ECE 324: Introduction to Machine Intelligence – Fall 2019

Basic Course Information

Welcome! We live in exciting 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 topics to be covered are: representation of data, classification, supervised learning,loss functions, activation functions, neural network optimization, learning rate, regularization, underfitting, overfitting, deep neural networks, convolutional neural networks, natural language processing and recurrent neural networks. 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

Karthik Bhaskar - karthikraja.kalaiselvibhaskar@mail.utoronto.ca
Michael Dimmick - m.dimmick@mail.utoronto.ca
Zining Zhu - zining.zhu@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)	25%
Midterm Test	30%
Project	45%

The midterm test will be held on Wednesday November 20th, from 2-4pm, in the Exam Centre Room 310. No aides or calculators are allowed in the test.

Course Web Sites

The web site for this course can be found on Quercus, accessible from `q.utoronto.ca`. All class announcements, and handouts - including assignments, will be accessible there. We will use Piazza.com to serve as the course discussion board where you can ask questions about course content, assignments and project. If you were not automatically signed up to the class Piazza page by the instructor, go to the sign up page here: `https://piazza.com/utoronto.ca/fall2019/ece324/home`,

Course Schedule

Type	Day	Time	Room
Lecture	Monday	10am-11	MS 2172
Lecture	Tuesday	11am-12	WB 130
Lecture	Thursday	9am-10	RS 211
Tutorial	Thursday	10am-11	RS 211

Assignments

A key part of the learning in this course is in the hands-on programming assignments.

There are five assignments in this course; the schedule is given in Table 1. Assignments are due at 9pm on the due date given in the Table. There is a penalty-free grace period of one hour past the time deadline. Work that is submitted between 1 hour and 24 hours past the deadline will receive a 20% grade deduction (that is 20% of the total achievable grade). No subsequent (past 24 hours) late work will accepted.

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

Number	Title	Date Assigned	Date Due	Grade Portion
1	Review of Python	September 5	September 12	1%
2	Single Neuron Classifier from Scratch	September 12	September 19	6%
3	MLP for Income Prediction	September 19	October 3	6%
4	CNN for Image Recognition	October 3	October 17	6%
5	NLP/Word2Vec and Sentence Classification	October 17	October 31	6%
Total				25%

Table 1: Assignment Schedule and Grade Portion

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%
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 3.7 and PyTorch as the main software environment. Assignment 1 and Assignment 3 give you the instructions on what to install on your computer. Briefly, these are the tools that will get you up and running:

- Get Python 3.7 and libraries from as described here: `https://docs.anaconda.com/anaconda/install`.
- Get the PyCharm python **community** 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 references may prove useful:

1. **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. `https://www.oreilly.com/library/view/deep-learning/9781491924570/` also available from Amazon.
2. **Deep Learning with PyTorch**, Author: Vishnu Subramanian. This book describes the pytorch-based code, but does not give a very view of the fundamentals.
3. **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.
4. **Machine Learning Yearning** by Andrew Ng - this is an online draft book that can is free, located at `https://www.deeplearning.ai/machine-learning-yearning/`
5. **Data Science from Scratch, First Principles with Python** by Joel Grus - this book covers much broader territory than this course, and it does it from the bottom up; it is relevant to the first few weeks of the course, but is also considered as a key long-term reference for the field. `https://joelgrus.com/2019/05/13/data-science-from-scratch-second-edition/`

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.