CSC 2541
Winter, 2005
Topics in Machine Learning: Natural Scene Statistics

Time:	Fridays 10–12
Location:	Pratt 378
Instructor:	Richard Zemel <zemel@cs.toronto.edu>
Office hours:	Wednesdays 10–12 Pratt 290D
Class URL:	www.cs.toronto.edu/~zemel/Courses/csc2541.html

Overview

This course is a graduate seminar devoted to research in natural scene statistics. Recent years have seen a growing interest in the statistics of natural scenes, such as images and sounds. These studies are important on two counts: both for understanding how these statistics influence and are reflected in our perceptual systems, and for developing algorithms that exploit these statistics. On the biological side, a long-standing assumption is that organisms are adapted to the statistical properties of the signals to which they are exposed. Recent developments in statistical modeling, along with powerful computational tools, have enabled researchers to study more sophisticated models for sensory inputs, to validate these models against large data sets, and to begin examining their relevance to representations in neurons.

On the applications front, identifying relevant scene statistics can be useful for solving perceptual inference problems. For example, the distribution of image gradients are different on and off object boundaries, a fact that can be exploited in building effective edge detectors. There are also important applications of statistical knowledge of images to computer graphics, and to the design of devices that interact with humans.

Note that the course differs from previous incarnations of CSC2541, Topics in Machine Learning. The rule for topics courses is that a graduate student can get credit for the same course number more than once, as long as the department can ensure that the courses are quite different and has a subtitle indicating the specific topic.

Course requirements

The course will be organized around a set of representative papers in this area. The format of the class will be a mix of background lectures and discussions of these papers. The primary requirement of this class is that you read and thoroughly think through the assigned papers.

The requirements, and grading in the class will be divided up as follows:

1. Presentations (40%): Each week, 1-2 people will be responsible for presenting the selected material and leading a discussion of it. The presentation should be approximately one hour, with the remainder of class reserved for discussion. The presenter(s) will be expected to produce:

   - a brief summary of the relevant background material
   - an in-depth presentation of the main ideas, techniques, and results of the paper
   - a few questions and issues for class discussion

   The presenter(s) should meet with me twice to discuss the presentation. Around 10-14 days before the class meeting, we will meet to discuss the main issues and relevant background; and 3-5 days before the class we will go through the material together and I will make suggestions about the presentation.

2. Class Participation (20%): This grade will be based on contributions to class discussions, and questions and comments about the material throughout the term. Each student will be required to bring to class two questions or comments about the paper we are discussing that day, and is expected to bring these up during or following the presentation.

3. Final Project (40%): There will also be a final project that will involve studying a particular method discussed in the course, or a different relevant one in the area of natural scene statistics. The project can be adapted to the background of the student. The project could be primarily empirical (implementing and studying the performance of some approach), or analytical (comparing alterrnative approaches to a particular problem), or some combination of these.

   On February 25th you will be required to submit a brief proposal outlining the final project you intend to carry out. This will be returned with feedback within a week and you will be expected to get started on your project immediately thereafter. You should be on the lookout throughout the earlier sections of the course for topics or issues that you find particularly interesting and would like to pursue in more detail in a project. I will hand out a list of suggested topics.

4. Final: There is no final exam for the course.

Prerequisites

The methods described in the course will rely on mathematics, and we will discuss mathematical derivations of the various methods. Many of the required concepts will be covered in class, and a list of texts containing relevant material will be supplied. However, some background in linear algebra, probability/statistics, and calculus will be useful.

Ideally each student will have taken at least one course in AI, and preferably in either machine learning and/or vision. But this is only a soft requirement.

Readings

There will be a paper or two assigned for each class meeting. It is expected that you will read these in detail, and come to class prepared with questions, comments, criticisms, etc. I will also supply pointers to supplemental readings, such as background material from machine learning, machine vision, or computational neuroscience.

Most of the papers are available on the web, and links are available. For any paper that is not, I will bring copies for everyone in the class one week before it is to be discussed.

For your presentations, please choose at least one week (preferably two) during the second week of classes so that we can fill out the schedule of presenters. You can change your selection by arranging a swap with another student, and letting me know at least two weeks in advance.

CLASS SCHEDULE
(subject to change)

Class	Topic	Paper Title	Authors
Jan 7	Introduction	*Statistics of natural images and models*	Huang, Mumford
Jan 14	Sparseness (generative models)	*Natural image statistics and efficient coding* *Independent component filters...*	Olshausen, Field v. Hateren, v.d. Schaaf
Jan 21	Applications (exp. family models)	*Learning low-level vision* *Filters, random fields and maximum entropy*	Freeman et al Zhu, Wu, Mumford
Jan 28	Intrinsic images	*Recovering intrinsic images...* *Deriving intrinsic images...*	Tappen et al Weiss
Feb 4	GSMs	*Scale mixtures of gaussian...*	Wainwright, Simoncelli
Feb 11	Higher-order ICA		Lewicki-Karklin

Other possible topics & papers:

- other higher-order models: Welling-Osinderio-Hinton
- neural data: Gallant; Rolls
- applications
  1. edge detection (Yuille)
  2. compression (Simoncelli)
  3. object recognition (Torralba)
  4. occluded structures, segmentation (Ruderman, Mumford, Zhu/Tu)
  5. tracking, optic flow
  6. matting
- time-varying data: Hyvarinen
- sounds: Lewicki
- active approach: Reinagel/Zador, cat videos