

Fundamentals of Neural Networks and Deep Learning (CSCI 4922/5922)

University of Colorado Boulder, Computer Science Department

Course Modes: The course can be taken synchronously or asynchronously. While we will meet in person at regularly scheduled times, all class sessions will be recorded and shared afterwards (in Canvas) for asynchronous viewing.

Lectures: Tuesdays and Thursdays at 12:30-1:45pm MT in MBE 155. Personal devices (e.g., laptop, phones) should only be used for our in-class activities. No food (including chewing gum) or drinks (aside from water) are allowed in the classroom.

Website: <https://dannagurari.colorado.edu/course/neural-networks-and-deep-learning-spring-2026/>

Instructor: Danna Gurari (she/her), danna.gurari@colorado.edu

Course Support:

- Teaching Assistant: Everley Tseng (she/her), yu-yun.tseng@colorado.edu
- Teaching Assistant: Nick Cooper (he/him), nicholas.cooper-1@colorado.edu
- Teaching Assistant: Jarek Reynolds (he/him), jarek.reynolds@colorado.edu
- Teaching Assistant: Nolan Brady (he/him), nolan.brady@colorado.edu

Ways to Contact Us:

- **Questions for Instructor:** The instructor will stay after lectures to answer questions.
- **Piazza:** We encourage you to post questions to Piazza. This approach can benefit other students who have similar questions, as they can view the answers on Piazza. It also provides an opportunity for peers to answer your questions.
- **Office Hours:** The TAs will host office hours every day of the week at times posted on the course website and Canvas. Instructions for how to attend are provided on the home page of Canvas.
- **Appointments:** If you have concerns that cannot be addressed during regular office hours, please email either Everley Tseng or Nick Cooper (course TAs) to request an appointment. They will check email Monday to Friday between 9am and 4pm. You can expect a response within one business day; if you do not receive a reply within that window, feel free to send a follow-up message. The TAs will involve the instructor for any items they are unable to address.

Course Overview

Summary

This course covers fundamentals of neural networks and deep learning and their use to address many societal problems. Students will learn to design and implement neural network architectures, train them, and evaluate their performance. Included will be examination of popular architectures such as fully connected, convolutional, recurrent, and transformer-based networks alongside learning strategies such as backpropagation, regularization, and transfer learning. Students will also gain practical, hands-on experience by applying learned skills to analyze visual, textual, and audio data.

Objectives

By the end of the course, the goals are for students to:

1. Recognize and analyze common neural network architectures. Towards this aim, students will:
 - Identify popular architectures — including fully connected, convolutional, recurrent, and transformer-based networks — and implement their core mathematical operations by hand
 - Compare and contrast their design principles as well as their strengths and limitations for various tasks
2. Recognize and analyze common techniques for training deep neural networks. Towards this aim, students will:
 - Identify key challenges with training neural networks of various depths and capacities
 - Compare and contrast techniques for overcoming learning challenges including training from scratch as well as performing transfer learning with frozen pretrained features, fine-tuning, and parameter-efficient methods (e.g., LoRA)
3. Gain hands-on experience with neural networks and deep learning for various artificial intelligence tasks. Towards this aim, students will:
 - Implement basic neural networks and learning algorithms from scratch
 - Design, train, deploy, and evaluate neural networks models with multiple data types using the deep learning library PyTorch and modern hardware (GPUs)
4. Explore responsible AI practices to both reduce environmental harm and promote ethical behavior in AI systems. Towards this aim, students will:
 - Experiment with sustainable techniques for model development including to speed up model inference and learning
 - Identify unethical behaviors of modern models, technical strategies to mitigate those behaviors, and societal frameworks to promote models' ethical use

5. Apply what is learned to propose and critique novel projects involving neural networks and deep learning. Towards this aim, students will:
 - Propose a project idea involving neural networks and deep learning
 - Design and execute experiments or analyses supporting the proposed idea
 - Create a presentation about the project
 - Write a report about the project
 - Review fellow students' project presentations and provide constructive feedback

Prerequisites

Machine learning (e.g., CSCI 4622 or 5622).

Readings

Readings will be assigned in the first part of the course. This will supplement the lectures by providing a more detailed and formal coverage of the topics.

Problem Sets

Three problem sets will be assigned in the first part of the course. Each assignment description will be posted in Canvas one week before the due date. These assignments will offer deeper examination of foundational concepts. Due to the large number of students in the course, we expect graded assignments will be returned 1 week after submission.

