CSC413H1S 20261 (All Lectures): Neural Networks and Deep Learning

Jump to Today ✎ Edit

CSC413H1S Neural Networks and Deep Learning

University of Toronto, St. George Campus, Winter 2026

All information on this syllabus is subject to change until Tuesday, January 6, 2026.

Table of Contents

Your Instructors

Alice Gao

Hi! I am Professor Alice Gao (she/her). Feel free to call me Professor, Professor Gao or Alice. I am an Assistant Professor, Teaching Stream in the Computer Science department. This is my fourth year at UofT. This year, I will be teaching Intro to ML (CSC311) and Deep

Learning (CSC413). Besides teaching, I also supervise undergraduate students for various projects. If you're interested in research, consider taking CSC494/495 with me. Visit **my website (https://www.cs.toronto.edu/~axgao/)** for more information on my teaching experience, research projects, and how to request a reference letter.

Please do not email me regarding this course. Instead, attend my office hours or email the course email account. My office is in Bahen 4240. If you want to chat, I recommend making an appointment beforehand.

Nikita Dhawan

Hello! My name is Nikita (she/her/hers), and I'll teach one of the sections of this course. I am 5th-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 about me can be found on my **website (https://www.cs.toronto.edu/~nikita/)**.

Seeking Help

↑ Contents ↑

The course staff are ready and eager to help you inside and outside of class. Please remember that seeking help is not a sign of weakness. To the contrary, we encourage you to seek help from us early and often so that we can help you succeed in this course. You can get in touch with us through the following channels.

Instructor Office Hours:

The instructors have set aside several weekly office hours. During office hours, feel free to ask us about course materials or anything else. If there are a lot of students and an upcoming course deadline, we will prioritize answering course-related questions.

Instructor Office Hours

	Day	Time	Location
Alice Gao	Monday	3pm - 4pm	Bahen 4290
Alice Gao	Tuesday	11am - noon	Bahen 4290
Nikita Dhawan	Tuesday	5pm - 6pm	Bahen 2272

Note that Alice will not be holding the office hour on Monday during the first week.

TA Office Hours

The TAs will hold weekly in-person and online office hours to answer your questions. We will announce the TA office hour information on Piazza.

Piazza

Piazza is the main platform for asking questions and receiving help from the course staff. It allows you to post questions anytime and benefit from shared discussions with peers.

Please post publicly whenever possible so others can learn from your questions. If your post contains sensitive information (e.g., assignment hints or personal details), make it private instead. We aim to respond within **2 business days**; questions posted outside business hours may be answered the next business day.

All **important announcements** will be made on Piazza (not Quercus), since Piazza allows for immediate follow-up questions. You are responsible for checking Piazza a few times each week to stay updated.

Course email address

csc413-2026-01@cs.toronto.edu (mailto:csc413-2026-01@cs.toronto.edu)

The course email address is for administrative requests only (e.g., enrolment, accommodations, extensions, exam conflicts). Questions related to course content (lectures, assignments, practice problems, etc.) should be posted on Piazza or during office hours. For special considerations requests and remark requests, it is sufficient to fill out the respective online form. Please only send an email if you haven't heard back about your request in a week or you need to provide information that cannot be included in the form responses.

Course Description

↑ Contents ↑

Deep learning is the branch of machine learning focused on training neural networks. Neural networks have proven to be powerful across a wide range of domains and tasks, including computer vision, natural language processing, speech recognition, and beyond. The success of these models is partially thanks to the fact that their performance tends to improve as more and more data is used to train them. Further, there have been many advances over the past few decades that have made it easier to attain good performance when using neural

networks. In this course, we will provide a thorough introduction to the field of deep learning. We will cover the basics of building and optimizing neural networks in addition to specifics of different model architectures and training schemes. See the Course Schedule table below for the topics and deadlines in the course.

There is no required textbook. The recommended textbook is **Dive into Deep Learning (http://d2l.ai/)** .

Course Schedule

↑ Contents ↑

This course schedule below outlines all the lectures, labs, assignment due dates, and term test dates for the whole term. I highly recommend that you take some time to **put all the important dates in your calendar right now**.

Course Schedule

Week	Lectures	Tests	Programming Assignments
Jan 5 - 9 Week 1	Introductions, Linear and Logistic Regression		**HW01** posted on Friday.
Jan 12 - 16 Week 2	Multi-layer Perceptrons, Backpropagation		**HW01** due 4pm Friday. **HW02** posted on Friday.
Jan 19 - 23 Week 3	Under/overfitting, Regularization, Numerical Stability, Autodiff		**HW02** due 4pm on Friday. **HW03** posted on Friday.
Jan 26 - 30 Week 4	Gradient descent, Adaptive Gradient Methods	**Test 1**, Date/Time: 6pm to 8:30pm, Thu, Jan 29, 2026. Location: EX 100	**HW03** due 4pm on Friday. **HW04** posted on Friday.

