  - We get to look at some simple example

By the end of this step, we know in theory

- What kind of problems we are dealing with in Reinforcement Learning
- What methods are available to solve these problems

What Do We Learn?

In nutshell: we learn *Reinforcement Learning!*

You may wonder how do we learn it? Well! in 3 steps

- *Step 2:* **Reinforcement Learning Methods**
  - Model-based Methods
    - *In some toy-scenarios, we can write the underlying problem analytically*
    - *This is not really the case in practice though!*
  - Model-free Methods
    - *In reality, we cannot write the problem analytically!*
    - *We need to solve the problem directly from data by efficient algorithms*

As we get over this part

- You have all background that you need on Reinforcement Learning
  - *You can formulate a Reinforcement Learning problem*
  - *You can specify at least one algorithms to solve it*

But! Your algorithm might take for ever to run! ☺

What Do We Learn?

In nutshell: we learn Reinforcement Learning!

You may wonder how do we learn it? Well! in 3 steps

- *Step 3:* **Deep Reinforcement Learning**
  - We now apply deep learning to solve those hard problems
    - *We use neural networks to learn the solution from few samples*
    - *We look into Deep Q-Learning and Policy Gradient Methods*
    - *We also learn actor-critic methods*

This is the major part of the course $\approx$ 50%

- You need good background on Deep Learning, i.e., to be fairly familiar with neural networks

High Level Chart

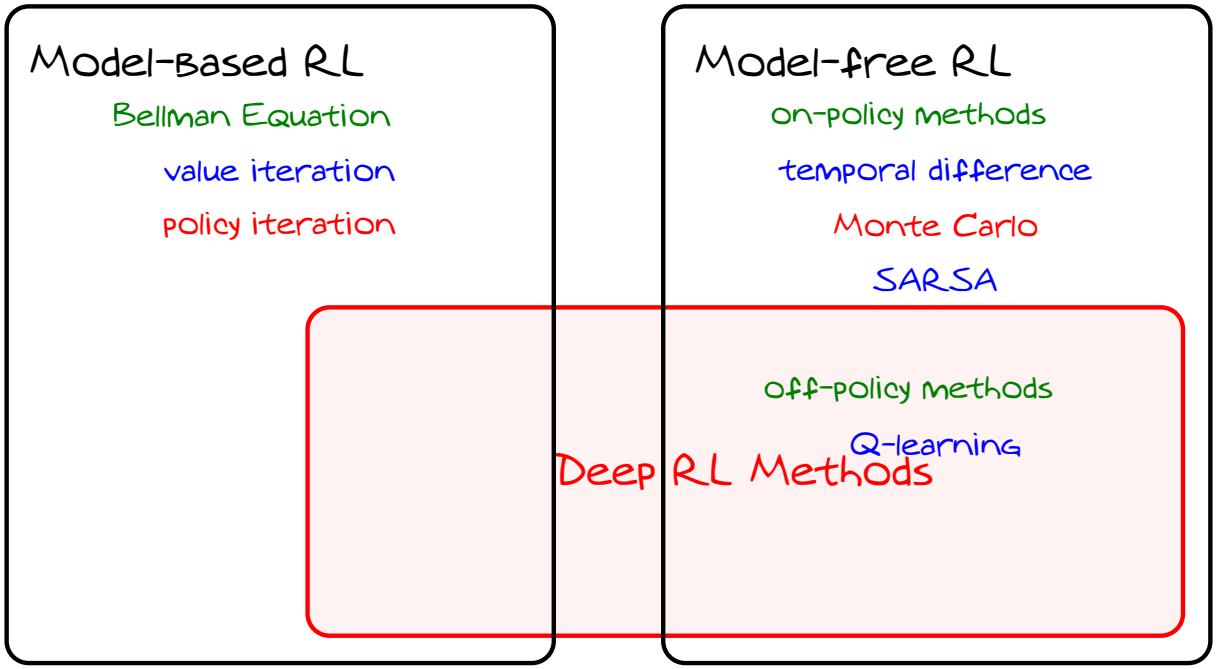

We will see this chart couple of times in this course!

