MIE 324: Introduction to Machine Intelligence – Fall 2018

Basic Course Information

Welcome! We live in exciting technological times in which computers have become essential in almost every human endeavour, whether it is music, design, healthcare, business or entertainment. The goal of this course is to provide a introduction to practical machine learning methods that have recently been shown to be very effective in many fields.

It will also give you an opportunity to improve your software skills and practices. You'll also get some practice at communication through a project proposal, reports and presentations. The key elements of the course are the hands-on practical programming assignments, and the course project.

The main topics to be covered are: representation of data, classification, supervised learning, natural language processing, deep neural networks, convolutional neural networks, recurrent neural networks, loss functions, activation functions, neural network optimization approaches from an intuitive stance, learning rate, regularization, underfitting, overfitting. There will be guest lectures from practitioners who present case studies of problems addressed and solved, and the subject of current research. We will also discuss and consider some of the ethical considerations that arise in this corner of AI.

The assignments will make use of the python language and the PyTorch deep learning framework.

Instructor

Jonathan Rose
Email: Jonathan.Rose@ece.utoronto.ca
Web: `http://www.eecg.toronto.edu/~jayar/`
Office: Engineering Annex Room 319
Office Hours: TBD

Teaching Assistants

Harris Chan - harris.chan@mail.utoronto.ca
Kevin Shen - tingke.shen@mail.utoronto.ca

Grading

The grade in the course will be based on five assignments, the midterm examination, and the large course project, which will be done in groups of two. The midterm examination will cover all material up to a point that will be specified in class. The project will include a proposal, report, in-class presentation and an individual interview.

Item	**Weight**
Assignments (5)	30%
Midterm	25%
Project	45%

The midterm test is tentatively set to be held on Friday October 26, 2018, during the 11am tutorial. Calculators are not allowed in the midterm.

Course Web Site

The web site for this course can be found on Quercus, accessible from `q.utoronto.ca`. All class announcements, and handouts - including assignments, and solutions will be accessible there. We will use Piazza.com, to serve as a home for a course discussion board where you can ask questions about course content and assignments. To sign up for Piazza go here: `piazza.com/utoronto.ca/fall2018/mie324`

Course Schedule

Lectures: Tuesday, Thursday and Friday from 2-3pm in BA 1240.
Tutorial: Friday at 11am in MY380 (the new engineering building at 55 St. George Street.)

Assignments

A key part of the learning in this course is in the hands-on programming assignments. The tutorials will be a time to get help with the assignments.

There are five labs in this course (including lab 0), and the lab schedule is given in Table 1. Labs are due at 11:59pm on the due date given in the Table. There will be a late penalty assessed of 10% of achievable grade per day late, up to a maximum of five days, after which the assigned grade will be zero. The first day of penalty occurs at the moment the assignment is late, and each subsequent penalty occurs exactly 24 hours later.

These assignments will likely take a substantial amount of your time.

Number	Title	Date Assigned	Date Due	Portion of Grade
0	Review of Python	September 6	September 13	2%
1	Image Classifier: Cats and Dogs	September 13	September 20	4%
2	MLP for Income Prediction	September 20	October 4	8%
3	CNN for Gesture Recognition	October 4	October 18	8%
4	NLP/Word2Vec and Sentence Classification	October 18	November 1	8%
Total				30%

Table 1: Assignment Schedule and Overall Grade Portion

There will be a late penalty of 10% of the grade of an assignment for every day it is late, where the first day is charged when the deadline passes.

Project

The project in this course will be a major software project that makes use of the material of the course to do something of your own choosing. The project in this course will be done in groups of two, and will account for 45% of your final grade. There are several phases and specific deadlines of the project, with several interim deliverables. Please see the associated document titled "Project Structure and Information."

Item	**Portion of Full Grade**
Proposal Document and Presentation	4%
In-Class Progress Report	4%
In-Class Final Presentation	7%
Final Report	15%
Individual Report	5%
Individual Interview	10%
Total	45%

Software Frameworks, Acceleration and Getting Started Tutorials

We will be using Python 3.6 and PyTorch as the main environment. Lab 0 and Lab 1 gives you the instructions on what to install on your computer. Briefly, these are the tools that will get you up and running:

- Get Python 3.6 and libraries from as described here: `https://docs.anaconda.com/anaconda/install`.
- Get the PyCharm python development environment here: `https://www.jetbrains.com/pycharm/download/`
- To brush up on your python, consider the following tutorials: `https://learnxinyminutes.com/docs/python3/` and `http://cs231n.github.io/python-numpy-tutorial/`

Please see the instructor if you do not have a reasonable computer on which to do these assignments and the project. The early assignments should not require very much computing power. For later assignments and the project, we will be making use of remote 'cloud' computing from Google.

Textbooks

There is no required textbook for this course, but the following text has reasonable coverage of the topics in the lanugage we will be using: **Deep Learning with PyTorch**, Author: Vishnu Subramanian, `https://www.packtpub.com/big-data-and-business-intelligence/deep-learning-pytorch` also available from Amazon. Here are some reference text that you may wish to review:

1. **Deep Learning** by Goodfellow, Bengio and Courville - free at `http://www.deeplearningbook.org`, but can also be purchased in hardcover. I do not cover the material in the way this book presents it, but it can make a useful reference.
2. **Deep Learning, A Practitioner's Approach** by Patterson & Gibson (O'Reilly) - this text doesn't use the same language we're using, but the front material is closer to how I will present neural networks.
3. **Machine Learning Yearning** by Andrew Ng - this is an online draft book that can is free, localted at `http://www.mlyearning.org`

In addition, there are many online resources that can teach aspects of the course, including the PyTorch tutorials located here: `https://pytorch.org/tutorials/`.

Academic Integrity

In developing solutions to assignments, you are free to discuss your approach to the solution with fellow students. This is fine as long as the final solution is yours alone. You should not post any of your assignment questions in a private or public online discussion forum or web site in order to solicit solutions from others. Note that, under the University of Toronto code of conduct, a person who supplies an assignment to be copied will be penalized in the same way as the one who makes the copy. We will use software to detect copying that is quite sophisticated and so is difficult to defeat.

Accessibility Needs

The University of Toronto is committed to accessibility. If you require accommodations or have any accessibility concerns, please visit `http://www.studentlife.utoronto.ca/as` as soon as possible.