INTRODUCTION TO MACHINE LEARNING

Syllabus: CSC 311 Winter 2020

1. Course Objective. Machine learning (ML) is a set of techniques that allow computers to learn from data and experience, rather than requiring humans to specify the desired behaviour manually. ML has become increasingly central both in AI as an academic field, and in industry. This course introduces the main concepts and ideas in machine learning, and provides an overview of many commonly used machine learning algorithms. It also serves as a foundation for more advanced ML courses.

By the end of this course, the students will learn about (roughly categorized)

- Machine Learning Problems: Supervised (regression and classification), Unsupervised (Clustering), Reinforcement Learning
- Models: Linear and Nonlinear (Basis Expansion and Neural Networks)
- Loss functions: Squared Loss, Cross Entropy, Hinge, Exponential, etc.
- Regularizers: ℓ_1 and ℓ_2
- Probabilistic viewpoint: Maximum Likelihood Estimation, Maximum A Posteriori, Bayesian inference
- Bias and Variance Tradeoff
- Ensemble methods: Bagging and Boosting
- Optimization technique in ML: Gradient Descent and Stochastic Gradient Descent

The students are expected to learn the intuition behind many machine learning algorithms and the mathematics behind them. Through homework assignments, they will also learn how to implement these methods and use them to solve simple machine learning problems.

2. Course webpages. The course webpage contains all course information, additional readings, assignments, announcements, office hours, etc. Please check it regularly!

- Homepage: `https://amfarahmand.github.io/csc311`
- Quercus: `https://q.utoronto.ca`
- MarkUs: Updated later
- Piazza: `https://piazza.com/utoronto.ca/winter2020/csc311/home`

3. Lectures and Tutorials. This course has two identical sections:

- L0101 (Farahmand)
  - Lecture: Tuesday 13:00–15:00 at MS 2172
  - Tutorial: Thursday 14:00–15:00 at MS 2172
- L5101 (Andrews)
  - Lecture: Thursday 18:00–20:00 at BA 1130
  - Tutorial: Thursday 20:00–21:00 at BA 1130

Even though attendance is not mandatory, it is highly recommended. We expect students to actively participate during the lectures. We may ask students to form small groups in the classroom and discuss questions asked by the instructor. We may have quizzes during the class to help them learn the content better. The quizzes will not be marked.

4. Teaching Team.

4.1. *Instructors.* The course is taught be the following instructors:

- Amir-massoud Farahmand
  Email: csc311-2020-01@cs.toronto.edu
  Office: BA2283
  Office Hours: Thursday, 13–14
- Emad A. M. Andrews
  Email: csc311-2020-01@cs.toronto.edu
  Office: BA2283
  Office Hours: Thursday, 20–22

4.2. *Teaching Assistants.* The following graduate students will serve as the TA for this course: Chunhao Chang, Rasa Hosseinzadeh, Julyan Keller-Baruch, Navid Korhani, Shun Liao, Ehsan Mehralian, Alexey Strokach, Jade Yu.

4.3. *Contacting the Instructors and TAs.* .
The best way to discuss the content of the course is by asking your questions in the classroom, in Piazza, or during the office hours. If you have a question related to the content of the course, ask it as soon as possible during the class. Do not wait until it is late (if there are too many questions during the class, we may ask you to postpone it to after the class, so we can finish the lecture, but we do not expect that to happen often).

If after reading slides and looking at resources, you still feel that you do not understand something, use office hours or Piazza. Using Piazza has the benefit that everyone in the class can see the answers, and they may learn something from it. But be careful not to ask any question that reveals a hint for your assignments.

Use emails for administrative requests that cannot be resolved during the office hours or on Piazza. In general, it would be considerate of you to reduce the number of emails you send us. We receive tons of emails per day, and answering all of them is time consuming. But if you have to ask something through email (maybe because it is a private matter or is because you require a special accommodation), please feel free to do so. We will answer them as soon as possible (but do not expect immediate answers; maybe 1 or 2 weekdays). Please do **not** send any course-related emails to our personal accounts. They will be ignored. If you have to send an email, use the official instructor email instead.

5. Course Evaluation. (Subject to minor changes)

- Four (4) homework assignments: 40%
- Five (5) reading assignments: 10%
- Midterm exam (1h): 20%
- Final exam (3h): 30%

You must achieve a minimum mark of 30% on the final exam to pass the course.

6. Course Outline. This course covers several commonly used machine learning algorithms and related methodological concepts. The tentative list of topics is:

1. Introduction to machine learning
2. Nearest neighbour methods
3. Decision trees
4. Bias-variance decomposition
5. Ensembles: Bagging and Random Forests
6. Linear models and feature expansion
7. Regression with linear models
8. Classification with linear models: Logistic Regression
9. Neural networks
10. Support Vector Machines
11. Ensembles: Boosting
12. Probabilistic Models: Maximum Likelihood Estimation, Bayesian inference, Maximum A Posteriori estimation
13. Principal Component Analysis
14. Clustering, K-Means, and EM
15. Matrix Factorization & Recommendation Systems
16. Reinforcement Learning
17. Other Recent Topics

7. Prerequisites. This is relatively a math-heavy course. It is expected that the student be comfortable using tools and concepts from (multivariate) calculus (e.g., gradient), linear algebra (e.g., norms, eigenvalues, SVD), probability theory (e.g., common discrete and continuous distributions, expectation, conditional expectation), and optimization (e.g., gradient descent, Newton method, constrained optimization and Lagrangian formulation). The student also needs to be comfortable writing medium-size codes in Python.

