




"A beginning is the time for taking the most delicate care that the balances are correct."

Frank Herbert, *Dune*

CMPUT 365

Introduction to RL


Plan

- Introduction
- Course logistics
 - Instruction team
 - Pre-requisites
 - Textbook
 - Coursera
 - Academic integrity
 - Evaluation
- What is reinforcement learning?

Note



Lectures may be audio recorded for the purpose of a student's individual study as part of an approved academic accommodation.

Please, interrupt me at any time!




About myself

- Name: Marlos C. Machado
- I was born in Brazil
- I have been living in Edmonton for 10+ years
- Ph.D. working on reinforcement learning
 - Interned at Microsoft Research, IBM Research, and DeepMind
- Worked 4 years at Google Brain and DeepMind
 - Among several other things, we deployed RL to fly balloons in the stratosphere



Course overview and logistics



- Canvas: [link](#)
- My website: [link](#)

Start here!



**UNIVERSITY
OF ALBERTA**
CMPUT 365 A1 (53468)
Fall 2026
REINFORCEMENT LEARNING

Territorial Acknowledgement

The University of Alberta acknowledges that we are located on Treaty 6 territory, and respects the histories, languages, and cultures of First Nations, Métis, Inuit, and all First Peoples of Canada, whose presence continues to enrich our vibrant community.

Instructor Information

Instructor
Marlos C. Machado
Contact Information
machado@ualberta.ca
Office Hours: 4x/week after class
Lecture Room 6 Time
ET/CT: E1 401, MWF 13:00 - 13:50
Course Page: <https://courses.ualberta.ca/course/96154>

Teaching Assistant(s)

TA Name: Ravish Kulkarni
Email: rkulkarni@ualberta.ca
Office:
TBD - Will be posted on Canvas

Key resources

- Syllabus
 - Canvas and my website.
- Teaching assistants



Bavish



Elton



Erik



Rick




Pedro



Shashank



Sharvani



Vaibhav

- TA email address: cmput365@ualberta.ca
- My email address: machado@ualberta.ca
- Slack invitation link: [link](#)

I want to make this course a **safe** and **inclusive** environment, for everyone.

It is ok to make mistakes.

We should all strive to be **respectful** to each other.

If you want me to address you by a **different name**, or if you want to tell me your **pronouns**, I'm more than happy to hear!

Office hours

- Slack and Canvas: Asynchronous
- Marlos: *After class* @ *here*
- TAs: *To be announced this week*


Pre-requisites

- CMPUT 175 or CMPUT 275
- CMPUT 267 or 466, or STAT 265
- Python
- Probability (e.g., expectations of random variables, conditional expectations)
- Calculus (e.g., partial derivatives)
- Linear algebra (e.g., vectors and matrices)

You should either be familiar with these topics or be ready to pick them up quickly as needed by consulting outside resources.

Pre-requisites

- CMPUT 175 or CMPUT 275
- CMPUT
- Python
- Probability
- Calculus
- Linear algebra



to pick them up quickly as needed by consulting outside resources.

This will **not** be a flipped classroom!

- In the past, this course used to be taught in a flipped classroom
 - Roughly, you are initially introduced to **new topics outside** the classroom, using classroom time to explore topics in greater depth
- The number of students in this class has been steadily increasing, though
 - I don't know how to scale a flipped classroom without relying more and more on TAs to teach you
- Some of the feedback I received revolved around it feeling too repetitive
 - First read the textbook, watch the recorded lectures, do exercises, and then come to class
 - You can (and **should**) still do some of that before coming to class
- All this to say this will be a regular class, for better or for worse 😊

Required textbook


Reinforcement Learning: An Introduction

Richard S. Sutton & Andrew G. Barto

MIT Press. 2nd Edition.

<http://www.incompleteideas.net/book/the-book-2nd.html>

- You will need to read the book!
That's how you study for this course!
- The book is really good!



Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)
Midterm 1	Midterm	20%	10/2/26
Midterm 2	Midterm	20%	10/30/26
Final Exam	Exam (final)	30%	TBD by CNAS

Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)
Midterm 1	Midterm	20%	10/2/26
Midterm 2	Midterm	20%	10/30/26
Final Exam	Exam (final)	30%	TBD by CNAS

Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)

Coursera, almost every week (starting next week, Sep 11): 31.5%

Midterm 2	Midterm	20%	10/30/26
-----------	---------	-----	----------

Final Exam	Exam (final)	30%	TBD by CNAS
------------	--------------	-----	-------------

Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)

Coursera, almost every week (starting next week, Sep 11): 31.5%

Late submissions will not be accepted. There are 11 quizzes and 11 graded assignments. You're expected to do all of them, but s**t happens, so you can miss 2 of each and still get full marks.

