PROBABILISTIC METHODS FOR MACHINE LEARNING

Syllabus: CSC412 Winter 2025

1. Instructors.

- Denys Linkov
  Email: csc412prof@cs.toronto.edu
  Office Hours: TH 5-6 BA2272

2. Lectures. This course has two identical sections

- CSC412H1-S-LEC5101
  - Lectures - MP134 - Mon 6-8pm
  - Tutorials - MP134 - Mon 8-9pm
- CSC412H1-S-LEC5201
  - Lectures - MP134 - Thurs 6-8pm
  - Tutorials - MP134 - Thurs 8-9pm

3. Teaching Assistants. Email: csc412ta@cs.toronto.edu

- Stephan Rabanser
- Vahid Balazadeh
- Kelly Zhu
- Alireza Mousavi
- Chen-Hao Chao
- Jerry Ji
- Alireza Keshavarzian

4. Course webpages. The course webpage contains all course information, additional readings, assignments, announcements, office hours, etc. You are expected to check the following sites regularly!

- diophontine.github.io/csc412/
- q.utoronto.ca
- https://piazza.com/utoronto.ca/winter2025/csc412
- https://markus.teach.cs.toronto.edu/markus

5. Course Evaluation.

- 4 assignments: 35%
  - A1 - 10% Due Feb 02
  - A2 - 10% Due Feb 16

  - A3 - 6% Due Mar 16
  - A4 - 9% Due Mar 30
- Midterm exam: 25%
  - Feb 24/27 Inclass
- Final exam: 40%

6. Course Outline. This course covers several commonly used machine learning algorithms and related methods. Topics may include:

1. Introduction
2. Probabilistic Models
3. Decision theory
4. Directed Graphical Models
5. Markov Random Fields
6. Exact inference
7. Message passing
8. Sampling & MCMC
9. Hidden Markov Models
10. Variational inference
11. Variational Autoencoders
12. Embeddings
13. Attention
14. Sparse Autoencoders
15. MoE - Mixture of Experts
16. Contrained Decoding
17. Speculative Decoding
18. Diffusion Models

7. Prerequisites & Exclusions.

- CSC311H1/ CSC411H1/ STA314H1/ ECE421H1/ ROB313H1/ CSCC11H3/CSC311H5
- Exclusion: STA414H1

8. Textbooks. There is no required course textbook. The following materials can be helpful.

- Kevin Murphy (2022). Probabilistic Machine Learning: An Introduction
- Kevin Murphy (2023). Probabilistic Machine Learning: Advanced Topics
- Trevor Hastie, Robert Tibshirani, Jerome Friedman (2009). The Elements of Statistical Learning
- Christopher M. Bishop (2006). Pattern Recognition and Machine Learning
- David MacKay (2003). Information Theory, Inference, and Learning Algorithms

9. Assignments. There will be 4 assignments in this course. The assignments will be released on the course webpage. Submission instructions will be provided with each assignment.

9.1. *Collaboration policy on the assignments.* Assignments must be your own individual work. After attempting the problems on an individual basis, you may discuss and work together on the homework assignments with up to two classmates. However, **you must write your own code and write up your own solutions individually and explicitly name any collaborators at the top of the homework.**

10. Exams. There will be an in-class midterm exam on the week of Feb 24th. Details will be announced in class and on the course webpage.You can use an optional A4 cheat sheet - double-sided. Final exam date is TBD. You can use two optional A4 cheat sheets - double-sided **Collaboration on the exams is not allowed!** Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.

11. Late policy. Ten percent of the value will be deducted for each late day (up to 3 days, then submission is blocked). No credit will be given for assignments submitted after 3 days.

12. Absence declaration. Students who are absent from academic participation for any reason (e.g., COVID, cold, flu and other illness or injury, family situation) and who require consideration for missed academic work have been asked to record their absence through the ACORN online absence declaration. The absence declaration is considered sufficient documentation to indicate an absence and no additional information or documentation should be required when seeking consideration from an instructor. Students should also advise their instructor of their absence. Instructors will not be automatically alerted when a student declares an absence. It is student's responsibility to let instructors know that they have used the Absence Declaration so that you can discuss any needed consideration, where appropriate.

13. 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.

14. 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. For some of the assignments, we will provide you a starter code in `Python` only.

15. 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 absence declaration form and let your instructor know immediately.
- The form will only be accepted as valid if the form is filled out according to the instructions on the form.
- If the midterm test is missed for a valid reason then the final test will be worth 65% 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.

16. Accommodation for Disability Policy. Please send your documented accessibility requests directly to the instructor, at least a week before the due date of each evaluation item. Extensions may be granted, and the duration will be determined based on the letter from the Accessibility Services at the University of Toronto.