COMP 4060/7950: Computer Vision (Fall 2012)

Course info

Instructor: Yang Wang(E2-404 EITC, 204-474-9740, ywang@cs.umanitoba.ca)

Office hour: Mon 2:00pm-3:00pm, Wed 12:00pm-1:00pm, or by appointment

Class time/location: Mon/Wed 3:45pm-5:00pm, E2-310 EITC

Course webpage: `http://www.cs.umanitoba.ca/~ywang/courses/comp4060/fall12/`

Overview

This is an advanced course as an introduction to computer vision – the discipline of "teaching machines to see". The goal of computer vision is to make sense of photographs, videos, and other imagery. In this course, we will cover both standard and advanced topics in computer vision. In the first half of the course, we will focus on basic techniques and algorithms in computer vision. In the second half of the course, we will explore more advanced topics. The course will strive to provide a foundation for students to understand the computer vision literature and to do independent research.

A related course is "Digital Image Processing" offered by Prof.Neil Bruce. Digital image processing is an area very closely related to computer vision. The differences between these two areas are as follows. Image processing mainly deals with low-level processing tasks (e.g. image enhancement, image filtering, denoising, etc). Both the input and output of an image processing system are usually images. In contrast, computer vision focuses on more high-level image understanding tasks. The output of a computer vision system is usually a high-level description (e.g. "what and where are the objects in the image?", "where are people doing in the video?", etc), i.e. it aims to tell the story of an image. Interested students are encouraged to take both courses, since the overlap of the topics covered in these two courses is minimal.

Prerequisites

Basic knowledge of linear algebra, calculus, probability. For undergraduate students, this means you should have taken all the MATH/STAT courses required for the CS major (MATH1300, MATH1500, STAT1000) with good grades. Ideally, students should also have taken a course in probability (e.g. STAT 2400). Previous exposures to image processing and/or machine learning are desirable, but not required. Some assignments involve programming in Matlab. Previous experience with Matlab is not required, although students are expected to learn the language. Students should also have enough programming and

algorithm skills to finish a project. For the final projects, students can use any programming languages they choose.

Grading scheme

Grading will be based on the following components. Please check the course website for their due dates and detailed instructions.

- Homework: there will be 3 homework assignments throughout the course.
- Midterm: there will be an in-class midterm exam.
- Reading summary/review: there are weekly paper reading assignments in the second half of the course. Students are required to write a critical summary for each assigned paper.
- Presentation: graduate students taking the class will be required to present a paper in the class.
- Course project: both undergraduate and graduate students in this course will do a course project. Projects can be done individually or in a group of two. For project involving more than two people, you need to get the consent from the instructor. Students are required to submit a project proposal, a final report, and do a project presentation.

Undergraduate and graduate students will be graded separately. Grading for undergrads/grads will be based on the following table:

	undergrads	grads
assignments	50%	30%
midterm	10%	10%
reading summary/review	10%	10%
paper presentation	0%	10%
course project	30%	40%

Late policy

Students are allowed 4 grace days to use at their own discretion over the semester for the assignments. Late days are counted from the time an assignment is due and is rounded up to the next whole day. For example, if an assignment is due on Friday 3:30pm, and is submitted on Saturday 6:00pm, 2 grace days will be used.

IMPORTANT: Late assignments will not be accepted after the grace days are exhausted. No extension will be given to other compoments of the grading scheme except in unusual circumstances at the discretion of the instructor (e.g. medical reasons with doctor's note).

IMPORTANT: Grace days are for homework assignments only. They may not be used for other components (e.g. projects, reading summary, etc.) of the grading scheme.

List of topics

This schedule is tentative and subject to change. Please check the course website on a regular basis for updated information.

Week 1: Introduction

Week 2: Image filters, edges, textures

Week 3: Interest points, descriptors, fitting and registration

Week 4: Alignment, instance recognition, grouping and segmentation

Week 5: Introduction to machine learning, face recognition

Week 6: Image classification

Week 7: Object detection, pictorial structure

Week 8: Part-based models

Week 9: Looking at people; student presentation

Week 10: Context, location, 3D; student presentation

Week 11: Student presentation

Week 12: Vision and language; student presentation

Week 13: Large-scale and internet vision; student presentation

Week 14: Guest lectures

Textbooks

No required texts. The following books are useful as reference:

- D. Forsyth and J. Ponce, "Computer Vision: A Modern Approach", 2nd edition, Prentice Hall, 2011.
- R. Szeliski, "Computer Vision: Algorithms and Applications", Springer, 2010.
- R. Hartley and A. Zisserman, "Multiple View Geometry in Computer Vision", Cambridge University Press, 2004.
- C. Bishop, "Pattern Recognition and Machine Learning", Springer, 2007.

Academic Dishonesty

Students are reminded that acts of academic dishonesty are serious offenses and are subject to disciplinary action by the university. Students should familiarize themselves with the policy information of the Faculty of Science (`http://umanitoba.ca/faculties/science/undergrad/resources/webdisciplinedocuments.html`).