Applied Machine Learning

Syllabus and logistics

Reihaneh Rabbany

Remote Class

- **Live Lectures:** Monday & Wednesdays, 10:05 am -11:25 (Montreal time)
  - Online Zoom meetings through Mycourses
  - Lectures will be **recorded** and automatically uploaded in Mycourses
- **Course Website:** http://www.reirab.com/comp551.html 

  - Syllabus, slides, deadlines, schedule, evaluation, etc.

Communications

- **Online Zoom Office Hours** **under zoom tab in Mycourses**
  - **Instructor:** Thursdays 10 am - 11 am
  - **TAs:** please check Mycourses's calendar
- **Course Email:** comp551mcgill@gmail.com 

- **Instructor Email**: rrabba@cs.mcgill.ca [for private communication "551 special" in title]
- **Course Slack:** will send invites soon

Prerequisites

- Strong linear algebra, probabilities, and Python programming is highly recommanded
- How can I refresh my background knowledge to follow the lectures better? a lot of excellent online materials, see which one you can follow easier, you can also refer to these reviews on probability and linear algebra.
- Tutorials next week on Jan 15th (Math) and Jan 19th (Python) 

- Two quizzes on main concepts needed for lectures [with unlimitted attmpts allowed], due Jan 23rd, released this Wed & next Mon 

Tutorials

Jan 15.	Safa	Probability & Linear Algebra	
Jan 19.	Manoj	Python	https://www.python.org/
Mid Feb.	Nishant	Scikit-learn	https://scikit-learn.org/
Early Mar.	Sumana	Pytorch	https://pytorch.org/

Pre-recorded, attend the corresponding TA's office hour for questions

Teaching Assistants

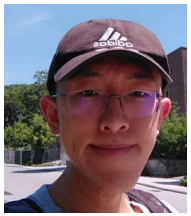

Tianyu Shi [tianyu.shi3]

Quiz TA

I am second year master student with Prof. Lijun Sun and a research intern at Mila, supervised by Prof. Laurent Charlin. I am working on reinforcement learning for intelligent transportation system.

Sumana Basu [sumana.basu]

3rd year PhD student at Mila, Supervisors: Doina Precup, Adriana Romero, Research Interest: Reinforcement Learning for Healthcare

Manoj Venkatesan [manoj.venkatesan]

2nd Year Master's student supervised by Dr.Hannah Michalska and working on Control Systems.

Safa Alver [safa.alver]

I am a 2nd year PhD student supervised by Doina Precup, working on Transfer in Reinforcement Learning

Haque Ishfaq [haque.ishfaq]

3rd year PhD student, supervised by Prof. Doina Precup, working on reinforcement learning theory.

Tianzi Yang [tianzi.yang]

Head TA

2nd year Ms student, working on machine learning and autonomous driving

Yan Miao [yan.miao]

1st year MSc student, working on deep learning and medical imaging under the supervision of Prof. Peter Savadjiev

Nishant Mishra [nishant.mishra]

Group TA

2nd year MS student, working on Deep Learning for Computer Vision

Teaching Assistants

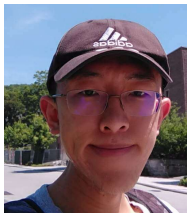

Tianyu Shi [tianyu.shi3]

Quiz TA

Office hours TBA, check Mycourses and course website

Helps you wrt questions for weekly quizzes

Team 0

Sumana Basu [sumana.basu]

Team 1

Manoj Venkatesan [manoj.venkatesan]

Team 2

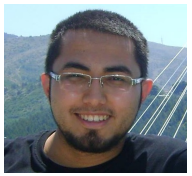

Safa Alver [safa.alver]

Team 3

Haque Ishfaq [haque.ishfaq]

Team 4

Tianzi Yang [tianzi.yang]

Head TA

Team 5

Yan Miao [yan.miao]

Team 6

Nishant Mishra [nishant.mishra]

