COMP 4360 Machine Learning

Winter 2012

Instructor

Jacky Baltes
EITC E2 402, University of Manitoba, Winnipeg, MB, Canada, R3T 2N2
Email: jacky@cs.umanitoba.ca
WWW: http://www.cs.umanitoba.ca/~jacky

Lecture Times: M,W,F: 15:30 - 16:20 in E2 304
Office Hours: M,W,F: 9:30 - 10:30 or by prior appointment via email

Prerequisites: Not to be held with the former 74.416, 74.319, or 74.426.

Introduction

This course covers methods that computer programs can use to adapt themselves to perform better on similar problems in the future, that is they learn from their experience. A selection of topics that may be covered in lecture are:

- **Concept Learning**: version spaces, inductive bias, learning of disjunctions, case-based meta learning.
- **Decision trees**: ID3 and C4.5, the overfitting problem.
- **Neural nets**: perceptrons, gradient descent, backpropagation.
- **Instance-based learning**: k-nearest neighbor algorithm, locally weighted regression, case-based reasoning.
- **Learning sets of rules**: sequential covering algorithm, learning first order rules, FOIL.
- **Genetic algorithms**: classification using genetic algorithms, genetic programming.
- **Reinforcement learning**: Dynamic programming, temporal difference learning, Q-learning.
- **Support Vector Machines**

Grading

The course mark is determined by: (a) a final exam (50%), (b) a midterm test (20%), and (c)

practical work (30% Assignments).

Assignments (30%)

Course assessment includes a large practical component. There are three assignments covering specific topics. Each assignment is worth 10%. The following list contains some sample assignment topics:

- Compare the performance of classification algorithms in learning to classify middle earth inhabitants from the Lord of the Rings Universe.
- Learn to predict indicators of annual income from Canadian census data.
- Evaluate the performance of different machine learning algorithms in the “homicidal chauffeur” game.
- Learn to control an inverted pendulum using reinforcement learning.
- Implement a system that can learn the orientation of a small robot based on a view of the robot. Some example views are shown below.

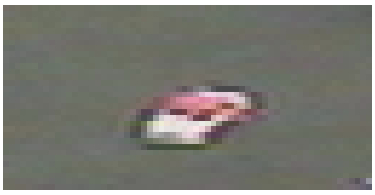
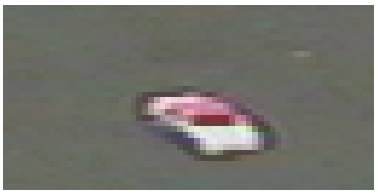


- Learn a path tracking controller for a small remote controlled toy car.

Midterm (20%)

There will be a midterm exam worth 20%. The midterm will be held in class.

Final Exam (50%)

There will be a final exam worth 50%. The exam will be three hours long. The final exam will be held during examination period at the end of the term. Exact time and location will be determined by Student Records.

Textbook

The textbook for this course is Tom. M. Mitchell, Machine Learning, 1st Edition, McGraw Hill, ISBN 0-07-042807-7, 1997. The material in the textbook will be supplemented by additional material covering the state of the art of machine learning in robotics. Students are expected to understand the material as well as being able to implement simple versions of the described techniques and algorithms.

Academic Dishonesty

Students are reminded that there are penalties for academic dishonesty. Academic dishonesty includes submitting assignments that are not entirely the student’s own work. See the UofM Calendar: Academic Dishonesty and Plagiarism and Cheating for more information.

A declaration sheet, which states that the work being submitted is completely your own, is available at University of Manitoba Honesty Declaration This sheet must be printed out, filled in, signed, and attached to every which is submitted. No assignment will be marked unless the declaration is attached.