Applied Machine Learning

Syllabus and logistics

Oumar Kaba
Reihaneh Rabbany

COMP 551 (Fall 2025)

In-person Class

- **Lectures:** Mon & Wed, 2:35 pm - 3:55 (Montreal time)
  - Leacock Building (LEA), room 26
  - Lectures will be **recorded** and uploaded in Mycourses
- **Course**

  Website: https://oumarkaba.github.io/comp551/comp551.html

  

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

Communications

- **Office Hours**
  - **Instructor & TAs:** please check Mycourses (TBD)
- **Course Email:** comp551-f25@mcgill.com
- **Instructor Email**: [for private communication "COMP551" in title]
  - Reihaneh.rabbany@mcgill.ca
  - sekou-oumar.kaba@mcgill.ca
- **Course Discussion:** Ed through Mycourses

bookmark

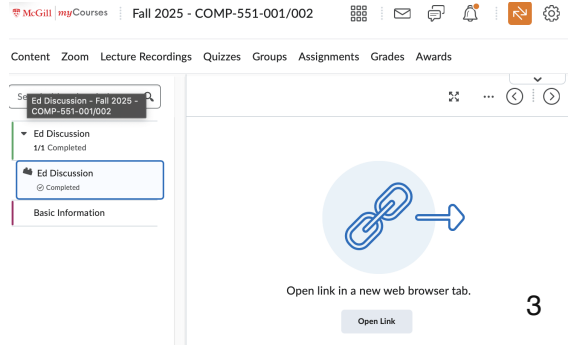

Prerequisites

- Strong linear algebra, probabilities, and Python programming is highly recommended: MATH 222, Math 113 and Comp 202
- 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.
- Two quizzes on main concepts needed for lectures [with unlimited attempts allowed], due September 9th {ADD/DROP deadline}, released

Tutorials

Early Sep.	Probability & Linear Algebra	
Early Sep.	Python	https://www.python.org/
Late Sep.	Scikit-learn	https://scikit-learn.org/
Late Oct.	Pytorch GPU Cluster Use	https://pytorch.org/

About this course

Tentative Outline

- Introduction
- Core concepts
- Linear regression
- Logistic and softmax regression
- Gradient descent methods
- Regularization
- Perceptrons & Multilayer Perceptrons
- Gradient computation and automatic differentiation
- Convolutional neural networks
- Nearest Neighbours
- Classification and regression trees
- Maximum likelihood and Bayesian Reasoning
- Naive Bayes
- Linear support vector machines
- Bagging & Boosting
- Unsupervised learning
- Dimensionality reduction
- Learning with graphs

