Computer Vision

CS 682 | Spring 2025 | Professor Gregory J. Stein

Course Information

3 Credits
Lectures: Async via Canvas
Office Hours:
Prof. Stein, TBD

Instructor Information

Prof. Gregory J. Stein
He, Him, His
gjstein@gmu.edu

Course Description

The aim of computer vision is to compute properties of the three-dimensional world from images so that it may be understood by machines. Topics in this class include how to create panoramic images, to build a 3D reconstruction of an environment from videos, and to recognize familiar people and objects, all through analysis of images and video clips.

This course is at the senior undergraduate level and aims to introduce students to the field of computer vision through hands-on programming projects. By the end of the course, students should be familiar with the standard tools used by the computer vision community, and be able to approach research papers in the field.

See the detailed class schedule below for more details on what will be covered during the course.

Course Learning Outcomes

1. Basic knowledge of the image formation process.
2. Basic knowledge of image processing techniques: edge detection, corner detection, segmentation, feature detection and matching.
3. Basics of video processing, motion computation, stereo scene understanding, and 3D vision and geometry
4. Ability to implement basic vision algorithms in Python using open source matrix processing tools: e.g., numpy.
5. Ability to apply the appropriate techniques to a problem, write a project report, and present the results in a video overview.

Course Structure

This course is broken modules, decomposed into one module per week released a couple weeks in advance of the expected watch time. Modules will typically be

broken into 20–40 minute blocks, containing a mostly slide-based lectures with the occasional live coding exercise, with which you are encouraged to follow along.

Interactive demos are good opportunities to follow along and test your understanding of the material. In the lectures where the demos exist, I will encourage you to pause the video and try some of the exercises on your own before I walk through them.

Prerequisite Knowledge for the Course

- **Algorithms and Data Structures** As many (if not all) of the computer vision techniques we discuss will be algorithmic in nature, understanding how to implement simple algorithms and understand object-oriented code will be an important skill in this course. Big-O notation will also come up in the course, though students will not be expected to dive deep into computational complexity.
- **Linear Algebra** This is the most important prerequisite, as many of the tools developed and derived in this course involve matrix manipulation. Understanding eigenvalues and eigenvectors will also be key skills. Other matrix factorization techniques will also be used in this course, though prior knowledge (e.g., of singular value decomposition) is not required.
- **Probability and Statistics for Engineers and Scientists** Probability will come up on occasion throughout the course, so a working knowledge of fundamentals (e.g., properties of Gaussian distributions, Bayes Rule) will be useful, though not strictly required. Specific concepts associated with hypothesis testing (and the Chi-squared test) will be useful but are also covered in the course.

In addition, students are expected to have some familiarity with the Python programming language (including `numpy`) in which all programming projects and assignments will be done.

Assignments & Grading

The grading in this course will come from three main components:

- 5 programming assignments, P1–P5 (50%; 10% each)
- 5 concept-building assignments, C1–C5 (20%; 5% each, since lowest grade dropped)
- 1 final project, FP (30%)

All assignments will be turned in via Canvas. Assignments will be assigned alongside that week's lecture content. You will have roughly 2.5 weeks to complete each assignment; see the specifics due dates for each on the course calendar.

Programming Assignments

In each programming assignment, students will be expected to implement some of the algorithms we will discuss during class. Assignments will be given in the form of Jupyter Notebooks and will often include some partially written code for students to complete.

Students will be expected to write up a report for each programming assignment, preferably in LaTeX, but MS Word is acceptable. As this is senior level course, some creativity on the solution will occasionally be required; some of the problems will have open-ended prompts and students should expect to explore the parameter space of the algorithms they implement and report on their findings. Students may collaborate on the programming assignments in small groups (no more than 2 or 3 students) but solutions must be written up independently.

Concept Building Assignments

To supplement the programming assignments, which test mostly practical knowledge, the concept building assignments will be shorter and ask two or three conceptual questions related to the course material. We prefer that you type your solutions (e.g., via LaTeX) but neatly handwritten and scanned submissions will also be accepted. Students may collaborate on the concept building assignments in small groups (no more than 2 or 3 students) but solutions must be written up independently.

