Course Outline
ECE271B – Statistical learning II
Department of Electrical and Computer Engineering
University of California, San Diego

Nuno Vasconcelos Spring 2007

Your responsibilities in this class fall into two main categories:

1. Class participation and homework 30%
2. Final project 70%

There will be 3 homework sets, handed-out roughly every other week. You are allowed to collaborate on homework as long as you write your solutions independently and acknowledge the collaboration in the problems where it was used.

Final project: Projects are individual. The ideal project would be a contribution publishable in a conference. This could be

- new theory or algorithms
- a creative application of existing techniques to, say, computer vision, speech processing, data mining, etc.
- a thorough experimental comparison of a set of existing techniques
- a tutorial review of a part of the literature not covered in class

The deadline to hand in your project report is the last day of classes. At this date you will have to hand in a 12-page project report. There will be 20 minute presentations of all projects on the last two classes. If this project is also being done for another class you have to state that. Preferably this should not happen. To allow me to keep track on your progress, you will have to hand in the following items.

April 19	2-page project proposal
May 3	meet with instructor (optional)
April 17	1-page intermediate project report
June 5/7	20 minute presentation
June 7	final report

Note that it is your responsibility to find a suitable project topic. We can, of course, talk about it. The project will be evaluated according to

1. the final paper 70% (50% for content, 20% for writing)
2. the presentation 30%

TA: Hamed Masnadi-Shirazi (hmasnadi@ucsd.edu)

Office hours: (for homework see TA first)

- TA: TBA
- instructor: Thursday, 7:00-8:30PM, EBU1-5602

Text: There is no adopted text. I will distribute handouts as we go. There are various books of interest. These are not required but can be used for alternative explanations of the material. The following list is organized by how easy the books are to read. If you want to pick one, I would recommend Hastie et al.

1. V. Cherkassky, *Learning from Data.* John Wiley, 1998.
2. T. Hastie, R. Tibshirani, J. H. Friedman, *The Elements of Statistical Learning.* Springer Verlag, 2001.
3. Ralph Herbrich, *Learning Kernel Classifiers*, MIT Press, 2002.
4. Vladimir Vapnik, *The Nature of Statistical Learning Theory.* Springer Verlag, 1999.
5. Bernhard Scholkpf and Alexander Smola, *Learning with Kernels*, MIT Press, 2002.
6. Vladimir Vapnik, *Statistical Learning Theory.* Springer Verlag, 1998.
7. Luc Devroye, Laszlo Gyorfi, Gabor Lugosi, *A Probabilistic Theory of Pattern Recognition.* Springer Verlag, 1998.

The following are some books on the general principles of machine learning, optimization, etc. Some of them touch most of the issues that we will cover, but only superficially.

1. Richard O. Duda, Peter E. Hart and David G. Stork *Pattern Classification.* New York, NY: John Wiley & Sons, 2001.
2. Tom Mitchell, *Machine Learning*, McGraw-Hill, 1997.
3. Christopher Bishop, *Neural Networks for Pattern Recognition.* Oxford University Press, 1996.
4. Stephen Boyd, Lieven Vandenberghe, *Convex Optimization.* Cambridge University Press, 2004.

There is a web page for the course, available from

`http://www.svcl.ucsd.edu/~nuno`

Instructor

Nuno Vasconcelos, EBU1 5602, 4-5550, e-mail: nuno@ece.ucsd.edu

LECTURE SUBJECT	Number of classes
Introduction	1
Linear learning machines and the perceptron	2
Kernels and feature spaces	2
Basics of Optimization	2
Support Vector Machines	2
Boosting	2
Other large margin methods	2
Learning theory	4