Machine Learning
CAS CS542 (Spring 2012)

Course Description: Introduction to modern machine learning concepts, techniques, and algorithms. Topics include regression, kernels, support vector machines, feature selection, boosting, clustering, hidden Markov models, and Bayesian networks. Programming assignments emphasize taking theory into practice, through applications on real-world data sets.

Prerequisites: CAS CS 112; or equivalent programming experience, and familiarity with linear algebra, probability, and statistics.

Lectures: MW 5:00p–6:30p, MCS B33

Instructor: Kevin Small

Office:	MCS 146
Office Hours:	Mondays, 6:30p-7:45p; Tuesdays, 2:30p-4:30p (or by appointment)
Email:	ksmall1@bu.edu

Teaching Fellow: Zhiqiang (Alex) Ren

Office:	TBA
Office Hours:	TBA
Email:	aren@cs.bu.edu

Course Readings: There is not an *explicitly* required text for this course. However, I will use [B06] will be the primary reference for the course and it is a worthwhile purchase if you plan on working in machine learning. All of these are on reserve in the Science/Engineering library. Selected readings from other books and papers will be distributed as electronic or hard copies.

Recommended Text
[B06] Bishop, C. M. (2006). Pattern Recognition and Machine Learning.

Optional Text
[M97] Mitchell, T. M. (1997). Machine Learning.

Reference Texts
[DHS00] Duda, R. O, Hart, P. E., and Stork, D. G. (2000). Pattern Classification.
[W10] Wasserman, L. (2010). All of Statistics: A Concise Course in Statistical Inference.
[HTF09] Hastie T., Tibshirani R., and Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction (Second Edition).
[KV94] Kearns, M. J. and Vazirani, U. (1994). An Introduction to Computational Learning Theory.
[KF09] Koller, D. and Friedman, N. (2009). Probabilistic Graphical Models: Principles and Techniques.
[SS01] Scholkopf, B. and Smola, A. J. (2001). Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond. [RN09] Russell, S. and Norvig, P. (2009). Artificial Intelligence: A Modern Approach (Third Edition).

Course Webage: `http://cs-people.bu.edu/ksmall1/cs542`

All lecture notes, course assignments and the course schedule can be found on these pages.

Course Mailing List: The course mailing list `cascs542a1-l@bu.edu` is managed by ITs majordomo engine. If you preregistered for CS542, then you are already on this list. If you drop the class, they you will be automatically removed from this list. Students are encouraged to post course-related messages and questions to the mailing list. To send e-mail to the course mailing list, use the e-mail address `cascs542a1-l@bu.edu` – the mailing list server only accepts email posted from a BU account.

Grading

Problem Sets (6)	35%
Project	15%
Midterm Exam	20%
Final Exam	30%

Late Assignments: Problem sets are due at the beginning of class. However, students are allowed 96 hours of "late credit" of which at most 48 hours may be used for any single assignment.

Examinations: Tests are closed-notes and closed-book. However, for the final examination you may bring a single page (both sides) crib sheet on normal-sized paper (8.5x11 inches). Your crib sheet should be handwritten (no mechanical or electronic reproductions are allowed). If you use a crib sheet, you will be asked to turn it in with your test.

Midterm Examination: Monday, March 5 (in class)
Final Examination: Monday, May 7 (6p-8p, location TBA)

Collaboration: You are allowed and encouraged to work together on the problem sets. However, after discussing the problem (and hopefully reaching a solution) each student must write down the solution independently, preferably with minimal reference to written notes from the joint session. In other words, you must understand the solution well enough in order to reconstruct it independently. Sharing source code for programming exercises is explicitly forbidden. Working together on exams, of course, is also not allowed.

Academic Honesty: All course participants must adhere to the College of Arts and Sciences Academic Conduct Code. Printed copies of the code are available from CAS 105.