SPRING 2023

EC ENGR 188: Special Course in Electrical Engineering

The Foundations of Computer Vision

Proposed New Course: EC ENGR 149

Course Description

This course covers the foundations of computer vision from both a theoretical and practical perspective. A particular emphasis is on “Classical” Computer Vision, which after taking this course, should be seen as complementary to deep learning. In fact, the pioneers of classical computer vision are the same community that developed convolutional neural networks (CNNs).

The course should be relevant for CS majors specializing in AI, ECE majors specializing in cyberphysical systems and information engineering, ME majors specializing in robotics, Statistics majors specializing in machine learning, CogSci and Psych majors specializing in perception, and Math majors looking for applications.

Instructors: Prof. Achuta Kadambi

TA: Zhen Wang, Howard Zhang

Grader: Ananay Garg, Christopher Hwang

Email Policy: The course has an email account: <vmgteaching@gmail.com> All class related concerns should go there. But before you email, please see the lecture 1 slides as they may answer some questions and shed light about the email policy. In particular, regrades are handled by the grader and email subject lines should say regrade so it goes to the grader’s attention. In general, if you do not agree with the graders assessment, you may escalate to the TA, and if that fails, escalate to the Professor during in-person meeting times.

Lecture Times: 8:05am - 9:50am

Lecture Location: KAPLAN 135 KNSY PV 1240B

Final Exam Time: Friday, June 16, 2023 8:00 AM - 11:00 AM

Prof. Kadambi Office hours: Wed 10-10:50am, go to Kadambi's office Engr IV 56-147J. If the office door is closed, then we had more people than could fit in my office, and we would have gone into the nearby Maxwell Room, and a post-it note will be on my office door. Wed 4/19 OH is rescheduled to Friday 4/21 at 4:30 - 5:20. Same location. This is due to a Grant Proposal Deadline.

TA Office Hours: Friday 11am-12pm

Piazza Link: https://piazza.com/ucla/spring2023/ecengr188

Prerequisites: Familiarity with Python programming, and command of Linear Algebra, Calculus, Signal Processing, and Probability. Familiarity with LaTeX (or willingness to spend 1-2 hours learning it). Familiarity with deep learning. If you have no exposure to deep learning, you will be asked to watch an additional videolecture (deep learning I) that replaces the Memorial Day lecture.

Grade Structure: The grading structure is different for undergraduate and graduate students. Graduate students can request the undergraduate grading scheme up until week 4.

Undergraduate

Lecture Quizzes	20%
Problem Sets	60%
Final Exam	20%
Participation (bonus: +1%)	

Graduate

Lecture Quizzes	20%
Final Project	80%
Participation (bonus: +1%)	

Grading thresholds are used as:

A+	at discretion	A	93-max	A-	90-93
B+	87-90	B	83-87	B-	80-83
C+	77-79	C	73-76	C-	70-72

Course Textbook: There is no required textbook, although the following textbooks can be seen as a helpful resource. These are all freely available.

- Computer Vision: Algorithms and Applications, 2nd ed.
- An Invitation to 3-D Vision
- Computational Imaging

Course Materials: The practical nature of the computer vision world is undergoing a revolution with GPT. For some of the problem sets it may be helpful to use a Github copilot subscription (which is powered by GPT). There is a 60 day free trial at the time of writing, and the cost of a monthly subscription over the 10 weeks is nominal and less than a textbook. Edit: a student has pointed out that this is free for students.

Details on Problem Sets / Homework Assignments: There will be four problem sets, released on the following schedule. Most problem sets are allotted 2 weeks for completion, with PSET 3 having an extra week as it covers more material. You will be allotted 3 late days to use for assignments as you choose. Beyond this 3 day grace, late assignments will not be accepted unless we have a medical note.

PSET 1 OUT: 4/7 DUE: 4/21
- A LaTeX file will be provided via overleaf: https://www.overleaf.com/read/djyhjpxjtftf;
- The accompanying Jupyter notebook and data needed for this problem set can be found in BruinLearn under Module Week 1;
- Edit your copy and turn in the PDF on Gradescope and Jupyter notebook.

