

Computer Science 6100/4100: Machine Learning

August 24, 2008

RPI, Fall 2008

Instructor: Sanmay Das

Class times: Mon, Thu 2:00-3:30 in Low 3130

Office: Lally 302

Phone: x2782

Office hours: Mondays from 3:30-4:30 PM and by appointment.

1 Course Description

1.1 Overview

This course is an introduction to many of the important themes in machine learning, including supervised and unsupervised learning, online learning, reinforcement learning, and learning in societies of agents. It will focus on statistical methods, how to evaluate success, and learning as a critical component of decision-making.

1.2 Objectives

This course is intended to be quite broad in its coverage. Students who take the class should develop a sufficiently good grasp of the fundamental ideas and terms in machine learning to understand the academic literature. Further, students should come away from the class capable of implementing machine learning algorithms and interpreting the outcomes of experiments they run using these implementations.

1.3 Prerequisites

This course requires a substantial level of mathematical maturity and comfort with probability, multivariable calculus, and basic linear algebra. Students are also expected to be able to write significant programs in a modern programming language. I encourage the use of statistical and mathematical languages such as R and Matlab. Please speak with me if you have any concerns about whether your background is sufficient for this class.

1.4 Format

Class sessions will mostly be lectures, but there may be some sessions that focus on discussion of papers that will be made available in advance. There will be two exams held during class hours, and three assignments, which will consist of a mixture of written and programming problems. I

may also sometimes ask for brief (1-2 single spaced pages in 11 or 12 point font) written summaries of papers we discuss. All written work **must** be handed in at or before the beginning of class.

1.5 Textbooks

I will not be using any particular textbook, and I do not recommend that you purchase a textbook just for this class. The following references may be useful:

1. *Artificial Intelligence: A Modern Approach* (2nd ed) by Stuart Russell and Peter Norvig, Prentice Hall, 2003.
2. *Statistical Inference* (2nd ed) by George Casella and Roger L. Berger, Duxbury Press, 2001.
3. *Reinforcement Learning: An Introduction* by Richard S. Sutton and Andrew G. Barto, MIT Press, 1999.
4. *Rational Herds: Economic Models of Social Learning* by Christophe P. Chamley, Cambridge University Press, 2004.

1.6 Preliminary Syllabus

Please note that this syllabus may change based on time constraints, the interests of the class, or other factors.

1. Introduction to the main themes of the course. Supervised, unsupervised, reinforcement, and “rational” (Bayesian) learning. Math review.
2. Estimators. Maximum-likelihood and Bayesian estimation.
3. Linear models for regression and the bias/variance tradeoff.
4. Supervised learning: generative and discriminative models. Naive Bayes and Logistic Regression.
5. Decision trees. Nearest-neighbor classification.
6. Evaluation – accuracy, precision, recall, ROC curves. Overfitting and Occam’s razor.
7. Hyperplanes. The perceptron algorithm. Support vector machines.
8. The kernel trick. Regularization.
9. Ensemble methods: Boosting and Bagging.
10. Utility theory and decision-making. Cost-sensitive classification.
11. Unsupervised learning. Expectation maximization and k-means.
12. Online (supervised) learning.
13. Temporal models: Hidden Markov Models and Kalman Filters.
14. Sequential decision-making. Markov Decision Problems.

15. Secretary problems.
16. Reinforcement learning. Exploration vs. exploitation.
17. Q-learning and Temporal Difference (TD)-learning. Generalization and function approximation in reinforcement learning.
18. Bayesian framework for incorporating information into decision-making. The Martingale Convergence Theorem.
19. Multi-armed bandits and Gittins Indices.
20. Social learning. Cascades and herd behavior. Learning in financial markets.

2 Policies

2.1 Announcements and Course Website

The course website is at <http://www.cs.rpi.edu/~sanmay/teaching/cs6100>. All announcements related to the class will be made either in lecture or on the website. **I will assume that any announcement made on the website is known to everyone in class within 24 hours of it being posted.** It is important to check the website regularly! You are responsible for all announcements made in lecture or on the website.

2.2 Assessment and Course Grade

Your overall course score will be determined (on a curve) using the following weights.

1. Assignment 1: 15%
2. Exam 1: 15%
3. Assignment 2: 20%
4. Course project: 25%
5. Exam 2: 15%
6. Class attendance and participation: 10%

Late assignments will be penalized by 20% per-day except for cases of (official) Institute-established illness or emergency. Since this is a graduate class, I expect students to attend regularly and participate actively in discussions – notice that 10% of your course grade is based on attendance and participation.

After each assignment or exam has been graded, I will present some summary statistics to give students a sense of how they are performing in the class. If you would like to appeal your grade on any of the assignments or exams, you may do so within 10 days of the assignment or exam being handed back. In order to appeal the grade, please provide a detailed written statement explaining why you believe the assigned grade is incorrect, in addition to the assignment or exam. I will regrade the entire assignment or exam and your grade may go up or down, or it may stay the same.

2.3 Collaboration and Academic Integrity

The statement below is based partially on a model statement from the provost.

Student-teacher relationships are based on trust. For example, students must trust that teachers have made appropriate decisions about the structure and content of the courses they teach, and teachers must trust that the assignments that students turn in are their own. Acts which violate this trust undermine the educational process. The Rensselaer Handbook of Student Rights and Responsibilities defines various forms of Academic Dishonesty and you should make yourself familiar with these.

In this class, you are allowed to collaborate on assignments to the following extent. You are welcome to discuss problems (both written and programming) with each other and to take your own notes during these discussions. However, you must write up solutions on your own. You must write, on the assignment, the names of students you discussed each problem with, and any external sources you used in a significant manner in solving the problem. Lack of citation of a source is a serious violation of this policy. In programming assignments, please do not look at the code of any other student in the class, or allow any other student to look at your code. There may be some assignments in which teamwork is permitted on programming projects. In these cases collaboration outside your team is not allowed.

Submitting an assignment that is in violation of this policy will automatically lead to receiving no credit for the assignment and a reduction of at least one grade modifier (e.g. from B to B-) beyond that in the overall course grade. However, depending on the circumstances, it could also lead to harsher penalties, for example, a failing grade in the class and initiation of the Institute-wide disciplinary process. If you have any questions about the level of collaboration permitted, or any other aspect of this policy, please speak with me about it before handing in the assignment!

2.4 General Note

I think it is a good idea to (and the Institute requires me to!) spell out class policies in detail so that everyone is on the same page as we start the semester. If you have questions on anything, either academic or policy-related, please do not hesitate to come and speak with me, either in office hours or by appointment.