Final Project

While the details for the final project are still TBD (and will be provided in the middle of the term), the final project will be an open-ended project, the aim of which will be for you to take initiative and explore some topic or application that goes beyond what we have done in the programming assignments. The final project proposal will be due about a month in advance of the project and be an opportunity to get feedback from me before you start working on it in more depth.

Participation & Lectures

Lectures will regularly include *breakout sessions* during which students will be expected to work collaboratively to tackle some small problem related to the lecture. However, I have chosen *not* to include participation or attendance as a part

of the overall grade. You are still strongly encouraged to come to lecture, as the breakout sessions are an integral part of the course experience.

Lateness Policy

Every assignment (except the Final Project) can be turned in 2 days late without penalty. I ask that you make a best effort to turn in the assignments on time and use the additional 2 days if you get stuck or if "life happens" (which it seems to do a lot these days). Any additional late days will result in a 10%/day penalty for all assignments, up to a maximum of 1 week after the original deadline, after which the assignment will not be accepted.

I should note that you are expected to be able to complete the assignments in two weeks. The lateness policy is to avoid last minute emails. You should aim to turn in the assignments by Friday and use the extra two days only if you need it.

If you feel you need an extension in addition to the 2 free late days, ask me *before the original deadline*, as I would like to avoid emails asking for extensions with only hours before late penalties start to accumulate.[1] You do not need to ask for permission on the "free" late days; they are automatically applied.

Assignments are (roughly) given on Mondays, due the Wednesday 2 weeks later, and so can be turned in until Friday at 11:59PM without penalty. If you turn in the assignments on time (on Wednesday), you will get a 2% bonus on that assignment's grade (up to the maximum).

Course Resources

Optional Textbooks & Readings

I have made my best effort to make the course as self-contained as possible. As such, there are no *required* textbooks for this course. If interested, you might consider purchasing *Concise Computer Vision: An Introduction into Theory and Algorithms* by Reinhard Klette, which I will occasionally refer to throughout the course for further reading. The book has nice illustrations of many of the concepts and has relatively easy-to-understand prose. I will also refer to *Computer Vision: Algorithms and Applications*, by Richard Szeliski (available as a PDF online for free), though that text is rather more advanced and can be difficult to follow if you do not

[1] I have tried my best to come up with a late policy that is as fair and flexible as possible. Deadlines are largely there for your protection: in the past, I granted some assignments extensions early on in the class to students that were then behind for much of the rest of the term. Please try to start on the assignments early and ask for extensions only if you think you really need them, so that this does not happen to you.

already know the concepts. Additional supplemental readings, usually in the form of papers will be included either as links in the lecture slides or as supplemental PDFs.

Lecture Slides, Resources and Recordings available through Canvas

All materials I use for the pre-recorded lectures, plus I will be providing lecture slides as PDF documents after each lecture via Blackboard. In addition to the lecture slides, many of the lectures in the first two-thirds of the course will have breakout sessions for you to code or follow along with during lecture; these will also be uploaded to the Blackboard in advance. Lecture materials for the week will be uploaded on Tuesdays.

Detailed Course Schedule

The course is largely broken into three Units:

1. **Images** (Modules 1–5) In which we will discuss the fundamental mathematical tools used to process images and identify common *features* between multiple images
2. **Structure** (Modules 6–10) In which we will study how we might understand and reconstruct the 3D world.
3. **Modern Applications** (Modules 11–14) Devoted to discussing state-of-the art research tools and applications, including Place Detection, Simultaneous Localization and Mapping, and Convolutional Neural Networks (and applications).

Dates for each module below are when you are roughly expected to keep up with the pace of the course. My intention is to release each module at least a week in advance so that you can a bit ahead if you would like.