You need them all for the certificate!

Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)
Midterm 1	Midterm	20%	10/2/26
Midterm 2	Midterm	20%	10/30/26

Two midterms, summing to 40%. Closed book.

Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)
Midterm 1	Midterm	20%	10/2/26
Midterm 2	Midterm	20%	10/30/26

Two midterms, summing to 40%. Closed book.
If you miss the midterm, you can can apply for an excused absence.
If granted, the weight of the missed midterm will be deferred to the final.

Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)
Midterm 1	Midterm	20%	10/2/26
Midterm 2	Midterm	20%	10/30/26
Final Exam	Exam (final)	30%	TBD by CNAS

Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)
Midterm 1	Midterm	20%	10/2/26
The final, worth 30%, will be about the whole course. If you miss the final, you can apply to a deferred final examination.			
Final Exam	Exam (final)	30%	TBD by CNAS

Title	Assessment Type	Weight	Date of Evalu...
Practice Quizzes	Other	9 x 1% = 9%	Day of the last class on the topic (see Course schedule)
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	Day of the last class on the topic (see Course schedule)
Midterm 1	Midterm	20%	10/2/26
Midterm 2	Midterm	20%	10/30/26
Final Exam	Exam (final)	30%	TBD by CNAS


Title	Assessment Type	Weight	D
Practice Quizzes	Other	9 x 1% = 9%	
Assessments (graded quizzes/notebooks on Coursera)	Other	9 x 2.5 = 22.5%	
Midterm 1	Midterm	20%	
Midterm 2	Midterm	20%	
Final Exam	Exam (final)	30%	100%

Total: **101.5%.**
You can **not-submit 2 quizzes** and **2 assessments.**

Grades **will not** be rounded at the end, and **no more extra marks** will be given.
No exceptions.

Coursera

- Coursera will be essential to CMPUT 365
- Today or tomorrow, you will be added to a private session of the RL courses (we will use your university's email)
 - If you don't have access you should let me know!
 - **IMPORTANT: If you don't use the private session you won't get credit for submitted work!**



The screenshot shows the Coursera website interface. At the top, the Coursera logo is on the left, followed by an 'Explore' button with a dropdown arrow. To the right is a search bar with the placeholder text 'What do you want to learn?' and a magnifying glass icon. Below the header, a breadcrumb trail reads 'Home > Browse > Data Science > Machine Learning'. A light blue banner below the breadcrumbs contains the text: 'This Labor Day, enjoy \$120 off Coursera Plus. Unlock access to 10,000+'. The main content area features the University of Alberta logo and name. Below this is the course title 'Reinforcement Learning Specialization' in a large, bold font. Under the title, a description reads: 'Master the Concepts of Reinforcement Learning. Implement a complete RL solution and understand how to apply AI tools to solve real-world problems.' Below the description are two small circular profile pictures of instructors, followed by the text 'Instructors: Adam White +1 more'. A prominent blue button with white text says 'Enroll for free' and 'Starts Sep 1'. To the right of the button, it says 'Try for free: Enroll to start your 7-day full access free trial'. Below the button, it states '60,167 already enrolled'. At the bottom, it says 'Included with coursera PLUS' with a blue 'PLUS' badge, followed by a link to 'Learn more'.

Coursera

Coursera

RL

Search in course

Search

Viewing:

CMPUT 365 - Winter 2026

Private

Upcoming

January 5, 2026 - April 27, 2026

🌐

Edit Course

Help

Fundamentals of Reinforcement Learning

< Course Material

Module 1

Module 2

Module 3

Module 4

Module 5

Grades

Notes

Discussion Forums

Messages

Live Events

Classmates

An Introduction to Sequential Decision-Making

46 min of videos left1h 10m of readings left2 graded assessments left

For the first week of this course, you will learn how to understand the exploration-exploitation trade-off in sequential decision-making, implement incremental algorithms for estimating action-values, and compare the strengths and weaknesses to...

Show Learning Objectives

The K-Armed Bandit Problem

Module 1 Learning Objectives

Reading • 10 min

Weekly Reading

Reading • 30 min

Let's play a game!

Ungraded Plugin • 15 min

Sequential Decision Making with Evaluative Feedback

Video • 5 min

Compare bandits to supervised learning

Discussion Prompt • 10 min

What to Learn? Estimating Action Values

Learning Action Values

Video • 4 min

What's underneath?