About this course

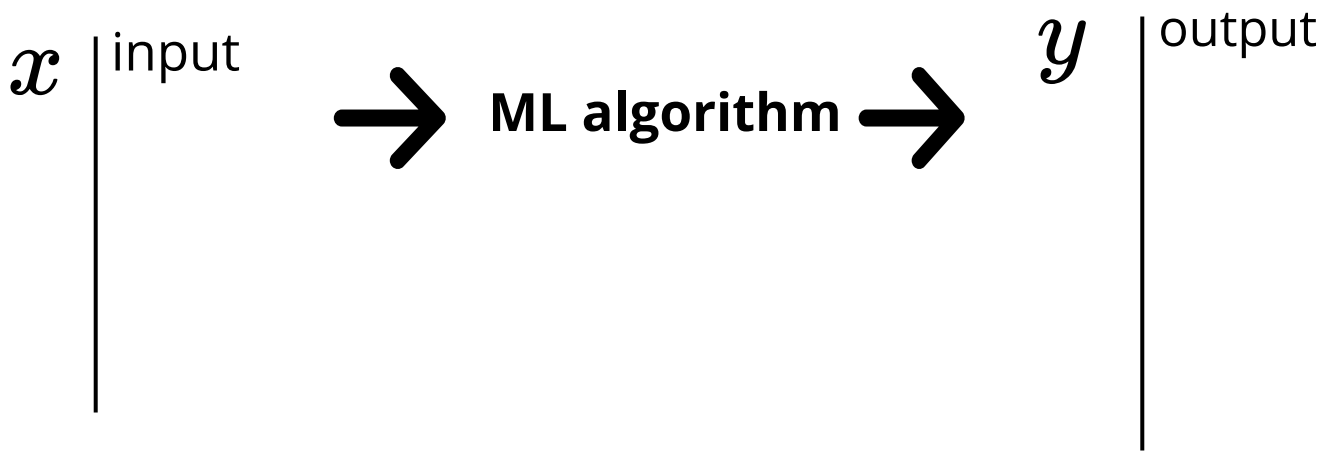

example

<tumorsize, texture, perimeter> = <18.2, 27.6, 117.5> → cancer = No

About this course

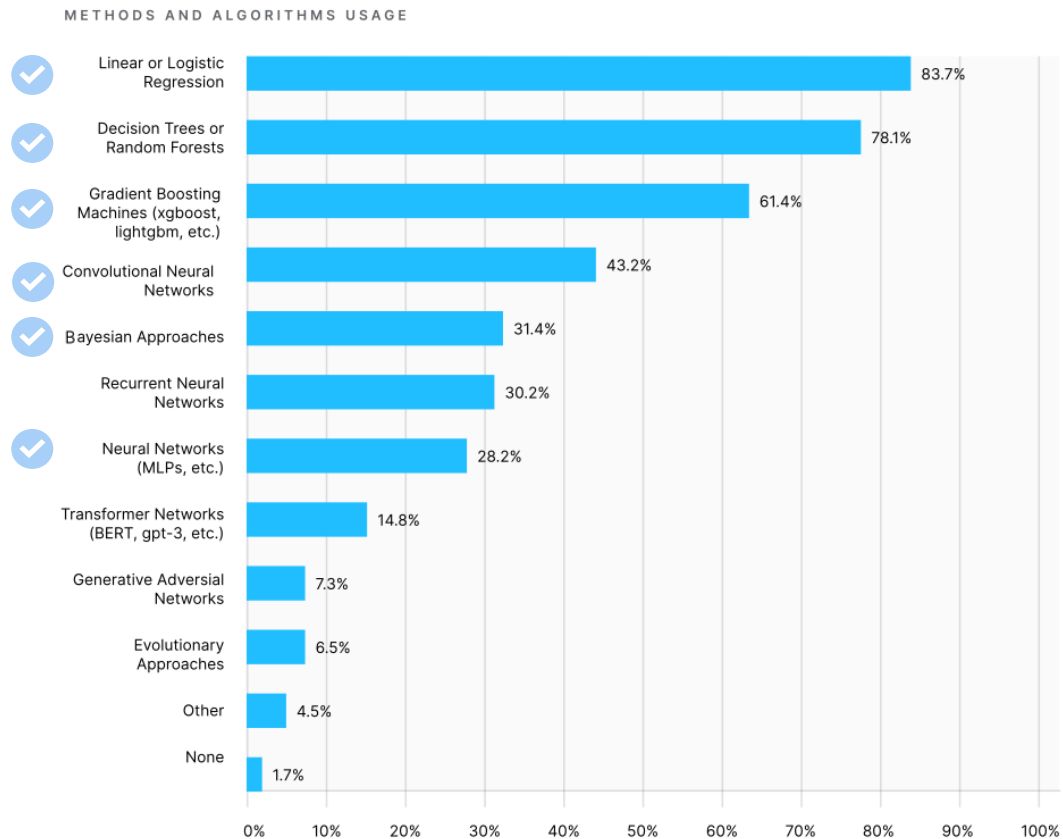

from 2020 Kaggle's survey on the state of Machine Learning and Data Science,
you can read the full version here

About this course

complementary components

Theory

Lectures
Weekly Practice Quizzes
Midterm Exam
Understand the theory behind learning algorithms

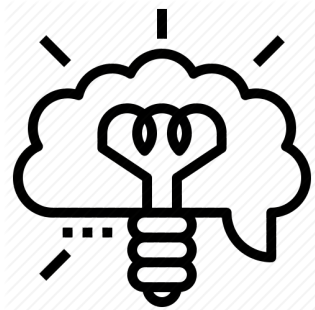

Application

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

About this course:

