Machine learning EE 445

Narayana Santhanam

Spring 2023

Welcome to EE445!

Material Topics will introduce the thinking and techniques behind modern statistics and machine learning. While we will use a text, the presentation in class will be different. Conceptually, we study regression, classification, principal components analysis, singular value decomposition, kernel methods and move on to more elaborate deep learning scenarios. Along the way, we will pick up basics of convex optimization, stochastic gradient descent and their accelerations in the context of neural networks. We will touch upon transformers, the big successes of today in natural language processing. The same principles go into a wide range of applications, and we will have looked at problems that will predict the efficacies of cancer treatments, work with face identification, and build recommender systems to name a few.

You will get jupyter notebooks to work on and these cover fundamentals of machine learning. They form 25% of the grade. These will help you pick up python if you are not already familiar, as well as give you basic templates you can reuse for your projects. These workbooks are help you ease into EE 445 without having taken a preliminary course highlighting the linear algebra fundamentals of machine Learning. Several initial notebooks would have already been covered in EE 345, and we will spend less time on these. They are not necessarily easy if you have just taken MATH 307, so you should take my help to complete them. A tentative list of the notebooks are as follows (we may increase or reduce or change the topics based on how we are progressing):

Workbook 1 Working with matrices in python, multiplication, inner and outer products (EE 345)

Workbook 2 Linear least squares, projections into linear spaces (EE 345)

Workbook 3 Computing significance values (t-values) of regression coefficients, uncovering "fake" features, LASSO (EE 345)

Workbook 4 Eigenvalues, eigenspaces, spectral decomposition theorem, Ridge regression

Workbook 5 Classification: Fisher LDA, LDA using linear least squares regression, QDA

Workbook 6 Classification: maximum entropy approaches (logistic regression), k-NN

Workbook 7 Single neuron networks

Workbook 8 Feedforward networks networks: Handwritten digit recognition

Workbook 9 Projecting into lower dimensions, SVD, PCA, autoencoder networks

Workbook 10 Generative Adverserial networks (GANs)

Projects, exams There will be no midterms, but you will have project/s to work on in addition to the python notebooks mentioned above. The preferred language is python, but R is a viable alternative as well (only if you are an expert already). The project/s will take up 25% of the final grade. There will be theory homeworks as well. The finals will be based on theory homeworks and account for another 25% of the final grade. At least 5% will be earmarked for class participation and presentations, and the remaining portion of the grade will be assigned to theory homeworks.

Credits 3

Class meetings/week 3 lecture hours

Textbook Patterns, predictions and actions (M. Hardt and B. Recht). The preprint of the book is available (legally) as a pdf download from the authors' website. The final publication version is available as a download from the UH Manoa library.

Optional references Neural networks and Deep Learning (Aggarwal), Introduction to statistical learning (James, Witten, Hastie, Tibshirani), Linear algebra and learning from data (Gilbert Strang).

Grading Python workbooks (25%), Final (25%), Project (25%) (25% TBD: split between theory homeworks and presentations)