ISyE 8803 – Topics in Statistical Machine Learning, Fall 2021

Tentative Syllabus

Logistics

The course aims to cover a list of research advances in statistical machine learning, to build literature and training to perform research in the area.

Class Time and Location: MW, 3:30pm-4:45pm, ISyE Main 126.

Instructor: Prof. Yao Xie, Groseclose #339, email: yao.xie@isye.gatech.edu
Instructor Office Hour: M 4:45pm-5:30pm.

Class Material: All course materials will be made available on Canvas.

Prerequisites: ISyE 6416, 6740, 6412 or permission of the instructor.

Honor Code: For any question involving Academic Honor Code issues, please consult www.honor.gatech.edu

Grading Policy: Course notes scribe - 50%, Project - 50%.

Scribe: Each student will sign up for one lecture that requires a scribe. Please use the template provided. The student also needs to come up with one numerical example for that lecture (a simulation or real-data example, written in Python). The quality of the scribe will be graded by the professor.

Project: by group, each group consists of 1-2 students. Proposal: At least 2 pages write-up using the provided template, including at least: project description, objective, planned approach and expected outcome. Final report: at least 4 pages report using the provided template.

Topics

(8/23) Introduction: Basics of ML and statistics

- estimation and hypothesis test, exponential family, supervised vs. unsupervised, model training and generalization

Theme 1: Advanced statistical model and methods

Low-rank models for high-dimensional data

- (8/25) Sparse, low-rank models
- (8/30) Advanced low-rank and tensor recovery (Guest Lecture: Mark Davenport, GT)
- (9/1) Manifold data learning (Guest Lecture: Xiuyuan Cheng, Duke)

Kernel methods

- (9/8) Nearest neighbor, knn
- (9/13) Kernel density estimation (KDE)
- (9/15) Kernel regression
- (9/20) Gaussian processes

Two-sample tests
- (9/22) Information divergence (KL, MMD, Wasserstein, sink-horn)
- (9/27) MMD two-sample test
- (9/29) Stein method

Neural networks theory - supervised
- (10/4) NN approximation (Guest Lecture: Alex Cloninger, UCSD)
- Training dynamics:
    - (10/6) Neural Tangent Kernel (NTK)
    - (10/13) Mean-field
- (10/18) Introduction to deep neural networks and training
- (10/20) Convolutional neural networks and graph neural networks

Sequential data
- (10/25) Spatio-temporal data modeling, point processes
- (11/1) Sequential data modeling, tracking, prediction
- (11/3) Sequential change detection

Causal inference (11/8)

<u>Theme 2: Algorithmic development</u>

Conformal inference and UQ
- (11/10) Conformal prediction

Modern optimization for learning
- (11/15, 11/17) Online convex optimization and first-order methods
- (11/22) Estimation of GLM using VI
- (11/29) Minimax and saddle point problem (Guest lecture: Jimmy Zhang/George Lan, GT)
- (12/1) Transfer Learning (Guest lecture: Tuo Zhao, GT)

Final presentations (12/6)

References

Review
http://www.stat.cmu.edu/~larry/=sml/Review.pdf

K-nearest neighbor
https://sebastianraschka.com/pdf/lecture-notes/stat479fs18/02_knn_notes.pdf

Causal inference
https://blog.ml.cmu.edu/2020/08/31/7-causality/
http://www.stat.cmu.edu/~larry/=sml/Causation.pdf

Kernel methods
http://www.gatsby.ucl.ac.uk/~gretton/coursefiles/rkhscourse.html

Conformal inference
https://36708.github.io/Lectures/feb16-class.pdf
http://www.stat.cmu.edu/~larry/=sml/Conformal

Optimal transport and wasserstein
http://www.stat.cmu.edu/~larry/=sml/Opt.pdf

Two-sample test
http://www.stat.cmu.edu/~larry/=sml/TwoSample.pdf

Manifold and dimensionality reduction
http://www.stat.cmu.edu/~larry/=sml/DimRed.pdf

Convex optimization for neural networks
http://web.stanford.edu/class/ee364b/lectures/convexNN.pdf

Neural network approximation
https://arxiv.org/pdf/2105.04026.pdf