

Neural Networks and Deep Learning (CSCI 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-2025/>

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: Neelima Prasad (she/her), neelima.prasad@colorado.edu
- Grader: Mohammad Qazim Bhat (he/him)

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 there are any outstanding concerns that cannot be addressed during regular office hours, please first email the TAs Everley Tseng and Nick Cooper to make an appointment or solicit an answer. Please note that at least 24 hours notice will typically be needed before a TA will be able to meet. The TAs will involve the instructor for any items they are unable to address.
- **Regrade requests:** All requests must be emailed to the TA, Nick Cooper, within 2 weeks of receiving the grade to be considered.

Course Overview

Summary

This course covers the fundamentals of neural networks and deep learning as well as how they are used to address many artificial intelligence problems in society. Students will learn to design and implement multi-layered neural network architectures, train them on large amounts of data, and evaluate their performance. Included will be examination of popular architectures such as fully connected networks, convolutional neural networks, recurrent neural networks, and transformers, alongside learning strategies such as backpropagation, initialization, and regularization. Students will also gain practical, hands-on experience by applying learned skills to analyze visual data (computer vision) and textual data (natural language processing).

Objectives

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

1. Understand the key concepts for designing deep learning models, a critical precursor to effective collaborations in industry or academia. Towards this aim, students will:
 - Characterize the key architectures used in designing neural networks
 - Recognize techniques used to train and evaluate deep learning models
 - Identify strengths and weaknesses of different neural network architectures and training approaches
2. Apply deep learning models to perform various artificial intelligence tasks. Towards this aim, students will:
 - Develop deep learning models from scratch
 - Experiment with established deep learning libraries
 - Evaluate deep learning models for tasks in various application domains, including for analyzing images and text
3. Conduct and communicate about a novel project. Towards this aim, students will:
 - Propose a novel project idea (this will be an iterative process)
 - Design and execute experiments to support the proposed idea
 - Create a presentation about the project
 - Review fellow students presented research and provide constructive feedback
 - Write a report about the project

Prerequisites

Programming competency and experience with probability, statistics, and linear algebra.

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 many of the topics.

Problem Sets

Four 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 lab 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 second part 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 "Deep Learning" by Ian Goodfellow, Yoshua Bengio, and Aaron Courville and "Deep Learning for NLP and Speech Recognition" by Uday Kamath, John Liu, and James Whitaker. Both texts are available online for free. Finally, we will also read publications from conference proceedings, all of which are available online for free.

Tentative Schedule

| Week | Topic(s) |
|-------------|--|
| 1 | Introduction, Artificial neurons |
| 2 | Feedforward neural networks, Objective functions, Gradient Descent |
| 3 | Training neural networks, Model Capacity, Regularization |
| 4 | Convolutional neural networks (CNN), Image classification |
| 5 | Datasets, Hardware, Hyperparameters, Dense prediction |
| 6 | Recurrent neural networks, Tokenization, Neural Embeddings |
| 7 | Attention, Transformers |
| 8 | Transformers, Multimodal models |
| 9 | Multimodal models, Foundation models, Prompts |
| 10 | Fine-tuning foundation models, Practical development challenges |
| 11 | Model compression, Efficient learning |
| 12 | Efficient learning, Future of deep learning |
| 13 | Responsible deep learning |
| 14 | Deep learning in industry (guest speakers) |
| 15 | Course summary |

Grading

Final course scores will be calculated as follows:

| <hr/> | |
|-------------------------------|-----|
| % of Final Class Grade | |
| <hr/> | |
| Problem Sets | 35% |
| Lab Assignments | 35% |
| Final Project | 30% |

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

| <hr/> | |
|--------------|-------------------------------|
| Grade | % of Final Class Grade |
| <hr/> | |
| 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 2 hours. After 2 hours, no credit will be given.

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 an assignment grade, when it is believed a mistake was made. After the allotted time frame, regrade requests will not be considered.

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.

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.