Group TA

Helps you for forming groups and resolving conflicts

Team 7

About this course

Tentative Outline

- Introduction
- Nearest Neighbours
- Basic concepts
- Maximum likelihood and Bayesian Reasoning
- Unsupervised learning
- Naive Bayes
- Expectation Maximization
- Dimensionality reduction
- Linear regression
- Logistic and softmax regression
- Gradient descent methods
- Regularization
- Classification and regression trees
- Bias-variance decomposition, bagging & random forests
- Perceptrons and linear support vector machines
- Multilayer Perceptrons
- Gradient computation and automatic differentiation
- Convolutional neural networks

About this course

complementary components

Theory

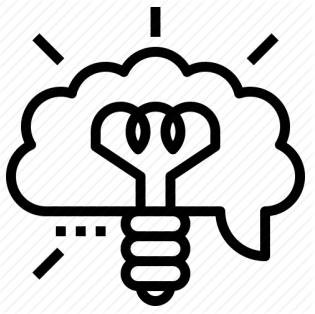

Lectures
Weekly Practice Quizzes
Checkpoint Quizzes
Understand the theory behind learning algorithms

Application

Codes in lectures
Mini-projects
Practice applying them in real-world

About this course

Tentative Guest Lectures

- April 14th: **Ethics in ML**
  - A guest lecture by Abhishek Gupta, founder of Montreal AI Ethics Institute
- April 15th: **RL in practice**
  - A guest lecture by Marc Bellemare, leads the reinforcement learning efforts of the Google Research team in Montréal

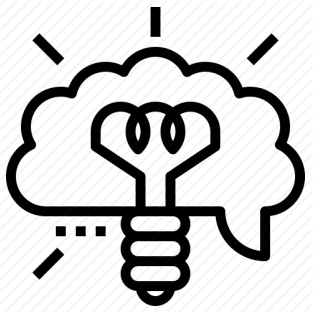

About this course:

Evaluation and grading

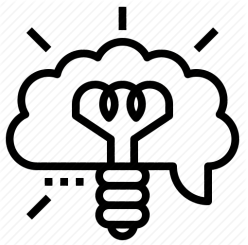

Regular Practice Quizzes - **20%** {from last lecture - short}

Knowledge Checkpoint Quizzes - **30%** {from all lectures to that point - longer}

Mini-projects - **50%** {group assignments}

About this course:

Evaluation and grading

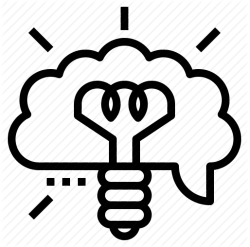

Regular Practice Quizzes - **20%**

{from last lecture - short}

- One per lecture to check the key concepts discussed in the last lecture
- Timed to be done in 1 hour after starting the quiz
- Available until the start of the next lecture
- Four lectures don't have practice quizzes and instead we have the checkpoint quizzes
- From the two subsequent quizzes, the best one will be considered for your final grade
- Starts from January 13th
- The first 2 practice quizzes check the prerequisites for the course and have a different setting than the regular practice quizzes. In particular, unlimited attempts are allowed and due date is set to after add/drop instead of the next lecture

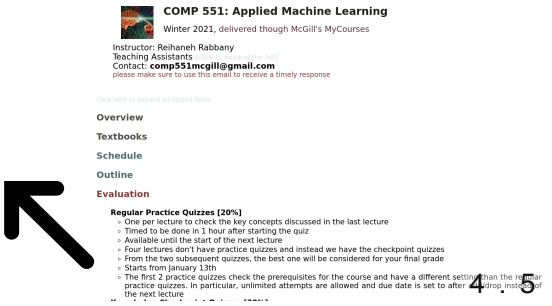

About this course:

