http://www.cs.unc.edu/~lazebnik/research/spring08/

COMP 776: Computer Vision

Spring 2009, T TH 9:30-10:45, SN 115

Instructor: **Svetlana Lazebnik** (lazebnik -at- cs.unc.edu)

Quick links: syllabus, schedule, useful resources

www.manzaramesh.in

Overview

In the simplest terms, computer vision is the discipline of "teaching machines how to see." This field dates back more than forty years, but the recent explosive growth of digital imaging technology makes the problems of automated image interpretation more exciting and relevant than ever. There are two major themes in the computer vision literature: 3D geometry and recognition. The first theme is about using vision as a source of *metric 3D information*: given one or more images of a scene taken by a camera with known or unknown parameters, how can we go from 2D to 3D, and how much can we tell about the 3D structure of the environment pictured in those images? The second theme, by contrast, is all about vision as a source of *semantic information*: can we recognize the objects, people, or activities pictured in the images, and understand the structure and relationships of different scene components just as a human would? This course will strive to provide a unified perspective on the different aspects of computer vision, and give students the ability to understand vision literature and implement components that are fundamental to many modern vision systems.

Prerequisites: Basic knowledge of probability, linear algebra, and calculus. MATLAB programming experience and previous exposure to image processing are desirable, but not required.

Textbook: **Computer Vision: A Modern Approach** by David Forsyth and Jean Ponce is the recommended textbook for the course. The instruction will follow this textbook very loosely. Many additional instructional materials will be used throughout the course.

Grading: Computer vision is a very hands-on subject. For this reason, the coursework will primarily consist of implementation (please make sure you have access to MATLAB with the

Image Processing Toolbox installed). There will be three or four minor programming assignments and a larger final assignment which will most likely consist of a recognition competition (details to follow). Class participation will be another important component of the grade. This involves coming to class regularly, asking questions, and answering review questions. **Without satisfactory participation, it will be impossible to get an "H" in the class.** The weights assigned to different course components will be as follows:

- Regular assignments: 50%
- Final assignment: 30%
- Participation: 20%

Syllabus

I. Image formation

- Camera models
- Light and color
- Linear filters and edges
- Feature extraction (corners and blobs)

II. Grouping and fitting

- Hough transform
- RANSAC
- Alignment

III. Geometric vision

- Camera calibration
- Epipolar geometry
- Two-view and multi-view stereo
- Structure from motion

IV. Recognition

- Bags of features
- Generative and discriminative models
- Face detection and recognition

V. "Miscellaneous"

- Segmentation
- Optical flow
- Tracking

Schedule

Date	Topic	Readings, assignments
January 13	What is computer vision? **PPT** (29MB), **PDF** (9MB)	**Resource:** MATLAB tutorial
January 15	Cameras **PPT** (15MB), **PDF** (3MB)	**Reading:** F&P ch. 1
January 20	Radiometry **PPT** (14MB), **PDF** (1MB)	**Reading:** F&P ch. 4, 5 **Homework:** Assignment 1 out
January 22	Shape from shading (see slides from Jan. 20), color: **PPT** (13MB), **PDF** (3MB)	**Reading:** F&P ch. 6
January 27	Color concluded (see slides from Jan. 22)	
January 29	Linear filtering **PPT** (4MB), **PDF** (3MB)	**Reading:** F&P ch. 7 **Assignment 1 due at 5 PM**
February 3	Edge detection: **PPT** (4MB), **PDF** (2MB); corner detection	**Reading:** F&P ch. 8
February 5	Corner and blob detection **PPT** (8MB), **PDF** (3MB)	**Resource:** Harris corner detector code
February 10	Feature extraction (see Feb. 5 slides); least squares (see Feb. 12 slides)	**Reading:** F&P sec. 3.1 **Homework:** Assignment 2 out
February 12	Robust fitting, RANSAC **PPT** (2MB), **PDF** (0.5MB)	**Reading:** F&P ch. 15
February 17	Hough transform: **PPT** (4MB), **PDF** (1MB); alignment	
February 19	Alignment concluded: **PPT** (9MB), **PDF** (4MB)	**Reading:** Distinctive image features from scale-invariant keypoints
February 24	Single-view geometry **PPT** (1MB), **PDF** (1MB)	**Reading:** F&P ch. 2, 3 **Assignment 2 due at 5 PM**
February 26	Epipolar geometry and stereo **PPT** (2MB), **PDF** (1MB)	**Reading:** F&P sec. 10.1, ch. 11 **Homework:** Assignment 3 out
March 3	Binocular stereo **PPT** (14MB), **PDF** (3MB)	F&P ch. 11
March 5	Multi-view stereo	

	PPT (32MB), **PDF** (5MB)	
March 17	Structure from motion **PPT** (4MB), **PDF** (1.5MB)	**Reading:** F&P sec. 12.3, 12.4, 13.3.1, 13.4, 13.5 **Assignment 3 due at 5 PM** **Homework:** Assignment 4 out
March 19	Intro to recognition **PPT** (20MB), **PDF** (6MB)	
March 24	Recognition: Concepts and issues **PPT** (14MB), **PDF** (4MB)	**Resource:** ICCV 2005/CVPR 2007 Short Course on Object Recognition
March 26	Bags of features: **PPT** (7MB), **PDF** (2MB)	
March 31	Discriminative models: **PPT** (1MB), **PDF** (0.5MB)	**Reading:** F&P sec. 22.1, 22.2, 22.5 **Assignment 4 due at 5 PM**
April 2	Generative models: **PPT** (7MB), **PDF** (1.5MB)	**Homework:** Assignment 5 out
April 7	Spatial models: **PPT** (6MB), **PDF** (1.5MB)	
April 9	Eigenfaces: **PPT** (3MB), **PDF** (1MB)	**Reading:** F&P sec. 22.3
April 14	Face detection: **PPT** (2.5MB), **PDF** (1MB)	**Reading:** Robust Real-Time Face Detection
April 16	Segmentation: **PPT** (11MB), **PDF** (3MB)	**Reading:** F&P ch. 14
April 21	Optical flow: **PPT** (5MB), **PDF** (2MB)	**Assignment 5 due at 5 PM -- FIRM DEADLINE!**
April 23	Tracking: **PPT** (20MB), **PDF** (4MB)	**Reading:** F&P ch. 17

Useful Resources

Tutorials, review materials

- **MATLAB tutorial** (via David Kriegman and Serge Belongie)
- **More MATLAB tutorials:** basic operations, programming, working with images (via Martial Hebert at CMU)
- **Linear algebra review** (via David Kriegman)
- **Random variables review** (via David Kriegman)

General reference

- **Draft chapters of Forsyth and Ponce**
- **Computer Vision: Algorithms and Applications** -- a textbook in progress by Richard Szeliski
- **Ballard and Brown's historic textbook**
- **CVOnline** -- Compendium of Computer Vision
- **ICCV 2005/CVPR 2007 Short Course on Object Recognition** -- by Fei-Fei Li, Rob Fergus, and Antonio Torralba

MATLAB reference

- **MATLAB guide from Mathworks**
- **MATLAB image processing toolbox**

The real world

- **The Computer Vision Industry** -- maintained by David Lowe

Acknowledgments

The course slides draw on materials generously made publicly available by D. Forsyth, J. Ponce, J. Koenderink, S. Seitz, R. Szeliski, B. Freeman, M. Pollefeys, D. Lowe, K. Grauman, A. Efros, F. Durand, L. Fei-Fei, A. Torralba, R. Fergus (and possibly others whose attributions I either couldn't find or omitted by my own negligence).