6.891 Machine Learning: Administriva

Contact Information

Lecturer:	Tommi Jaakkola	tommi@ai.mit.edu, NE43-735, 253-0440
TA:	Jason Rennie	jrennie@ai.mit.edu, NE43-733, 253-5339
Web Page:	http://www.ai.mit.edu/courses/6.891/	

Classes

Lectures:	Tuesday and Thursday 2:30-4pm	room 4-149
Recitations:	Wednesday and Friday 1-2pm	room 34-302
Office Hours:	Tuesday 4:00-5:30pm, Friday 9:30am-10:30am	room NE43-733

Homework

Homework will be due every other Thursday at the beginning of class. There will be a total of six homeworks. Late homeworks will not be accepted. Collaboration is encouraged while solving problems, however every student must hand in his or her own homework. Please list your collaborators on every homework that you turn in. Copying will not be tolerated.

Exams

We will be holding one mid-term and one final. The mid-term will scheduled for a date approximately half way through the course and the final will be held during the usual MIT finals period.

Grading

Your final grade will be based on homework, exams and the final project in the following proportions: Final 30%, Mid-term 15%, Homework 30%, Project 25%.

Text

We will be using the second edition of *Pattern Classification* by Duda, Hart & Stork as our book. The book is not yet in print; an electronic copy will be made available on-line via the course web page.

6.891 Machine Learning: Syllabus

This introductory course on machine learning will give an overview of many techniques and algorithms in machine learning, beginning with topics such as simple perceptrons and ending up with more recent topics such as boosting, support vector machines, hidden Markov models, and Bayesian networks. The course will give the student the basic ideas and intuition behind modern machine learning methods as well as a bit more formal understanding of how and why they work. The underlying theme in the course is statistical inference as this provides the foundation for most of the methods covered.

Topics to be covered:

- Classification, regression
- Neural networks, back-propagation
- Generalization, regularization
- Feature selection, feature induction
- Combining multiple models, additive models, boosting
- Support vector machines, kernels
- Complexity, VC-dimension
- Model selection, structural risk minimization
- Density estimation, parametric, non-parametric
- Clustering, agglomerative, mixture models
- Hidden markov models, forward-backward algorithm, EM
- Bayesian networks, graphs, estimation, model selection
- Markov random fields, dynamic Bayesian networks
- Time series prediction, independent component analysis

A more detailed lecture schedule will be available soon from the course web page.