MOD	DATE	TOPICS	PROJECTS & QUIZZES
1	1/21	Course Introduction Course Logistics Image Fundamentals Linear Image Filtering	P1 Out 1/21
2	1/27	Resampling & Image Pyramids Fourier Transforms	C1 Out 1/27
3	2/03	Feature & Corner Detection Feature Invariance	P1 Due 2/05 P2 Out 2/03
4	2/10	Image Transformations Feature Descriptors	C1 Due 2/12 C2 Out 2/10
5	2/17	Image Alignment RANSAC	P2 Due 2/19 P3 Out 2/17

6	2/24	Camera Models Blender Tutorial	C2 Due 2/26 C3 Out 2/24 Final Description Out
7	3/03	Panoramas Single-View Modeling Nonideal Cameras	P3 Due 3/05 P4 Out 3/03
8	3/17	Two-view Geometry Stereo Vision	C3 Due 3/19 C4 Out 3/17
9	3/24	Light Photometric Stereo	P4 Due 3/26 P5 Out 3/24
10	3/31	Multi-view Stereo Applications Structure from Motion	C4 Due 4/02 C5 Out 3/31
11	4/07	Deep Learning Basics Neural Networks Convolutional Neural Networks	P5 Due 4/09
12	4/14	Image Classification Bag of Words Methods	
13	4/21	Simultaneous Localization & Mapping (SLAM) Modern SLAM	C5 Due 4/23
14	4/28	Ethical and Practical Considerations Generative Image Methods Course Summary	

Inclusion & Integrity

I stand by Mason's commitment to diversity and inclusion and hope to foster an inclusive environment in which all feel welcome in my class.

> *True diversity is defined not only as differences in individual backgrounds, personal identities, intellectual approaches, and demographics; it is also the removal of barriers and the creation of space that allow individuals to fully engage in the life of the university.*

Every student in this class is exactly where they belong and it is my honor to welcome each of you to join us in learning throughout this semester. Every student in this class, regardless of background, sex, gender, race, ethnicity, class, political affiliation, physical or mental ability, veteran status, nationality, or any other identity category, is an equal member of our class.

You have the right to be called by whatever name you wish, to be referred to by whatever pronoun you identify, and to adjust these at any point. If you feel uncomfortable in any aspect of our instruction that results in any barrier to your inclusion in this course, please contact me (your professor) directly.

Honor Code Statement

The GMU Honor Code is in effect at all times. In addition, the CS Department has further honor code policies regarding programming projects, which are detailed here. Any deviation from the GMU or the CS department Honor Code is considered an Honor Code violation.

Disability Accommodation

If you have a documented learning disability or other condition which may affect academic performance, make sure this documentation is on file with the Office of Disability Services and then discuss with the professor about accommodations. Submitting the paperwork at the deadline for a project or quiz is far too late! Even if you don't know if you plan on utilizing the accommodations ahead of time, it's in your best interest to prepare them ahead of time.

Mental Wellness

Graduate School can be a stressful environment and the realities of remote work can amplify these stresses. My “door” is always open; if you are struggling with the course work or would like someone to talk to, feel free to reach out to me. GMU also provides many mental health resources that I encourage you to look at.

Sexual Harassment and Interpersonal Violence

As a faculty member and designated "Responsible Employee," I am required to report all disclosures of sexual assault, interpersonal violence, and stalking to Mason's Title IX Coordinator per university policy 1412. If you wish to speak with someone confidentially, please contact the Student Support and Advocacy Center (703-380-1434), Counseling and Psychological Services (703-993-2380), Student Health Services, or Mason's Title IX Coordinator (703-993-8730; cde@gmu.edu).

Other Course Policies from GMU’s official guidelines

These four policies affect students in all courses at George Mason University. This Course Policy Addendum must be made available to students in all courses (see Catalog Policy AP.2.5).

Academic Standards

Academic Standards exist to promote authentic scholarship, support the institution’s goal of maintaining high standards of academic excellence, and encourage continued ethical behavior of faculty and students to cultivate an

educational community which values integrity and produces graduates who carry this commitment forward into professional practice.