Evaluation and grading

Regular Practice Quizzes - **10%** {1% each}

Midterm exams - **40%** {15% - 25%}

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

Lecture summaries - **1% bonus** {caped at 5}

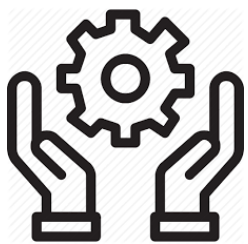

About this course:

Evaluation and grading

Regular Practice Quizzes - **10%**

- One per week to check the key concepts discussed each week
- Available until the start of the next Mon lecture
- Unlimited attempts are allowed

About this course:

Evaluation and grading

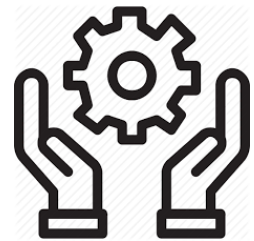

Mini-projects - **50%** {group assignments}

- 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
- The goal is not to divide and conquer but to collaborate, do not wait for others to complete their tasks, help eachother do all the parts in the assignment
- All group members receive the same mark unless there are major complaints on not contributing, responding, etc. from group-mates, which will be resolved on 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

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.

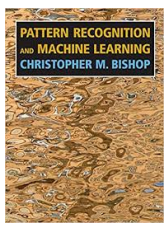

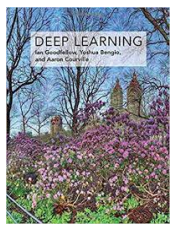

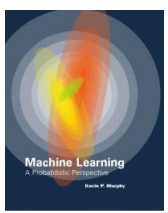

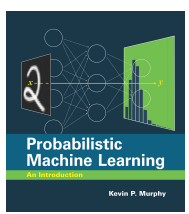

- [Bishop] Pattern Recognition and Machine Learning *by Christopher Bishop (2007)*
- [Goodfellow] Deep Learning *by Ian Goodfellow, Yoshua Bengio, and Aaron Courville (2016)*
- [Murphy] Machine Learning: A Probabilistic Perspective *by Kevin Murphy (2012)*
- Murphy'22] Probabilistic Machine Learning: An Introduction, by Kevin P. Murphy (2022)

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.

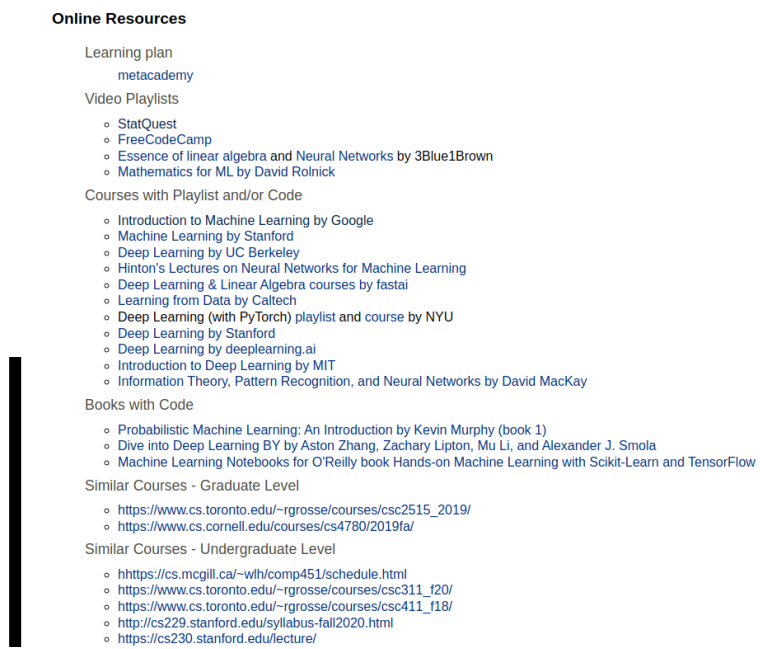

Online Resources

Learning plan

- metacademy

Video Playlists

- StatQuest
- FreeCodeCamp
- Essence of linear algebra and Neural Networks by 3Blue1Brown
- Mathematics for ML by David Rolnick

Courses with Playlist and/or Code

- Introduction to Machine Learning by Google
- Machine Learning by Stanford
- 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

