EE-UY 4423: Introduction to Machine Learning

Overview

This course provides a hands on approach to machine learning and statistical pattern recognition. The course will provide an introduction to fundamental algorithms for linear regression, classification, model selection, support vector machines, dimensionality reduction and clustering. The material will be developed with hands on python-based exercises on real and synthetic data. Applications will be demonstrated in audio and image processing, robotic control, and text and web analysis.

This class will also be cross-listed with CSE (CSE number TBD)

- Prof: Sundeep Rangan, TA: TBD
- Lectures TuTh 11-12:30, Room TBD
- Recitation Fr 11:30-12:50, Room TBD
    - Used for weekly quiz, help with Python labs and homework.
- Text: James, Witten, Hastie and Tibshirani, "An Introduction to Statistical Learning", https://web.stanford.edu/~hastie/local.ftp/Springer/ISLR_print1.pdf

  Note: While this text uses R, the class will be in Python.

- Supplementary texts:
    - Raschka, "Python Machine Learning", 2013
    - Bishop, "Pattern Recognition and Machine Learning"
- Grading:
    - Labs and final project: 40%, Homework: 15%, Midterm 1: 15%, Midterm 2: 15%, Final: 15%
    - Labs will involve approximately six python-based exercises.
- Pre-requisites:
    - Undergraduate probability and linear algebra
    - 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
    - What is statistical learning?
    - Introduction to python
- Single Variable Linear regression
    - Single variable regression
    - Assessing model accuracy, R^2 and RSE
    - Probability review: joint distributions, correlations of two random variables
    - Probabilistic analysis of regression
- Multiple variable regression
    - Problem formulation, feature vectors
    - Matrix description and linear algebra review
    - Correlation matrices
    - Non-linear transformations and feature spaces
- Linear classification
    - Overview of classification
    - Linear classifiers
    - Logistic regression
    - Review of optimization, gradient descent
    - Linear discriminant analysis (LDA)
    - Assessing model accuracy. ROC and AUC
- Cross validation
    - Overfitting and underfitting, bias and variance
    - Leave one out CV
    - K-fold cross validation
- Model selection
    - Subset selection
    - Ridge regression
    - Sparsity and LASSO
- Neural networks:
    - Perceptron and the perceptron algorithm
    - Multi-layer perceptrons
    - Back propagation
    - An intro to “deep” networks
- PCA
    - Principal components and dimensionality reduction
    - Computations via the Singular Value Decomposition
- K-Means
    - Concepts of clustering
    - K-means algorithm