Machine Learning

CS 2750 (ISSP 2170) – Spring 2007

Lecture meeting time: Monday, Wednesday**:** 1:00 PM-2:15 PM
Classroom: 5313 Sennott Square (SENSQ)

Instructor:	**Milos Hauskrecht**	**TA:**	**Chenhai Xi**
Office:	5329 Sennott Square Building	**Office:**	5503 Sennott Square
Office Hours:	TBA	**Office Hours:**	TBA
Phone:	(412) 624–8845	**Phone:**	(412) 624-8832
Email:	milos@cs.pitt.edu	**Email:**	chenhai@cs.pitt.edu
Course Web page:	http://www.cs.pitt.edu/~milos/courses/cs2750/	**Web Page:**	http://www.cs.pitt.edu/~chenhai/

Course Description:

The goal of the field of machine learning is to build computer systems that learn from experience and that are capable to adapt to their environments. Learning techniques and methods developed by researchers in this field have been successfully applied to a variety of learning tasks in a broad range of areas, including, for example, text classification, gene discovery, financial forecasting, credit card fraud detection, collaborative filtering, design of adaptive web agents and others.

This introductory machine learning course will give an overview of many models and algorithms used in modern machine learning, including linear models, multi-layer neural networks, support vector machines, density estimation methods, Bayesian belief networks, mixture models, clustering, ensamble methods, and reinforcement learning. The course will give the student the basic ideas and intuition behind these methods, as well as, a more formal understanding of how and why they work. Students will have an opportunity to experiment with machine learning techniques and apply them a selected problem in the context of a term project.

Prerequisites: Knowledge of matrices and linear algebra (CS 0280), probability (CS 1151), statistics (CS 1000), programming (CS 1501) or equivalent, or the permission of the instructor.

Textbook:

Chris Bishop. *Pattern recognition and Machine Learning.* Springer, 2006.

Homework assignments

Homework assignments will have mostly a character of projects and will require you to implement some of the learning algorithms covered during lectures. Programming assignments will be implemented in Matlab. The assignments (both written and programming parts) are due at the beginning of the class on the day specified on the assignment. In general, no extensions will be granted.

Collaborations: You may discuss material with your fellow students, but the report and programs should be written individually.

Term projects

The term project is due at the end of the semester and accounts for a significant portion of your grade. You can choose your own problem topic. You will be asked to write a short proposal for the purpose of approval and feedback. The project must have a distinctive and non-trivial learning or adaptive component. In general, a project may consist of a replication of previously published results, design of new learning methods and their testing, or application of machine learning to a domain or a problem of your interest.

Policy on Cheating

Cheating and any other anti-intellectual behavior, including giving your work to someone else, will be dealt with severely. If you feel you may have violated the rules speak to us as soon as possible. Please make sure you read, understand and abide by the Academic Integrity Code for the Faculty and College of Arts and Sciences (http://www.fcas.pitt.edu/academicintegrity.html).

Students with Disabilities

If you have a disability for which you are or may be requesting an accommodation, you are encouraged to contact both your instructor and Disability Resources and Services, 216 William Pitt Union, (412) 648-7890/(412) 383-7355 (TTY), as early as possible in the term. DRS will verify your disability and determine reasonable accommodations for this course.

Tentative syllabus:

- Machine learning introduction
- Density estimation
- Supervised learning:
  - Linear and logistic regression
  - Generative classification models
  - Multi-layer neural networks
  - Support vector machines
- Unsupervised learning
  - Bayesian belief networks (BBNs)
  - Learning parameters and structure of BBNs
  - Expectation maximization
  - Clustering
- Dimensionality reduction/feature selection
  - Feature filtering
  - Wrapper methods
  - PCA
- Ensemble methods (mixtures of experts, bagging and boosting)
- Reinforcement Learning