CS395T: Special Topics in Computer Vision, Spring 2010
Object Recognition

Meets: Wednesdays 3:30-6:30 pm
ACES 3.408
Unique # 54470

Instructor: Kristen Grauman
Email: grauman@cs.utexas.edu
Office: CSA 114

TA: Sudheendra Vijayanarasimhan
Email: svnaras@cs.utexas.edu
Office: CSA 106

Course overview:

Topics: This is a graduate seminar course in computer vision. We will survey and discuss current vision papers relating to object recognition, auto-annotation of images, and scene understanding. The goals of the course will be to understand current approaches to some important problems, to actively analyze their strengths and weaknesses, and to identify interesting open questions and possible directions for future research.

See the syllabus for an outline of the main topics we'll be covering.

Requirements: Students will be responsible for writing paper reviews each week, participating in discussions, completing one programming assignment, presenting once or twice in class (depending on enrollment, and possibly done in teams), and completing a project (done in pairs).

Note that presentations are due **one week before** the slot your presentation is scheduled. This means you will need to read the papers, prepare experiments, make plans with your partner, create slides, etc. more than one week before the date you are signed up for. The idea is to meet and discuss ahead of time, so that we can iterate as needed the week leading up to your presentation.

More details on the requirements and grading breakdown are here.

Prereqs: Courses in computer vision and/or machine learning (378 Computer Vision and/or 446 Machine Learning, or similar); ability to understand and analyze conference papers in this area; programming required for experiment presentations and projects.

Please talk to me if you are unsure if the course is a good match for your background. I generally recommend scanning through a few papers on the syllabus to gauge what kind of background is expected. I don't assume you are already familiar with every single algorithm/tool/image feature a given paper mentions, but you should feel comfortable following the key ideas.

Syllabus overview:

I. Single-object recognition fundamentals: representation, matching, and classification
   A. Specific objects
   B. Classification and global models
   C. Objects composed of parts
   D. Region-based methods

II. Beyond single objects: recognizing categories in context and learning their properties
   A. Context
   B. Attributes
   C. Actions and objects/scenes

III. Scalability issues in category learning, detection, and search
   A. Too many pixels!
   B. Too many categories!
   C. Too many images!

IV. Recognition and "everyday" visual data
   A. Landmarks, locations, and tourists
   B. Alignment with text
   C. Pictures of people

** The course website lists the papers that are required reading for each topic (usually three papers per week), as well as pointers to associated web pages and code packages. Please check out these pointers for useful material when preparing presentations for class. The “extra” papers listed for each topic are meant to provide additional optional background, and a good starting point for a literature review if you were to pursue a class project in that area.

Important dates:

Mon Jan 25:	Topic preferences for presentations due via email to Sudheendra.
Fri Feb 12:	Implementation assignment due (available now on webpage)
Mon Mar 8:	Project proposal abstract due
Mon Mar 22:	Project update and extended outline due
Week of Mar 29:	Individual project update meetings (out-of-class, by appt)
Wed May 5:	Project presentations due
Thurs May 6:	Extra class day for presentations part II (tentative)
Thurs May 13:	Final project papers due

Requirements:

Students are expected to do the assigned reading, participate in class discussions, write two paper reviews each week, and complete a final project and warmup coding assignment. There are two types of presentations: a paper presentation that involves doing background research on a topic (using the papers from the provided list), and an experiment presentation relevant to one of the selected papers. Details of our schedule will depend on total enrollment, meaning these will likely be done in partners.

Note that presentations are due one week before the slot your presentation is scheduled. This means you will need to read the papers, make plans with your partner, prepare experiments, slides, etc. more than one week before the date you are signed up for. The idea is to meet and discuss ahead of time, so that we can iterate as needed the week leading up to your presentation.

See **Blackboard->CourseDocuments** for examples of good paper reviews, presentations, and experiment examples to get a sense of what's expected.

Paper reviews

Submit two paper reviews per week for the assigned papers. (We'll usually read 3 papers each week; just choose any 2 to review.) Each review should address the following (in any order):

- Give a summary of the paper in your own words (very brief, 2-3 sentences)
- What is the main contribution of the paper?
- What are the primary strengths and weaknesses of the paper?
- How convincing are the experiments? If something specific is lacking, what should have been tested?
- Describe one specific way in which the work could be extended.
- Additional comments, including unclear points

Reviews are due by **10 PM on the night before** class (Tuesday). Email reviews to Kristen and Sudheendra, pasting the text directly into your mail (no attachments, please). Include [CS395] in the subject header. In weeks that you are presenting, skip writing the reviews.

Paper presentations

Each student (or team, depending on the total enrollment) will give a presentation in class covering 3 papers on a topic selected from the course syllabus list. This 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 **30 minutes**. Please run through it beforehand and check the time (a good rule of thumb: generally 30 minutes ~ target 30-35 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.

Experiment presentations

For each topic one person will present the results of some experimental evaluation of some main idea in a paper we read. Basically 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 may be able to find code or binaries provided by the authors online (see links on our course page schedule alongside the papers). The goal is to 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 that highlights an expected strength or weakness of the approach

Note that 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 what we have read.

To avoid redundancy with the paper presentation, this presentation should not spend time explaining the methods of the papers; you can assume this is covered already.

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 needed 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 leverage your experience later when working on the project. Be sure to explain the rationale for the outcomes, and conclude with a summary of the message(s) your example illustrates.

An experiment presentation should take about 20-30 minutes. In addition to the presentation, make a simple webpage to outline the demo and include links to any existing code, data, etc. you may 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 (by Wed): email slides to Kristen and Sudheendra, 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 instructors, practice, and check time length of the talk.

The day of the presentation: send final slides (and web page URL for experiment presentation) to Kristen.

Projects

A project could be built around any of the following, and should be done with a partner:

- an extension to a technique we study in class
- an in-depth analysis and empirical evaluation of one or two existing related techniques
- design a novel approach and perform accompanying experiments
- a thorough survey/review paper summarizing and analyzing the work in a sub-area from the syllabus (this one only should not be done in partners)

Initial project proposals will be due before the middle of the term. Details on the format will be given in class beforehand.

Grading:

Grades will be determined roughly as follows:

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