Syllabus

Faculty - Click to Adopt a Textbook

Welcome and Getting Started

The syllabus and a course checklist are provided.

Meet the Faculty

Please review the syllabus and download the checklist for your own use throughout the course.

About this Course

Syllabus

Textbook Information for Students

Add your syllabus and a checklist/course schedule as files below by clicking the + and Choosing Upload from Computer.

Course Expectations

Then delete these instructions.

Grading Policy

Minimum Technical Requirements

University Policy and Support

Technology Support Information

Meet Your Classmates

Getting Started Quiz

Basic Information

Meeting Location: T/R 09:40-10:55 **in-person (Gavett Hall Room 202)**

Instructor: Prof. Jeff Luo (https://www.rochester.edu/u/jluo/)

Instructor's Office Hours: T/R 2-3 pm in Wegmans Hall 3101

TAs and Office Hours:

- Ziyun Zeng (https://scholar.google.com/citations?user=b2DIlscAAAAJ), *Graduate TA* (Fridays 1-3 pm in Wegmans Hall 3504)
- İsmail Can Yağmur (https://github.com/canyagmur), *Graduate TA* (Wednesdays 1:20-3:20 pm in Wegmans Hall 3002)

Information Flow: We will use **Blackboard** to make announcements, post lecture notes/videos and assignments, and submit student work (**never submit any work via email**).

Credit Hours: This is a four-credit-hour course. The credit hours are met by weekly lectures (150 minutes), office hours, homework assignments, and the final course project.

Description

Fundamentals of computer vision, including image formation, elements of human vision, low-level image processing, and pattern recognition techniques. Advanced topics include modern visual features, graphical models, model-based and data-driven approaches, and contextual inference, as well as illustrative examples of successes and challenges in applications. CSC 449, a graduate-level course, typically requires additional readings and assignments. Based on past classes, CSC 249, an undergraduate-level course, has the same requirements with a 5% bonus for completing the same requirements as 449.

Rough Topic Outline

- **Image Formation and Multiple View Geometry:** camera models, light and color, camera calibration, image alignment, stereo vision
- **Features and Filters:** image transforms, linear filtering, edges, corners, feature extraction, feature matching, optical flow
- **Segmentation and Fitting:** segmentation by clustering, interactive segmentation, robust fitting, and RANSAC, Hough transform
- **Recognition:** feature learning, convolutional and recurrent neural networks, advanced CNN architectures, attention mechanisms, feature visualization, image classification, object detection, instance segmentation, face recognition, shape analysis, video analysis, vision-language models
- **Description and Generation:** image captioning, VQA, generative adversarial networks, diffusion models, image and video generation

Course Objective

The course is designed as an upper-level elective for Computer Science undergraduate students and an AI-area breadth course for graduate students. Students will develop a strong understanding of formulating and solving problems in computer vision.

Prerequisites

Basic knowledge of probability, linear algebra (MTH 165 strongly recommended); data structures, algorithms; programming experience. For assignments and projects, the most common programming language is Python or MATLAB.

Hardware/Software Requirements

You need a computer and a reliable network to access online lectures and course materials. Besides, many of the course assignments require you to have access to a powerful GPU system. This can usually be resolved by:

1. Using our provided CSUG account to ssh to shared GPU resources, or
2. Owning an NVIDIA GPU system with at least 6GB GPU memory, or
3. Renting a cloud GPU via cloud service providers, e.g., Google Cloud.

By taking this course, you confirm that you have met the system requirements.

Textbook and Other Resources

Machine/Computer Vision has been under rapid development over the past few years. There is no perfect textbook that covers everything needed for an introductory computer vision course in today's context. However, we will use the following textbook, which is widely used in many other computer vision courses.

- Computer Vision: A Modern Approach, **2nd Edition**, Forsyth and Ponce, Pearson, 2012

In addition, I recommend the following textbooks as supplementary reading.

- Computer Vision: Algorithms and Applications, Szeliski, Springer, 2011 [link (http://szeliski.org/Book/)]
- Deep Learning, Goodfellow, Bengio, and Courville, The MIT Press, 2016 [link (https://www.deeplearningbook.org/)]
- Multiple View Geometry in Computer Vision, 2nd Edition, Hartley and Zisserman, Cambridge University Press, 2004

Schedule

Course Overview - Start Here

This is a **tentative** schedule of the course and is subject to **modifications over the term**.

Lectures and exams will be given **in-person only**.

Dates	Thursday (R)	Tuesday (T)	Chapter	Project
1/20, 22	Introduction	Image Formation / Camera Models	1,2	
1/27, 29	Color Vision	Low-level Image Processing	3,4,5	
2/3, 5	Texture	Review of Probability	6, notes	
2/10, 12	Pattern Recognition Techniques I	Pattern Recognition Techniques II	15, notes	
2/17, 19	Modern Visual Features I	Modern Visual Features II	notes	**Teaming Starts**
2/24, 26	CNN and Feature Visualization I	CNN and Feature Visualization II	notes	