Week	Lectures	Tests	Programming Assignments
		Coverage: weeks 1 to 3.	
Feb 2 - 6 Week 5	Convolutional Neural Networks, Batch/Layer Normalization, Residual Connections		**HW04** due 4pm on Friday. **HW05** posted on Friday.
Feb 9 - 13 Week 6	Sequences, Recurrent Neural Networks, Sequence-to-Sequence Learning		**HW05** due 4pm on Friday. **HW06** posted on Friday.
Feb 16 - 20	Reading Week		
Feb 23 - 27 Week 7	Attention		**HW06** due 4pm on Friday. **HW07** posted on Friday.
Mar 2 - 6 Week 8	Transformers		**HW07** due 4pm on Friday. **HW08** posted on Friday.

Week	Lectures	Tests	Programming Assignments
Mar 9 - 13 Week 9	Large Language Models	**Test 2**, Date/Time: 6pm to 8:30pm, Thu, Mar 12, 2026. Location: EX 100 Coverage: weeks 4 - 8.	**HW08** due 4pm on Friday. **HW09** posted on Friday.
Mar 16 - 20 Week 10	Architecture Grab Bag: Low Rank Adaptation, Graph Neural Networks, Mixture of Experts		**HW09** due 4pm on Friday. **HW10** posted on Friday.
Mar 23 - 27 Week 11	Architecture Grab Bag: Variational Autoencoders, Transposed Convolution, UNet.		**HW10** due 4pm on Friday. **HW11** posted on Friday.
Mar 30 - Apr 3 Week 12	Deep Learning Engineering. Fairness, Accountability, and Transparency		**HW11** due 4pm on Friday.

Grading Scheme

↑ Contents ↑

Please take a careful look at the grading scheme chart below. You 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%.

Grading Scheme

Component	Weight
Theoretical	
Test 1	15%
Test 2	15%
Final Exam	40%
Applied	
Programming Assignments	30%

Lectures

There are 2 hours of lectures every week. The times and locations are in the table below. Please refer to the course schedule above for the detailed weekly outline.

All lectures will be recorded and posted automatically. *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 personal academic use, but *you may not copy, share, or otherwise distribute them* without explicit permission from the instructor.

There will be no lecture during the third scheduled time slot per week. Instead, we will use these hours for TA office hours as much as possible. If a TA is holding an office hour during the time slot, we will announce it on Piazza a few days in advance.

Lecture and Tutorial Time Slots and Locations

Section	Lectures (2 hours)	TA Office Hours
101	Tuesday, 3-5pm, BA 1190	Thursday, 3-4pm, GB 220
201	Wednesday, 1-3pm, GB 221	Friday, 1-2pm, BA 1160
301	Tuesday, 1-3pm, GB 220	Thursday, 1-2pm, SS 1073

Programming Assignments

Deep learning is fundamentally a very applied subject. As such, the best way for you to learn the materials is to convert the theory learned in class into code. You will complete **11 programming assignments**, approximately one assignment each week. These assignments provide hands-on experience implementing deep learning models with the support of TAs.

The programming assignments will be completed **individually**. They are due at **4 PM ET on Fridays via MarkUs (https://markus.teach.cs.toronto.edu/markus/courses)** . Your **final assignment grade** will be the average of the **best 9 out of 11 assignments** since we will automatically **drop your 2 lowest assignment grades**.

Programming assignments will be **graded using automated unit tests**. We will provide some of the unit tests to you on MarkUs. You can run these unit tests on MarkUs as long as you have tokens available. Each token, once used, will regenerate in one hour. After the deadline, your assignment grade will be determined by running all the unit tests.

Two Tests

There will be **two tests**, each approximately **80 minutes long.** See the course schedule for the dates and times of the two tests. You can bring **one aid sheet (two-sided 8.5" by 11")** to each test. The term tests will be graded and returned via **Crowdmark**. You can access Crowdmark through the Crowdmark LTI link in the Modules section of Quercus.

The tests will contain **theoretical questions only;** there will be no programming questions. You can expect two types of questions.

