Syllabus for CS 677: Advanced Computer Vision, Fall 2020

Instructor

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

Teaching Assistant: TBD

ONLINE Course

Due to the prevalence of COVID-19 in Los Angeles, this class will be offered entirely in online format only. All lectures, office hours and interactions with TAs will be online. Exams will also be administered online though the modality has not yet been finalized. This is the first time that the course is being offered entirely online, without live feedback, which may introduce unanticipated variations to the content and conduct of the course.

Brief Course Description

The course will provide an overview of the challenges of vision, the common approaches and current techniques. While specific examples and applications may be used to illustrate, the focus will be on fundamental techniques and algorithms. We assume no prior knowledge of computer vision but still aim to study many modern, state-of-art techniques.

Course Availability

CS677 is available for CS PhD credit but is also open to CS and EE MS students. In exceptional cases, undergraduate students may also be admitted. PhD students will be given priority in enrollment; however, MS students are expected to be the large majority of enrollment. As the class is being offered online, we are not restricted by physical space but enrollment will still be limited by the amount of high-quality support we can provide to each student.

Prerequisites

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

2. **Programming**: Ability to program in *Python*.

3. **Other Courses**: There are no specific pre-requisite courses. In particular, courses in AI, Machine Learning, Deep Learning, Computer Vision and Image Processing are *not required*.

4. Entrance Exam: *No exam* will be given to assess pre-requisites. However, *GPA* may be used to screen students for preparedness.

Textbooks

There is, unfortunately, not a single, modern textbook available to cover the topics in this course. We will use published papers and tutorials extensively. Nonetheless, following books will be helpful for study.

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

"Deep Learning: Algorithms and Applications", I. Goodfellow, Y. Bengio and A. Courville, 2017 (online version available at no cost for personal use).

"A Guide to Convolutional Neural Networks for Computer Vision", S. Khan, H. Rahmani, S. Shah and M. Bennamoun, 2018 (online version available from a USC account).

Recommended:

"Computer Vision: Algorithms and Applications", Richard Szeliski, 2010 (online version available at no cost for personal use).

Grading Breakdown

There will be two exams: Exam1 and Exam2, each counting for 25% of the grade (for a total of 50%). There will be one mathematical assignments and five or six programming assignments. Large scale "projects" are not planned. Total assignment will count for 30% of the grade. 10% of the grade will be assigned to attendance (does not apply to remote DEN students). Lastly, an end of the term, "term paper", will count for the remaining 10% of the grade. We aim to waive the last requirement but this requires university-level approval; if granted, no term paper will be required and the weights given to the two exams will increase to 30% each.

Programming Assignments

The assignments must be completed using the Python language. We will use OpenCV library for the traditional part of the course and PyTorch for the deep learning component. It is expected that some cloud resources will be made available for assignments requiring use of GPUs; students are not required to have GPU-enabled personal computers of their own.

Detailed Course Syllabus:

The topic of computer vision is evolving very rapidly. Recent advances have come largely from "data-driven" deep learning and neural networks. However, traditional, "model-based" methods continue to be of interest and use in practice. This course will cover both traditional and deep-learning approaches.

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

1. **Introduction (1 week)**
   Background, requirements and issues, human vision
2. **Image formation: geometry and photometry (1.5 weeks)**
   Geometry, photometry (brightness and color), quantization, camera calibration
3. **Image segmentation and Feature Extraction (1 week)**
   Various methods of image segmentation, edge detection, object proposals, SIFT features
4. **Multi-view Geometry (2 weeks)**
   Shape from stereo and motion, feature matching, surface fitting, Active ranging
5. **Object Recognition: Traditional Methods (1 weeks)**
   HoG/SIFT features, Bayes classifiers, SVM classifiers
6. **Introduction to Neural Networks (1 week)**
   Artificial neural networks, loss functions, backpropagation and SGD, Batch Normalization
7. **Object Recognition: Deep Learning Methods (2.5 weeks)**
   Image classification, object detection and semantic segmentation, adversarial attacks. Various neural network architectures, visualization techniques.
8. **Motion analysis and Activity Recognition (1 week)**
   Motion detection and tracking, Inference of human activity from image sequences
9. **Selected Topics (1 week)**
   Examples: Face recognition, Image grounding, Visual question answering