- Probabilistic Machine Learning: An Introduction by Kevin Murphy (book 1)
- Dive into Deep Learning BY by Aston Zhang, Zachary Lipton, Mu Li, and Alexander J. Smola
- Machine Learning Notebooks for O'Reilly book Hands-on Machine Learning with Scikit-Learn and TensorFlow

Similar Courses - Graduate Level

- https://www.cs.toronto.edu/~rgrosse/courses/csc2515_2019/
- https://www.cs.cornell.edu/courses/cs4780/2019fa/

Similar Courses - Undergraduate Level

- hhttps://cs.mcgill.ca/~wlh/comp451/schedule.html
- https://www.cs.toronto.edu/~rgrosse/courses/csc311_f20/
- https://www.cs.toronto.edu/~rgrosse/courses/csc411_f18/
- http://cs229.stanford.edu/syllabus-fall2020.html
- https://cs230.stanford.edu/lecture/
- Cheatsheets: https://stanford.edu/~shervine/teaching/

Who is in this class? You

200 registered

Mostly undergraduates year 3

Mostly with Computer background

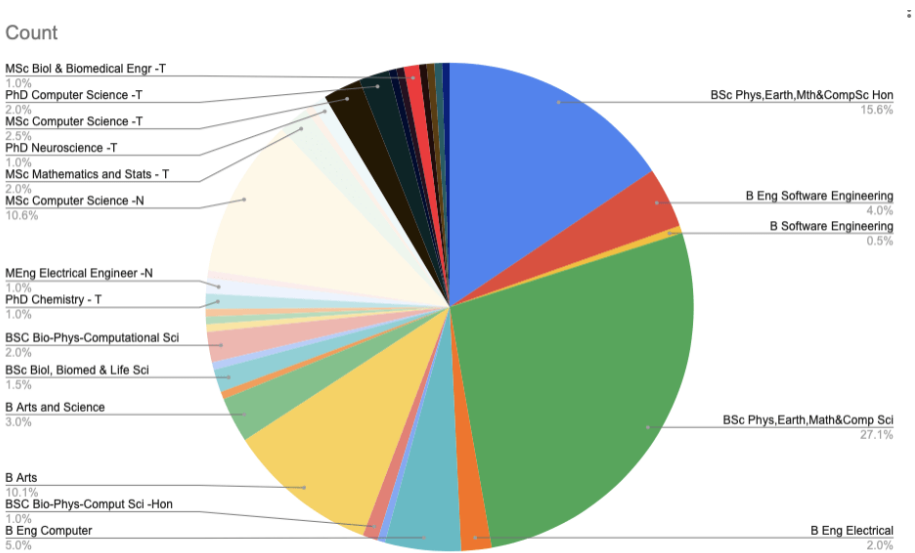

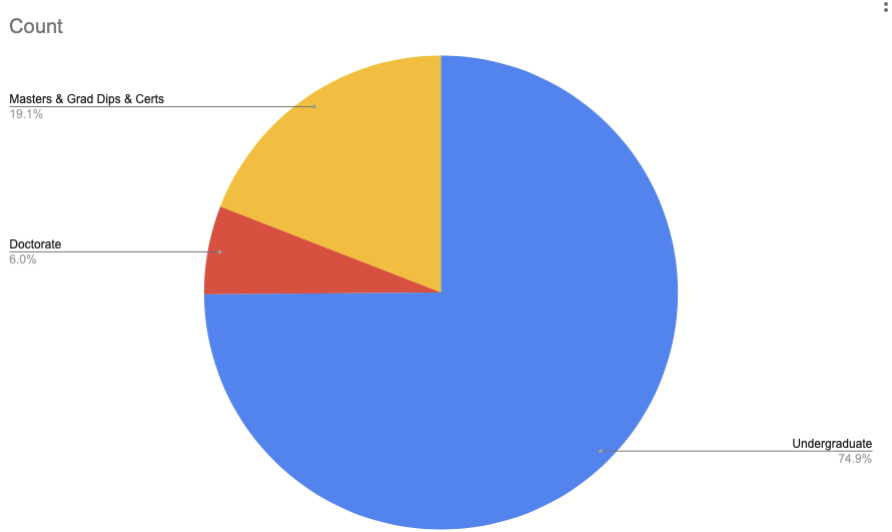

Who is in this class? Teaching Team

Instructors: **Oumar Kaba** & **Reihaneh Rabbany**

Teaching Assistants

Charlotte Volk	My name is Charlotte, I've just finished up the first year of my masters and am planning to fast track to PhD in the winter. I'm studying biologically plausible learning algorithms with Blake Richards. This is my 3rd time TAing for COMP 551, I'm excited to get started!
Rafid Saif	I'm Rafid, I am also finishing the first year of my masters in Electrical Engineering. I'm currently working on skill-based learning for robotics and am advised by Hsiu-Chin Lin
Mohamad Danesh	I'm a 3rd year PhD student at CS working on robot learning under Hsiu-Chin LIn.
Sebastian Sabry	My name is Sebastian, I've finished my first year of my masters switching to a Thesis with Dr. Rabbany! I'm excited to TA for COMP 551.
Zahra Tehraninasab	My name is Zahra and I'm a master's student working on Computer Vision for medical imqge analysis under professor Tal Arbel. This is my second time TAing for COMP 551. Looking forward to meeting you all!

Who is in this class? RR

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!

Who is in this class? RR

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)

