ECE699-005: Learning From Data - 2017
Fall 2017

Syllabus and Course Information

Course Overview

This is an introductory course in machine learning (ML) and pattern recognition that covers basic theory, algorithms, and applications. ML is a key technology in Big Data, and in many financial, medical, commercial, and scientific applications.

Machine learning is the science of getting computers to act without being explicitly programmed. In the past decade, machine learning has given us self-driving cars, practical speech recognition, effective web search, and a vastly improved understanding of the human genome. Machine learning is so pervasive today that you probably use it dozens of times a day without knowing it.

ML has become one of the hottest fields of study today, and is one that is taken by undergraduate and graduate students in many universities throughout the country. This course balances theory and practice, and covers the mathematical as well as the heuristic aspects. Some of the questions that will be asked include:

What is learning?
Can a machine learn?
How is machine learning done?
How can machine learning be done well?
What are the traps and what are the things to look out for?

In this class, you will learn about the most effective machine learning techniques, and gain practice implementing them and getting them to work on real data. You will learn not only about the theoretical underpinnings of learning, but also gain the practical know-how needed to quickly and powerfully apply these techniques to new problems.

This course provides a broad introduction to machine learning, data mining, and statistical pattern recognition. Topics include: (i) Supervised learning (parametric/non-parametric algorithms, support vector machines, kernels, neural networks). (ii) Unsupervised learning (clustering, dimensionality reduction, recommender systems, deep learning). (iii) Learning theory (bias/variance tradeoffs, VC theory, large margins). The course will draw from numerous case studies and applications.

Instructor: Professor Monson H. Hayes
Nguyen Engineering Building, Suite 3100
(703) 993-3969
hayes@gmu.edu

Class: Thursday, 4:30 pm – 7:10

Office Hours: Thursday 7:00 – 8:30 and by appointment.

Textbook: Y. Abu-Mostafa, M. Magdon-Ismail and H-T. Lin
Learning from Data, AML Press, 2012

Course Webpage

All course materials will be available from a link to the course that may be found on the instructor's web page, http://ece.gmu.edu/~hayes

Course Assignments and Exams

Homework and computer exercises
Course project with a presentation

Prerequisites

The student is expected to have an understanding of basic probability. The student should also have a strong background in mathematics, particularly linear algebra. An understanding of optimization theory would be beneficial, but not necessary.

List of Topics may Include

1. Introduction to Learning
2. Bayesian Estimation
3. Maximum Likelihood
4. The Linear Model
5. The Perceptron Algorithm and the Pocket Algorithm
6. Training versus Testing
7. The Theory of Generalization
8. The VC Dimension and the Bias-Variance Tradeoff
9. Linear Regression and Classification
10. Trees and Random Forests
11. Least Squares
12. Neural networks
13. Overfitting, Regularization, and Validation
14. Support Vector Machines and Kernel Methods
15. Linear dimensionality reduction (PCA and ICA)
16. Mixture models
17. Autoencoders
18. Energy models – Boltzman Machines and RBMs
19. Clustering
20. Deep Learning
21. Diversity and Unbalanced Data Problems