Ungraded Plugin • 15 min

Estimating Action Values Incrementally

Video • 5 min

Academic integrity

- [Code of Student Behaviour](#)
- [Student Conduct Policy](#)
- [Academic Integrity website](#)
- **Appropriate collaboration:** You are allowed to discuss the quizzes and assignments with your classmates. Note, however, that you are not allowed to exchange any written text, code, or to give and/or receive detailed step-by-step instructions on how to solve the proposed problems.
- **Cell phones:** Cell phones are to be turned off during lectures, labs and seminars.
- **Recording and/or Distribution of Course Materials:** Audio or video recording, digital or otherwise, by students is allowed only with my prior written consent as a part of an approved accommodation plan.

Academic integrity – **Expectations for AI use**

The primary goal of this course is to foster *individual* critical, creative thinking, and problem-solving skills related to reinforcement learning. Thus, in order to achieve such learning outcomes, you can submit each practice quiz and graded assignment multiple times, which allows for many learning opportunities.

The use of advanced AI-tools based on large-language models such as ChatGPT is **strictly prohibited** for all quizzes and graded assignments. The only exception is their use for Python-related queries (but the use of such tools to help with the programming assignments themselves is still strictly prohibited).

As stated in the university's [AI-Squared - Artificial Intelligence and Academic Integrity](#) webpage, “*learning is not only about the product; learning is also about the process of acquiring new knowledge or learning ways to think and reason.*”

Schedule

- The course will be structured in “weeks”. **Not every week starts on Monday**
- We have 12 weeks of content classes and we’ll cover 13 weeks of the MOOC
 - This corresponds to 9 chapters of the textbook


Schedule

- The course will be structured in “weeks”. **Not every week starts on Monday**
- We have 12 weeks of content classes and we’ll cover 13 weeks of the MOOC
 - This corresponds to 9 chapters of the textbook
- A practice quiz and a graded assignment are due at the end of each “week” in terms of content – You should look at the syllabus / schedule all the time
- The deadline for submitting assignments and quizzes is 23:59:59

Schedule

Course Schedule & Assigned Readings


Week	Date	Topic	Deadlines (all due at 23:59:59)	Readings
0	Wed, Sep 2	Course overview Discussion about what is reinf. learning		
0	Fri, Sep 4	Background review: Probability, statistics, linear algebra, and calculus		
Mon, Sep 7		Labour Day		
1	Wed, Sep 9	Fundamentals of RL: An introduction to sequential decision-making		Chapter 2, up to §2.7 (pp. 25-36), and §2.10 (pp. 42-44)
1	Fri, Sep 11	Fundamentals of RL: An introduction to sequential decision-making	Practice quiz and Progr. assignment (Bandits & exploration / exploitation)	
2	Mon, Sep 14	Fundamentals of RL: Markov decision processes (MDPs)		Chapter 3, up to §3.3 (pp. 47-56)
2	Wed, Sep 16	Fundamentals of RL: Markov decision processes (MDPs)	Practice quiz (MDPs)	
3	Fri, Sep 18	Fundamentals of RL: Value functions & Bellman equations		Chapter 3, §3.5-§3.8 (pp. 58-69)
3	Mon, Sep 21	Fundamentals of RL: Value functions & Bellman equations		
3	Wed, Sep 23	Fundamentals of RL: Value functions & Bellman equations	Practice and Graded quiz (Value functions & Bellman equations)	
4	Fri, Sep 25	Fundamentals of RL: Dynamic programming		Chapter 4, §4.1-§4.4 (pp. 73-84); §4.6-§4.7 (pp. 86-89)
4	Mon, Sep 28	Fundamentals of RL: Dynamic programming	Practice quiz and Progr. assignment (Optimal policies with dyn. progr.)	
Wed, Sep 30		National Day for Truth and Reconciliation		



What is reinforcement learning?

Artificial intelligence

“AI is the ability of machines to perform tasks that are typically associated with human intelligence, such as learning and problem-solving.” –Wikipedia



(Russell & Norvig; 2010)

Artificial intelligence

“AI is the ability of machines to perform tasks that are typically associated with human intelligence, such as learning and problem-solving.” –Wikipedia

The less a science has advanced, the more its terminology tends to rest on an uncritical assumption of mutual understanding.

– W. V. Quine




Machine learning

Machine learning is a subfield of AI in which the system's desired behavior is not explicitly programmed, instead it is *learned* from data

Machine learning

Machine learning is a subfield of AI in which the system's desired behavior is not explicitly programmed, instead it is *learned* from data

- “*Supervised learning* is learning from a training set of labeled examples provided by a knowledgeable external supervisor” (Sutton & Barto; 2018)




Cat



Cat



Not cat



Cat or not cat?