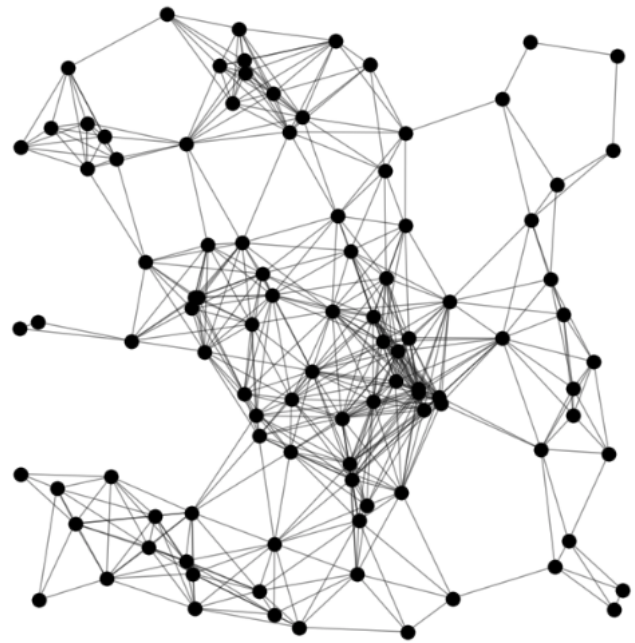

Who is in this class? RR

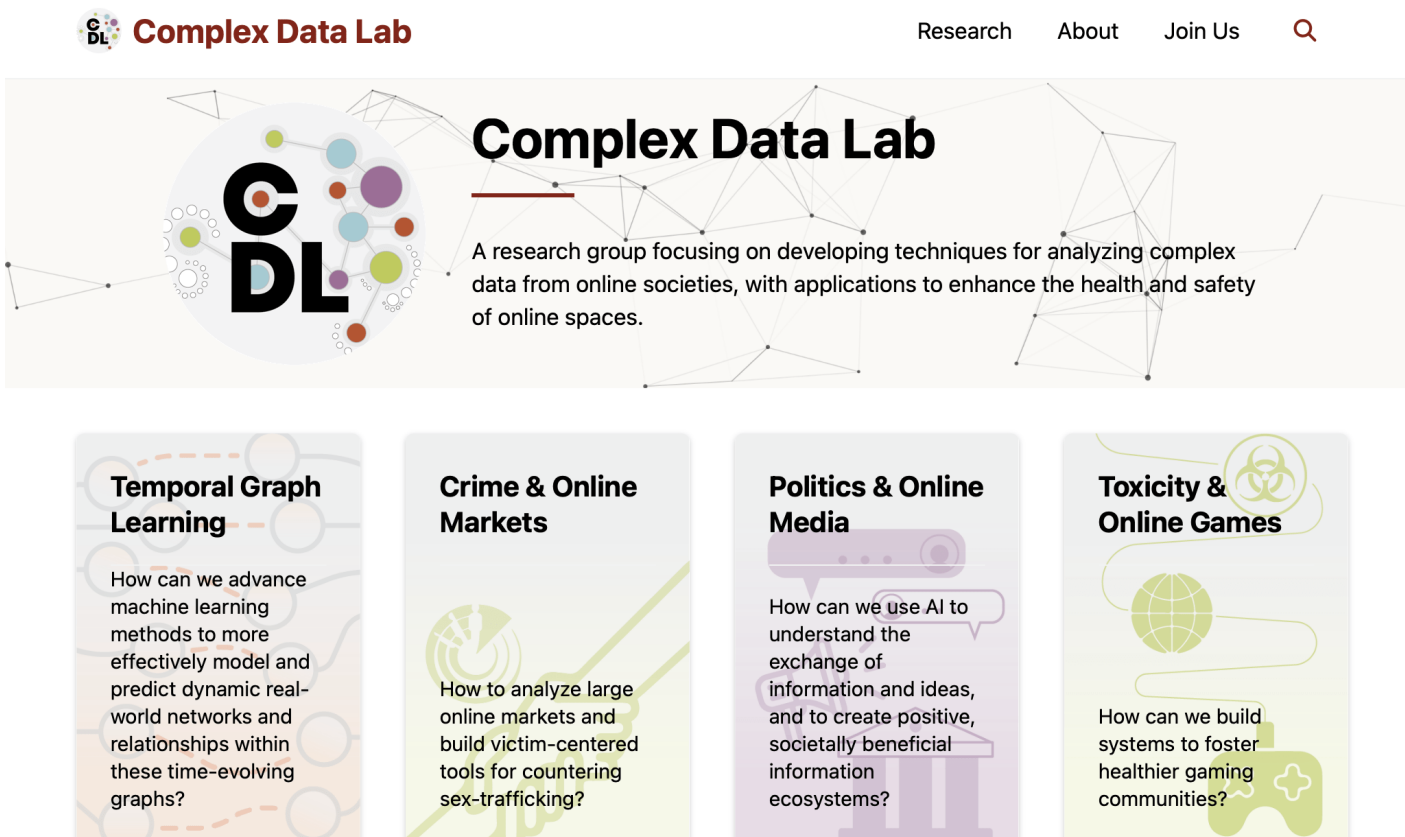

https://www.complexdatalab.com/

Who is in this class? OK

Sékou-Oumar Kaba (call me Oumar)

Last year PhD candidate

Mila & McGill University

https://oumarkaba.github.io/

Previously

Masters: Theoretical physics, with applications to materials

Industry: Data science, Montreal-based companies

Research internship: Huawei Research, ML for networks

Research internship: Microsoft Research, ML for chemistry

Who is in this class? OK

My research: **Physics for machine learning**

Bring the idea of symmetry from physics to machine learning systems. Essentially make them understand that there are transformations that preserve the identity of objects.