How Do We Get Trained?

There are three learning components in the course

- *Assignments*
  - ↳ You get three sets of assignments
  - ↳ No need to say that they are the best thing to understand the course!
  - ↳ And, of course we do lots of programming in there!
    - ↳ *Your code should not return error!*
  - ↳ Each assignment will be solved in Tutorial after the deadline
    - ↳ Submission by deadline at 11:59 PM: full mark
    - ↳ Up to 3 days delay is allowed: each day deducts 10%
- Midterm Exam
- Final Project

How Do We Get Trained?

There are three learning components in the course

- Assignments
- *Midterm Exam*
  - ↳ We will have one written exam in the middle of semester
    - ↳ Questions that can be solved by hand, so no programming in the exam
    - ↳ It is on ***Tuesday, October 21, 2025***
    - ↳ It counts for 25% of the whole mark
- Final Project

How Do We Get Trained?

There are three learning components in the course

- Assignments
- Midterm Exam
- *Final Project*
  - ↳ The most interesting part of the course
  - ↳ We build groups by **Week 5**
    - ↳ A list of predefined project topics is provided ≈ Week 3

      Topics are all on Deep Reinforcement Learning
    - ↳ Each group chooses a topic
    - ↳ Each group member submits a progress report by **Week 10**
    - ↳ Each group submits a final paper
    - ↳ Each group submits its implementation
    - ↳ Each group presents its poster in a poster session on last week

Lots of Programming in *Python*

We are going to do lots of programming in Python

- Basic knowledge in Python is necessary
- We use *PyTorch* and *NumPy* a lot
  - *Don't run away if you haven't used them too much*
  - *If you know Python and have good programming skills, you're fine!*
- It's important to mention *knowing Deep Learning is a must*!
  - You may follow without Deep Learning till midterm
  - After midterm we need to use deep learning
- We will learn a new library
  - *Gymnasium:* standard API widely used for Reinforcement Learning

Course Calendar

Week #	Date	Notes	Posted	Deadline
1	*Sep 01 - Sep 05*			
2	*Sep 08 - Sep 12*		Assignment 1:*Basics*	
3	*Sep 15 - Sep 19*			
4	*Sep 22 - Sep 26*			Assignment 1: *Basics*
5	*Sep 29 - Oct 03*		Assignment 2: *Tabular RL*	**Project:** *Flyer and Proposal*
6	*Oct 06 - Oct 10*			
7	*Oct 13 - Oct 17*			Assignment 2: *Tabular RL*
8	*Oct 20 - Oct 24*	***Midterm Exam -- Oct 21***		
9	*Oct 27 - Oct 31*	***Reading Week-- No Lectures***		
10	*Nov 03 - Nov 07*		Assignment 3: *Deep RL*	**Project:** *Progress Briefing*
11	*Nov 10 - Nov 14*			
12	*Nov 17 - Nov 21*			Assignment 4: *Deep RL*
13	*Nov 24 - Nov 28*			
14	*Dec 01 - Dec 05*	***Last Lecture on Dec 2***		***Project: Final Submission and Presentation***

- *You could check it out here*

No Major Prerequisites

Except Deep Learning, the course is self-containing meaning that

you will learn all other background!

We assume that we all have some basic math in mind

↳ linear algebra, calculus, and probability theory

But we review whatever we need from these topics whenever needed!