Machine learning

Machine learning is a subfield of AI in which the system's desired behavior is not explicitly programmed, instead it is *learned* from data

- “*Supervised learning* is learning from a training set of labeled examples provided by a knowledgeable external supervisor” (Sutton & Barto; 2018)
- “*Unsupervised learning* is typically about finding structure hidden in collections of unlabeled data” (Sutton & Barto; 2018)



Machine learning

Machine learning is a subfield of AI in which the system's desired behavior is not explicitly programmed, instead it is *learned* from data

- “*Supervised learning* is learning from a training set of labeled examples provided by a knowledgeable external supervisor” (Sutton & Barto; 2018)
- “*Unsupervised learning* is typically about finding structure hidden in collections of unlabeled data” (Sutton & Barto; 2018)



... and *reinforcement learning*!

Artificial intelligence

Machine learning

Reinforcement learning

Reinforcement learning

Reinforcement learning is a computational approach to learning from interaction to maximize a numerical reward signal (Sutton & Barto; 2018)



Reinforcement learning

Reinforcement learning is a computational approach to learning from interaction to maximize a numerical reward signal (Sutton & Barto; 2018)

- The idea of learning by interacting with our environment is very natural
- It is based on the idea of a learning system that wants something, and that adapts its behavior to get that



Artificial intelligence

Machine learning

Reinforcement learning

Reinforcement learning

Reinforcement learning is a computational approach to learning from interaction to maximize a numerical reward signal (Sutton & Barto; 2018)

- The idea of learning by interacting with our environment is very natural
- It is based on the idea of a learning system that wants something, and that adapts its behavior to get that



Some features are unique to reinforcement learning:

- Trial-and-error
- The trade-off between exploration and exploitation
- The delayed credit assignment / delayed reward problem

Artificial intelligence

Machine learning

Reinforcement learning

Reinforcement learning

Reinforcement learning is a computational paradigm for learning from interaction to maximize a numerical reward signal (Sutton & Barto, 2018)

- The idea of learning by interacting with an environment is very natural
- It is based on the idea of a *learning agent* that wants something, and that acts to get that

Some features are unique to reinforcement learning:

- Trial-and-error learning
- The trade-off between exploration and exploitation
- The delayed reinforcement / delayed reward problem


Problem or solution?



RL is now commonly deployed in the real-world

- **Recommendation systems**
 - Ads, news articles, videos, etc
- **General game playing**
 - Go, Chess, Shogi, Atari 2600, Starcraft, Minecraft, Gran Turismo
- **Industrial automation**
 - Cooling commercial buildings
 - Inventory management
 - Gas turbine optimization
 - Optimizing combustion in coal-fired power plants
- **Algorithms**
 - Video compression on YouTube
 - Faster matrix multiplication
 - Faster sorting algorithms
- **Control / Robotics**
 - Navigating stratospheric balloons
 - Plasm control for nuclear fusion
- **And more (see Csaba's [slides](#))**
 - COVID-19 border testing
 - Conversational agents
 - ...

The 2024 ACM A.M. Turing Award Winners “Created” RL




Association for
Computing Machinery

On intelligence, AGI, ASI, etc etc...

- People in the field have different, non-competing, perspectives and motivations
 - Some study RL to learn about / develop tools for solving sequential decision-making problems
 - Some look at RL as a computational model of intelligence

On intelligence, AGI, etc etc...


- People in the field have different, non-competing, perspectives and motivations
 - Some study RL to learn about / develop tools for solving sequential decision-making problems
 - Some look at RL as a computational model of intelligence
- RL was originally developed to understand intelligence/the brain
 - We should develop a critical view around these topics, and an ability to recognize hype / PR pieces




A. Barto (2024)

On intelligence, AGI, etc etc...


- People in the field have different, non-competing, perspectives and motivations
 - Some study RL to learn about / develop tools for solving sequential decision-making problems
 - Some look at RL as a computational model of intelligence
- RL was originally developed to understand intelligence/the brain
 - We should develop a critical view around these topics, and an ability to recognize hype / PR pieces
- Both perspectives are valid and both had had successes in the past **But they are different!!**



A. Barto (2024)



(Schultz, Dayan, & Montague; 1997)




(Stachenfeld, Botvinick, & Gershman; 2017)


(Silver et al.; 2016)



(Degraeve et al.; 2022)



(Bellemare et al.; 2020)



Next class

- What **I** plan to do: A reminder about the required theoretical background
 - Probability (e.g., expectations of random variables, conditional expectations)
 - Calculus (e.g., partial derivatives)
 - Linear algebra (e.g., vectors and matrices)
 - I won't remind / teach you Python.
- What I recommend **YOU** to do for next class:
 - Make sure you have access to Coursera, eClass, and Slack
 - Brush up whatever you feel you are rusty on in terms of background
 - Read Chapter 1 of the textbook (not mandatory)