Introduction to Statistical Inference and Learning

Instructors: Yaniv Tenzer, 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. Hypothesis testing, Neyman-Pearson lemma, multiple hypothesis testing, FDR.
8. Principal Component Analysis, dimensionality reduction
9. Regression in low and high dimensions
10. Sparsity and compressed sensing
11. 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. C. Giraud, *Introduction to high-dimensional statistics*, CRC.
4. K. Knight, *Mathematical Statistics.*
5. C. Bishop, *Pattern Recognition and Machine Learning.*
6. Thomas and Cover, *Elements of Information Theory.*

Grading:

Throughout the course we will give 4 homework exercises. It is *mandatory* to submit at least 75% of them.

There will also be a final exam.

The grade in the course consists of a weighted average: 60% exam + 40% HW exercises.