Evaluation and grading

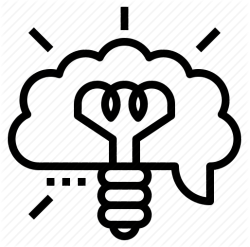

Knowledge Checkpoint Quizzes - **30%**

{from all lectures to that point - longer}

- One per month released on a Thursday
- From all topics discussed up to the date of the quiz
- Timed to be done in 3 hours, within a 3 days availability period
- Mark the due dates: Jan 21-23rd [3%], Feb 12-14th [7%], March 12-14th [10%], April 15-17th [10%]

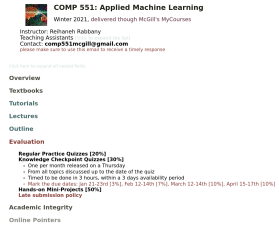

bookmark

About this course:

Evaluation and grading

Mini-projects - **50%** {group assignments}

- Four programming assignments to be done in groups of three*, *no exception to this given the grading load on TAs
- Groups can stay the same between projects, you can also regroup when needed
- All group members receive the same mark unless there are major complains on not contributing, responding, etc. from group-mates, which will be resolved in a case by case basis. If a significant difficulty/conflict arises, please send an email to the course email, cc the group-TA and put 'Group-TA' in the title
- Mark the due dates: Feb 5th [10%], Feb 28th [15%], March 25th [15%], April 25th [10%]

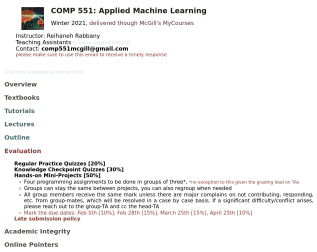

bookmark

Late submissions

All due dates are **11:59 pm** in Montreal unless stated otherwise. **No make-up quizzes** will be given. For mini-projects, 2^k% percent will be deducted per k days of delay.

If you experience barriers (including a covid related issue) to learning in this course, submitting the projects, etc., please do not hesitate to discuss them with me directly, and please make sure to put "551 special" in the header to make sure I see your email [for general course correspondence, please use the course email: comp551mcgill@gmail.com].

As a point of reference, you can reach the Office for Students with Disabilities at 514-398-6009

Code of Conduct

- Do not share or (re)post any of the course materials online. This includes: video lectures, codes, quizzes, zoom links, etc.
- Be respectful in the course forums and other communications
- Submit your own work for projects and quizzes

Academic Integrity

The ``McGill University values academic integrity. Therefore, all students must understand the meaning and consequences of cheating, plagiarism and other academic offenses under the Code of Student Conduct and Disciplinary Procedures'' (see McGill's webpage for more information). (Approved by Senate on 29 January 2003)

Relevant Textbooks

No required textbook but slides will cover chapters from the following books, all available online, which can be used as reference materials.

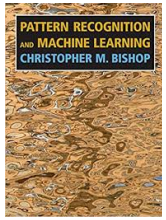

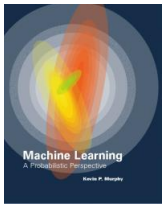

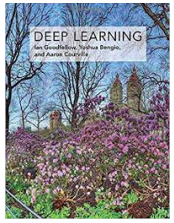

[Bishop] Pattern Recognition and Machine Learning by Christopher Bishop (2007), available online

[Murphy] Machine Learning: A Probabilistic Perspective by Kevin Murphy (2012), available online through the library

[GBC] Deep Learning (2016) by Ian Goodfellow, Yoshua Bengio, and Aaron Courville, available online

Resources

Numerous great online resources at different levels, a selection is listed on the course website

Some may be more accessible than this course since they are designed for a different audience, but please note that this is a course designed for graduate students in computer science without ML background, with a heavy theory component. Consider following alternative courses at McGill:

- COMP 451: fundamentals of ML
  - Good ML entry course for CS undergraduates
- ECSE 551: ML for engineers

*Both antirequisite to 551

COMP 551: Applied Machine Learning

Winter 2021, delivered though McGill's MyCourses

Instructor: Reihaneh Rabbany
Teaching Assistants [click to expand the list]
Contact: **comp551mcgill@gmail.com**
please make sure to use this email to receive a timely response

