Machine learning EE 645

Narayana Santhanam

Fall 2018

Welcome to EE645!

This course is a graduate level course in machine learning.

Pre-requisites The Spring 2024 edition welcomes students from multiple departments: in addition to the traditional ECE, Computer Science, and other STEM departments (ME, CEE, Math, Astronomy), we have students from social sciences, business and urban planning. Basic familiarity with statistics, probability and some linear algebra, supplemented with a willingness to learn fundamentals, is all that is required. Ideally, students should have had some exposure to using machine learning/AI (even if it is only YouTube videos or playing with them as black boxes). The goal is for all students to pick up the central tenets of this field in a way that can guide them to use AI/machine learning in a sophisticated and nuanced fashion.

Material The course will be divided into several modules. There is no single text for the course, but there are some resources that may serve you well beyond the course:

- "Understanding machine learning: From Theory to Algorithms" by Shai Shalev-Shwartz and Shai Ben-David
- "Patterns, Predictions and Actions" by Moritz Hardt and Benjamin Recht
- "Dive into Deep Learning": online book by Zhang et. al.
- "Neural Network Learning: Theoretical Foundation" by Martin Anthony and Peter Bartlett (an old but really good book)

The following is a brief outline of the topics we will cover. Please note that this is subject to extensive change and will be modified based on class.

Soft start, Weeks 1-2 We begin with topics you may already be familiar with. This is to refresh some fundamentals, set notation, and provide some context to things we will learn, but not exactly to teach the material again. If you are not already familiar with them, please talk to me. If you are missing a couple of topics in the preliminary list, do not worry—we will help you catch up. Topics covered include linear regression, brief mentions of ridge regression, LASSO and neural networks, linear classification (Fisher discriminant), principal components analysis. New topics singular value decomposition, linear projections and the Lindenstrauss-Johnson Lemma.

Kernel methods Weeks 3-5 Kernel methods tend to be state of the art when the amount of data is limited. We will look at kernel methods in general, including support vector machines, kernel PCA and Gaussian processes. We will round this topic up with the concept of random Fourier features. There will be a mini-project on this topic.

Regularization, Week 6 ℓ_1 and ℓ_2 regularization, the Matrix norm, ResNet architecture in neural networks.

Language models/NLP Weeks 7-9 Topic models, Non-negative Matrix factorization, attention and the Transformer architecture, Large Language models. There will be a mini-project in using some of the architectures in this section.

Theory of Learning, Weeks 10-12 This covers VC dimension, PAC learnability, and PAC Bayes methods, which provide state of the art guarantees on neural networks. While this is a theoretically intense section, it is also possible to imbibe a qualitative understand of the topics.

Convolution Networks, Week 13 Modern convolution networks.

Reinforcement Learning, Week 14-15 Dynamic Programming and reinforcement learning, value iteration, deep RL.

Projects There will be no midterms, but you will have projects to work on for the duration of the course. The preferred languages are R or python. In addition to the mini-projects above, you will also do a longer term project of your choice in teams. We will also have presentations that we will do in the finals week, based partly on the project your team has done, and partly on a choice of topic you will discuss with me.

Office hours Office hours (tentatively): Fri 9-10am, Mon 10:30-11:30am.