CS 260: Machine Learning Theory
Fall 2011

Course Overview

This course will provide a broad overview of the theoretical foundations underlying common machine learning algorithms. Topics covered will include:

- An introduction to the Probably Approximately Correct (PAC) model of learning
- Uniform convergence bounds
- VC theory
- Online learning, mistake bounds, the Perceptron algorithm, and Winnow
- Learning from expert advice, no-regret learning, and the Weighted Majority algorithm
- Weak learners, ensemble methods and boosting
- Support vector machines and kernels
- Connections to game theory

Who should take this course? This course is primarily intended for students with background and interest in machine learning who would like to know more about the theoretical foundations, and students with background and interest in theoretical computer science who would like to know more about machine learning. Although this course has no official prerequisites, students who enroll must have a basic knowledge of probability (roughly at the level of the first half of CS 112) and must be comfortable reading and writing formal mathematical proofs.

Who should not take this course? This course is not intended for students who are primarily interested in applications of machine learning. There is no programming required, and we won't spend much time talking about implementation. If you would prefer a more applied view of machine learning, consider taking CS 276A.

This course counts for credit for the AI major and minor field requirements for computer science Ph.D. students.

Meeting Times

Lectures: Mondays & Wednesdays, 2:00-3:50pm, Boelter 5422
Attendance at lectures is required.

Staff Contact Info and Office Hours

Instructor: Prof. Jenn Wortman Vaughan
Office Hours: Fridays 2–3pm, 4532H Boelter Hall
Contact: jenn at cs

Grader: Jacob Mathew
Contact: jacobgmathew at gmail

Breakdown of Grades

Grades will be based on the following components:

- *Homework Assignments* (60%): There will be four homework assignments due in class. Assignments submitted up to 24 hours late will be penalized 25%. **No assignments will be**

accepted for credit more than 24 hours after the deadline. All solutions must be typed and printed out. Solutions that are not typed will be penalized 25%, and unreadable answers will not be graded. Your solutions will be graded on both **correctness and clarity**. If you cannot solve a problem completely, you will get more partial credit if you identify the gaps in your argument than you will if you try to cover them up. For some assignments, it might be the case that only a subset of the problems will be graded in detail. You may or may not be told which problems will be graded in advance.

- *Final Projects* (40%): Projects consist of a written report and in-class presentation. All projects must be completed in groups of 2–3 students; there are **no exceptions** to this rule. 25% of your project grade will be based on your proposal and status report, 50% on your write-up, and 25% on your presentation, which will be scored by both the instructor and your peers in class.

Regrade Policy

If you believe that a mistake was made in the grading of a homework assignment, you may submit a request for a regrade. This request must be submitted in writing directly to the grader, and must include a clear explanation of the reason you believe you should have received more points. All regrade requests must be received within one week.

Textbook & Readings

There is no required textbook for this course, but several recommended textbooks are listed on the course website. These texts are on reserve at the Science and Engineering Library. Links to suggested reading material will be posted on the course website throughout the quarter.

Academic Honesty Policy

Collaboration on the homework assignments is encouraged! Students are free to discuss the homework problems with anyone in the class under the following conditions:

1. Each student must write down his or her solutions **independently**, and must understand the solutions he or she writes down. Talking over solutions is fine, but reading or copying another student's answers is never acceptable!
2. Each student must write a list of all of his or her collaborators at the top of each assignment. This list should include **anyone with whom the assignment was discussed**.

Additionally, any sources used must be properly credited. These policies are described in more detail in the Academic Honesty Policy that must be signed by every student in the class. When in doubt, consult this policy or ask the instructor.

Course Website

All of this information and more is available on the course website:

`http://www.cs.ucla.edu/~jenn/courses/F11.html`

Please check this website regularly for reading assignments, problem sets, and announcements.