Course Description

ECE 445: Machine Learning for Engineers (Topics in ECE)
Fall 2019

Schedule — TTh 5:00 – 6:20 PM
Place — RWH-206

Instructor

Waheed U. Bajwa
723 CoRE, Tel. 848-445-8541
`http://inspirelab.us`

Prerequisites

Enrolled students must have taken undergraduate courses in probability theory and linear algebra. The course will also require extensive programming, especially as part of the term project. In keeping with the industry standards, all programming will need to be done in Jupyter notebooks (`http://jupyter.org/`) using either Julia, Python, or R (individual students will get to pick any one of these languages in most assignments). Submission of all programming assignments/projects will take place through GitHub (`http://github.com/`). In many instances, students will be forbidden from using popular machine learning packages such as `scikit-learn` for assignments.

Learning Outcomes

- Mastery of the basic terminology and concepts in machine learning
- Understanding of the basic building blocks of practical machine learning systems
- Mathematical understanding of commonly used machine learning algorithms
- Ability to develop basic machine learning systems from scratch in Jupyter
- Recognition of common pitfalls that come with machine learning systems

Who should take this course?

- Students interested in machine learning and data science careers in the industry
- Students interested in applying machine learning techniques in their own disciplines
- Students interested in graduate school with a focus on machine learning

Who should not take this course?

- Students interested only in getting an easy ‘A’ grade
- Students who are uncomfortable with abstract and/or rigorous mathematics
- Students who are afraid of putting in four to eight hours per week for this course
- Students who are uncomfortable with the concept of programming in Jupyter

Required Texts

Much of the material taught in this class will come from the following two texts:

G. James, D. Witten, T. Hastie and R. Tibshirani
An Introduction to Statistical Learning with Applications in R
Springer; 1st ed. 2013, Corr. 7th printing 2017 edition (September 1, 2017)
Available at: `http://www-bcf.usc.edu/~gareth/ISL/`

T. Hastie, R. Tibshirani, and J. Friedman
The Elements of Statistical Learning: Data Mining, Inference, and Prediction
Springer; 2nd edition (2016)
Available at: `https://web.stanford.edu/~hastie/ElemStatLearn/`

Students should also maintain detailed notes of the material taught in class. In addition, students will be occasionally provided material and internet links for further referencing.

Tentative Course Outline

Weeks 1–3:

- Review of basic probability theory and linear algebra concepts
- Introduction to machine learning and its basic terminology
- Understanding the machine learning pipeline
- Feature engineering and representation learning
- Introduction to the role of numerical computations in machine learning

Week 4:

- Introduction to basic machine learning algorithms
- Theoretical description of machine learning algorithms
  - Statistical risk minimization and empirical risk minimization
  - Role of geometry in machine learning algorithms
  - Optimization and computational aspects of machine learning algorithms

Weeks 5–7:[1] Classification

- Parametric models
  - Naive Bayes classification

[1] First in-class exam will take place during this period.

- Linear classification (including support vector machines)

- Non-parametric models
  - Nearest-neighbor classification
  - Kernel support vector machines
- Frequentist versus Bayesian philosophy
- Brief introduction to neural networks

Weeks 8–9: Practical machine learning systems

- Testing, training, and validation data, and cross-validation
- Distribution mismatch (transfer learning) and missing data
- Bias–variance tradeoff, overfitting, and Occam's razor
- (Stochastic) Gradient descent

Weeks 10–11: Regression

- Linear regression
- Nonlinear regression
- Penalized regression

Weeks 12–14:[2] Clustering

- Gaussian mixture model
- Expectation–maximization algorithm
- k-means and related algorithms

Week 14: Introduction to advanced topics (subject to availability of time)

- Hidden Markov models and graphical models
- Random forests and ensemble learning
- Ranking and recommendation systems
- Density estimation

[2] Second in-class exam will take place during this period.