- **Conceptual** questions aim to test your understanding of the properties of various architectures. These questions will be primarily short-answer questions.
- **Calculation** questions aim to test your understanding of the mathematical aspects of the architectures. For example, these questions may ask you to calculate the output dimensions or the number of parameters for a given architecture. These questions will typically be longer questions with multiple parts.

Research in education shows that the best way to prepare for these types of assessments is to **practice solving problems regularly, ideally under time constraints**.

Final Exam

The final exam will cover all the topics in the course, with a slight emphasis on the materials not covered in the two tests. The final exam will be 3 hours long. You can bring **one aid sheet (two-sided 8.5" by 11")** to the final exam. The final exam schedules will be available on **the A&S page (https://www.artsci.utoronto.ca/current/faculty-registrar/exams-assessments/exam-assessment-schedule)** .

Practice Problems

To help you prepare for the supervised assessments (two tests and the final exam), the instructors will be drafting practice problems for each topic during the term and release these problems whenever they become available. We will not be posting solutions to these problems. We encourage you to use these problems as resources to prepare for the tests and the final exam.

Course Policies

↑ Contents ↑

Generative AI Policies and Guidelines

In this course, you may use generative artificial intelligence (AI) tools to support your learning on any **unsupervised** assessments, namely the programming assignments. You may **not** use generative AI on **supervised assessments**, including the two term tests, and the final exam. Generative AI tools include (but are not limited to) ChatGPT, Copilot, and open-source models that you have trained or deployed yourself.

While generative AI can be a powerful learning aid, we caution against relying on it to complete coursework. The majority of your final grade (70%) is based on supervised assessments, where you are not allowed to use generative AI. For this reason, we encourage you to use generative AI primarily as a tool to **support understanding and practice**, rather than as a substitute for your own problem-solving.

When working on **practice problems**, you are encouraged to use generative AI to clarify the problem statement, explore high-level solution strategies, or review a sample solution. After consulting a sample solution, however, we strongly recommend stepping away from the material and later attempting to reproduce the solution on your own, without generative AI and under time constraints. This process—known as *retrieval practice*—is one of the most effective ways to build durable understanding. If you do not practice reproducing solutions independently, you may find that you cannot recall or apply them during supervised assessments, even if they seemed clear when first reviewed. In this way, over-reliance on generative AI can create a false sense of mastery that ultimately undermines learning and performance.

The programming assignments can often be completed quickly with the help of generative AI. However, if your goal is to develop the ability to write and adapt code independently, we strongly recommend attempting the labs on your own first. You are, of course, encouraged to use generative AI when you are genuinely stuck. Completing labs with minimal AI assistance helps you practice translating

machine learning theory into working code—a non-trivial skill that reinforces conceptual understanding and is essential for success on term tests, the final exam, and future coursework.

Academic Integrity

↑ Contents ↑

Academic offences are taken very seriously and can have serious consequences. To help you understand expectations in this course, we outline common examples of academic offences and clarify how collaboration is permitted.

Submitting work that is not your own is an academic offence. For the programming assignments, this includes submitting code written by another student, submitting the same or substantially similar code as another student, or sharing code for others to submit. You may discuss ideas and strategies with others, but all submitted code must be written by you.

Using unauthorized aid during a supervised assessment is also an academic offence. In this course, supervised assessments include the term tests and the final exam. During these assessments, students may not use generative AI tools, external resources, or assistance from others unless explicitly stated otherwise.

At the same time, collaboration is encouraged when done appropriately. For programming labs, you may discuss problems, approaches, and debugging strategies with other students and consult external resources. However, you must write and submit your own code and clearly attribute all sources you consulted, including websites, tools, or individuals. Similarly, for written work, you may discuss ideas and receive feedback from others, but the final writing must be your own and any sources must be properly cited.

Students are encouraged to review the **Academic Integrity at U of T (https://www.academicintegrity.utoronto.ca/)** website for additional guidance on collaboration, attribution, and strategies for your learning while following academic integrity principles.

Special Consideration Policies

If you miss a deadline 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.

Policies for **Programming Assignments**:

- Recall that we will be dropping the 2 lowest assignment grade when calculating your final assignment grade.
- You can request an extension of up to 7 days by filling out the special considerations form below. The extension will be granted automatically.
- If you need special considerations beyond the extension of 7 days (and dropping the 2 lowest assignment grades), please send us an email at the course email account **csc413-2026-01@cs.toronto.edu (mailto:csc413-2026-01@cs.toronto.edu)** .