Textbooks

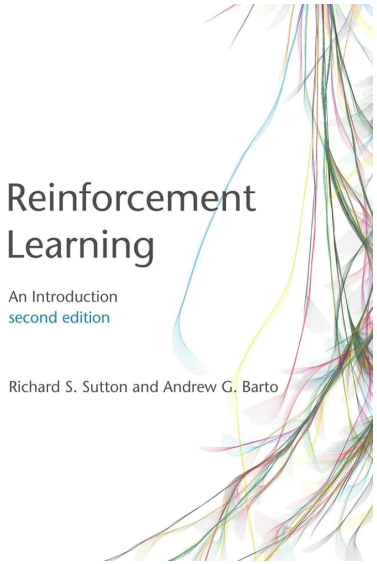

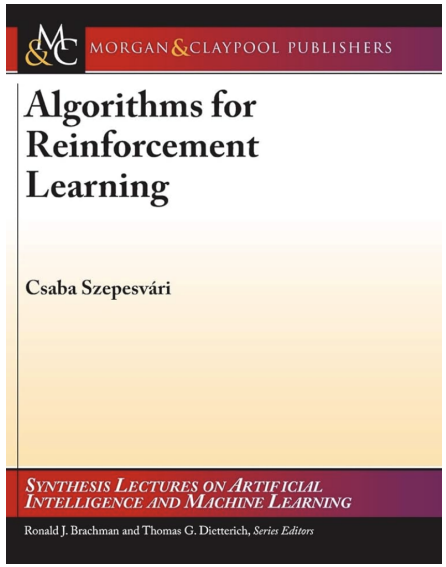

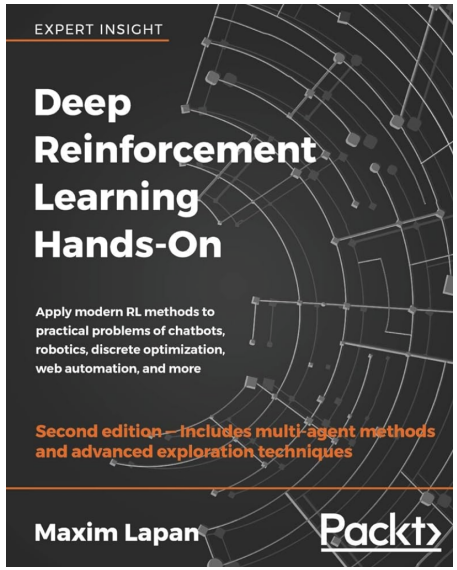

All materials are provided in the course. It's however good to know some texts!

- Sutton abd Brato can be accessed online at this link
- Szepesvári is available online here
- Lapan is a good source for Part 3

Terms and Conditions!

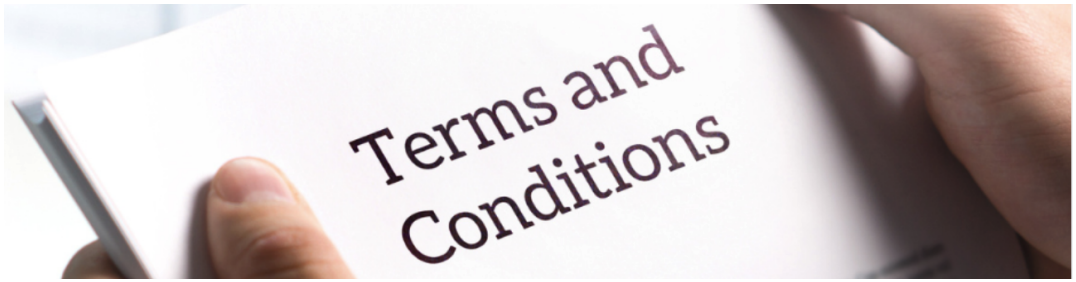

The instructor keeps the right reserved for himself to modify the slides

- *last minute before the lecture* ☺
- *after the lecture has been given*
  - *Typically happens due to typos*

The instructor keeps the right reserved for himself to deliver the lecture-notes

- *in form of mini-batches* ☺

Date and Signature .....

Introducing Glum

Glum does not buy my words! e.g.,

+ *Well! I know Deep Learning, you think you can teach me Reinforcement Learning!*
– Sure! Let's try!

So, please excuse me if I explain things sometimes in too much detail! I need to convince Glum!

No such thing as a stupid question!

Did you know that we got a Wikipedia page on this?

- ↳ Trust me! Your question will never sound stupid!
- ↳ If you don't ask; then, I need to ask!
  - ↳ Interaction is the best tool to avoid getting bored!

Any Questions? ☺