

CS480/680: Introduction to Machine Learning

Spring 2024

Time and location: MW 10:00-11:20AM, 11:30-12:50, MC 2034

Web: <https://cs.uwaterloo.ca/~y328yu/mycourses/480>

Learn: <https://learn.uwaterloo.ca/d21/home/1021993>

Piazza: <https://piazza.com/uwaterloo.ca/spring2024/cs480cs680>

Syllabus

Instructor

Yao-Liang Yu
DC 3617
School of Computer Science
University of Waterloo

Web: <https://cs.uwaterloo.ca/~y328yu>
Email: yaoliang.yu@uwaterloo.ca
Office hour: W 1:00-2:00pm or by email appointment

TAs:

- Lu, Haoye (h229lu)
- Lu, Yiwei (y485lu)
- Malekmohammadi, Saber (s3malekm) x 2
- Mouzakis, Argyris (amouzaki)
- Szabados, Spencer (sszabado)

Course Description

This course introduces students to the design of algorithms that enable machines to “learn,” as well as some of their significant impacts on society. In contrast to the classic paradigm where machines are programmed by specifying a set of instructions that dictate what exactly a machine should do, a new paradigm is developed whereby machines are presented with examples from which they learn what to do. This is especially useful in complex tasks such as natural language processing, information retrieval, data mining, computer vision and robotics where it is not practical for a programmer to enumerate all possible situations in order to specify suitable instructions for all situations. Instead, a machine is fed with a large collection of examples from which it automatically learns suitable rules to follow.

The course will introduce the basics of machine learning and data analysis, and demonstrate how machine learning techniques have impacted life and society. The course starts with some classical tasks and algorithms in machine learning, then transitions to modern advances and successful applications in real life, and finally ends with discussions on the societal impact that machine learning has generated. As machine learning becomes increasingly popular in practice and starts to penetrate into critical areas such as law, health care, scarce resource allocation, autonomous driving, elections, etc., various ethical and safety issues have been discovered. A significant part of the course will be devoted to raising students’ awareness on such ethical issues, including how to make machine learning

algorithms fair and unbiased, how to protect individual privacy without compromising learning, how machine learning algorithms might be attacked, how to measure and improve the security of machine learning algorithms, and how to make machine learning algorithms explainable and interpretable.

Course Objectives

At the end of the course, students should have the ability to:

- Recognize and formalize a task as a machine learning problem;
- Identify suitable algorithms to tackle different machine learning problems;
- Understand and implement a plethora of foundational machine learning algorithms;
- Apply machine learning algorithms to real datasets;
- Become aware of potential ethical and safety issues of machine learning on society.

Course Overview

- Introduction (first lecture: May 6, 2024)
- Perceptron
- Linear Regression
- Logistic Regression
- Support Vector Machines
- Reproducing Kernels
- Multi-layer Perceptron
- Deep Neural Networks
- Convolutional Neural Networks
- Graph Neural Networks
- Attention
- Decision Trees
- Bagging, Boosting and Random Forest
- Generative Adversarial Networks
- Flow Models
- Variational Autoencoders
- Optimal Transport
- Diffusion Models
- Contrastive Learning
- Hidden Markov Models
- Calibration
- Algorithmic Fairness
- Adversarial Machine Learning
- Differential Privacy

Prerequisites

Good knowledge of linear algebra (vector space, eigenvalue, matrix multiplication, etc.) and basic probability (random variable, distribution, expectation, conditional probability, Bayes rule, etc.). Exposure to numerical computing and optimization is a plus but not required. Basic programming language such as Python, or Julia (can learn in a few hours).

Textbooks

We will pose lecture notes or slides before class. The following fine books are recommended:

- *Mathematical Analysis of Machine Learning Algorithms*. Tong Zhang. Cambridge University Press, 2023.
- *Patterns, Predictions, and Actions*. Moritz Hardt and Benjamin Recht. Princeton University Press, 2022.
- *Probabilistic Machine Learning: Advanced Topics*. Kevin Patrick Murphy. MIT Press, 2023.
- *Probabilistic Machine Learning: An Introduction*. Kevin Patrick Murphy. MIT Press, 2022.
- *Dive into Deep Learning*. Aston Zhang, Zack C. Lipton, Mu Li and Alex J. Smola. 2019.
- *Deep Learning*. Ian Goodfellow, Yoshua Bengio and Aaron Courville. MIT Press, 2016.

