GR 8201(3), (Spring 2019) Topics in Theoretical Statistics
Recent Statistical Insights in Kernel Machines

Schedule

Time: Tuesdays 2:10-4pm, **Location:** 1025 SSW.

Instructor: Samory Kpotufe. *email:* skk2175@columbia.edu
Office hours: Monday 1:45 pm to 3pm, Department of Statistics, Office 911.

Description:

Kernel Machines (e.g., SVMs, Kernel Ridge, Gaussian Processes, Kernel PCA, ...) remain one of the most successful and better understood methods in Machine Learning. Much recent research effort has gone into improving their ability to scale to large amounts of high-dimensional data, while retaining their statistical accuracy whenever possible. This has called for refined understanding of their statistical properties, rooted in mathematical tools from functional analysis and operator theory (towards capturing the statistical *complexity* of the functions spaces and linear operations involved). In contrast the initial statistical understanding from Machine Learning Theory relies on so-called *margin conditions* which seem to only capture part of the picture, especially given the apparent complexity of modern data.

The first few weeks (3-6) will serve to develop the mathematical tools used in much recent work (initiated by folks such as S. Smale, I. Steinwart, M. Pontil, T. Poggio, and co-authors), followed by paper presentations (by students) and invited speakers (2-3) working at the boundary of this field.

Basic background

While we'll try to have self-contained discussions, familiarity with the following will be helpful.

- Basic probability concepts, e.g., measurability, integration, characterictic functions, $\mathcal{L}_p(\mu)$ spaces, ...
- Basic Linear Algebra, e.g., vector spaces, Spectral and Singular Value theorems, ...
- Basic Real Analysis, e.g. completeness, compactness, forms of continuity, ...
- Basic Statistical concepts, e.g., ℓ_p Risks, Regularization, basic concentration inequalities such as Chernoff, Bernstein, ...

Grading

Paper Presentations (99%)**.** Papers will be assigned early, to be presented by one person or a group. I will be available to help guide reading and eventual presentation. For maximum benefit, everyone will be expected to read all papers, and get into the details of those papers they are presenting.

Projects (1%)**.** Roughly one or two group projects whose results are to be presented at the end of the semester. These would mostly consists of simple programming assignments that test the theoretical insights developed in the rest of the course.