Policies for **Tests**:

- If you have a scheduling conflict with the test date/time, you must notify the course staff at the course email account **at least one week before the test date**. We will schedule one make-up test within one week of the original test date and notify you of the date and time of the make-up test. Students who report conflicts later than one week before the original test date will **not** be eligible for the make-up test, as it must be scheduled in advance.
- If you **miss the test and the make-up test** for approved reasons, we will shift the test's weight to the final exam.

Policies for students 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 appropriate special consideration request form below. We will aim to respond to you within 48 business hours.

Special consideration request form (https://forms.office.com/r/GD9VGpbwt3)

Make sure to include supporting documentation with your request. Please read the new **Student Absences (https://www.artsci.utoronto.ca/current/academics/student-absences)** page from the Faculty of Arts & Science carefully. It contains detailed information on the recognized forms of documentation and the circumstances under which you should use the Absence Declaration tool.

Below are some **invalid** reasons for applying for a special consideration request. These reasons are invalid because the circumstances are **neither unexpected nor outside of your control**.

- Heavy course load
- Multiple assignments are due in the same week.
- Multiple tests are scheduled in the same week.

- I need to catch up on missed work.

If you have difficulty managing stress and time, don't hesitate to contact your College Registrar, who can suggest wellness counselling, academic advising, and/or learning strategist services.

Remark Requests

↑ Contents ↑

If you found an error in the **marking of a test**, you may submit a remark request. We do **NOT accept remark requests in the first 24 hours after the grades are released**. You should spend this time reading and understanding your assessment results, the sample solutions and the marking scheme.

Below are some examples of **valid** reasons for requesting a remark.

- My answer was marked incorrectly based on the marking scheme.
- There was an error when adding up the marks.
- I should get more marks based on the marking scheme.
- I interpreted the question differently, which caused my answer to differ from the sample solutions.

After 24 hours, you will have **one week** to submit a remark request by filling out the appropriate form below. Please provide a detailed justification --- this will help us process your request efficiently.

Remark request form for tests (TBD)

We will process the remark requests after one week.

Student Support Resources

↑ Contents ↑

UofT Mental Health Portal (https://mentalhealth.utoronto.ca/)

If you go into Explore Our Care Model, One-On-One Sessions, View Services and Resources, you will get to **Mental Health Clinical Services (https://studentlife.utoronto.ca/service/mental-health-clinical-services/)** . The first option allows you to book **same-day counselling appointments (https://studentlife.utoronto.ca/service/same-day-counselling-appointment/)** . Book an appointment by calling Health & Wellness at 416-978-8030 (select option 5).

If you are in a crisis, UofT Telus Health Student Support (formerly U of T My SSP) provides *real-time, confidential*, *24-hour* support for any school, health, or general life concern at no cost to you. Call 1-844-451-9700 or 001-416-380-6578 (if outside of North America).

You can also contact the Good2Talk Student Helpline (Call 1-866-925-5454 or text GOOD2TALKON to 686868).

Accessibility Services (https://studentlife.utoronto.ca/service/accessibility-services-registration-and-documentation-requirements/)

Check out the **Accessibility Services registration & documentation requirements – St. George Campus** (https://studentlife.utoronto.ca/service/accessibility-services-registration-and-documentation-requirements/) . Register with Accessibility Services before 5 p.m. on Friday, October 13, 2023, for final assessment accommodations.

Student Life Portal (https://studentlife.utoronto.ca/)

Check under **Health and Wellness** (https://studentlife.utoronto.ca/department/health-wellness/) for a range of programs and services.

Take a look at the **Academic Success** (https://studentlife.utoronto.ca/task_levels/academic-success/) section. There are many resources on topics such as **Better note-taking** (https://studentlife.utoronto.ca/task/better-note-taking/) , **Studying, concentration and memory** (https://studentlife.utoronto.ca/task/studying-concentration-and-memory/) , **Goal setting and motivation** (https://studentlife.utoronto.ca/task/goal-setting-and-motivation/) , **Reboot after an academic setback** (https://studentlife.utoronto.ca/task/reboot-after-an-academic-setback/) , etc. They also offer a free online course on **5 Keys to Succeed at UofT** (https://studentlife.utoronto.ca/program/5-keys-to-succeed-at-u-of-t-online-course/) .

Recognized Study Groups (https://sidneysmithcommons.artsci.utoronto.ca/recognized-study-groups/)

Recognized Study Groups (RSG) are student-led study groups of up to eight classmates enrolled in the same Faculty of Arts & Science course. RSGs can meet online or in person on the St. George Campus. You can apply to lead or join an RSG at the start of each academic term.

Course Summary:

Date	Details	Due