PSET 2 OUT: 4/21 DUE: 5/5
- A LaTeX file will be provided via overleaf: TBD
- Edit your copy and turn in the PDF on Gradescope and Jupyter notebook.

PSET 3 OUT: 5/5 DUE: 5/26
- A LaTeX file will be provided via overleaf
- Edit your copy and turn in the PDF on Gradescope and Jupyter notebook.

PSET 4 OUT: 5/26 DUE: 6/9
- A LaTeX file will be provided via overleaf
- Edit your copy and turn in the PDF on Gradescope and Jupyter notebook.

Details on Final Project: The default final project will be curated and involve each of the three R's. The goal is to "mimic" one function of a Tesla Autopilot or similar driving system, and in particular, obtain the 3D semantic labeling of a scene. The Three "R"'s will be applied in sequence. First, students will begin with images taken at different viewpoints that need to be properly represented. Features need to be obtained. From this, students will apply the second "R", which is reconstruction of a 3D scene. Finally, students will apply a third "R" which is recognition of scenes. The final project accumulates knowledge from the lecture in a sequential manner. Recommended checkpoints for code freeze are provided in the syllabus. The final project guidelines document will be updated with class progress and can be found at this link. The final project is to be done individually.

***Tentative* Topics are Below.**

Legend

Red:	**R**epresentation	(Draws a lot from Signal Processing)
Green:	**R**econstruction	(Draws a lot from Linear Algebra)
Blue:	**R**ecognition	(Draws a lot from Deep Learning)

Date	Lecture Topic	#
4/3	Introduction, Course Logistics, Introducing the 3 R's, Introducing our "First R"**: R**epresentation. What is an image?	1
4/5	Image Processing I. 2D Signals, Systems, and Properties	2
4/10	Image Processing II. 2D convolution,	3
4/12	2D Fourier Transform. Intro to Deconvolution and Noise,	4
4/17	Deconvolution and Blind Deconvolution. More Filters: Gaussian, Median, Bilateral, Sobel, Laplacian.	5
4/19	From Edges to Corners, Blobs, and Characteristic Scale	6
4/24	**R**epresenting Local Correspondence I (Descriptors + Detectors)	7
4/26	**[REMOTE LECTURE] R**epresenting Local Correspondence II (Matching) Recommend code freeze on **R**epresentation portion of the final project.	8
5/1	Introducing our "Second R": **R**econstruction. 2D and 3D transformations, Euclidean, affine	9
5/3	**R**econstruction on the Planar Domain I: 2D-2D transformations, projection, and the DLT algorithm.	10
5/8	**R**econstruction on the Planar Domain II: (Homographies, Perspective projections, projective/homographic transformations, homogeneous DLT)	11
5/10	**R**econstruction on the Planar Domain III: Combining SIFT + RANSAC + Homography.	12
5/15	**R**econstruction in 3D space. Camera Models Part I.	13

5/17	Reconstruction in 3D space. Camera Models Part II. Introduction to Two-view Geometry	14
5/22	Stereo	15
5/24	Stereo Continued	16
5/29	**No Class - Memorial Day Holiday.** Recommend code freeze on Reconstruction portion of the final project. Videolecture for Deep learning I: Mandatory for students new to deep learning. Optional if you have taken a deep learning class (anywhere)	17
5/31	Deep Learning II: Self-supervised Learning, Contrastive Learning	18
6/5	Deep Learning III: Vision and Language	19
6/7	Deep Learning IV: Revisiting Representation, Reconstruction, and Recognition in a joint algorithmic framework. Recommend code freeze on Recognition portion of the final project.	20
6/16	**Undergraduate Final Exam** (Room TBD) Friday, June 16, 2023 8:00 AM - 11:00 AM	
6/16	**Graduate Final Project Writeup and Code Due** (No Late Days for Final Project)	

Follow-Up or Concurrent Courses: Students who enjoy this course, may consider EE.114 (Speech and Image Processing), EE.146 (Introduction to Machine Learning), EE.211 (Digital Image Processing), CS.269 (Deformable Models), CS.275 (Artificial Life), and EE.239 (Computational Imaging) as examples of synergistic courses at UCLA.