CS7670: Special Topics in Computer Vision, Fall 2011, Cornell University

Time: T/Th 2:55-4:10pm
Place: Upson 315
Homepage: http://www.cs.cornell.edu/courses/cs7670/2011fa/

Instructor: Noah Snavely
Email: snavely@cs.cornell.edu
Office: Upson 4157
Office hours: TBA

Course overview

Course description:
In the past decade computer vision has made incredible progress across the board, in geometry, recognition, image processing, computational photography, and other areas. In this graduate seminar in computer vision, we will survey and discuss state-of-the-art research in this quickly moving field, with a focus on 3D geometry, image matching and retrieval, use of the Internet to gather and annotate data, and scene understanding. This will draw on papers from both computer vision and computer graphics venues. We also hope to have a few guest lectures.

Requirements:
Students will be responsible for reading and reviewing a subset of papers discussed each week (submitting their reviews to CMS), participating in class discussions, giving one or two presentations to the class (depending on enrollment, and possibly done in pairs), and completing a research project. In preparing a presentation, students will also be expected to experiment with the methods they describe, either by trying out existing code online or implementing a simplified version of the method. Students will have leeway to decide what methods and papers interest them in presenting a topic, and should feel free to suggest other topics that interest them.

Presentations will be due **one week** before they are to be presented, meaning that you will need to read the papers, do any necessary experiments, and prepare your presentation at least one week before your presentation slot. This will allow time for the instructor to give you feedback on your presentation.

Prerequisites:
Students are expected to have a working knowledge of computer vision at the level of CS6670 (Computer Vision) or equivalent, and should be willing and able to understand and analyze recent conference papers in this area. If you are unsure if this course is right for you, please come talk to me. Perusing a few papers on the syllabus is a good way to gauge what kind of background is necessary. This course is expected to be interactive, relevant to the latest research, and (most of all), fun.

List of Suggested Topics

1. 3D geometry
    a. Structure from motion
    b. Multi-view stereo
    c. Interactive 3D reconstruction
    d. New sensors (e.g., Kinect)
2. Image matching and retrieval
    a. Feature detection and matching
    b. Bag-of-words models
    c. Large-scale image collections
    d. Machine learning methods for image matching
3. The Internet
    a. Location recognition
    b. Image editing
    c. Human-aided computer vision
4. Object recognition and scene understanding
    a. Context
    b. Attributes
    c. Materials
    d. Illumination
    e. Faces
5. Computational photography

If you are interested in a topic not listed here, please contact the instructor.

Important Dates (subject to change)

Thu Aug 25:	First class meeting
Tue Aug 30:	Topic preferences due (via CMS) – please specify your top three topics
Tue Sep 20:	Project proposal abstract due
Fri Oct 14:	Project update due
Week of Oct 17:	Individual project update meetings (by appointment)
Mon Dec 5:	Project presentations due
Fri Dec 9:	Final project papers due

Requirements

Students are expected to do the assigned readings, participate in class discussions, write two research paper reviews per week, and complete a final research project. In addition, each student will prepare approximately two presentations on a current topic in computer vision, drawing on a set of suggested research papers. The presentation will summarize the topic, as well as present the results of at least one of the student's own experiments with this topic, using existing code or the student's own implementation. **Note that presentations are due one week before you are scheduled to present**, in order to leave time for feedback from the instructor. Noah will kick things off by presenting in the first one or two lecture slots.

Paper reviews:
You will submit two paper reviews per week for the assigned papers. Since we'll usually be reading more than two papers per week, you can choose which papers to review. Reviews are due by 10pm the night before the paper will be covered in class (i.e., either Monday night or Wednesday night), and can be submitted via CMS. Each review should address the following:

- Give a brief summary of the paper in your own words, including the main contribution
- What are the strengths and weaknesses of the paper?
- How convincing are the experiments? What, if anything, is lacking?
- Describe one specific way in which the work could be extended.
- Additional comments, including points that were unclear

If you are presenting a paper in a given week, you need not submit reviews that week.

Paper presentations:
Each student (or possibly team) will give a presentation in class covering two papers on a topic selected from the list of suggested topics, or another topic of interest (with instructor approval). Each presentation should overview the papers and explain key technical details, and synthesize any underlying commonalities or highlight interesting distinctions. The talk should be well-organized and polished, sticking to about 45 minutes. Please run through it beforehand and check the time (a good rule of thumb is to stick to 35-40 slides total). Include these components in the presentation:

- Clear statement of the problem
- Why the problem is interesting, important, difficult
- Key technical ideas, how they work, main contributions, strengths and weaknesses
- Evaluation, summary of key experiments and data
- How the technical approaches agree/differ
- Open issues raised in the papers, likely extensions

Try to use applications to motivate the work when possible, and look for visual elements (images, videos) to put in the presentation. Check out the links on the class webpage, and also look at authors' webpages for supplementary materials. It's ok to grab a few slides from

conference talks etc. when available, but be sure to clearly cite the source on each slide that is not your own.

In addition, in your presentation you will present the results of some experimental evaluation of some main idea in a presented paper. The goal is to implement a distilled version of an essential technical idea in the paper, and show us some toy example of how this works in practice. For many papers, you will be able to find code or binaries provided by the authors online (see links on the course page alongside the papers). Your experiment should help us gain a more complete intuition about the work we are studying. You might:

- Experiment with different types of datasets
- Examine the method's sensitivity to relevant parameters
- Show a simplified example highlighting an strength or weakness of the method

The goal here is not to recreate published results or to build systems as described in the paper. Instead, you are looking to make a small illustrative demo that will let us more deeply understand the papers we are reading.

Spend some time playing with your implementation, and put thought into what would be an instructive toy example to show the class. The demo should allow us to learn something about the method, not just see it. If you had to implement something yourself, explain how you did it, and especially point out any details or choices that weren't straightforward, in case others in the class can build on your experience later when working on the final project. Be sure to explain the rationale for the outcomes, and conclude with a summary of the messages your example illustrates.

In addition to the presentation, make a simple webpage to outline the demo and include links to any existing code, data, etc. you have used. We'll point to that page for the rest of the class to reference.

Timetable for presenters:

- One week before your presentation is to be given (i.e., by Tuesday or Thursday): email slides to the instructor, and schedule a time to meet and go through a practice run.
- The week of your presentation: refine the slides based on input from the instructor, practice, and check time length of the talk.
- The day of the presentation: send final slides (and URL for experiment summary) to the instructor.

Projects:

Students will do a final research project (individually, or possibly in pairs with instructor approval). The final project should be related to a topic we cover in class, but should involve new research, either by extending a topic discussed, doing a thorough experimental evaluation of two or more techniques in a given area, or posing a new approach and performing appropriate experiments. The final projects will involve a proposal, a midterm report (possibly

involving a short presentation), and a final presentation to the class. The timeline will roughly follow the one under **Important Dates**.

Grading

Grading for this class will roughly follow these guidelines:

- 20% participation (includes attendance, in-class discussions, and paper reviews)
- 40% paper presentations (includes draft submitted one week prior, and in-class presentation)
- 40% final project (includes proposal, updates, presentation, final paper)

Special thanks to Kristin Grauman for the template for this document.