Machine Learning

CS 2750 (ISSP 2170) – Spring 2020

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

Instructor:	**Milos Hauskrecht**	**TA:**	**TBA**
Office:	**5329 SENSQ**	**Office:**	
Office Hours:	TBA	**Office Hours:**	
Phone:	(412) 624–8845	**Phone:**	
Email:	milos@cs.pitt.edu	**Email:**	
Course Web page:	http://www.cs.pitt.edu/~milos/courses/cs2750/	**Web Page:**	

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, 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 machine learning, including linear classification and regression models, multi-layer neural networks, support vector machines, Bayesian belief networks, mixture models, clustering, ensemble methods, and reinforcement learning. The aim of the course is to 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 implement and experiment with the different machine learning techniques on various datasets during homework assignments and a term project.

Prerequisites: Knowledge of calculus, 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 consist of a mix of theoretical problems and programming tasks. The programming tasks will require you to implement in Matlab some of the learning algorithms covered during lectures. Please visit http://technology.pitt.edu/software/browse/matlab.html to obtain a Matlab license for students. 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: No collaboration on homework assignments, programs, term projects and exams unless you are specifically instructed to work in groups, is permitted.

Grading

The final grade for the course will be determined based on homework assignments, exams, the term project and your lecture attendance and activity. The midterm exam will be held on March 5, 2019 during the class. The final exam will be the week of April 13-17. The term project presentations will be held the week of April 20-24.

Policy on Cheating

Cheating and any other anti-intellectual behavior, including giving your work to someone else, will be dealt with severely and will result in the Fail (F) grade. 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 University of Pittsburgh and School of Computing and Information (SCI) at: http://sci.pitt.edu/current-students/policies/academic-integrity-policy/

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:
  - basic parametric distributions
  - exponential family models
  - non-parametric density estimation methods
- Supervised learning:
  - Linear and logistic regression
  - Generative classification models
  - Multi-layer neural networks
  - Support vector machines
  - Decision trees
  - Non-parametric classification models
- Probabilistic graphical models
  - Bayesian belief networks (BBNs)
  - Learning parameters of BBNs
  - Expectation-maximization
- Clustering:
  - K-means clustering
  - Soft-clustering
  - Hierarchical clustering
- Dimensionality reduction/feature selection
  - Feature filtering
  - Wrapper methods
  - Principal Component Analysis
- Ensemble methods (mixtures of experts, bagging and boosting)
- Reinforcement Learning