ELEC 400M : Machine Learning Fundamentals for Engineers Fall 2022

Lecture Syllabus

Instructor: Prof. Xiaoxiao Li *Scribe: Xiaoxiao Li*

1 Course Description

- This is a Special Topics course focusing on foundations and concepts of machine learning and its applications to engineering problems. Students are expected to have obtained a solid background in probability and random variables, as demonstrated by successfully completing one of the following courses: ELEC/STAT 321, MATH/STAT 302, MATH 318.
- This course can be applied towards the advanced electives requirement of the BASc in Electrical Engineering program and the BASc in Computer Engineering program.
- Further, credit will be granted for only one of: ELEC 400M, CPSC 330, CPSC 340.

2 Contact Information

- Instructor: Xiaoxiao Li
- Email: xiaoxiao.li@ece.ubc.ca

3 Time and Location

- Class Meets: Tue & Thu || 5:00 pm – 6:30 pm || ORCH 1011
- TA Office Hours: TBA
  - Wenlong Deng dwenlong@student.ubc.ca
  - Sadegh Mahdavi smahdavi@ece.ubc.ca
  - Chun-Yin Huang chunyinhuang17@gmail.com
- Instructor Office Hours: Thursday 4-5pm (by appointment only)

4 Prerequisites

- Proficiency in Python
  All class assignments will be in Python.
- College Calculus, Linear Algebra
  You should be comfortable taking derivatives and understanding matrix vector operations and notation.

- Basic Probability and Statistics
  You should know basics of probabilities, Gaussian distributions, mean, standard deviation, etc.

5 Course Goals

The course aims to provide an introductory level exposure to machine learning concepts with a balance between practical and theoretical aspects and hands-on experience suitable for engineering students. At the end of the course, students will be able to: apply the concept of learning and machine learning to real-world problems; identify the machine learning tasks and select suitable machine learning models; execute training and validation of models; apply techniques to control overfitting and assess the success of learning; use and modify available software for machine learning models and apply to new problems; realize the ongoing challenges and problems in machine learning; continue with specialized and advance machine learning courses.

6 Computational Resources

GPU computing is required for this class. I strongly recommend to Google Colab or use your own/lab's GPU since that is the most convenient way of writing and testing code with GUI. Click here to try out the Colab tutorial.

7 Course Content

This course will cover the following topics:

1. Introduction to Machine Learning (Sep 6)
2. Machine Learning Basics
   - Assignment 1 Submission (Sep 13)
   - Concepts and Basic Math (Sep 8)
   - Linear Regression and Shrinkage(Sep 13)
   - Logistic Regression (Sep 15)
   - Intro to Machine Learning Practice (Python, Pytorch, Co-lab, etc.) (Sep 20)
   - Model Training and Evaluation (Sep 27)
   - Assignment 1 Submission (Sep 29)
3. Supervised Learning
   - Announce Assignment 2 (Oct 4)
   - Introduction to Supervised Learning and K-Nearest Neighbors (Oct 4)
   - Support Vector Machines (Oct 6, Oct 11)

- Decision Tree and Random Forest (Oct 13, Oct 18)

4. Unsupervised Learning
   - Clustering and Gaussian Mixture Model (Oct 20)
   - Principal Components Analysis (Oct 25)
   - Assignment 2 Submission (Oct 28)
5. Overview of Deep Neural Networks
   - Announce Assignment 3 (Oct 28)
   - Background and Introduction to Multilayer Perceptrons (Oct 27)
     - Fully Connected Layers
     - Activation Functions
     - Objective Functions
   - Backpropogation and Optimization (Nov 1)
   - Assignment 3 Submission (Nov 11)
6. Introduction to Deep Learning Models
   - Convolutional Neural Networks (Nov 8, Nov 15)
   - Recurrent Neural Networks (Nov 17)
   - AutoEncoder (Nov 22)
   - Generative Adversarial Network (Nov 24)
7. Machine Learning in Real Applications
   - Real Problems and Solutions (Nov 29)
   - Self-Supervised Learning for NLP, Speech and Image (Dec 1)
   - System for ML and ML for System (Dec 6)
8. Report Submission (Nov 30)
9. Final Project Report Submission (Dec 16)

8 Grading, Assignments, and Final Project

- 4 Assignments: 60% = 3*20%
  - Conceptual and practical questions
  - Programming questions
- Article reading and report:10%
  - Comment on the recent AI topics: fairness, privacy, ...
  - Comment on the recent AI products: Alexa, Apple keyboard, ...

  - For Teamwork, no more than 3 people.

- Final project: 30% [1]
  - A machine learning project including data collection, data preprocessing, data analysis using machine learning models. You need to submit codes together with a well structured report (at least 2 pages and no more than 10 pages). **No Teamwork allowed**.
  - *Passing the course does on conditional on if you pass the final project*

- Late submission will result in *0.8 decay per day. Extension is only accepted via applying for Academic Concession.

9 Suggested Reading Materials

- Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. The elements of statistical learning. Vol. 1. No. 10. New York: Springer series in statistics, 2001.
- Müller, Andreas C., and Sarah Guido. Introduction to machine learning with Python: a guide for data scientists. " O'Reilly Media, Inc.", 2016.
- Goodfellow, Ian, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. Deep learning. Vol. 1, no. 2. Cambridge: MIT press, 2016.
- Torfi, Amirsina. Deep Learning Roadmap. https://www.machinelearningmindset.com/books/

10 Acknowledgment

* Our course materials and design are referred to the the following resources, thanks for the great work done by the smart people!

- https://speech.ee.ntu.edu.tw/ tlkagk/courses.html
- http://cs231n.stanford.edu/
- http://deeplearning.cs.cmu.edu/
- https://www.deeplearningbook.org/lecture_slides.html
- https://www.cs.princeton.edu/courses/archive/spring16/cos495/
- http://ttic.uchicago.edu/ shubhendu/Pages/CMSC35246.html
- https://www.cc.gatech.edu/classes/AY2018/cs7643_fall
- http://introtodeeplearning.com/
- https://hrlblab.github.io/cs3891.html
- Prof. Lutz Lampe's teaching materials
- Prof. Qi Dou's teaching materials

[1] You need to pass the final project to pass the course.