

**CS478 – Machine Learning
Cornell University
Spring 2000**

Instructor: Assistant Professor Claire Cardie
Office: Upson 4142
Email: cardie@cs.cornell.edu
WWW: <http://www.cs.cornell.edu/Info/People/cardie/>
Office Hours: Weds 3:30-4:30; Thurs 1-2
Required Text: *Machine Learning*, Tom Mitchell

Course Description:

Machine learning is concerned with the question of how to construct computer programs that automatically improve their performance through experience. This course introduces students to the primary approaches to machine learning including inductive inference of decision trees, case-based learning, neural network learning, statistical learning methods, genetic algorithms, bayesian learning, and clustering. In addition this course will introduce theoretical issues including inductive bias, PAC learning, the mistake bound framework, and boosting. Evaluation will be based on homeworks, research paper critiques, a course project, and a final.

Course Goals:

The goal of this course is to introduce students to current machine learning and data mining methods. It is intended to provide enough background to allow students to apply machine learning and data mining techniques to learning problems in a variety of application areas. A course project will be required. Skills required to understand, critique, and extend existing research in machine learning will be emphasized.

Communication:

Email: You can reach me at cardie@cs.cornell.edu.

WWW: The course home page will contain postscript and pdf versions of handouts, homeworks, and view-graphs (when available).

Evaluation:

Homework (35%): There will be four homework assignments. They will consist of a combination of written problems and programming assignments. Your lowest score on the assignments will be dropped. The assignments will be clumped mostly in the first half of the semester.

Paper critiques (10%): During the semester, you'll be asked to write one-page critiques of selected research papers. There will be about six of these.

Final project and paper (25%): You will also complete a machine learning project of your choosing. The results of the project will be written up in a short paper (seven pages maximum). It is due during the last week of classes. Your grade will be based on the quality of the idea investigated, the project report and the code that you write. There will be a separate handout on the course project and paper.

Exam (20%): Yes, there will be a final exam.

Class participation / Other (10%): You'll be expected to participate in class discussion or otherwise demonstrate an interest in the material studied in the course.

Course Policies

Academic integrity policy: You are responsible for knowing and following Cornell's academic integrity policy. Highlights of the policy are available for your perusal at the course web page. For CS478, there are no collaborative assignments; therefore, you are not allowed to discuss homeworks or programming assignments or paper critiques with anyone except me or the TA's. All code that you turn in must be entirely yours.

Late assignment policy: Barring extenuating circumstances, all homeworks, critiques, and projects must be turned in on the date specified, AT THE START OF CLASS. Assignments turned in within 24 hours of the due date will be penalized one full grade (e.g. A→B). Assignments turned in within 48 hours of the due date will be penalized two full grades (e.g. A→B). Assignments more than 48 hours late will not be accepted. Sorry.