Introduction to Statistical Inference and Learning

Instructors: Yiqing Wang, Boaz Nadler

In this course we will cover the basic concepts underlying modern data analysis, machine learning and statistical inference. Subject to time constraints, topics covered will include

1. Basic probability, inequalities, classical distributions.
2. Basic information theory, entropy, KL divergence.
3. Parameter estimation, maximum likelihood, Bayesian approaches, MAP, MMSE, Empirical risk minimization, consistency, Cramer-Rao lower bound (information inequality).
4. Parametric and non-parametric models.
5. density estimation, kernel smoothing.
6. The bias-variance tradeoff and the curse of dimensionality in high dimensional problems.
7. Stein's phenomenon.
8. Hypothesis testing and the Neyman-Pearson lemma.
9. Principal Component Analysis, dimensionality reduction.
10. Markov Models and Hidden Markov Models.
11. Some statistical challenges related to big data.
12. Supervised and unsupervised ensemble learning.

Suggested Reading:

1. L. Wasserman, *All of statistics*, Springer. (also take a look at *all of non-parametric statistics* by the same author.)
2. T. Hastie, R. Tibshirani and J. Friedman, *The elements of statistical learning,* Springer.
3. V. Vapnik, *The nature of statistical learning theory*, Springer, 2nd edition.
4. K. Knight, *Mathematical Statistics.*
5. C. Bishop, *Pattern Recognition and Machine Learning.*