CSC413/2516: Neural Networks and Deep Learning

University of Toronto Computer Science Winter 2025

All course information (except place/time) is subject to change until January 8, 2025. After that point, any significant changes will be announced to the class mailing list.

Course Description

Machine learning is a powerful set of techniques that allow computers to learn from data rather than having a human expert program behaviour by hand. Neural networks are a class of machine learning algorithms originally inspired by the brain but have recently seen much success in practical applications. They're at the heart of production systems at companies like Google and Meta for face recognition, speech-to-text, and language understanding.

This course gives an overview of 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 many labelled examples of that behaviour. The last 1/3 focuses on unsupervised learning and reinforcement learning.

Prerequisites

- Introduction to Machine Learning:
  - CSC311H1/CSC311H5/ CSCC11H3/ CSC411H1/ STA314H1/ ECE421H1/ ROB313H1
- Multivariable Calculus:
  - MAT235Y1/MAT237Y1/MAT257Y1/ MAT257Y5/ MAT291H1/ MAT294H1/ AER210H1/ (MAT232H5, MAT236H5)/ (MAT233H5, MAT236H5)/ (MATB41H3, MATB42H3)
- Linear Algebra:
  - MAT223H1/ MAT240H1/ MAT185H1/ MAT188H1/ MAT223H5/ MATA23H3

Where and When

Section	Instructor	Lecture Time	Lecture Room	Tutorial Time	Tutorial Room
LEC 101/2001	Alice Gao	Tue 3-5pm	SS 1083	Thu 3-4pm	MP 137
LEC 201/2101	Roger Grosse	Wed 9-11am	WI 1016	Fri 10-11am	WI 1016
LEC 301/2201	Nikita Dhawan	Wed 1-3pm	ES B149	Fri 1-2pm	ES B149

Teaching Staff

Instructor: Alice Gao

Hi! I'm Professor Alice Gao (she/her/hers), one of your instructors for this course. Feel free to call me Professor, Professor Gao, or Alice. I'm an assistant professor in the computer science department, and this is my third year at UofT. I've taught Intro to AI (CSC384) and Intro to ML (CSC311) regularly. This is my first time teaching Deep Learning (CSC413), and I'm very excited! If you enjoy this course, consider taking another one with me. I'd love to get to know you better! I also work on research projects in computer science education. If you're interested in research, consider taking CSC494/495 with me. We'll work on a research project and learn together! For more information on my teaching experience, research projects, and how to request a reference letter, visit my website.

Please do not email me regarding this course. Instead, attend my office hours or send a message to the course email account (see the "Seeking Help" section below for details). My office is in Bahen 4240. If you want to chat, I recommend making an appointment beforehand.

Office Location: Bahen 4240

E-mail address: ax.gao@utoronto.ca

Instructor: Roger Grosse

E-mail Address: rgrosse@cs.toronto.edu

Hi! I'm Roger Grosse (he/him/his), another instructor for the course. You can call me Roger. I'm as Associate Professor of Computer Science, focused on deep learning and AI alignment. I last taught this course in 2019, and am looking forward to getting back into it! For more information, please check out my website.

Office location: Pratt 290F (that's the D. L. Pratt building, not the E. J. Pratt Library!)

Email: rgrosse@cs.toronto.edu

Instructor: Nikita Dhawan

Hello! My name is Nikita (she/her/hers), and I'll teach one of the sections of this course. I am 4th-year PhD student at UofT. My research focuses on machine learning and building or using neural nets, which means it uses or has used a lot of the techniques in this course! Please feel free to come chat about topics in the course, research at the intersection of ML and healthcare, or anything else, during my office hours. More information can be found on my website.

Office Location: Bahen 2272 (through the Help Centre at 2270)

E-mail Address: nikita@cs.toronto.edu

Schedule

Lectures

Lecture recordings will be generated and posted automatically on the OCCS Student App. *Remember that course videos and materials belong to your instructor and the University and are protected by copyright.* You are permitted to download videos and materials for your own academic use, but *you may not copy, share, or otherwise distribute them* without explicit permission from the instructor.

Click on the lecture topic for additional materials (readings, slides, etc.).

Note: slides will differ slightly between instructors, but will be equivalent in terms of core content which will appear on tests.

