Computer Science and Engineering 417A: Introduction to Machine Learning

Washington University in St. Louis, Fall 2014

Instructor: Sanmay Das
email: sanmay at seas dot wustl dot edu
Class times: Tue, Thu 10:00-11:30 in McDonnell 162
Office: Jolley 510
Office hours: TBA and by appointment.

1 Course Description

1.1 Overview

This course is a broad introduction to machine learning, covering supervised learning, unsupervised learning, decision-making under uncertainty, and reinforcement learning. Topics that will be covered include generative and discriminative techniques for classification (likely including regression, Naive Bayes, decision trees, neural networks, nearest-neighbor methods, support vector machines, and boosting), clustering and dimensionality reduction, dynamic programming, and temporal difference methods. Note that there is some overlap with topics in the 500-level courses on Artifificial Intelligence and Machine Learning, but the material covered in this class will be at a more elementary level.

1.2 Prerequisites

CSE 241 and ESE 326 (or Math 320) or equivalents; Linear algebra and multi-variable calculus. If you do not have a basic background in CS through data structures and algorithms, or if you are not comfortable with calculus and probability, you may have a hard time in this class.

1.3 Format

Class sessions will be lectures. There will be two in-class exams, one in mid-October, and one on the last day of class, December 4. There will be no separate final exam. There will be 6-8 homework assignments that will involve a mix of programming/computational exercises and pencil-and-paper problems.

1.4 Textbooks

Most of the time I will not post lecture notes. Instead I will give references to the parts of the textbooks that correspond to the material covered in class on a given day. I may also post additional required or recommended reading to the website.

We will use the following two textbooks:

1. *Learning From Data*, Y. Abu-Mostafa, M. Magdon-Ismail, and H-T Lin.
   `http://amlbook.com/`

2. *Artificial Intelligence: A Modern Approach*, S. Russell and Peter Norvig.
   `http://aima.cs.berkeley.edu/`

1.5 Preliminary List of Topics

This preliminary list of topics may change based on time constraints and the progress of the class.

1. What is machine learning? Types of learning.

2. Generalization in finite and infinite hypothesis spaces. Training versus testing, model complexity, the VC bound, the bias-variance tradeoff.

3. Linear models: the perceptron, regression, logistic regression.

4. Nonlinear transformations of data.

5. The problem of overfitting. Regularization and validation as ways of preventing overfitting.

6. Modern supervised learning techniques, including Naive Bayes, decision trees, neural networks, nearest-neighbor methods, support vector machines, and boosting.

7. Unsupervised learning: clustering (k-means) and dimensionality reduction (principal component analysis)

8. Sequential decision-making: The Bellman equation and dynamic programming.

9. Reinforcement learning: temporal difference methods and Q-learning.

2 Policies

2.1 Announcements and Course Website

The main course website is at `http://www.cse.wustl.edu/~sanmay/teaching/cse417`. All announcements related to the class will be made either in lecture or on the website. **I will assume that any announcement made on the website is known to everyone in class within 24 hours of it being posted.** It is important to check the website regularly! You are responsible for all announcements made in lecture or on the website.

We will use Piazza for all questions and discussions related to the class. Please post questions on Piazza – they will reach the professor and all the TAs, and you will get a quicker response. Individual emails about class issues will typically be met with a response saying "Please post your question to Piazza (anonymously if you so desire)." A link to the Piazza site will be on the main course website.

2.2 Assessment and Course Grade

Your overall course score will be determined (on a curve) using the following weights.

1. Homework assignments: 50%
2. Each in-class exam: 25%

Assignments will typically be due **at the beginning of lecture**. I will **not grant any extension requests** unless a staff member from Engineering Student Services emails me directly with information on why you need an extension, in which case I may grant one. If you would like to appeal your grade on any work, you may do so within 10 days of the work being handed back or the grade being received. In order to appeal the grade, please provide a detailed written statement explaining why you believe the assigned grade is incorrect, in addition to the work itself. We will regrade the entire piece of work, and your grade may go up or down, or it may stay the same.

2.3 Collaboration and Academic Integrity

In this class, you are allowed to collaborate on assignments to the following extent. You are welcome to discuss problems with each other and to take your own notes during these discussions. However, you must write up solutions on your own. You must write, on the assignment, the names of students you discussed each problem with, and any external sources you used in a significant manner in solving the problem. Lack of citation of a source is a serious violation of this policy. You may not give or receive help from other students in the class on exams.

Submitting an assignment or exam that is in violation of this policy will automatically lead to receiving no credit for the assignment and a reduction of at least one grade modifier (e.g. from B to B-) beyond that in the overall course grade. However, depending on the circumstances, it could also lead to harsher penalties, for example, a failing grade in the class and initiation of the school's formal academic integrity review process. If you have any questions about the level of collaboration permitted, or any other aspect of this policy, please speak with the instructor or one of the TAs about it before handing in the assignment!