Click here to expand all nested fields

Overview

Textbooks

Tutorials

Lectures

Outline

Evaluation

Academic Integrity

Online Pointers

Learning plan
Video Playlists
- StatQuest
- FreeCodeCamp
- Essence of linear algebra and Neural Networks by 3Blue1Brown

Courses with Playlist and/or Code
- Introduction to Machine Learning by Google
- Deep Learning by UC Berkeley
- Hinton's Lectures on Neural Networks for Machine Learning
- Deep Learning & Linear Algebra courses by fastai
- Learning from Data by Caltech
- Deep Learning (with PyTorch) playlist and course by NYU
- Deep Learning by Stanford
- Deep Learning by deeplearning.ai
- Introduction to Deep Learning by MIT
- Information Theory, Pattern Recognition, and Neural Networks by David MacKay

Books with Code
Similar Courses - Graduate Level

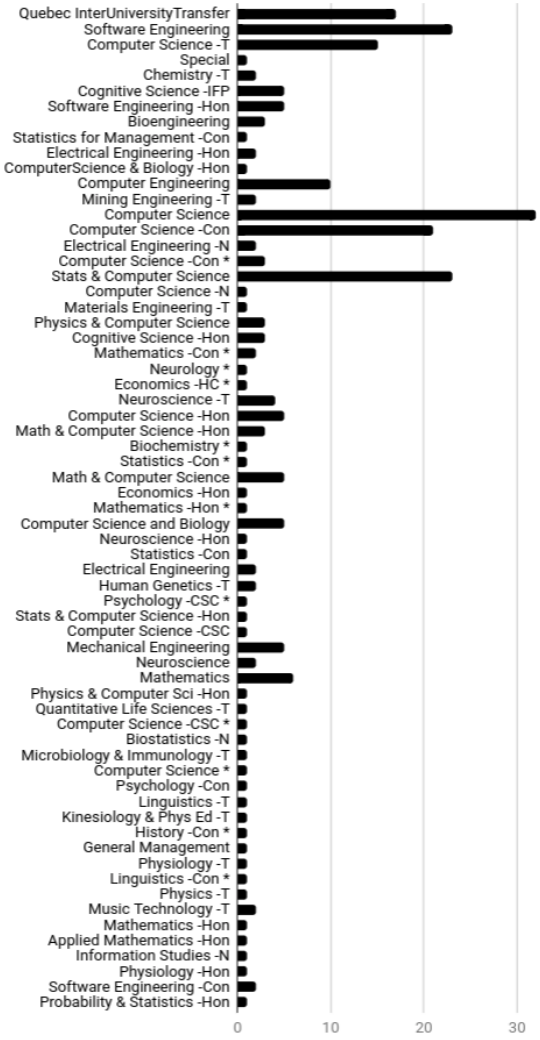

Who is in this class? You

250 registered

Mostly undergraduates year 3

Mostly with Computer background

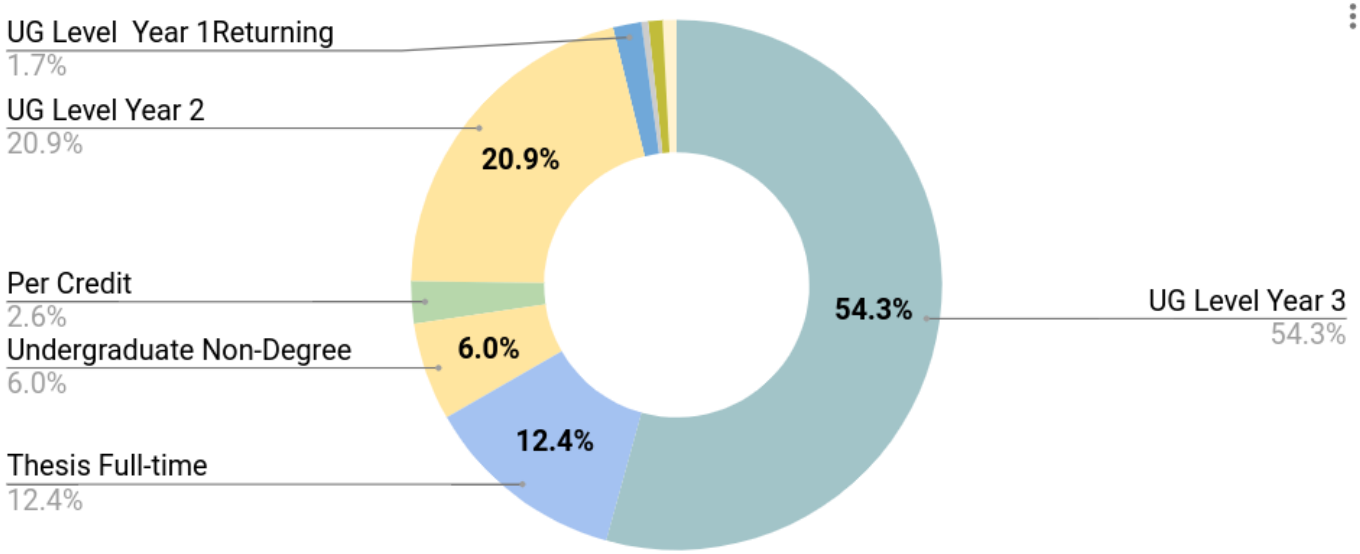

Who is in this class? Me

Reihaneh Rabbany

