

Recent Advances in Computer Vision (CSCI 7000)

University of Colorado Boulder, Computer Science Department

Course Lectures: Mondays and Wednesdays, 4:10-5:25pm MT

Course Mode: synchronously remote, which means classes will be held virtually during the listed days and times.

Instructor: Danna (pronounced similar to “Donna”) Gurari (rhymes with Ferrari)

Instructor Nickname: for simplicity, you can call me Dr. G.

Pronouns: she/her

Email: danna.gurari@colorado.edu

Appointments: If you wish to make an appointment, the best way to do this is by email. Please note that I typically need at least 24 hours notice. If you prefer to email me a question, I typically respond to emails on weekday mornings.

Course Overview

Summary

Computer vision is increasingly becoming part of the fabric of our society, enabling new technologies to analyze images and videos in order to support our daily lives. Students in this graduate-level course will explore both fundamental and state-of-art problems in the field. This will entail examination of the types of algorithms commonly used to solve computer vision problems along with the benchmarks used to track progress on those problems. The course will be taught in a seminar style format, with students taking a very active role in reading, presenting, and critiquing recent research papers from premiere computer vision conferences (such as CVPR, ICCV, and ECCV). Early career graduate students who are considering research in computer vision or related topics are encouraged to participate.

Objectives

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

1. Understand core computer vision problems and typical solutions, a critical precursor to effective collaborations in industry and academia. Towards this aim, students will:
 - Recognize and define core computer vision problems
 - Identify benchmarks used by the research community to track progress on each problem (i.e., data source, data annotation process, evaluation metrics)
 - Identify types of algorithms commonly used to solve each problem alongside their general properties that make them well-suited for the problem
2. Analyze and present cutting-edge research. Towards this aim, students will:
 - Identify in research papers the novelty claims, mechanisms used to validate the claims (e.g., theories and experiments), and why the papers’ contributions matter to society

- Deliver oral presentations that explain research papers
 - Discuss the merits and limitations of research papers
 - Prepare and moderate discussions about research papers
3. Conduct research in computer vision to enhance their expertise on a topic of their choice. Towards this aim, students will:
- Design and execute a research project on a novel idea involving computer vision, such as analyzing an existing method, implementing/evaluating a new method, surveying the state-of-art for a specific problem, or outlining how to tackle a new computer vision problem
 - Deliver an oral presentation that explains the research
 - Review fellow students' presented research and provide constructive feedback
 - Communicate the research through a final report

Prerequisites

Machine learning experience (e.g., 4622, 5622, or 5922).

Website

<https://www.cs.colorado.edu/~DrG/Courses/RecentAdvancesInComputerVision/>

Class Participation

Students are expected to attend every class. Every student should demonstrate ongoing engagement in class discussions. Beyond two unexcused absences will lower your final grade.

Readings with Assignments

In the first part of the course, students will be expected each week to read two research papers and complete assignments about those readings. Assigned readings will be posted on the course website and assignments will be posted on Canvas. The assignments will offer training in understanding and thinking critically about computer vision research papers and concepts. The assignments will require a brief summary of select aspects of the papers and “discussion points” to explore during class. Discussion points can be in the form of questions, critiques, connections to other readings, or plausible future work that students think are interesting to investigate in greater detail.

Student-Led Lectures

In the first part of the course, students will also be expected to present research papers on a select topic. Each student/team will lead one week's lectures on one computer vision topic. For the first few weeks, the course instructor will lead all lectures.

Topics will be selected and assigned during the first week of the semester based on students' preferences (the number of topics and number of students leading each lecture may change slightly during the semester depending on the number of students in the course). Students are expected to begin preparing their lecture as soon as they receive the topic.

Each week's lectures should consist of two parts. The first lecture should include a 45-minute presentation that: (i) defines the problem, (ii) motivates the practical importance of solving this problem with a computer vision solution (i.e., applications that can/do benefit society), (iii) describes 1-2 existing datasets used to track progress on this problem, and (iv) describes metric(s) used to evaluate the performance of computer vision models. Then, the lecture will conclude with a facilitated class discussion about the merits and limitations of existing community-shared datasets and evaluation metrics, organized by the instructor around the questions and discussion points submitted by all students. The second lecture of the week should include a 45-minute presentation that describes at least two papers that each introduce a computer vision model covering: (i) the novelty claims of each paper, (ii) mechanisms used to validate the claims, and (iii) open technical questions/problems. Optionally, this lecture can also include a programming tutorial and/or demo. Then, the lecture will conclude with a facilitated class discussion about the merits and limitations of existing computer vision models, organized by the instructor around the questions and discussion points submitted by all students. Students can incorporate materials from outside sources in their presentations (for example, content from the paper's authors or slides), but proper credit MUST be given.

