

Neural Networks and Deep Learning (CSCI 5922)

University of Colorado Boulder, Computer Science Department

Course Lectures: Mondays and Wednesdays, 1:25-2:40pm MT

Course Mode: Synchronously remote, which means classes will be held virtually during the listed days and times. Classes will also be recorded for students who can only join asynchronously.

Course Website: <https://dannagurari.colorado.edu/course/neural-networks-and-deep-learning-spring-2024/>

Instructor: Danna Gurari

Instructor Nickname: if easier, feel free to call me Dr. G.

Pronouns: she/her

Email: danna.gurari@colorado.edu

Teaching Assistant: Everley Tseng

Pronouns: she/her

Email: yu-yun.tseng@colorado.edu

Teaching Assistant: Nick Cooper

Pronouns: he/him

Email: nicholas.cooper-1@colorado.edu

Grader: Sruthi Sampath Kumar

Pronouns: she/her

Ways to Contact Us:

- **Questions for Instructor:** Instructor will stay after each lecture to answer questions.
- **Piazza:** We encourage you to post questions to Piazza. This can benefit other students, who may have similar questions to see the answers posted on Piazza, and it also allows peers to answer your questions. You can post questions to Piazza from Canvas.
- **Office Hours:** The TAs will host office hours every day of the week at times posted on the course website. To sign-up, please enter your name on the office hours spreadsheet shared on Canvas and then join the video meeting link shared on Canvas.
- **Regrade requests:** All requests must be emailed to the TA, Everley Tseng, within 2 weeks of receiving the grade to be considered.
- **Appointments:** For any other matters, please first email the TA Everley Tseng 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.

Course Overview

Summary

This course will cover machine learning approaches that are based on neural networks, also referred to as ‘deep learning’, and how these approaches are applied to solve artificial intelligence problems. This will include examination of the history of deep learning and basic architectures used in modern deep learning algorithms. We also will practice applying deep learning algorithms to visual data (computer vision) and textual data (natural language processing). The course homework will consist of problem sets and lab assignments for the first part of the semester followed by a self-designed project the second part of the semester. Readings will come both from textbooks and research publications.

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 fundamental architectures used in designing neural networks
 - Identify techniques used to train and evaluate deep learning algorithms
 - Recognize 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:
 - Experiment with deep learning libraries, especially Keras
 - Evaluate deep learning algorithms for tasks in various application domains, including for analyzing images and text
3. Conduct and communicate 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
 - 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 the foundational concepts in deep learning. You will be given one week to complete each assignment. 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 one week and a half before the due date. These lab assignments will give students hands-on practice in developing deep learning models. You will be given one and a half weeks to complete each assignment. 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 in conducting and communicating original work by applying their deep learning knowledge in a novel way. It will consist of three milestones to help students define/refine the 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
2	Artificial neurons, Feedforward neural networks (NN)
3	Gradient descent, NN training
4	NN training, Regularization
5	Convolutional neural networks (CNN), Introduction to CV, Image classification
6	Pretrained CNN features, Model fine-tuning, Object detection
7	Image-to-image translation, Semantic segmentation, Recurrent neural networks
8	Introduction to NLP, Neural word embeddings, Introduction to attention
9	Transformers
10	Foundation models, Prompts, Multi-task learning, Multimodal learning
11	Practical systems-level development challenges, Self-supervised learning
12	Model compression, Efficient 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.

Regrade Requests

Students may submit a regrade request within 2 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 me 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](#).

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 me 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*

All of us benefit from support during times of struggle. You are not alone. There are many helpful resources available on campus and an important part of the University experience is learning how to ask for help. Asking for support sooner rather than later is often helpful. 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.