Anstatt der Vorlesung am Dienstag, den 10. Dezember, sollen alle Studenten bitte in den Vortrag von Mikhail Belkin gehen.

Ankündigung auf der Rückseite.

UNIVERSITÄT
LEIPZIG

Prof. Dr. Mikhail Belkin
The Ohio State University

From classical statistics to modern AI: the lessons of deep learning

Abstract

For years, Empirical Risk Minimization (ERM) served as the main theoretical foundation for machine learning. Yet, even early on, certain phenomena, such as lack of over-fitting in boosting, posed a challenge to the ERM view.

In recent years these foundational cracks have widened precipitously, as best practices of deep learning were in direct contradiction to the methodologies suggested by analyses of ERM or bias-variance trade-off. It has even been suggested that empirical successes of deep learning obviate the need for mathematical analyses in understanding data. This is not the case. As I will discuss, modern methods operate in regimes, mostly overlooked in the past that can be identified and usefully analyzed with existing mathematical tools.

These regimes are best described by interpolation, where the model is selected from the space of predictors fitting the training data exactly, according to a certain inductive bias. Furthermore, I will show how classical (capacity based) models and modern models based on inductive biases can be unified within a single "double descent" risk curve, which subsumes the classical U-shaped bias-variance trade-off.

Finally, I will discuss some important implications for optimization, and comment on the nature of the inductive biases observed in deep learning.

10. December 2019

11:00 Uhr

Lecture Hall: Felix-Klein-Hörsaal, Paulinum P 501, Augustusplatz 10

Followed by a reception