The student/team leading the lecture is expected to select research papers to cover that were published recently at a premiere computer vision conference (e.g., CVPR, ICCV, ECCV). The student/team must meet with the instructor at least two weeks prior to the first lecture to share 4-6 candidate papers to cover (including one about a specific dataset challenge that will be assigned to the class) and resolve any open questions. The student/team must meet with the instructor again at least one week prior to the first lecture to review a draft of the slide decks for both lectures and resolve any open questions. The instructor's goal for the meetings is to help make the presentations as awesome as possible and build students' knowledge and confidence in presenting the material.

Final Project

The second part of the course will center on a research project. The goal of this project is for students to develop their skills in conducting and communicating original research on a topic of their choice. If choosing to conduct a survey paper, it is expected to be comprehensive and critical of the literature. If proposing a new model or evaluating existing models, experiments must be conducted to demonstrate the models' strengths and weaknesses. This project will consist of four milestones to help students define/refine the scope so that ultimate success is possible: a proposal, outline, four-minute live presentation, and final report (at least six pages, single spaced, 12 point font). An additional part of this effort will be to review peers' final presentations and provide constructive feedback, such that the feedback can then be used to strengthen the final report.

Tentative Schedule

Week	Topic(s)	Assignments Due
1	Introduction	Lecture Topic Selection
2	Object Recognition	Reading Assignment
3	Object Detection	Reading Assignment
4	Semantic Segmentation, Object Tracking	Reading Assignment
5	Tentative: Scene Classification	Reading Assignment
6	Tentative: Attribute Recognition	Reading Assignment
7	Tentative: Salient Object Detection	Reading Assignment
8	Tentative: Panoptic Segmentation	Reading Assignment
9	Tentative: Video Classification	Reading Assignment
10	Tentative: Video Localization	Reading Assignment, Project Proposal
11	Tentative: Image Captioning	Reading Assignment
12	Tentative: Visual Question Answering	Reading Assignment
13	Synthesis	Project Outline
14	Efficient Computer Vision	
15	Responsible Computer Vision	Project Presentation
16	<i>No Class</i>	Final Project Report

Resources

Links to required readings will be posted on the course website for each class meeting. Readings will draw heavily from conference proceedings at top computer vision publication venues such as CVPR, ECCV, and ICCV.

Grading

Final course scores will be calculated as follows:

<hr/>	
% of Final Class Grade	
<hr/>	
Class Participation	10%
Reading Assignments	30%
Student-Led Lecture	30%
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 5 hours. After 5 hours, no credit will be given.

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

Both 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](#).

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:

- *Coping with Stress and Personal Hardships*

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 or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, please consider taking advantage of the support available to you from the [Counseling and Psychiatric Services \(CAPS\)](#).

- *Accommodations for Disability*

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. Disability Services determines accommodations based on documented disabilities in the academic environment. Information on requesting accommodations is located on the [Disability Services website](#). Contact Disability Services at 303-492-8671 or dsinfo@colorado.edu for further assistance. If you have

a temporary medical condition, see [Temporary Medical Conditions](#) on the Disability Services website.

- *Writing Center*

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

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](#).

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.

Requirements for COVID-19

As a matter of public health and safety due to the pandemic, all members of the CU Boulder community and all visitors to campus must follow university, department and building requirements and all public health orders in place to reduce the risk of spreading infectious disease. Students who fail to adhere to these requirements will be asked to leave class, and students who do not leave class when asked or who refuse to comply with these requirements will be referred to [Student Conduct and Conflict Resolution](#). For more information, see the policy on [classroom behavior](#) and the [Student Code of Conduct](#). If you require accommodation because a disability prevents you from fulfilling these safety measures, please follow the steps in the Accommodation for Disabilities statement on this syllabus.

As of Aug. 13, 2021, CU Boulder has returned to requiring masks in classrooms and laboratories regardless of vaccination status. This requirement is a temporary precaution during the delta surge to supplement CU Boulders COVID-19 vaccine requirement. Exemptions include individuals who cannot medically tolerate a face covering, as well as those who are hearing-impaired or otherwise disabled or who are communicating with someone who is hearing-impaired or otherwise disabled and where the ability to see the mouth is essential to communication. If you qualify for a mask-related accommodation, please follow the steps

in the Accommodation for Disabilities statement on this syllabus. In addition, vaccinated instructional faculty who are engaged in an indoor instructional activity and are separated by at least 6 feet from the nearest person are exempt from wearing masks if they so choose.