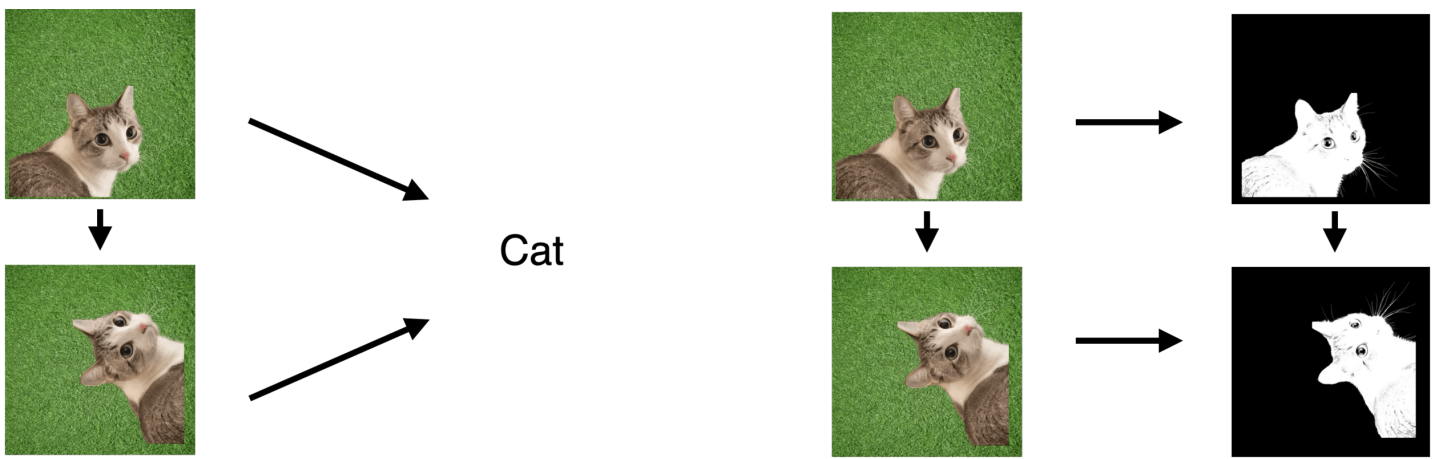

Can we prove if a machine learning system will generalize to transformations or not?

Who is in this class? OK

My research: **Machine learning for physics and chemistry**

I am especially interesting in investigating how we could model materials at the atomic level and discover new ones using these methods.

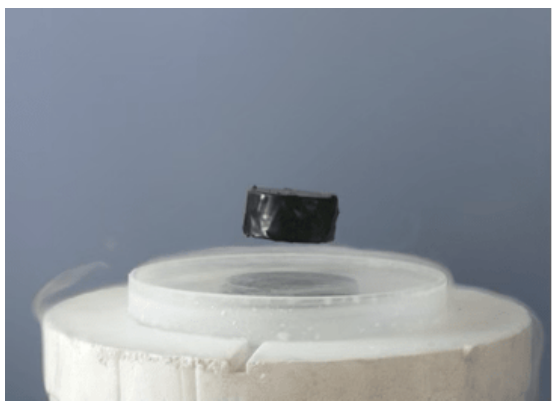

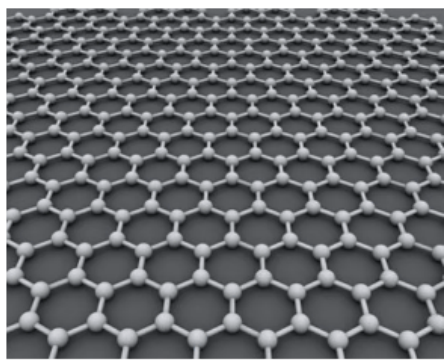

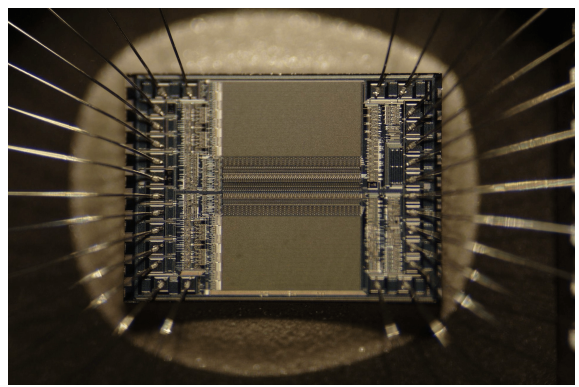

Who is in this class? OK

My research: **Machine learning for physics and chemistry**

I am especially interesting in investigating how we could model materials at the atomic level and discover new ones using these methods.

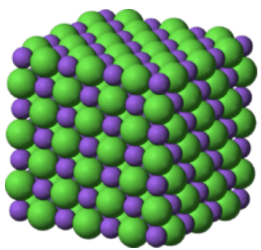

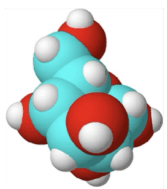

Neural net

Property of interest

Property of interest

Neural net

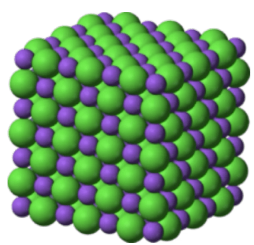

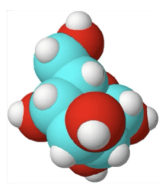

Questions?