The official prerequisites for this course are: CSC207H1, Mat235Y1/Mat237Y1/Mat257Y1/(minimum of 77% in Mat135H1 and Mat136H1)/(minimum of 73% in Mat137Y1)/(minimum of 67% in Mat157Y1), Mat221H1/ Mat223H1/ Mat240H1, STA247H1/STA255H1/STA257H1

Exclusion: CSC411H1, STA314H1, Ece421H1. Note: Students not enrolled in the Computer Science Major or Specialist program at the UTSG, UTM, or UTSC are limited to a maximum of three 300-/400-level CSC/ECE half-courses.

8. Textbooks. There is no required course textbook. The following materials can be helpful and some of them will be referred to in the class. Most of them are freely available online.

- Trevor Hastie, Robert Tibshirani, and Jerome Friedman, The Elements of Statistical Learning, Second Edition, 2009.
- Christopher M. Bishop, Pattern Recognition and Machine Learning, 2006
- Richard S. Sutton and Andrew G. Barto, Reinforcement Learning: An Introduction, Second Edition, 2018.
- Ian Goodfellow, Yoshua Bengio and Aaron Courville, Deep Learning, 2016
- Kevin Murphy, Machine Learning: A Probabilistic Perspective, 2012.

- Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani, An Introduction to Statistical Learning, 2017.
- Shai Shalev-Shwartz and Shai Ben-David, Understanding Machine Learning: From Theory to Algorithms, 2014.
- David MacKay, Information Theory, Inference, and Learning Algorithms, 2003.

9. Assignments. There will be four (4) homework assignments in this course. These assignments will have both mathematical derivations and programming components in Python. These assignments must be submitted through MarkUs. We encourage typesetting the written part of the homework assignments using LaTeX, but scans of handwritten solutions are also acceptable as long as they are legible. If they are not, we may deduct marks. For the programming part, your Python code should execute correctly and without any errors. It should also be written clearly and be up to acceptable standards of software engineering.

There will be five (5) reading assignments. These are selected from seminal papers in machine learning or closely related fields such as statistics. These papers complement the topics we cover in the course, or show you how a research paper in ML is written. You have to read them and try to understand as much as possible. Some papers are easy and some may be difficult. It is not important that you completely understand a paper, but you should put some effort into it.

We ask you to summarize each paper in a short paragraph (say, 100-200 words) and try to come up with two suggestions on how the method described in the paper can be used or extended. These five assignments contribute 10% to your final mark.

The reading assignments are only lightly evaluated: You should submit your summary before the end of the class. We randomly check some of your summaries to see whether or not they are in fact relevant to the assigned paper. If they are, you get the points. If they are not, you will get 0% out of 10% (even if it is only one of your submissions).

9.1. *Collaboration policy.* Collaboration on the assignments is not allowed. Each student is responsible for their own work. Discussion of assignments should be limited to clarification of the handout itself, and should not involve any sharing of pseudocode or code or simulation results. Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.

10. Exams. There will be a midterm exam (tentatively) on the week of February 24th and a final exam. Detail will be announced in class and on the course webpage.

11. Late policy. Each assignment can be submitted until three (3) days after the due date. Afterwards, the submission is blocked. For each late day, 10% will be deducted.

Extensions will be granted only in special situations, and you will need a Student Medical Certificate or a written request approved by the course coordinator at least one week before the due date.

12. Grading concerns. Any requests to have graded work re-evaluated must be made within one week of the date the grade is released. Re-evaluation may result in a decrease in the grade.

13. Computing. In the assignments and project, you may need to write your own programs, debug them, and use them to conduct various experiments, plot curves, etc. You may use any programming language, but `Python` is preferable. On the midterm and final exams, you will not be asked to understand a particular languages syntax and will not need to provide code in any particular language. You may be asked to provide pseudo-code.

14. Missed Tests.

- If a test is missed for a valid reason, you must submit documentation to the course instructor.
- If a test is missed for a valid medical reason, you must submit the University of Toronto Verification of Student Illness or Injury form to your instructor within one week of the test.
- The form will only be accepted as valid if the form is filled out according to the instructions on the form.
- Important: The form must indicate that the degree of incapacitation on academic functioning is moderate, serious, or severe in order to be considered a valid medical reason for missing the term test. If the form indicates that the degree of incapacitation on academic functioning is negligible or mild then this will not be considered a valid medical reason.
- If the midterm test is missed for a valid reason then the final test will be worth 50% of your final grade. Other reasons for missing a test will require prior approval by your instructor. If prior approval is not received for non-medical reasons then you will receive a term test grade of zero.

15. Accommodations for Disability. Students with diverse learning styles and needs are very welcome in this course. In particular, if you have a disability/health consideration that may require accommodations, please feel free to approach us and/or Accessibility Services at (416) 978 8060; `studentlife.utoronto.ca/as`.

16. Academic Integrity Elements. Academic integrity is a fundamental value for us. We take cases of academic misconduct very seriously. Plagiarism, cheating, making up facts, unauthorized assistance, misrepresenting your identity, etc. are all considered academic offences. Please familiarize yourself with the University of Toronto's Code of Behaviour on Academic Matters.