Course Syllabus

ECE 445: Machine Learning for Engineers
(14:332:445: Topics in ECE)
Fall 2020

Table of Contents

1 Instructional Staff

Instructor

Waheed U. Bajwa (waheed.bajwa@rutgers.edu)
723 CoRE, Tel. 848-445-8541
http://inspirelab.us

Teaching Assistant

To Be Assigned (will be announced on Canvas)

Means of Communication

Instructor E-Mail: `waheed.bajwa@rutgers.edu` **(all personal questions)**
Canvas LMS Site: `https://bit.ly/ECE445f20` **(overall course management)**
Canvas Discussions: `https://bit.ly/ECE445f20discussion` **(all non-personal questions)**
Twitter: `@SigProcessing (#RUECE445)` **(light commentary related to course topics)**

2 Lecture Timings and Office Hours

The course is being offered as a *synchronous, remote* one with video recordings of each lecture being made available on Canvas for those who cannot attend the synchronous lectures. Based on the pre-semester survey that was taken by 89% of the students enrolled in the course as of August 19, 2020, the following timings have been finalized for office hours of the instructor and synchronous, remote lectures (which, incidently, coincide with the original in-person course timings prior to the July 6th Rutgers decision of remote instruction).

Synchronous Zoom Lectures

Tuesdays and Thursdays: 5:00 – 6:20 PM
(Zoom links are available within the Canvas site)

Instructor Zoom Office Hours

Wednesdays: 10:30 – 11:00 AM
Fridays: 5:30 – 6:00 PM
(Zoom links are available within the Canvas site)

TA Office Hours

To Be Decided (will be announced on Canvas)

3 Course Prerequisites

While the course will motivate the covered material through the use of various engineering applications, being an engineering student (or a particular major within engineering) is *not* a pre-requisite for enrollment. However, enrolled students must have taken undergraduate courses in **probability theory** and **linear algebra**. The course will also require extensive **programming** for reinforcement of concepts introduced in the course. In keeping with the industry standards, all programming will need to be done in notebooks [e.g., Jupyter (`http://jupyter.org/`) and Google Colab (`https://colab.research.google.com/`)] using either Julia, Python, or R (individual students will get to pick any one of these languages in most assignments). In many instances, students will be forbidden from using popular machine learning packages such as `scikit-learn` for assignments.

4 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 small-scale machine learning systems from scratch
- Recognition of common pitfalls that come with machine learning systems

5 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://faculty.marshall.usc.edu/gareth-james/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.

6 Grading Policy

The final course grade will be based upon:

1. Homework (5%)
2. Mini Jupyter/Colab exercises (25%)
3. In-class exam #1 (25%)
4. In-class exam #2 (20%)
5. Term project (25%)

Grades will be assigned on a relative basis. The relative scale though will vary based upon the performance of the overall class. In an ideal setting, students above class average will get B and higher and students at or below class average will get C+ and lower, respectively. If the class performs really well, however, then the B will turn into a B+. Similarly, if the class performs really bad then the B will turn into a C+ (or even C).

7 Late Submission Policy

Every student gets an automatic submission grace period of up to 3 days for a maximum of two homeworks as well as for a maximum of two programming exercises. Utilization of the first grace period for a homework (resp., an exercise) is without any penalty. Utilization of the second grace period comes with a 30% penalty. No late submissions will be accepted from a student who has utilized both these grace periods, regardless of the emergency or unique circumstances. It is therefore advised that students avail themselves of these grace periods in true emergencies.

8 Academic Integrity Policy

It is important that the students enrolled in this class familiarize themselves with the Rutgers Academic Integrity Policy, `http://nbacademicintegrity.rutgers.edu/home/academic-integrity-policy/`, and the definition of plagiarism (`www.plagiarism.org/plagiarism-101/what-is-plagiarism/`), which includes code plagiarism. It is also important for the students to realize that pseudo-tutoring from platforms such as *Chegg* and *Course Hero* that result in solutions to homeworks, assignments, exams, etc., is serious academic misconduct. Note that all cases of academic misconduct in the course, whether minor or major, will not only be reported to the Office of Student Conduct, but will, in most cases, also result in loss of one or more letter grades.

9 General Advice for the Students

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 programming in Julia, Python, or R

Some tips for making the learning of class material easier!

- If you feel lost during the class, please make a point of seeing me in the office hours. You will be surprised to know that I do not turn into a monster during office hours `:)`.
- Because of the mathematically intensive nature of the course, one cannot learn it by forgetting about it till it is time for an exam. It is therefore important that you try to keep up with the class material on a regular basis.
- Class lectures are not enough to learn everything about the course. Reading material and homework problems will be assigned on a regular basis to help you learn all the important aspects of the course. Please make sure you keep up with these things, which will be communicated via email and via the course website.
- While the percentage of the grade assigned to homeworks and Jupyter exercises is not too large, these two categories are going to teach you the most and ensure that you do well on the exams and the term project. The purpose of keeping the percentage relatively small is that you don't feel pressured to blindly cheat from other students. You are encouraged to discuss things with others, but you will be doing yourself a big favor by doing the homeworks and exercises in the end by yourself.

10 A Tentative Course Outline

Week 0 (Pre-requisite Review)

- Review of basic probability theory and linear algebra concepts

Weeks 1–2 (Big Picture Concepts)

- Introduction to machine learning and its basic terminology
- Understanding the machine learning pipeline in production systems

Weeks 2–3 (Features)

- Feature engineering and feature/representation learning
- Principal component analysis (PCA) for feature learning

Week 4 (Basics of Machine Learning Algorithms)

- Basic building blocks of machine learning algorithms
- Mathematical basis of machine learning algorithms

Weeks 5–6 (Basic Classification Algorithms)

- Bayes' classification
- Naive Bayes' classification
- Linear discriminant analysis (LDA)
- Quadratic discriminant analysis (QDA)
- Nearest-neighbor classification

Week 7 (Evaluation of Classification Algorithms)

- Commonly used performance metrics
- Testing, training, and cross-validation

Weeks 8–10 (Additional Classification Algorithms)[1]

- Logistic regression
- Perceptron
- Linear support vector machine
- Kernel support vector machines

Weeks 10–11 (Practical Aspects of Machine Learning Systems)

- Bias–variance tradeoff, overfitting, and Occam's razor
- Basic principles of numerical optimization and gradient descent
- Privacy, ethics, and biases within machine learning systems

Weeks 11–12 (Basic Regression Algorithms)

- Least-squares regression
- Ridge regression
- Lasso regression

Weeks 13–14 (Basic Clustering Algorithms)[2]

- K-means clustering
- Gaussian mixture model
- Expectation–maximization algorithm

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