Introduction to Machine Learning

CS 1675 – Spring 2019

Lecture meeting time: Tuesday, Thursday**:** 9:30AM – 10:45 AM
Classroom: 5129 Sennott Square (SENSQ)
Recitations:

- Section 1: Wednesday: 11:00am-11:50am, SENSQ 6110
- Section 2: Wednesday: 3:00pm - 3:50pm, SENSQ 6110

Instructor:	**Milos Hauskrecht**	**TA:**	**Jeongmin Lee**
Office:	5329 SENSQ	**Office:**	5324 SENSQ
Office Hours:	TBA	**Office Hours:**	TBA
Phone:	(412) 624–8845	**Phone:**	4-8455
Email:	milos@pitt.edu	**Email:**	jlee@cs.pitt.edu
Course Web page:	http://people.cs.pitt.edu/~milos/courses/cs1675/		

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 models and algorithms used in machine learning including linear regression and classification models, multi-layer neural networks, support vector machines, Bayesian belief networks, clustering, ensemble methods, and reinforcement of learning. The course will give the student the basic ideas and intuition behind these methods, and a more formal understanding of how and why they work. Through homework assignments students will have an opportunity to implement and experiment with many machine learning techniques and apply them to various real-world datasets.

Prerequisites: STAT 1000, 1100, or 1151 (or equivalent), and CS 1501, or the permission of the instructor.

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

Additional readings/text:
Hal Daume III. A Course in Machine Learning. http://ciml.info

Homework assignments
Homework assignments will be a mix of theoretical questions and programming assignments. The programming will require you to implement some of the learning algorithms covered during the lectures in

Matlab. Please visit http://technology.pitt.edu/software/matlab-for-students to see how to get a free copy of Matlab license for students. The assignments (both reports 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.

Policy on collaboration: No collaboration on homework assignments, programs, 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, and your lecture attendance and activity. The midterm exam will be held on March 7, 2018 during the class. The final exam will be scheduled by the university/ school during the final exams week.

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 (DRS – www.drs.pitt.edu), 140 William Pitt Union, (412) 648-7890, drsrecep@pitt.edu, (412) 228-5347 for P3 ASL users, 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
  - 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
- Clustering:
  - K-means clustering
  - Hierarchical clustering
- Dimensionality reduction/feature selection
  - Feature filtering
  - Wrapper methods
  - Principal Component Analysis
- Ensemble methods (mixtures of experts, bagging and boosting)
- Reinforcement Learning