Overview

Machine learning is a powerful set of techniques that allow computers to learn from data rather than having a human expert program a behaviour by hand. Neural networks are a class of machine learning algorithm originally inspired by the brain, but which have recently seen a lot of success at practical applications. They're at the heart of production systems at companies like Google and Facebook for face recognition, speech-to-text, and language understanding. This course gives an overview of both the foundational ideas and the recent advances in neural net algorithms. Roughly the first 2/3 of the course focuses on supervised learning: training the network to produce a specified behaviour when one has lots of labelled examples of that behaviour. The last 1/3 focuses on unsupervised learning.

Course Staff

LEC0101	Wednesday 9am-11am, DV 2082
Instructor	Pouria Fewzee
Office Hours	Wednesday 12pm-2pm MN5107
Email	pouria.fewzee (at) utoronto.ca Please include CSC321 in your email subject For logistic related questions, please email Lisa
LEC0101	Wednesday 11am-1pm, IB 335
Instructor	Lisa Zhang *Course Coordinator
Office Hours	Monday 12pm-2pm DH3078 and by appointment
Email	lczhang (at) cs.toronto.edu Please include CSC321 in your email subject For non-personal, course related questions, please use Piazza

Website

The course website is required reading and is located here:

`https://www.cs.toronto.edu/~lczhang/321/`

The website contains important information and links: Quercus (for announcements), Markus (for submitting assignments), Piazza (course discussion board), lecture materials, assignmnet handouts, and other links.

Textbook

We will be using the CSC321 Course Notes written by Prof. Roger Grosse, to be posted on the course website.

Tentative Schedule

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

Week	Lecture	Tutorial	Deadline
1	Introduction; Linear Regression	Pre-requisite review (optional)	
2	Gradient Descent; Classification	Linear Regression	Homework 1
3	Nearest Neighbour; Generalization	Linear Classification	Homework 2
4	Multi-Layer Perceptrons	PyTorch	Project 1
5	Distributed Representation	Optimization	Homework 3
6	Convolutional Neural Network (CNN)	CNN in PyTorch	Project 2
	Reading Week		
7	Midterm; CNN Application	No tutorials	Project 3 Data Collection
8	Recurrent Neural Network (RNN)	Midterm Takeup	
9	Word Embeddings	RNN	Homework 4
10	Unsupervised Learning and Autoencoders	Fancy RNNs	Project 3
11	Residual Network	Autoencoder	Homework 5
12	Generative Adversarial Networks		Project 4

Marking Scheme

The marking scheme for this course is as follows:

Work	Weight	Comments
Math Homework	20%	There are 5 in totoal, completed individually.
Programming Projects	20%	There are 4 in total, completed individually or in pairs.
Midterm	15%	Feb 26th
Exam	40%	
Floating (added to midterm or exam)	5%	We choose the one that gives you the highest grade

Assessments

There are five written math homework assignments, to be completed individually. These are small, regular checkups to help you keep up with the course.

There are also four programming projects, to be completed on Google Colab either individually or in a group of 2. In addition to correctness, you are graded on design, code quality, and the quality of your written explanations.

Students must earn 40% or above on the exam to pass the course; otherwise, the final course mark will be set no higher than 47%.

Late Work Policy

All assignments (homework and projects) are to be submitted electronically using MarkUs by 9pm on the due date. Assignments submitted electronically will be timestamped based on the server time, not the student's local/PC time.

There will be a one hour grace period for late assignments. In addition, each student will receive six grace tokens; each grace token can be used for a four-hour extension for an assignment. For example, you may choose to use all six grace tokens on the first assignment, extending its deadline by 24 hours. Or, you may wish to use three tokens for each of two assignments, extending each deadline by 12 hours. MarkUs automatically tracks and deducts grace tokens.

If you work with a partner on the programming project, grace tokens are deducted from every team member, not just one of you. For example, if Alice and Bob are working together, and wish to submit a project 6 hours late, they must both have at least two grace tokens remaining.

Missed Term Work

Students should immediately contact the instructor via email by no later than the due date if a deadline cannot be met. In the case of medical problems, you are required to have a doctor complete a UTM medical certificate, which must say "I saw [student's name] on [date] and it is my medical opinion based on my examination that the student is medically unfit to complete his work at this time. He/she should be able to continue with their studies by [date]." Medical notes saying that the Doctor saw you after you recovered are not adequate and will not be accepted.

Midterm

The midterm test takes place during class, and covers all material from weeks 1 to 6. No aids will be permitted. If the university is closed on February 26th, the midterm will be rescheduled for the next lecture on March 4th.

Final Exam

The final term test is comprehensive. No aids will be permitted.

Academic Offenses

Honesty and fairness are fundamental to the University of Torontos mission. Plagiarism is a form of academic fraud and is treated very seriously. The work that you submit must be your own and cannot contain anyone elses work or ideas without proper attribution. You are expected to read the handout How not to plagiarize (http://www.writing.utoronto.ca/advice/using-sources/how-not-to-plagiarize) and to be familiar with the Code of behaviour on academic matters, which is linked from the UTM calendar under the link Codes and policies.
You may not discuss the homework with anyone other than your instructors and teaching assistants. You may not discuss the programming problems with anyone other than your instructors, teaching assistants, and your partner. The only exceptions are:

- Asking and answering questions on Piazza. Please do not share partial solutions or code on Piazza.
- During office hours, under the supervision of an instructor or teaching assistant.
- During certain tutorials, under the supervision of an instructor or teaching assistant.