Course Overview - Start Here

3/3, 5	Image Segmentation	Fitting	9,10	
3/10, 12	**Spring Break (3/7-3/15)**	**Spring Break (3/7-3/15)**		Spring break 3/7-3/15
3/17, 19	Face and People Recognition I	Face and People Recognition II (3/14)	notes	**Proposal Due 3/19**
3/24, 26	**Midterm Exam**	Deep Learning and CNN		**Midterm 3/24**
3/31, 4/2	Deep Learning and CNN II	*Special Lecture: Vision + Language*	notes	
4/7, 9	*Special Lecture: Vision + Language II*	Deformable Shape Models	notes	
4/14, 16	Video Analysis	GAN and Image Synthesis	notes	
4/21, 23	*Special Lecture: Medical Image Computing*	*Special Lecture: Text-to-Visual Generation*	notes	

4/28, 30	Guest Lecture: Dr. Xiang Li, Harvard/MGH, **Medical Imaging and chatGPT**	**Project Presentation I (4/30)**		reading period (5/2-5/4)
5/5, 7	**Project Presentation II (5/5)**	**Project Presentation III (5/7)**		final exam (5/5-5/10) **Report Due 5/9**
5/12, 14	No Class: Term Finished (5/10)			**Grades due 5/14**

Grading (total 102%)

- **midterm exam**: 30%
- **homework assignments**: 35% (including a small programming project worth 10%). Each (worth 5%) will contain a mix of analytical questions and programming questions. They will be done independently by each student. **All due dates are according to what is posted in Blackboard** (not in the lecture notes, which could be outdated).
- **research project & presentation**: 35% (proposal counts 5%, presentation counts 10%, report counts 20%). Students will form teams of up to two members to implement a research project. The project will generate a minimum 4-page (excluding references) double-column report and a presentation at the end of the term. The expectation for the final project is something "new", e.g., (1) solving a

new problem, (2) applying an existing algorithm to new data, or (3) applying a new algorithm (or a modified version of an existing algorithm) to the same data.
- new findings from a comparative study of using different algorithms for the same problem
- **class participation and effort**: 2%

Students sometimes ask about the nature of the midterm exam. The 10 questions in the midterm exam will test students' understanding of major concepts and procedures. There is no need to memorize complicated equations (except the fundamental ones) or make complicated calculations (you can score 100% even if you cannot do multiplication/division on paper - just leave the right numbers in the right places). Each question should not take more than 5 minutes if you know the concepts and procedures. A pencil is all you need. No other device (**including VR lens**) is allowed. For the few with special permission to take the exam via Zoom, you should print the exam on paper and write on the printed paper. There is sufficient space under each question to write down the correct answer. In fact, if you need more space than given, you are most likely on the wrong track. You can use the backside of the paper, but you must draw an arrow on the front to show where it is. Some questions require you to draw illustrations (in such cases, do not use words to describe what you can draw). I never gave practice questions in my 20+ years of teaching. The exam is a means rather than the end. The best preparation is to go over the lecture notes and make sure you understand, again, the major concepts and procedures. So what are they? Well, I always mention them more than once in the classes and/or lecture notes.

Late Work Policy: Deduct 10% of the total score of that assignment for each day past the deadline for each assignment. For example, if you submit HW1 one day past the deadline, then the max score you will get is 90% of that assignment. We will not accept/grade late submissions after one week past the deadline. All assignments should be submitted on Blackboard.

Regrading: Any questions about the grading of a piece of work must be raised within one week of the date that the work was returned. In other words, if you do not pick up your work in a timely fashion, you may forfeit your right to question the grading of your work.

Policy on Incomplete: A student will receive an Incomplete if s/he misses either the final project or any two or more of the assignments.

Other Policies

Code of Conduct: Check here for classroom "etiquette" expectations.

Academic Honesty: All assignments and activities associated with this course must be performed in accordance with the University of Rochester's Academic Honesty Policy. More information is available here (http://www.rochester.edu/college/honesty/). You are expected to read, understand, and follow the policy. Additionally:

- *Homework Collaboration:* In general, homework is to be completed independently, and any forms of collaboration are strongly discouraged. You may discuss homework problems with others (in this case, you are required to report the **names** of those with whom you discussed an assignment in your submission; failing to do so is a violation of academic honesty), but you must **not retain** written notes from your conversations with other students, or share data via computer files to be used in completing your homework. Your written work must be completed without reference to such notes, with the exception of class and recitation notes, which may be retained in written form.
- *Online Materials:* You may look for the documentation, manuals, and tutorials for a programming language or library. You may study a systematic introduction and tutorial for a particular research topic or problem space. In all cases, you should cite the information sources. However, you **should not** intentionally search for homework solutions or solutions to any similar problems posted by past students or others online. You **should not** intentionally use ChatGPT or similar tools to write the key content of the course project report automatically. In general, the use of AI tools is not prohibited in the final course project as long as you do something creative. **When in doubt, you should disclose the use of such an agent/resource and/or consult the instructor for approval**. Failing to do so is a violation of academic honesty.
- *General Rule:* When in doubt, cite the source.

Posting homework and project solutions to public repositories on sites like GitHub is a violation of the College's Academic Honesty Policy, Section V.B.2 Giving Unauthorized Aid.

Disabilities: Please see me about your required accommodations as early as possible in the term. The University of Rochester respects and welcomes students of all backgrounds and abilities. In the event you encounter any barrier(s) to full participation in this course due to the impact of a disability, please contact the Office of Disability Resources. The access coordinators in the Office of Disability Resources