

CS 195: Computer Vision (Fall 2024)

Instructor: Md Alimoor Reza

Assistant Professor
Department of Mathematics and Computer Science
Drake University
Office: Collier-Scripps Room#323

Logistics: The lectures and most course content will be hosted at <https://drake.blackboard.com>. lecture sessions, and completing individual (and collaborative) activities: programming assignments, in-class activities, quizzes, participation in polls, and a final project.

Class Meeting Time and Place: We will meet in person twice weekly for lecture sessions. Section-specific information is as follows:

- Location: Collier-Scripps Room#335
- Time: Monday/Wednesday 3:30 pm - 04:45 pm

Office hours: The instructor will hold weekly office hours, either in person or online via Zoom. Please see Blackboard for up-to-date times and URLs. If you cannot make office hours, contact the instructor to make a separate appointment.

- Office hour#1: Tuesday: 12:00 pm - 2:30 pm
- Office hour#2: Thursday 12:00 pm - 2:30 pm
- Location: Collier-Scripps#323

Course overview: This is an introductory course in computer vision. I will give a broad overview of the field. The emphasis will be on algorithms, mathematical models, and classical and deep learning techniques that are broadly applicable to many problems not only in computer vision but also in other areas of AI and CS. Topics will include (tentatively and not necessarily in this order):

- *Basic image manipulation:* digital image representation, image filtering, morphology.
- *Classification:* classical (eg, bags-of-words), and deep learning based (ResNet, ViT, Swin Transformer)
- *Detection:* classical (eg, sliding window based detection), and deep learning based (Mask R-CNN, YOLO)
- *Segmentation:* classical (eg, clustering-based, graph-based), and deep learning based segmentation (SAM)
- *Probabilistic modeling:* Gaussian distributions, maximum likelihood estimate, sampling from a distribution
- *Image generation:* deep learning-based generative models (GAN, VAE, and diffusion probabilistic model)
- *Feature detection:* edge detection, corner detection, invariant interest point detection
- *Geometry and image transformations:* camera models, image transformations, 2D & 3D geometry.
- *3D Reconstruction:* stereo, structure from motion, shape from silhouette.

Learning outcome: After successfully taking this course, you will be able to:

- Study and apply a variety of computer vision **recognition techniques** *multi-label image classification, object detection, object detection, and semantic segmentation* using different datasets and state-of-the-art models.
- Study and apply a variety of computer vision **reorganization techniques**, such as *clustering based, graph-based, and vision foundation model-based* segmentation.
- Study and apply a variety content generation techniques (or **reproduction techniques**) using *generative modeling such as Generative Adversarial Network (GAN), Variational Autoencoder (VAE), and Diffusion Models*.
- Study and apply a variety 3D **reconstruction techniques** using images.
- Develop software that utilizes libraries in order to conduct computer-vision-based data analysis.
- Feel comfortable with using industry-standard tools such as Google Colab, GitHub, PyTorch, Huggingface, etc
- Gain an understanding of the pros and cons of different paradigms (**R**ecognition, **R**eorganization, **R**econstruction, and **R**eproduction) so that students can choose appropriate solutions given a problem description
- Receive hands-on experience with commonly used algorithms and software tools within computer vision.

Prerequisites: CS167 (ML) or CS143 (AI) or equivalent is recommended but not absolutely necessary. You will need to be proficient in a general-purpose programming language; as a quick rule-of-thumb, you should be able to easily implement basic matrix operations using basic data structures of the language (e.g. matrix multiplication using arrays). I recommend Python. Computer vision draws from many other fields including linear algebra, basic calculus, machine learning, graph theory, probability theory, geometry, and statistics. Exposure to these topics is helpful but not required, although you may need to do some background reading.

Textbooks and materials: Readings will come from papers and several textbooks, primarily:

- Szeliski, *Computer Vision: Algorithms and Applications, 2nd Edition (2022)*, available via PDF online at <http://szeliski.org/Book/>.
- Kevin Murphy, *Probabilistic Machine Learning: An Introduction*, <http://probml.github.io/pml-book/book1.html#toc>.

Schedule, readings, and resources will be available via Blackboard, <https://drake.blackboard.com>.

