

Overview

An introduction to the programming and application of artificial intelligence, with a focus on deep learning. By the end of the course, students will be able to implement neural networks to perform classification on image, text, and other types of data. Students will also have a high-level understanding of neural network models used to generate images, such as autoencoders and GAN. We will focus on implementations using Python, Numpy, and PyTorch.

Meeting Time

Lecture: Monday 6pm-8pm MY330

Lecture: Thursday 6pm-7pm MY330

Lab: Thursday 7pm-8pm SF1013

Course Staff

| | |
|----------------------------|---|
| Instructor | Lisa Zhang |
| Office Hours | Thursday 4pm-5pm BA2197 and by appointment |
| Email | lczhang (at) cs.toronto.edu Please include APS360 in your email subject For non-personal, course related questions, please use Piazza |
| Teaching Assistants | Andrew Jung, Huan Ling, Farzaneh Mahdisoltani, Jake Snell |

Website

The course website is required reading and is located here:

<https://www.cs.toronto.edu/~lczhang/360/>

The website contains important information and links: Quercus (for submitting work and receiving grades/feedback), Piazza (course discussion board), lecture materials, lab handouts, project information, and other links.

Textbook

None. Lecture notes and recommended readings will be posted each week.

If you prefer reading textbook, the following text might be helpful:

“Deep Learning with PyTorch” by Vishnu Subramanian

“Introduction to Deep Learning” by Eugene Charniak (uses tensorflow instead of PyTorch)

Tentative Schedule

Here is the rough week-by-week schedule, subject to change.

| Weeks | Monday (lec) | Thursday (lec) | Thursday (lab) |
|--------------|----------------------------------|----------------------------|-----------------------|
| Week 1 | Introduction, Artificial Neurons | Pigeons to Neural Networks | Lab 1 |
| Week 2 | Neural Network Training | Neural Network Terminology | Lab 2 |
| Week 3 | (Victoria Day) | Multi-class Classification | Lab 3a |
| Week 4 | Convolutional Neural Networks | CNN Architecture | Lab 3b |
| Week 5 | Deconvolutions and Autoencoders | Word Embeddings | Lab 4 |
| Week 6 | Recurrent Neural Networks | Recurrent Neural Networks | Lab 5 |
| Week 7 | Midterm Review | Midterm Test | Midterm Test |
| Reading week | | | |
| Week 8 | (Canada Day) | Text Generation | Text Generation (lec) |
| Week 9 | Generative Adversarial Networks | Guest Lecture (TBD) | Project |
| Week 10 | Reinforcement Learning | Guest Lecture (TBD) | Project |
| Week 11 | AI Ethics | Guest Lecture (TBD) | Project |
| Week 12 | Final Term Test Review | Final Term Test | Final Term Test |
| Week 13 | (Civic Holiday) | Project | Project |
| Week 14 | Presentations | Presentations | Presentations |

Marking Scheme

The marking scheme for this course is as follows:

| Work | Weight | Comments |
|-----------------|--------|---|
| Assignments | 15% | |
| Midterm | 20% | Thursday June 20, 6pm-8pm, Room TBD |
| Project | 30% | |
| Final Term Test | 35% | Thursday August 1, 6pm-8:30pm, Room TBD |

Labs

A key part of the learning in this course is the hands-on programming labs. The scheduled lab time will provide time to work on and receive TA help on the lab exercises. There are 5 labs, all due in the first half of the course. All labs should be completed individually. The weighting of the labs are as following:

| Lab | Weight | Deadline |
|----------|--------|--|
| Lab 1 | 2% | Due May 15, 9pm |
| Lab 2 | 3% | Due May 22, 9pm |
| Lab 3(a) | 1% | Due May 24, 11:59pm (*note the irregularity) |
| Lab 3(b) | 3% | Due June 5, 9pm |
| Lab 4 | 3% | Due June 12, 9pm |
| Lab 5 | 3% | Due June 19, 9pm |

Project

The project in this course will require students to implement a major piece of software that makes use of the material of the course to do something of your own choosing. It is a substantial focus of the second half of this course. The project will be done in groups of three, and will account for 30% of your final grade. There are several phases and specific deadlines of the project, with several interim deliverables:

| Deliverable | Weight | Comment |
|--|--------|---------------------------------|
| Team Formation and Uniqueness Approval | | Due June 21, 9pm |
| Project Proposal | 3% | Due June 27, 9pm |
| Progress Meeting with a TA Mentor | 3% | Between July 8 and July 15, 9pm |
| Progress Report | 4% | Due July 24, 9pm |
| Presentation Slides | | Due August 12th, 3pm |
| Project Presentation | 10% | August 12th / 15th, 6pm-8pm |
| Project Repository | 10% | Due August 15, 9pm |
| Project Addendum | | Due August 15, 9pm |

Late Work Policy

There is a penalty-free grace period of **one hour** past the deadline. Any work that is submitted between 1 hour and 24 hours past the deadline will receive a **20%** grade deduction. No other late work is accepted. Quercus submission time will be used, not your local computer time. You can submit your labs as many times as you want before the deadline, so please submit often and early.

Midterm

The midterm test takes place during class, and covers all material from weeks 1 to 6. No aids will be permitted. The midterm room location will not be the regular lecture location, and is to be announced. Last term's midterm will be posted on the course website.

Final Term Test

The final term test is comprehensive, and will take place August 1, 6pm-8:30pm. No aids will be permitted. Last term's exam will be posted on the course website. This final term test replaces the final exam.

Use of TurnItIn

Turnitin.com will be used to assist in the evaluation of the originality of some of the term work. Turnitin.com is only a tool which will assist in detecting plagiarism. Normally, students will be required to submit their course essays to Turnitin.com for a review of textual similarity and detection of possible plagiarism. In doing so, students will allow their essays to be included as source documents in the Turnitin.com reference database, where they will be used solely for the purpose of detecting plagiarism. The terms that apply to the University's use of the Turnitin.com service are described on the Turnitin.com web site - <http://turnitin.com/>. **If you prefer not to submit your work to Turnitin.com, please let the instructor know by May 14.**

Accessibility

The University of Toronto and your instructors are committed to accessibility. If you require accommodations, or there is anything course-related we can do to help, please get in touch.

**Academic
Offenses**

The University of Toronto expects you to be a full member of the academic community and to observe the rules and conventions of academic discourse. In particular, all of the work you submit must be your own and your work must not be submitted by someone else. Plagiarism is a form of academic fraud, and the department uses software that compares submissions for evidence of similarity. The full text of the policy that governs Academic Integrity at U of T (the “Code of Behaviour on Academic Matters”) can be found at:

www.governingcouncil.utoronto.ca/policies/behaveac.htm

Please don't cheat. It is unpleasant for everyone involved, including us. Here are a couple of general guidelines to help you avoid plagiarism:

- Never look at another student's homework. Never show another student your solution. This applies to all drafts of a solution and to incomplete and even incorrect solutions.
- Keep discussions with other students focused on *concepts* and *examples*. Any code or solutions that you submit should be your alone.
- Do not post any of your assignment questions in a private or public online discussion forum or website 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.