Computer Vision

CS 682 | Fall 2021 | Professor Gregory J. Stein

Course Information

CS 682
3 Credits

Instructor Information

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

Note: this syllabus may change!

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 introductory graduate 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 read and understand research papers in the field.

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

Course Structure

Ensuring that the students are engaging with the material and with one another can be difficult. As such, most lectures will be broken into subparts: each class will consist of two ~50 minute lectures, each followed by a 20 minute breakout session during which students will collaborate to solve short programming assignments related to the course material. Remaining class time will consist of a shorter lecture covering a related topic or research area.

Prerequisites

- Algorithms and Data Structures
- Artificial Intelligence

- Linear Algebra and Calculus

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 (50%)
- 8 take home "quizzes" (24%)
- 1 final project (26%)

All assignments will be turned in via Blackboard.

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 in LaTeX. As this is an introductory graduate course, some creativity on the solutions is expected; 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 3 or 4 students) but solutions must be written up independently.

Quizzes

To supplement the programming assignments, which test mostly practical knowledge, the quizzes will be shorter and ask two or three conceptual questions related to the course material. Solutions must be typed using LaTeX. Students may collaborate on the quizzes in small groups (no more than 3 or 4 students) but solutions must be written up independently.

There will be a total of 8 take home quizzes assigned during the semester. Quizzes **cannot be turned in late** without prior discussion with the instructor. At the end of the term the lowest two quiz scores (including missed quizzes) will be dropped.

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, due to COVID-19, I have chosen *not* to include participation as a part of the overall grade. You are still strongly encouraged to come to lecture, as the breakout sessions are ana integral part of the course experience.

Course Resources

Textbook & Readings

Textbook readings will be assigned from Computer Vision: Algorithms and Applications, by Richard Szeliski (available as a PDF online for free). Additional supplemental readings, usually in the form of papers will be included either as links in the lecture slides or as PDFs uploaded to Blackboard or (if it exists) the course Piazza.

Lecture Slides

I will be providing lecture slides as PDF documents after each lecture via Blackboard or Piazza.

Discussion

TBD: we will be using either Piazza or Blackboard for class discussion.

Detailed Course Schedule

The following schedule is *tentative* and subject to change.

The course is largely broken into three Units:

1. **Images** (Lectures 1–5) In which we will discuss the fundamental mathematical tools used to process images and identify common *features* between multiple images
2. **Structure** (Lectures 6–10) In which we will study how we might understand and reconstruct the 3D world.
3. **Modern Applications** (Lectures 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).

LEC	TOPICS	PROJECTS	QUIZES
1	Introduction	P1 Out	

	Fundamentals & Image Filtering		
2	Fourier Transforms Resampling & Image Pyramids		Q01 Out
3	Feature & Corner Detection Feature Invariance The Hough Transform (optional)	P1 Due P2 Out	Q01 Due Q02 Out
4	Image Transformations Feature Descriptors		Q02 Due
5	Image Alignment RANSAC	P2 Due P3 Out	Q03 Out
6	Camera Models Blender Tutorial Camera Models cont.		Q03 Due Q04 Out
7	Panoramas Single-view Modeling	P3 Due P4 Out	Q04 Due
8	Two-view Geometry Stereo Vision		Q05 Out
9	Light Photometric Stereo	P4 Due P5 Out	Q05 Due
10	Multi-View Stereo Applications Structure from Motion		Q06 Out
11	Neural Networks Convolutional Neural Networks	P5 Due Proposal Due	Q06 Due Q07 Out
12	Image Classification Bag of Words Methods		Q07 Due Q08 Out
13	Simultaneous Localization and Mapping Modern SLAM (Maybe: Optical Flow)		Q08 Due
14	Generative Image Techniques (Maybe: Image Segmentation) Course Summary	Final Due	

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 our 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 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. 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).

Privacy and Email

Students must use their Masonlive email account to receive important University information, including communications related to this class. I cannot respond to messages sent from or send messages to a non-Mason email address.

To protect your privacy, I also cannot list your Masonlive email address on any public forum or provide it to any other students. You may, of course, give your email address to any other students.

Video recordings of class meetings that are shared only with the instructors and students officially enrolled in a class do not violate FERPA or any other privacy expectation.

All course materials posted to Blackboard or other course site are private; by federal law, any materials that identify specific students (via their name, voice, or image) must not be shared with anyone not enrolled in this class.