As members of the George Mason University community, we are committed to fostering an environment of trust, respect, and scholarly excellence. Our academic standards are the foundation of this commitment, guiding our behavior and interactions within this academic community. The practices for implementing these standards adapt to modern practices, disciplinary contexts, and technological advancements. Our standards are embodied in our courses, policies, and scholarship, and are upheld in the following principles:

- *Honesty:* Providing accurate information in all academic endeavors, including communications, assignments, and examinations.
- *Acknowledgement:* Giving proper credit for all contributions to one's work. This involves the use of accurate citations and references for any ideas, words, or materials created by others in the style appropriate to the discipline. It also includes acknowledging shared authorship in group projects, co-authored pieces, and project reports.
- *Uniqueness of Work:* Ensuring that all submitted work is the result of one's own effort and is original, including free from self-plagiarism. This principle extends to written assignments, code, presentations, exams, and all other forms of academic work.

Violations of these standards—including but not limited to plagiarism, fabrication, and cheating—are taken seriously and will be addressed in accordance with university policies. The process for reporting, investigating, and adjudicating violations is outlined in the university's procedures. Consequences of violations may include academic sanctions, disciplinary actions, and other measures necessary to uphold the integrity of our academic community.

The principles outlined in these academic standards reflect our collective commitment to upholding the highest standards of honesty, acknowledgement, and uniqueness of work. By adhering to these principles, we ensure the continued excellence and integrity of George Mason University's academic community.

Student responsibility: Students are responsible for understanding how these general expectations regarding academic standards apply to each course, assignment, or exam they participate in; students should ask their instructor for clarification on any aspect that is not clear to them.

Accommodations for Students with Disabilities

Disability Services at George Mason University is committed to upholding the letter and spirit of the laws that ensure equal treatment of people with disabilities. Under the administration of University Life, Disability Services implements and coordinates reasonable accommodations and disability-related services that afford equal access to university programs and activities. Students can begin the registration process with Disability Services at any time during their enrollment at George Mason University. If you are seeking accommodations, please visit https://ds.gmu.edu/ for detailed information about the Disability Services registration process. Disability Services is located in Student Union Building I (SUB I), Suite 2500. Email: ods@gmu.edu. Phone: (703) 993-2474.

Student responsibility: Students are responsible for registering with Disability Services and communicating about their approved accommodations with their instructor *in advance* of any relevant class meeting, assignment, or exam.

FERPA and Use of GMU Email Addresses for Course Communication

The Family Educational Rights and Privacy Act (FERPA) governs the disclosure of education records for eligible students and is an essential aspect of any course. *Students must use their GMU email account* to receive important University information, including communications related to this class. Instructors will not respond to messages sent from or send messages regarding course content to a non-GMU email address.

Student responsibility: Students are responsible for checking their GMU email regularly for course-related information, and/or ensuring that GMU email messages are forwarded to an account they do check.

Title IX Resources and Required Reporting

As a part of George Mason University's commitment to providing a safe and non-discriminatory learning, living, and working environment for all members of the University community, the University does not discriminate on the basis of sex or gender in any of its education or employment programs and activities. Accordingly, **all non-confidential employees, including your faculty instructor, have a legal requirement to report to the Title IX Coordinator, all relevant details obtained directly or indirectly about any incident of Prohibited Conduct** (such as sexual harassment, sexual assault, gender-based stalking, dating/domestic violence). Upon notifying the Title IX Coordinator of possible Prohibited Conduct, the Title IX Coordinator will assess the report and determine if outreach is required. If outreach is required, the individual the report is about (the "Complainant") will

receive a communication, likely in the form of an email, offering that person the option to meet with a representative of the Title IX office.

For more information about non-confidential employees, resources, and Prohibited Conduct, please see University Policy 1202: Sexual and Gender-Based Misconduct and Other Forms of Interpersonal Violence. Questions regarding Title IX can be directed to the Title IX Coordinator via email to TitleIX@gmu.edu, by phone at 703-993-8730, or in person on the Fairfax campus in Aquia 373.

Student opportunity: If you prefer to speak to someone *confidentially*, please contact one of Mason's confidential employees in Student Support and Advocacy (SSAC), Counseling and Psychological Services (CAPS), Student Health Services (SHS), and/or the Office of the University Ombudsperson.