Course Schedule

Week	Week Of	Lecture Topic	Deadlines
1	@January 6, 2025	Introduction & Linear Models	
2	@January 13, 2025	Multilayer Perceptrons and Backpropagation	
3	@January 20, 2025	Autodiff, Distributed Representations, and Optimization	
4	@January 27, 2025	Convolutional Neural Networks and Image Classification	HW1 due Fri, Jan 31, at 1pm
5	@February 3, 2025	Optimization and Generalization	
6	@February 10, 2025	Interpretability and Large-Scale Generative Models	Test 1 on Mon, Feb 10 6-7pm
	@February 17, 2025	Reading Week. No lectures or tutorials.	
7	@February 24, 2025	Recurrent Neural Networks and Attention	
8	@March 3, 2025	Transformers and Autoregressive Models	HW2 due Fri, Mar 7, at 1pm
9	@March 10, 2025	Large Language Models	
10	@March 17, 2025	Generative Models of Images and Graph Neural Networks	Test 2 on Tue, Mar 18 6-7pm
11	@March 24, 2025	Deep RL	
12	@March 31, 2025	TBD	HW3 due Fri, Apr 4, at 1pm

Tutorials

Week	Week Of	Tutorial Topic
1	@January 6, 2025	Multivariable Calculus Review
2	@January 13, 2025	PyTorch

Homework Assignments

There will be three homework assignments throughout the course. Please refer to the schedule above for the due dates.

Each assignment will include both written and programming exercises:

- **Written Exercises**: These questions are designed to give you additional practice with the course material. Each assignment will contain 2-3 conceptual questions, which should take a few hours to complete.
- **Programming Exercises**: You may be asked to write, modify, or use Python code to implement a simple version of a learning procedure recently covered in the course. Along with your code, you will need to submit a brief write-up describing the results you obtained.

Midterm Tests

There will be two term tests. Test 1 will cover the materials in weeks 1-4. Test 2 will cover the materials in weeks 5-8. Think of the two tests as equivalent to 1 midterm.

Check the schedule above for the dates and times for the term tests. The tests will be graded and returned through Crowdmark.

Based on educational research, the best way to prepare for a test is to do practice problems. If you need help creating practice problems, please ask us during office hours or on Piazza.

Final Exam (for the undergrad students, i.e. CSC413)

The final exam will cover all the topics in the course, with more emphasis on the topics after term test 2. You can bring one aid sheet (two-sided 8.5'' by 11') to the final exam. The final exam schedules will be available at the end of February on the A&S page.

Final Project (for the graduate students, i.e. CSC2516)

🏗️ Final Project

Resources and Seeking Help

- Quercus:
  - We will post non-public materials, e.g. homework and test solutions, on Quercus.
- Piazza:
  - You should have been signed up on Piazza automatically. The link is through Quercus.
  - We will send announcements through Piazza, which will reach you via email. Doing so will allow you to ask follow-up questions on Piazza.
- Course email address: csc413-2025-01@cs.toronto.edu
  - Please email us using the course email account for any administrative issues.

- Instructor Office Hours
  - Students can attend the OH of whichever instructor they prefer.
  - Nikita Dhawan: Wednesdays 3-4PM @ Bahen 2272
  - Alice Gao: (These will be combined office hours for CSC311 and CSC413)
    - Wednesdays 3:30 - 5 pm @ Bahen 4240
    - Thursdays 1:30 - 3 pm @ Bahen 4240
  - Roger Grosse: Wednesdays 4-5pm @ Pratt 290F
- TA Office Hours: We will be scheduling TA office hours and announcing them on Piazza.

Course Policies

📜 Course Policies

Course Policies

Grading Scheme

Undergraduate students must obtain a minimum grade of 40% on the final exam to pass this course. If your final exam grade is less than 40%, the maximum final grade you can obtain for the course will be 47%.

Component	Weight
Three Assignments	30%
Two Midterm Tests	30%
Undergraduate Students: Final Exam Graduate Students: Final Project	40%

The higher of the two midterm test grades will count for 20%, and the lower will count for 10%, for a total of 30%.

Special Consideration Policies

If you need an assignment deadline extension, cannot complete an assignment, or cannot write a test due to extraordinary circumstances beyond your control, please submit a special consideration request as soon as possible. Special consideration will be evaluated on a case-by-case basis and is not granted automatically. Sometimes, we cannot grant you exactly the special consideration you seek.

