

## Overview

A basic introduction to the history, technology, programming and applications of artificial intelligence, with emphasis on fast evolving field of machine learning. The main focus of the course will be neural networks, deep learning, and their applications. An applied approach will be taken, where students get hands-on exposure to AI techniques through the use of state-of-the-art machine learning software frameworks.

## Course Staff

<b>Instructor</b>	Lisa Zhang
<b>Office Hours</b>	TBD
<b>Email</b>	lczhang (at) cs.toronto.edu Please include APS360 in your email subject For non-personal, course related questions, please use Piazza
<b>Head Teaching Assistant</b>	Hojjat Salehinejad
<b>Teaching Assistant</b>	Andrew Jung
<b>Teaching Assistant</b>	Bibin Sebastian
<b>Teaching Assistant</b>	Kingsley Chang

## 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, assignment handouts, project information, and other links.

## Textbook

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

## Tentative Schedule

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

Weeks	Monday	Thursday
Week 1	Introduction	Pigeons to Neural Networks
Week 2	Terminology	Neural Network Training
Week 3	Define Neural Networks	Convolutional Neural Networks (CNN)
Week 4	Training CNN	CNN Architecture
Week 5	Autoencoders	Unsupervised Learning; word2vec
Week 6	Language Models	Recurrent Neural Networks
Reading week		
Week 7	Distillation	Midterm Test
Week 8	Guest Lecture (TBD)	Generative Adversarial Networks
Week 9	Guest Lecture (TBD)	Reinforcement Learning
Week 10	Guest Lecture (TBD)	Transfer Learning
Week 11	Fairness in AI	Ethics of AI
Week 12	Presentations	Presentations

## Marking Scheme

The marking scheme for this course is as follows:

Work	Weight	Comments
Assignments	20%	
Midterm	15%	Thursday Feb 28, RW117
Project	30%	
Final exam	35%	

Assignments

A key part of the learning in this course is the hands-on programming assignments. The tutorials will be a time to get help on these assignments. There are 5 assignments, each worth 4% of your course grade, all due in the first half of the course. All assignments should be completed individually. Assignments are due at 9pm (not midnight!)

Assignment	Weight	Deadline
Assignment 1	4%	Due Jan 20, 9pm
Assignment 2	4%	Due Jan 27, 9pm
Assignment 3	4%	Due Feb 3, 9pm
Assignment 4	4%	Due Feb 10, 9pm
Assignment 5	4%	Due Feb 17, 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
Project Proposal	3%	Due Feb 24, 9pm
Meeting with a TA Mentor	3%	Between Mar 4 and Mar 11, 9pm
Progress Report	4%	Due Mar 17, 9pm
Presentation Slides		Due Mar 29, 9pm
Project Presentation	10%	Apr 1 and Apr 4 during class
Project Report	10%	Apr 5

Late Work Policy  
Midterm Exam

There will be a late penalty of 20% for any late work. Late work is accepted up to 24 hours past the deadline.

Our midterm test takes place during class, and covers all material from weeks 1 to 6.

The final exam is comprehensive and takes place after classes are over. **You must earn 40% or above on the exam to pass the course; otherwise, your final course mark will be set no higher than 47%.**

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](http://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.