Lab Assignments

Three lab assignments will be assigned in the first part of the course. Each assignment description will be posted in Canvas 1.5-2 weeks before the due date. These assignments will give students hands-on practice in developing and analyzing deep learning models. Due to the large number of students in the course, we expect graded assignments will be returned 1 week after submission.

Final Project

The final third of the course will center on strengthening students' skills with applying their deep learning knowledge to create original work and communicating about it. This effort will consist of three milestones to help students define/refine the project scope so that ultimate success is possible: an outline, presentation, and final report. Details about each project milestone will be posted on the course website prior to each milestone's deadline.

Resources

Assigned readings will be posted on the course website for each class meeting. The primary texts will be "Neural Networks and Deep Learning" by Michael Nielsen; "Deep Learning for NLP and Speech Recognition" by Uday Kamath, John Liu, and James Whitaker; and original research publications that established foundational concepts in neural networks and deep learning. All texts are available online for free.

Distinction Between Undergraduate and Graduate Sections

Assignments are shared across the undergraduate and graduate versions of this course. However, the standards of assessment differ: undergraduate students are expected to demonstrate foundational expertise of the course concepts, while graduate students are expected to demonstrate intermediate expertise through more rigorous work. Specifically, the graduate-level requirements differ in two key ways:

- Final Lab Assignment: Graduate students will train models using a substantially larger dataset, which introduces greater development complexity and so requires more sophisticated implementations.
- Final Project: Graduate students will submit a longer and more comprehensive final report, providing deeper analysis, critical evaluation, and stronger connections to the broader research or application context.

Tentative Schedule

Week	Topic(s)
1	Course introduction
2	Artificial neurons; Fully connected neural networks
3	Convolutional and recurrent neural networks
4	Fundamentals of shallow and deep learning
5	Deep learning for deeper models
6	Dense prediction; Multimodal vision-language learning
7	Attention; Transformer basics
8	Pioneering transformers
9	Power of scaling; Foundation models; Prompts
10	Fine-tuning foundation models; Practical development challenges
11	Conversational, reasoning, and agentic models; Efficient learning and inference
12	Responsible neural networks and deep learning
13	Cutting edge and future of deep learning
14	Deep learning in industry (guest speakers)
15	Project presentations and course summary

Grading

Final course scores will be calculated as follows:

% of Final Class Grade	
Problem Sets	25%
Lab Assignments	45%
Final Project	30%

Final course scores represent the following grades (scores are rounded to the nearest integer):

Grade	% of Final Class Grade
A	94-100%
A-	90-93%
B+	87-89%
B	84-86%
B-	80-83%
C+	77-79%
C	74-76%

Late Policy

Late submissions will be penalized 1% of the grade per hour up to 12 hours, rounded up to the nearest hour; then 20% for up to 24 hours late; and then 50% for up to 48 hours late. After 48 hours, no credit will be given. Exceptions will be made only for exceptional circumstances.

Extra Credit Policy

No extra credit opportunities will be provided beyond those already in the assignments.

Regrade Requests

Students may submit a regrade request within two weeks of receiving their assignment grade if they believe a mistake was made. Regrade requests submitted after this period will not be considered. All requests must be emailed to the TA, Nick Cooper.

Policies

Statement on Learning Success

Your success in this course is important to me. We will all need accommodations because we all learn differently. If there are aspects of this course that prevent you from learning or exclude you, please let the Instructor know as soon as possible. Together we'll develop strategies to meet both your needs and the requirements of the course.

Classroom Behavior

Students and faculty are responsible for maintaining an appropriate learning environment in all instructional settings, whether in person, remote or online. Those who fail to adhere

to such behavioral standards may be subject to discipline. Professional courtesy and sensitivity are especially important with respect to individuals and topics dealing with race, color, national origin, sex, pregnancy, age, disability, creed, religion, sexual orientation, gender identity, gender expression, veteran status, political affiliation or political philosophy. For more information, see the policies on [classroom behavior](#) and the [Student Conduct & Conflict Resolution policies](#).

Honor Code

All students enrolled in a University of Colorado Boulder course are responsible for knowing and adhering to the Honor Code academic integrity policy. Violations of the Honor Code may include, but are not limited to: plagiarism, cheating, fabrication, lying, bribery, threat, unauthorized access to academic materials, clicker fraud, submitting the same or similar work in more than one course without permission from all course instructors involved, and aiding academic dishonesty. All incidents of academic misconduct will be reported to the Honor Code (honor@colorado.edu; 303-492-5550). Students found responsible for violating the academic integrity policy will be subject to nonacademic sanctions from the Honor Code as well as academic sanctions from the faculty member. Additional information regarding the Honor Code academic integrity policy can be found on the [Honor Code website](#).