Special Consideration Policies on **Midterm Tests:**

If you **miss one or both midterm tests** for approved reasons, you will have an opportunity to take make-up tests. Please send your supporting documentation to the course email account as soon as you miss a test. We will schedule one time

slot for both make-up tests sometime after test 2. This means that if you miss both midterm tests, you will have to write both make-up tests back to back.

Lateness and Special Consideration Policies on Assignments:

If you need an extension on an assignment deadline due to unexpected physical or mental health issues, disability-related barriers, or similar circumstances, please contact us at the course email account with supporting documentation. We can grant extensions of up to 7 days in such cases.

For other reasons, such as a heavy workload or unexpected personal issues, we will accept late assignment submissions for up to 3 days past the deadline. However, a 10% grade penalty will be applied for each day the assignment is late.

Special Consideration Policies if you are registered with Accessibility Services:

Your accommodation letter will allow for an extension of up to 7 full days. However, due to the incremental nature of CS courses, granting such a long extension from the onset may cause you to fall behind and be disadvantaged. As such, we will start by suggesting an initial 3-day extension. We will grant the 7-day extension later if necessary.

Please complete and submit the special consideration request form (TBD) to request special consideration. We aim to respond within 48 business hours.

Include Supporting Documentation: Ensure you attach all necessary documentation with your request. For detailed information on acceptable documentation and when to use the Absence Declaration tool, please read the Student Absences page from the Faculty of Arts & Science.

Invalid Reasons for Special Consideration Requests: The following reasons are not valid because they are neither unexpected nor outside of your control:

- Heavy course load
- Multiple assignments due in the same week

- Multiple tests scheduled in the same week
- Needing to catch up on missed work

If you are struggling with stress or time management, please contact your College Registrar. They can provide support through wellness counselling, academic advising, and learning strategist services.

Remark Requests

If you believe there was an error in the marking of an assignment or a test, or if you have questions about how we marked your work, you may submit a remark request.

Please note: We do **NOT** accept remark requests in the first 24 hours after the grades are released. This time should be spent reading and understanding your assessment results, the sample solutions, and the marking scheme.

Here are some examples of **valid** reasons for requesting a remark:

- Your answer was marked incorrectly based on the marking scheme.
- There was an error in adding up the marks.
- You believe you should receive more marks based on the marking scheme.
- You interpreted the question differently, which caused your answer to differ from the sample solutions.

After the initial 24 hours, you will have **one week** to submit a remark request by filling out the appropriate form below. Please provide a detailed justification to help us process your request efficiently. We will review all remark requests after the one-week period is over.

Remark Request Form for Assignments

Remark Request Form for Term Tests

Academic Integrity

By the time you get to an advanced course like csc413, you've heard this lots of times, so we'll keep it brief: avoid academic offences (a.k.a. cheating).

Collaboration policy: You are welcome to work together with other students on the homework. You are also welcome to use any resources you find (online tutorials, textbooks, papers, etc.) to help you complete the homework. However, you must write up your submission by yourself and not use any content generated by someone else or generative AI. You must cite all the collaboration and resources you used to complete each assignment. If you hand in homework that makes use of content that you did not create or you do not disclose the collaboration or resources, you will get a 0 for that homework. Note also that if you rely too much on outside resources, you will likely not learn the material and will do poorly on the exams, during which such resources will not be available.

Generative AI Policy

General questions: You are welcome (and encouraged) to ask AI chatbots (such as ChatGPT, Claude, Gemini, or Grok) general questions about the course content.

Access to Microsoft Copilot: As a UofT student, you can access a protected version of Microsoft Copilot. You can access Microsoft Copilot by navigating to microsoft365.com/chat and follow the prompts to login to your University of Toronto Microsoft 365 account. We encourage you to explore it and use it to enhance your learning.

Homework Assignments: You may ask general questions about concepts related to the homework questions. However, you may not directly ask them for hints on homework assignment questions; e.g., you should not be copying and pasting directly from the assignment handout into the chat interface. Also, you may not directly use the outputs of AI chatbots in your homework solutions (even paraphrased) unless instructed to do so.

You must include any chat transcripts related to a homework assignment along with your submission for the assignment. We will interpret the above policy leniently when judging whether the GenAI use is appropriate, though we reserve the right to update the GenAI policy if we find that students are using it in a way that reduces the educational value of the homeworks.

Final Project: For the CSC2516 final project, you may use GenAI however you like. However, you must properly attribute any work (e.g. generated code). You must attach all chat transcripts related to the project along with your submission.