

AdaBoost and robust one-bit compressed sensing

Matthias Löffler

ETH Zürich, Seminar for Statistics, Switzerland

AdaBoost is a popular and successful classification algorithm from the machine learning literature. Empirically, it often performs best, when it is overparameterized and run long after it interpolates the data. Previous attempts to explain the success of AdaBoost do not take these features into account or do not yield consistent estimators. By contrast, we prove error bounds by explicitly exploiting overparameterization and data interpolation by relating AdaBoost to the max-l1-margin estimator and studying the latter through the lens of one-bit robust compressed sensing. In particular, this provides an explanation why interpolating adversarial noise can be harmless for classification problems. This is joint work with G. Chinot, F. Kuchelmeister and S. van de Geer.