Syllabus: EEE 498/591 Foundations of Machine Learning: Theory to Practice (Fall 2020)

Instructor: Lalitha Sankar

Contact Info: email: `lsankar@asu.edu` Office: ERC 585

Meeting Info: Tue, Thu 3:00pm–4:15pm, Location: Hayden Library C34

Lab Hours: TBD

(1) Are you a student who wants a rigorous and detailed introduction to machine learning?

(2) Are you also eager to see how to apply ML methods to datasets using Python?

If you answered yes to these questions then the four-credit course EEE498/EEE591: Foundations of Machine Learning: From Theory to Practice is the course for you. EEE498 is a course for graduating seniors as well as juniors who have a background in linear algebra and some understanding of probability theory. EEE591 is for graduate students with the same background requirements. The 4-credit course also involves a lab component with implementation and testing of the algorithms taught in class. Knowledge of Python is useful but not necessary and will be introduced in the lab.

Important note: The one-credit lab **will not** include additional work; the lab and course projects are the **same** – the lab provides opportunities to help solidify understanding of the algorithms and mathematical models taught in the class **while including access to a TA**.

Course Description: Machine learning explores the design, analysis, and construction of algorithms that can learn from data and make inferences or predictions about future outcomes. The focus is on a **methodical approach** that will highlight the role of statistical and computational methods in analysis of data. This course includes a **near equal dose of theory and practice** with the goal of providing a thorough grounding in the fundamental methodologies and algorithms in machine learning. The focus will be a methodical way of learning that begins from the theoretical underpinnings of machine learning focused broadly on two distinct types of learning methods, namely supervised and unsupervised learning. Within each type, various well-studied and formulated approaches will be studied. The one-credit lab focuses on the implementation aspects and is Python-based.

Course outcome: students should be able to apply ML techniques to a variety of engineering problems.

Course Topics:

- Introduction to machine learning, review of probability and linear algebra
- Supervised learning; linear regression
- Weighted least squares; logistic regression
- Perceptron and general linear models
- Support Vector Machines, Kernels
- Unsupervised learning: k-means, expectation maximization
- Principal and independent component analysis
- Decision trees, boosting, bagging
- Bias and variance, learning theory, deep learning

Prerequisites: Linear algebra and **some** knowledge of probability theory (all relevant probability concepts will be covered in class)

Textbook and Reference Materials:

[Murphy] *Machine Learning: A Probabilistic Perspective*, Kevin Murphy.

Supplementary/Reference material: For a gentler introduction to machine learning; free online:

- A Course in Machine Learning by Hal Daume III
- Pattern Recognition and Machine Learning, Christopher Bishop.
- The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Trevor Hastie, Robert Tibshirani, Jerome Friedman.
- Computer Age Statistical Inference: Algorithms, Evidence and Data Science, Bradley Efron, Trevor Hastie.

You may also find these reference materials useful throughout the semester.

Machine Learning (and related topics)

- Crib sheet of math for ML by Iain Murray
- Understanding Machine Learning: From Theory to Algorithms, Shai Shalev-Shwartz, Shai Ben-David. An introduction to theoretical machine learning.
- Foundations of Data Science, by Avrim Blum, John Hopcroft and Ravi Kannan. This freely available pdf has nice chapters on machine learning (chapter 5), clustering (chapter 7) and SVD (chapter 3).

Linear Algebra and Matrix Analysis

- These wonderful videos by 3blue1brown provide a gentle and highly intuitive overview of linear algebra.
- Linear Algebra Review and Reference by Zico Kolter and Chuong Do (free). Light refresher for linear algebra and matrix calculus if you're a bit rusty.
- Linear Algebra, David Cherney, Tom Denton, Rohit Thomas and Andrew Waldron (free). Introductory linear algebra text.
- Matrix Analysis Horn and Johnson. A great reference from elementary to advanced material.

Probability and Statistics

- Probability Review by Arian Maleki and Tom Do. (From Andrew Ng's machine learning class.)
- Sankar's notes (on canvas) on probability (please request access)
- All of Statistics, Larry Wasserman. Chapters 1-5 are a great probability refresher and the book is a good reference for statistics.
- A First Course in Probability, Sheldon Ross. Elementary concepts (previous editions are a couple bucks on Amazon)

Optimization

- Numerical Optimization, Nocedal, Wright. Practical algorithms and advice for general optimization problems.
- Convex Optimization: Algorithms and Complexity, Sbastien Bubeck. Elegant proofs for the most popular optimization procedures used in machine learning.

Python

- www.learnpython.org "Whether you are an experienced programmer or not, this website is intended for everyone who wishes to learn the Python programming language."
- Convex Optimization: Algorithms and Complexity, Sbastien Bubeck. Elegant proofs for the most popular optimization procedures used in machine learning.
- NumPy for Matlab users

Latex

- Learn Latex in 30 minutes
- Overleaf. An online Latex editor.
- Standalone Latex editor on your local machine
- Detexify LaTeX handwritten symbol recognition
- Latex Math symbols