

Learnability Beyond Uniform Convergence

Shai Shalev

Hebrew University

The problem of characterizing learnability is the most basic question of statistical learning theory. A fundamental result is that learnability is equivalent to uniform convergence of the empirical risk to the population risk, and that if a problem is learnable, it is learnable via empirical risk minimization. The equivalence of uniform convergence and learnability was formally established only in the supervised classification and regression setting. We show that in (even slightly) more complex prediction problems learnability does not imply uniform convergence. We discuss several alternative attempts to characterize learnability.

The presentation is based on a joint research with Ohad Shamir, Nati Srebro, Karthik Sridharan, and with Amit Daniely, Sivan Sabato, and Shai Ben-David.