Acceptable Use of AI in This Class

While you may collaborate with generative AI (gen AI) tools throughout this course, any submitted work must be your own. Ways that Gen AI can contribute include for brainstorming and editing write-ups. If you have any questions about whether a specific use of gen AI is permitted for an assignment, please consult with the instructors before proceeding and/or cite how you leveraged the resource.

Requirements for Infectious Diseases

Members of the CU Boulder community and visitors to campus must follow university, department, and building health and safety requirements and all public health orders to reduce the risk of spreading infectious diseases.

The CU Boulder campus is currently mask optional. However, if masks are again required in classrooms, students who fail to adhere to masking requirements will be asked to leave class. Students who do not leave class when asked or who refuse to comply with these requirements will be referred to Student Conduct & Conflict Resolution. Students who require accommodation because a disability prevents them from fulfilling safety measures related to infectious disease will be asked to follow the steps in the “Accommodation for Disabilities” statement on this syllabus.

For those who feel ill and think you might have COVID-19 or if you have tested positive for COVID-19, please stay home and follow the [further guidance of the Public Health Office](#). For those who have been in close contact with someone who has COVID-19 but do not have any symptoms and have not tested positive for COVID-19, you do not need to stay home.

Excused Absences

A student will be given an opportunity to complete any work missed due to absences in observance of a religious holy day or military service. Campus policy regarding religious

observances requires that faculty make every effort to deal reasonably and fairly with all students who, because of religious obligations, have conflicts with scheduled exams, assignments or required attendance. The student must notify the Instructor at least two weeks in advance of the absence. The student will not be penalized for excused absences, but must complete the missed material within a reasonable time after the excused absence. Please see the [campus policy regarding religious observances](#) for full details.

University Resources for Students

There are a range of resources available on campus to support you and your academic success:

- *Mental Health and Wellness*

If you are struggling with personal stressors, mental health or substance use concerns that are impacting academic or daily life, please contact [Counseling and Psychiatric Services \(CAPS\)](#) located in C4C or call (303) 492-2277 24/7. Free and unlimited telehealth is also available through Academic Live Care. The Academic Live Care site also provides information about additional wellness services on campus that are available to students.

- *Accommodations for Disability*

[Disability Services](#) determines accommodations based on documented disabilities in the academic environment. If you qualify for accommodations because of a disability, please submit your accommodation letter from Disability Services to your faculty member in a timely manner so that your needs can be addressed. Contact Disability Services at 303-492-8671 or dsinfo@colorado.edu for further assistance. If you have a temporary medical condition or required medical isolation for which you require accommodation, see [Temporary Medical Conditions](#) on the Disability Services website.

- *Writing Center*

All students are encouraged to consult the [University Writing Center](#) for support.

- If you are experiencing an emergency or unanticipated life event, you can apply for financial assistance through the [Student Emergency Fund](#).

Personal Pronouns

CU Boulder recognizes that students' legal information doesn't always align with how they identify. Students may update their preferred names and pronouns via the student portal; those preferred names and pronouns are listed on instructors' class rosters. In the absence of such updates, the name that appears on the class roster is the student's legal name.

Sexual Misconduct, Discrimination, Harassment and/or Related Retaliation

The University of Colorado Boulder (CU Boulder) is committed to fostering an inclusive and welcoming learning, working, and living environment. CU Boulder will not tolerate acts of sexual misconduct (harassment, exploitation, and assault), intimate partner violence (dating or domestic violence), stalking, or protected-class discrimination or harassment by or against members of our community. Individuals who believe they have been subject to misconduct or retaliatory actions for reporting a concern should contact the Office of Institutional Equity and Compliance (OIEC) at 303-492-2127 or email cureport@colorado.edu. Information

about OIEC, university policies, [reporting options](#), and the campus resources can be found on the [OIEC website](#).

Please know that faculty and graduate instructors have a responsibility to inform OIEC when they are made aware of incidents of sexual misconduct, dating and domestic violence, stalking, discrimination, harassment and/or related retaliation, to ensure that individuals impacted receive information about their rights, support resources, and reporting options. To learn more about reporting and support options for a variety of concerns, visit [Dont Ignore It](#).

Sharing of Course Materials is Prohibited

Class recordings are reserved only for students in this class for educational purposes. The recordings should not be shared outside the class in any form. Violation of this restriction by a student could lead to Student Misconduct proceedings.