CS 677 Advanced Computer Vision

TTh 9:00AM-10:50 AM.

Instructor

Prof Ram Nevatia
Phone : (213) 740-6427
Email : nevatia AT usc.edu
Office Hours : TBD
Office Location : PHE 204

Teaching Assistant: TBD

Background:

This course is an enhanced version of the earlier computer vision course (cs574) that has been offered for many years. CS677 is now a 4 unit course so contains proportionately larger amounts of material compared to cs574 and has been updated to keep pace with fast changes in the field. CS677 is available for CS PhD credit but is also open to MS students.

Note that 574 is NOT a pre-requisite to 677 and it is not our plan to offer cs574 in the near future. As neither 574 nor 677 were offered in AY 16-17, demand for the course may be high in Fall 2017. We have limited capacity but do plan to repeat the offering in Spring 2018.

Brief Course Description

The course will provide an overview of the challenges of vision, the commonly used techniques and the current approaches. While specific examples and applications may be used to illustrate, the focus will be on fundamental techniques and algorithms. It aims to prepare students to start academic research or challenging careers in industry.

Prerequisites

1. Mathematics: Knowledge of and ability to use calculus, analytical geometry, linear algebra and probability theory.

2. Programming: Ability to program in C++ or Python (later preferred).

Textbooks

Required:
Computer Vision: A Modern Approach, D. Forsyth and J. Ponce, 2010

Recommended:
Computer Vision: Algorithms and Applications, Richard Szeliski, 2010.

Recommended:
Deep Learning: Algorithms and Applications, I. Goodfellow, Y. Bengio and A. Courville, 2017.

Grading Breakdown:

There will be two exams: Exam1 and Exam2, each counting for 30% of the grade (for a total of 60%).

There will be two mathematical assignments and four to six programming assignments. Large scale "projects" are not planned. Total assignment will count for 30% of the grade. Remaining 10% will be assigned to attendance and participation.

Detailed Course Syllabus:

Following is a list of topics expected to be covered, in anticipated order, and with expected time to be spent on them. However, as this is a new course, this list should be taken as being only indicative and actual topics, the order and the time may vary depending on various factors including student interests and preparation.

1. **Introduction (1 week)**
   Background, requirements and issues, human vision.
2. **Image formation: geometry and photometry (2 weeks)**
   Geometry, brightness, quantization, camera calibration, photometry (brightness and color)
3. **Image segmentation (1 weeks)**
   Various methods of image segmentation
4. **Multi-view Geometry (3 weeks)**
   Shape from stereo and motion, feature matching, surface fitting, Active ranging
5. **Object Recognition: Traditional Methods (2 weeks)**
   HoG/SIFT features, Bayes classifiers, SVM classifiers
6. **Object Recognition: Deep Learning Methods (2-3 weeks)**
   Deep neural networks, classification networks, object proposal networks
7. **Motion analysis and Activity Recognition (1-2 weeks)**
   Motion detection and tracking, Inference of human activity from image sequences