Canada CIFAR AI Chair and core member at Mila

Assistant Professor in the School of Computer Science

http://www.reirab.com/

Had CMPUT 551 Winter 2009 with Enrl of 9!

We might have some background noise

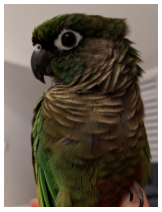

Who is in this class? Me

My research is on Network science, data mining and machine learning, with a focus on analyzing real-world interconnected data, and social good applications.

- Physics (complex systems)
- Sociology (social networks)
- Mathematics (graph theory)
- Data Mining (graph mining)
- Machine Learning (relational learning, graph neural networks)

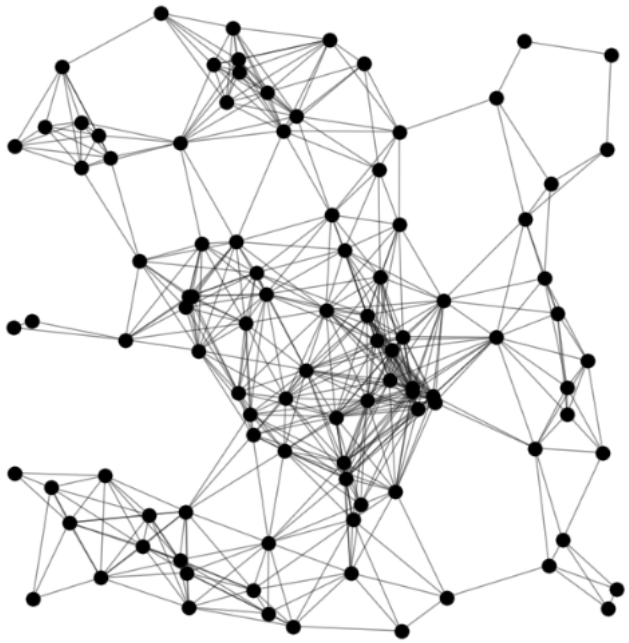

Who is in this class? Me

AI4Good applications where mining connections is the key

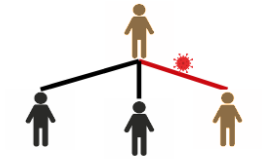

- Interconnected populations
  - Modelling covid-19 by incorporating contact graphs and flight networks
- Interconnected crime
  - Detecting organized human trafficking in online escort markets
- Interconnected discussions
  - Detecting coordinated groups in online political discourse

Questions?