

COMP 652: Machine Learning - Fall 2008 - Syllabus

General Info

Location: McConnell Engineering building, room 103
Times: Monday and Wednesday, 2:35-3:55pm
Instructor: Prof. Ted Perkins
Office: McGill Centre for Bioinformatics (Duff; see web page for directions)
Phone: (514) 398-5018
Email: perkins@mcb.mcgill.ca
Office hours: To be determined
Teaching assistant: To be determined
Class web page: http://www.cs.mcgill.ca/~perkins/COMP652_Fall2008/index.html

Course Description

The field of machine learning is concerned with the question of how to construct computer programs that improve automatically with experience. In recent years, many successful applications of machine learning have been developed, ranging from data-mining programs that learn to detect fraudulent credit card transactions, to autonomous vehicles that learn to drive on public highways. At the same time, there have been important advances in the theory and algorithms that form the foundation of this field. The goal of this class is to provide an overview of the state-of-art algorithms used in machine learning. We will discuss both the theoretical properties of these algorithms and their practical applications.

Prerequisites

Students should have a basic knowledge of computer science, including ability to program in some reasonable, high-level language (such as C/C++, Java, Matlab, etc.). Students may be required to use existing machine learning software packages, necessitating either knowledge of an appropriate data analysis environment that incorporates those tools or enough computer savvy to download and run such software. Students should have knowledge of multivariate calculus (e.g., MATH 222 at McGill), linear algebra (e.g., MATH 223), and probability theory (e.g., MATH 323).

Course materials

There is no required course textbook. Lecture will be prepared from a variety of sources, and lecture notes will be posted on the web page. Good textbooks on machine learning include:

- T. Mitchell, "Machine Learning", McGraw-Hill, 1997.
- R. O. Duda, P. E. Hart & D. G. Stork, "Pattern Classification. Second Edition", Wiley & Sons, 2001.
- C. M. Bishop, "Pattern Recognition and Machine Learning", Springer, 2006.

- C. M. Bishop, "Neural Networks for Pattern Recognition", Oxford University Press, 1996.
- T. Hastie, R. Tibshirani and J. Friedman, "The elements of statistical learning", Springer, 2001.
- R. S. Sutton and A. G. Barto, "Reinforcement learning: An introduction", MIT Press, 1998.
- E. Alpaydin, "Introduction to Machine Learning", MIT Press, 2004.

Recommended readings from the books above will be specified for each lecture.

Lecture notes will be posted on the web page. (Exactly who produces those notes is a matter of discussion.)

Evaluation

Students' marks will depend on two (or three) components:

- 6-9 homework assignments, which will have "paper and pencil" as well as programming questions. These will generally be assigned on a class day and you will have 1 week to complete them. They will be accepted late, at -20% penalty, on the day of the next class period.
- A course project, which could involve a complete (but non-complicated) application for ML algorithms--or something else.

The relative weight of the different components is as follows. Each homework carries equal weight. The course project will have the same weight as one homework. Class participation is encouraged, and, while not formally graded, may be used to push borderline grades up.

Rough course outline (*guaranteed to change!*)

- Introduction (1 lecture)
- Linear and logistic regression, Bayesian interpretation and Bayesian learning, model selection, artificial neural networks (4 lectures)
- Nearest-neighbor algorithms (1 lecture)
- Generative learning algorithms, Naive Bayes (1 lecture)
- Support vector machines (3 lectures)
- Decision Trees (1 lecture)
- Experimental design and evaluation (1 lecture)
- Computational Learning Theory (2 lectures)
- Ensemble methods (2 lectures)
- Reinforcement learning (4 lectures)
- Unsupervised learning (5 lectures)
- Wrap-up (1 lecture)