Grading and requirements:

- *Notebook assignments (30%)*: Take home programming assignments (submit on CodePost).
- *In-class activities (15%)*: Simple coding activities or paper-based tasks (submit on Blackboard).
- *Quizzes (30%)*: 3 quizzes, 10% each.
- *Final project (20%)*: one final project (see below the details).
- *Participation (05%)*: physical presence, participation in polls (not based on correctness).

Notebooks Assignments (30%) We will explore various computer vision techniques using Python programming language. You'll regularly submit notebook files (.ipynb) to showcase your skills. Since I won't typically run your code due to long computation times, please ensure your results are saved in the notebook. You can expect to submit **6-7 assignments** throughout the course.

In-class Activities (15%) You will also submit notebook files with simple coding or paper-based activities. These are easier, as they follow the lecture and are usually done during or right after class. Since I won't run your code, be sure to save the results in the notebook. You can expect to submit about **12-13 in-class activities** throughout the course.

Quizzes (30%) There will be **3 quizzes** that will be administered via Blackboard. They will not be timed, and you will have a few days to complete them. Quizzes should be completed individually. There is no time limit on these quizzes. As in the real world, you will be allowed to use external resources like the class notes and the internet. You will be required to cite any sources that you used while completing these quizzes other than the class notes.

Final project: 1 final project will consist of research or implementation on a topic of your choice (in consultation with the instructor). Projects may be done individually or in groups of up to 3 students. There will be three deliverables:

- *Project proposal*: a brief (1-2 pages) proposal, due approximately 9th week into the semester.
- *Project report*: a summary report along with an electronic copy of the source code developed during the project.
- *Project presentation*: a class presentation. students are encouraged but not required to submit their work to a regional conference (eg, CCSC) or workshop.

Attendance/Participation (05%): This class is highly interactive, meaning that active participation is both expected and the norm. You will receive credit for your participation, and it will be counted towards your final grade. I will keep track of your involvement using a signature sheet. Throughout the course, I will pose questions using polling software and conduct in-class Q&A sessions to better understand how the class is grasping the content. These responses will not be evaluated for correctness but rather for completion.

I respect your privacy. If you encounter challenges (physical health, mental health, or life in general) that interfere with your ability to participate in the course or complete your work, I will not require any kind of documentation. You also do not need to explain; you can simply inform me that you are experiencing problems and we will work together to figure out a plan that will enable you to complete the course if you want to. For example, if for some reason, you are unable to make the in-person class session, please email me and I will provide you with the Zoom link for the day so you can attend class virtually. If you are unable to participate in the course for a prolonged period, we will discuss whether an incomplete is the best option.

Academic Integrity Policy: *We take academic integrity very seriously.* You are required to abide by the Drake University policy on academic integrity, as described in the Statement on Academic Dishonesty: Cheating and Plagiarism (<https://www.drake.edu/studentlife/handbook-resources/handbook/academic/>). It is your responsibility to understand these policies. Students agree that by taking this course, papers and source code submitted to us may be subject to textual similarity review, for example, by Turnitin.com. These submissions may be included as source documents in reference databases solely for the purpose of detecting plagiarism of such papers or codes.

Drake Academic Affairs Guidelines for use of Generative AI: Drake University students should adapt to using generative AI tools responsibly in class, research, and work, while exploring their opportunities and challenges in teaching and research. Visit the *Center for Teaching & Learning Internal Site* for guidelines and resources related to the academic and professional use of Generative AI: <https://tinyurl.com/yhnsfnw7>.

Accommodation for students with disabilities: Drake University is committed to providing equitable access to learning opportunities for all students. The Disability Services office (107 Old Main) collaborates with students who have disabilities to provide and/or arrange reasonable accommodations. If you have, or think you may have, a disability (e.g., mental health, attentional, learning, autism spectrum disorders, chronic health, traumatic brain injury and concussions, vision, hearing, mobility, or speech impairments), please contact:

- Michelle Laughlin, Student Disability Services Coordinator (x1835)
- michelle.laughlin@drake.edu

to arrange a confidential discussion regarding equitable access and reasonable accommodations.

Holiday Observance: If you miss class because of a holiday or observance, you can fill out the form to automatically notify me (faculty). You can find the form on the <https://www.drake.edu/diversity/initiatives> for Initiatives and Programs, or click here to directly access the https://drake.qualtrics.com/jfe/form/SV_d5qfVUKtuTQdg7b.