CSCI 5561 Computer Vision (Fall 2023)

11:15 A.M. - 12:15 P.M., MW (01/18/2023 - 05/01/2023), ME 108

Instructor: Volkan Isler (Office hour: Monday 12:15-1:15 on Zoom, or by appointment on Wed 1-2pm).
TAs: Jingfan Guo (guo00109@umn.edu. Office hour: Tue 1pm on Zoom / Fri 1pm at Shepherd Lab 237)

Please use canvas for course related communication. In particular, we will set up discussion boards for homeworks assignments and the project so the entire class can benefit from the answers and the discussion. If you need to reach the instructor or the TAs, please include [CSCI 5561] at the beginning of your subject line.

Course Description

The primary goal of this course is to introduce students to computer vision -- how to make computers make sense of images! Traditionally, computer vision has been studied in a bottom up fashion. At the lowest level, images are processed to obtain simple features (e.g. corners, edges). Then comes geometry in which these features are used to build 3D models of the environment. Next up is mid-level vision in which features are grouped into segments or other salient regions. At the highest level is the problem of object recognition which could be considered as a gateway to general intelligence. In the last decade, the field has been transforming at a mind boggling rate where learning based architectures have become increasingly dominant. It is impossible to cover all aspects of computer vision in depth in a single course. We will focus on fundamentals and let you explore state of the art methods through the course project The course will be organized around three modules: Geometry, Image Features and Deep Learning. After taking this course, you will:

- Learn about basic tools and techniques for designing (and analyzing) computer vision algorithms and architectures
- Become familiar with active research challenges in vision
- Gain hands-on experience in building vision systems

Important Note about Required Background -- Please read carefully!

While this is an introductory course in computer vision, anyone who is interested in taking it needs to understand that it is a graduate level computer science course!
Graduate level. Computer Science.

This means that you need a certain degree of mathematics and algorithms background in order to take this course. In particular:

Prerequisites: CSci 5511 and 5521 are required prerequisites. This requirement will be enforced. These prerequisites are there to ensure that you have the sufficient background in

- Linear algebra
- Probability and statistics
- Optimization, including algorithmic methods such as dynamic programming
- Programming: in-class exercises and projects are in Python

Topics Covered and Text Book

There is no required textbook for the course. We will take a hands-on approach and combine lectures with programming assignments. There are excellent resources online, and feel free to use them. For classical textbooks,

- For signal and image processing, there are excellent textbooks such as
  Oppenheim's Signal's and Systems
  Gonzales and Wood's Digital Image Processing are standard references.
  Another good resource is Hany Farid's manuscript.
  You might want to refer to these books when we cover concepts such as convolution, filtering and sampling (which are also the foundations of all convolutional neural network architectures)
- We will rely on Understanding Deep Learning by Simon Prince for the deep learning component
- Multiple View Geometry in Computer Vision by Hartley and Zisserman is the standard reference for geometry.
- For general computer vision texts, the two recommended textbooks are Computer Vision: A Modern Approach by Forsyth and Ponce (2nd Edition) , and Rick Szeliski's book.
- I will make my notes available, which will supplement Prof Park's slides.

Topics:

- Images as signals; Convolutions and the Fourier Transform; Sampling; Corners, Edges, Invariance, Image Features
- Projective and multi-view geometry; Stereo reconstruction
- Detection, classification, segmentation
- Selected topics: transformers, diffusion models, NeRF

Expectations

We will cover a relatively large range of topics which rely on various background material. As a result, at times, you might find that the current topic is quite challenging (or rather basic). For example, dynamic programming may be easy to grasp for a student who has taken advanced algorithms courses. The same student may find projective geometry quite challenging. Here are some tips that you might helpful:

- Do not rely on a single source to learn the material. If the topic in one book does not make sense, do not hesitate to review other books, or even basic texts such as your linear algebra textbook. If you need guidance, do not hesitate to talk to your instructor or the TA. There are excellent resources online.
- Learn from each other.
- Work regularly. As the list of topics is quite diverse, make sure that you review the material in a timely fashion.

Evaluation

Evaluation will be composed of the following components:

- Homework (60%): Up to five assignments. Important: The assignments will require writing (lots of!) python code.
- Course Project (40%): More about the project will be posted on the project page.

Grades: 93.0% or above yields an A, 90.0% A-, 86% = B+, 82% = B, 78% = B-, 74% = C+, 70% = C, 67% = C-, 63% = D+, 60% = D, and less than 60% yields an F.

Other grade related issues: We will give ample time for each homework assignment. It is critical that you start early. The assignments will be due Fridays. If you submit late, but before the following Monday, you get 75% of the grade. Afterwards, 0%. Questions about a specific grade should be raised within a week after the grade is given. Incompletes will, in general, not be given. Exceptions will be considered only when a provably serious family or personal emergency arises, proof is presented, and the student has already completed all but a small portion of the work.

Scholastic Misconduct

Scholastic misconduct is broadly defined as "any act that violates the right of another student in academic work or that involves misrepresentation of your own work. Scholastic dishonesty includes, (but is not necessarily limited to): cheating on assignments or examinations; plagiarizing, which means misrepresenting as your own work any part of work done by another; submitting the same paper, or substantially similar papers, to meet the requirements of more than one course without the approval and consent of all instructors concerned; depriving another student of necessary course materials; or interfering with another student's work." **Important:** Computer Vision libraries should make your life easy but should not do the homework for you! You are welcome to use standard matrix functions, or utilities to read and

write images, but when, e.g., the problem asks you to implement an edge detector, you should not use built-in edge detection functions. When in doubt, just ask. Using implementations off-the-web or elsewhere in homeworks yields an immediate F. **If you use code from anyone or anywhere, you must give reference. Otherwise, this constitutes cheating.** While it is allowed to discuss homework with other students at a high level, you should write your own code independently. Do NOT use other students' code (including students who have taken this class previously). An automatic plagiarism system will be run to determine the similarity of code.

A Note from DRC

The University of Minnesota is committed to providing equitable access to learning opportunities for all students. Disability Resource Center (DRC) is the campus office that 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, chronic health, sensory, or physical), please contact DRC at 612-626-1333 to arrange a confidential discussion regarding equitable access and reasonable accommodations.
- If you are registered with DRC and have a current letter requesting reasonable accommodations, we encourage you to contact Prof. Isler early in the semester to review how the accommodations will be applied in the course.

Additional information is available on the DRC website: https://diversity.umn.edu/disability/

As a student you may experience a range of issues that can cause barriers to learning, such as strained relationships, increased anxiety. alcohol/drug problems, feeling down, difficulty concentrating and/or lack of motivation. These mental health concerns or stressful events may lead to diminished academic performance or reduce your ability to participate in daily activities. University of Minnesota services are available to assist you with addressing these and other concerns you may be experiencing. You can learn more about the broad range of confidential mental health services available on campus via http://www.mentalhealth.umn.edu

If you require disability or other accommodations, please reach out to Prof. Isler to discuss them as soon as possible.