

CS 6243 Syllabus – Spring 2005

Machine Learning

This course studies machine learning techniques in the area of artificial intelligence. Major topics covered include inductive learning, unsupervised learning, speedup learning, and computational learning theory. Prerequisites: CS 5233 (Artificial Intelligence) or CS 5633 (Analysis of Algorithms).

Instructor

Name: Tom Bylander
Office: SB 3.02.05B
Phone: 458-5693
Email: bylander@cs.utsa.edu
URL: <http://www.cs.utsa.edu/~bylander/cs6243/index.html>
Office Hours: Monday 1-2pm, Wednesday 10-11am, Friday 11am-12, or by appointment.

Books

Required: T. Mitchell, *Machine Learning*, McGraw-Hill, 1997.

Grading

Homework & Labs	20%
Midterm	20%
Presentation	10%
Project	20%
Final Exam	30%

Assignments may be handed in late by the next class session with a 10% penalty. Thereafter, a 20% penalty will be assessed.

Software

We will be using Weka (<http://www.cs.waikato.ac.nz/~ml/weka/>) which is a collection of machine learning algorithms implemented in Java. You are encouraged to download and install Weka in your own account. If you are unfamiliar with Java, consider Sun's own *Java Tutorial* (<http://java.sun.com/docs/books/tutorial/>). One of many online books on Java is *Thinking in Java* by Bruce Eckel (<http://www.mindview.net/Books/TIJ/>).

Assignments

Homeworks will occasionally be assigned throughout the semester. Many homeworks will involve using, modifying, and/or developing programs. 1 homework will be dropped from the grade.

Each student will present one paper in the last two weeks of the course. A list of papers to select from will be provided.

There will be one programming project in this course. Projects suggested by the instructor will involve the development of a learning algorithm in Weka. Students may propose and work on alternatives, subject to the instructor's approval.

Attendance and Participation

Regular class attendance per se is not required, but note that assignments and exams will be partly based on the lectures. Make-up exams are permitted as long as it's a reasonable excuse, you inform me in a timely fashion, and you document the excuse.

Tentative Schedule

<u>Week</u>	<u>Topic</u>	<u>Reading</u>
Jan. 18,20	Introduction, Concept Learning, Inductive Bias	§1, §2
25,27	Nearest Neighbor	§8
Feb. 1,3	Decision Trees	§3
8,10	Artificial Neural Networks (ANNs)	§4
15,17	Bayesian Learning	§6
25,27	Learning Sets Of Rules	§10
March 1,3	Support Vector Machines (SVMs)	Handout
8,10	Midterm on March 8	
22,24	Bagging and Boosting	Handout
29,31	Evaluating Hypotheses	§5
Apr. 5,7	Computational Learning Theory	§7
12,14	Reinforcement Learning	§13
19,21	Unsupervised Learning	Handout
26,28	Paper Presentations	
May 2	Paper Presentations	
May 13	Final, 10:30am-1:15pm	