CAP5610 Machine Learning

Term: Fall 2014

Time: Tu Th 12-1:15pm

Room: HEC0111

Instructor: Guo-Jun Qi, guojun.qi@ucf.edu

Office hours: Tu Th 2 – 3pm at HEC318

Pre-requisites

Calculus, linear algebra and probability theory

Background

In this course, we will begin with the basic concepts in machine learning, in the context of several classic topics from supervised learning (classification), unsupervised learning (model fitting, clustering) to feature learning, dimensionality reduction etc. We will also discuss advanced topics on learning theory, graphical models, and dynamic Bayesian model if time allows.

Announcement: Please send me your project title (temporary) and your team members to cap5610ucf@gmail.com. This is only for the purpose of composing program for project presentation scheduled for Nov. 20 and Nov. 25 in class. Every group will be given 10 min talk + 5 min Q&A. Project and presentation will be scored by peers (like a reviewer) and the instructor.

Books:

Required textbook: Pattern Recognition and Machine Learning (PRML), C. Bishop, Springer, 2006.

Other suggested books:

Machine Learning, T. Mitchell, McGraw-Hill, 1997.

The Elements of Statistical Learning: Data Mining, Inference, and Prediction, 2nd Edition, T. Hastie, R. Tibshirani, J. Friedman, Springer, 2011.

Probabilistic Graphical Models: Principles and Techniques, D. Koller, and N. Friedman, The MIT Press, 2009.

Covered topics (subject to change)**:**

1. Introduction to machine learning: feature extraction, examples of machine learning problems, training and test protocol, datasets, KNN
2. Naive Bayesian classifier
3. Linear regression and classification
4. Support vector machine and kernel method
5. Neural Networks and back propagation
6. Unsupervised learning problems: clustering, PCA, LDA, CCA etc.
7. Modeling fitting and EM algorithm
8. Graphical model: Baysian networks, Markov random fields, approximation inference, variational method, sampling, loopy belief propagation
9. Advanced topics: matrix factorization, metric learning, latent models, online learning, active learning, latent models, sparse coding, nonparametric Bayesian model etc. (will cover as many as possible if time allows)

Grading policy:

Homework and machine problems will account for 60% of the grade. There will be a mid-term exam (20%) and a final project (20%). There is no final exam.

Final Project:

Students will form several groups in a size of 3-4 people to collaborate on the final project. A final project can be a survey (paper review) on a certain machine learning topic, or a proposed project to solve a machine learning problem. Instructor will also suggest some topics for the project. A project proposal shall be submitted to instructor for review before it is approved. The proposal must specify who will be attending the project, the roles of each member, and description of the proposed survey topic or problem. Project report will be graded. Project presentation will be arranged if time allows.

Machine Problems:

The course is not intended to teach a programming language, so students can use any languages to solve the machine problems (e.g., C, C++, Java, Matlab, Python, Octave). You are required to submit a report, which shall include a brief description of how you implement the algorithm in the language you choose, the parameter setting, your test protocol, as well as the result you obtain. Source codes shall be submitted along with the report. Submit reports and source codes to cap5610ucf@gmail.com.

Lecture Nodes:

- August 19, Lecture 01: Introduction to Machine Learning [**pdf**]
  Reading assignment Sec 1.2, Sec 1.4 of PRML.
  Link to MNIST dataset in different formats: [**MNIST a**] [**MNIST b**] (two datasets are the same)
- August 21, Lecture 02: Review of Probability Theory [**pdf**]
  Reading assignment Chap 2 and Appendix B of PRML.
- August 26, Lecture 03: Bayes Classifier [**pdf**]
  Reading assignment Section 1.5 of PRML.
- August 28, Lecture 04: Bayes Classifier II [**pdf**]
  Note that MP 2 has been assigned, which will be due by 11:59am, Sep 18.
- September 3, Lecture 05: Logistic Regression [**pdf**]
- September 5, Lecture 06: Linear Regression [**pdf**]
- September 9, Lecture 07: Support Vector Machines [**pdf**]
- September 11, Lecture 08: Support Vector Machines II [**pdf**]
- September 16, Lecture 09: Neural Networks [**pdf**]
- September 18, Lecture 10: Deep Learning [**pdf**]
- September 23, Lecture 11: Dimensionality Reduction [**pdf**]
- September 25, Lecture 12: Dimensionality Reduction II [**pdf**]
- September 30, Lecture 13: Independence Component Analysis [**pdf**]
- October 2, Lecture 14: Canonical Component Analysis [**pdf**]
- October 7, Lecture 15: Clustering Analysis [**pdf**]
- October 9, Lecture 16: Expectation-Maximization Algorithms [**pdf**]
- October 14, Lecture 17: Hidden Markov Model [**pdf**]
- October 16, Lecture 18: Hidden Markov Model II [**pdf**]
- October 21, Lecture 19: Decision Trees [**pdf**]
- October 23, Lecture 20: Midterm Review [**pdf**]
- October 28, Lecture 21: Boosting [**pdf**]