Statistical and Learning Techniques for Computer Vision

Fall Semester, 2006

Course Syllabus

Instructor:	Dr. Jens Rittscher	Prof. Chuck Stewart
Office:	303 MRC	302 MRC
Email:	`rittj@cs.rpi.edu`	`stewart@cs.rpi.edu`
Hours:	M/R 3:30-4:30	M/R 3:30-4:30, Tu 1-3
Web:	`http://www.cs.rpi.edu/academics/courses/fall06/sltcv`	

Aim of the Course

Over the past decade, statistical modeling and machine learning techniques have moved into a central role in the field of computer vision. In part this is due to the advent of fast computational methods such as graph cuts and belief propagation, and in part this is due to a new understanding of the limits of physical modeling in addressing computer vision problems. These methods have been used in a wide range of problems, including but not limited to image restoration, segmentation, motion estimation, tracking, texture modeling and synthesis, and object recognition. The goal of this course is to provide students with the background necessary to understand the most important and widely-used techniques, and to learn to use them in their research.

Coverage

The course is roughly divided into three sections: random variables, Markov random fields (MRFs), and learning and classification methods. The section on random variables provides background on estimation, kernel density estimations and the EM algorithm. The section on MRFs covers modeling using MRFs, efficient computational approximations, learning the structure and parameters of MRFs, and recent generalizations. This leads into a study of graphical models in general and a comparison between graphical and discriminative models. This provides the transition into the last section of the course, which covers learning and classification methods, including support vector machines, neural networks, boosting, as well as subspace methods and learning of manifolds. Applications will be discussed throughout the course, sometimes as a small part of lecture, sometimes as an entire lecture, and sometimes through homework exercises.

The following is a tentative lecture-by-lecture summary of the course topics.

Random Variables	
Lecture 1	Random variables
Lecture 2	Probability density estimation
Lecture 3	Kernel density estimation
Lecture 4	Mixture models and EM
Markov Random Fields	
Lecture 5	Markov random fields
Lecture 6	Markov chain Monte Carlo and Gibbs sampling
Lecture 7	Variational methods
Lecture 8	Belief propagation
Lecture 9	Graph cuts
Lecture 10	Summary of MRF inference methods
Lecture 11	Learning MRF parameters
Lecture 12	From MRFs to general graphical models
Lecture 13	Conditional random fields
Learning and Classification	
Lecture 14	Discriminative vs. generative methods
Lecture 15	Support vector machines I
Lecture 16	Support vector machines II
Lecture 17	Neural networks
Lecture 18	Clustering methods
Lecture 19	Weak classifiers & boosting
Lecture 20	Classification & learning
Lecture 21	Projection methods
Lecture 22	Independent component analysis
Lecture 23	Continuous latent variables
Lecture 24	Concellation models

Required Background

Students should have taken at least one course in computer vision (or a closely-related field) and have a mathematical background that includes linear algebra and introductory probability and statistics.

Course Requirements

Course requirements will be divided into three parts: 50% for weekly homework assignments, 30% for three two-week programming projects, and 20% for a final literature review project. The weekly homework problems will vary from mathematics and theory to short programming problems (Matlab) to reading and analyzing papers from the research literature. The programming projects will be distributed throughout the semester. The final literature review project will involve a summary, comparison and analysis of two or more papers from the research literature. Details will be provided later in the semester.

Unless otherwise specified, students may work together in developing solutions to homework problems, although each student must write and submit his/her own solutions independently. On occasion students will be asked to solve a particular problem independently. These problems will typically be worth more than others. Similar rules apply to programming projects. The literature review project must be independent work.

Assignments must be turned in at the start of class on the due date. Any assignment turned in after the start of class is considered late. Students may use up to 4 late days (partial or whole) during the semester, but at most 2 on any particular assignment.

Readings

There is no required textbook for this course. Each lecture will be based on a different set of resources. Pointers to additional background reading will be provided with each lecture, and sometimes readings of specific book chapters or papers will be assigned. Despite the lack of a textbook, we do recommend the purchase of one or more companion books. Most important is the newly available book

> Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.

Another important resource is

> Duda, Hart and Stork, *Pattern Classification*, John Wiley and Sons, 2001.

Both books were developed for the pattern recognition research community, which is closely-aligned with computer vision, but has a longer and more consistent history in the use of statistical methods.