Grading (**Tentative**)

There will be 4 homework assignments, each worth 18% of your final grade. Assignments will be posted on the course webpage and announced on course piazza page. Expect to have 1 assignment every 3 weeks (roughly). For programming questions, [choose one of the following two languages: Python \(recommended\) or Julia](#).

[Completed assignments will be submitted through LEARN](#). **Submit early and often!**

As usual, it is OK to seek for help, but you must write your solutions independently and individually, and you should always acknowledge any help you get (book, friend, internet, etc.).

Mark appeals should be requested within two weeks of receiving the mark. The appeal could go either ways, so request only if you truly believe something is wrong.

Late Policy

We do **NOT** accept any late homework submissions (even if you are 1 second late), unless you have a legitimate reason with formal proof (e.g. hospitalization, family urgency, etc.) and notify the instructor in time. Traveling, busy with other stuff, internet connections, having a bad day, or simply forgetting to submit, are not considered legitimate.

Project (28% of the final grade) (**Tentative**)

Students are expected to compete in a Kaggle competition and submit the results there. More details can be found on the course webpage later. At the end of the competition, you need to submit a report that outlines your approach, along with your source code to reproduce your results. Please use the provided template and limit yourself to 4 pages (excluding references).

For CS680 students, per the instructor's approval, you may replace the above competition with a research project. Talk to the instructor with a tentative plan.

Academic Integrity

In order to maintain a culture of academic integrity, members of the University of Waterloo community are expected to promote honesty, trust, fairness, respect and responsibility. Check the university website for more information.

Grievance

A student who believes that a decision affecting some aspect of his/her university life has been unfair or unreasonable may have grounds for initiating a grievance. Read Student Petitions and Grievances, Section 4. When in doubt please be certain to contact the department's administrative assistant who will provide further assistance.

Discipline

A student is expected to know what constitutes academic integrity to avoid committing an academic offence, and to take responsibility for his/her actions. A student who is unsure whether an action constitutes an offence, or who needs help in learning how to avoid offences (e.g., plagiarism, cheating) or about "rules" for group work/collaboration should seek guidance from the course instructor, academic advisor, or the undergraduate Associate Dean. For information on categories of offences and types of penalties, students should refer to Policy 71, Student Discipline. For typical penalties check Guidelines for the Assessment of Penalties.

Appeals

A decision made or penalty imposed under Policy 70 (Student Petitions and Grievances) (other than a petition) or Policy 71 (Student Discipline) may be appealed if there is a ground. A student who believes he/she has a ground for an appeal should refer to Policy 72 (Student Appeals).

Students with Disabilities

AccessAbility Services collaborates with all academic departments to arrange appropriate accommodations for students with temporary or permanent disabilities without compromising the academic integrity of the curriculum. If you require academic accommodations, please register with the AccessAbility Services at the beginning of each academic term.

Intellectual Property

Students should be aware that this course contains the intellectual property of their instructor, TA, and/or the University of Waterloo. Intellectual property includes items such as:

- Lecture content, spoken and written (and any audio/video recording thereof);
- Lecture handouts, presentations, and other materials prepared for the course (e.g., PowerPoint slides);
- Questions or solution sets from various types of assessments (e.g., assignments, quizzes, tests, final exams); and
- Work protected by copyright (e.g., any work authored by the instructor or TA or used by the instructor or TA with permission of the copyright owner).

Course materials and the intellectual property contained therein, are used to enhance a student's educational experience. However, sharing this intellectual property without the intellectual property owner's permission is a violation of intellectual property rights. For this reason, it is necessary to ask the instructor, TA and/or the University of Waterloo for permission before uploading and sharing the intellectual property of others online (e.g., to an online repository).

Permission from an instructor, TA or the University is also necessary before sharing the intellectual property of others from completed courses with students taking the same/similar courses in subsequent terms/years. In many cases, instructors might be happy to allow distribution of certain materials. However, doing so without expressed permission is considered a violation of intellectual property rights.

Please alert the instructor if you become aware of intellectual property belonging to others (past or present) circulating, either through the student body or online. The intellectual property rights owner deserves to know (and may have already given their consent).