EL-GY 9123: Introduction to Machine Learning

Overview

This course is an introductory graduate-level machine learning for electrical and computer engineering students. The class covers fundamental algorithms in machine learning including linear regression, classification, model selection, support vector machines, dimensionality reduction and clustering. The material will be developed with computer exercises on real and synthetic data. Applications are demonstrated in audio and image processing, robotic control, and text and web analysis. No prior machine learning experience is required.

Prof: Sundeep Rangan, srangan@nyu.edu

- Lectures Tu 6:30-9:00, Room JAB 475
- Texts:
  - Hastie, Tibshirani, Friedman, "Elements of Statistical Learning". http://statweb.stanford.edu/~tibs/ElemStatLearn/printings/ESLII_print10.pdf
  - Raschka, "Python Machine Learning", 2015. http://file.allitebooks.com/20151017/Python%20Machine%20Learning.pdf
- Supplementary texts and resources
  - Bishop, "Pattern Recognition and Machine Learning"
  - Installing python (need to do this before first recitation): http://jupyter-notebook-beginner-guide.readthedocs.io/en/latest/index.html
  - Python tutorial: https://docs.python.org/3/tutorial/
- Grading:
  - Midterm 30%, Final 30%, Labs and homework 30%, Quizzes 10%
  - Optional project for up to 20% of grade.
  - Labs will involve approximately six python-based exercises.
  - Midterm and final exams are closed book with cheat sheets. Students will need to be able to write simple python in the exams.
- Pre-requisites:
  - Undergraduate probability and linear algebra. No ML experience is expected for this class
  - Students may NOT enroll in this class if they have taken any one of: CSE-GY 6923 (Intro grad ML), EE-UY 4423 (Intro UG ML), EL-GY 9133 (Advanced ML)
  - Students with ML experience are encouraged to take graduate-level Probability (EL-GY 6303) in the Fall and advanced ML in the Spring.
  - Programming experience is essential, including some exposure or willingness to learn object-oriented programming. No experience in python is required as python will be taught as part of the class.

Tentative Outline

- Introduction to ML
- Linear regression
- Model selection and assessment
- Linear classification
- Numerical optimization
- Support vector machines
- Midterm
- Neural networks and Tensorflow
- Convolutional neural networks
- Non-parametric methods
- PCA and dimensionality reduction
- Clustering and Expectation Maximization
- Tree-based methods
- Final exam