Syllabus: EEE 598 Statistical Machine Learning: Theory to Practice (Fall 2022)

Instructor: Lalitha Sankar

Contact Info: email: lsankar@asu.edu   Office: GWC 436

Office Hours: Tue 9-10 am, Thu 11:00 am –12:00 pm: https://asu.zoom.us/j/7529435385 (online via zoom)

Meeting Info: Tue, Thu 3:00pm–4:15pm, Location: STPV-324

(Unofficial) TA Info:

TA: Nathan Stromberg; Monica Welfert; email: {nstrombe,mwelfert}@asu.edu

EEE598: Statistical Machine Learning: From Theory to Practice is a three-credit course introducing the theory and practice of machine learning (ML). EEE598 is a course for graduate students who have some background in linear algebra and a basic understanding of probability theory (all relevant probability concepts will be covered in class). This 3-credit course focuses on both theoretical concepts and practical implementation and testing of machine learning methods and algorithms. Knowledge of Python is useful but not necessary; while it won't be taught formally, students are expected to pick it up as a part of homework assignments (some office hours will be dedicated every week to help with programming for the homework).

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 aim of this course is to introduce students across engineering to basic data science concepts and algorithms in a rigorous manner. A desired outcome is for students to be able to learn to distinguish between different algorithms and determine which methods are ideal for a problem setting at hand.

Evaluations will be based on a combination of homework assignments and a term paper project. Homework assignments are a combination of applied and basic mathematical concepts as well as programming problems. The culminating experience of this course is via a term paper project. This is where students can leverage their own academic backgrounds to identify projects relevant to their study/research and determine algorithms (learned in class or building upon concepts learned in the course) that may work in those settings. Understanding and implementing published papers and using public datasets is a big part of this learning process.

Grading:	Homework Assignments (5 in total)	60 %
	Class/Slack Participation	5 %
	Homework 0 (review; not in reg. 5)	10 % **bonus**
	Final Project (two Phases)	35 %

Final project will be broken into two phases, an intent phase where students (in groups of 2-3) do preliminary research and submit a brief report on the topic they wish to focus on. This is due mid-October (can be viewed as midterm) and counts for 25% of the final project grade. Graded review from this will help with the final project report which will be due at the end of the semester (during exam week) which will count for 75% of the final project grade.

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; Bias and variance
- Weighted least squares; logistic regression
- Perceptron and general linear models
- Support Vector Machines, Kernels
- p-values, hypothesis testing, Bayesian inference
- Unsupervised learning: k-means, expectation maximization
- Principal and independent component analysis
- Decision trees, boosting, bagging
- learning theory, deep learning

Prerequisites: Linear algebra and **an elementary** knowledge of probability/statistics (all relevant probability concepts will be covered in class) **Co-requisite:** EEE554 or equivalent (email Professor for clarification, if needed)

Textbook and Reference Materials: (both are free online)

[HTF] *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Trevor Hastie, Robert Tibshirani, Jerome Friedman.

[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.
- Computer Age Statistical Inference: Algorithms, Evidence and Data Science, Bradley Efron, Trevor Hastie.

Watt *Machine Learning Refined: Foundations, Algorithms, Applications*, Watt, Borghani, Katsaggelos, Cambridge University Press

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.

- TexPad for both OSX and Windows
- Standalone Latex editor on your local machine
- Detexify LaTeX handwritten symbol recognition
- Latex Math symbols