

Learning with Information Constraints

Ohad Shamir, Weizmann Institute

Abstract:

Many machine learning approaches are characterized by information constraints on how they interact with training data. One prominent example is communication constraints in distributed learning, while other examples include memory and sequential access constraints in stochastic optimization and online algorithms, and partial access to the underlying data in sequential decision making under uncertainty (e.g. multi-armed bandits). While the statistical effect of limited training data is well-understood, the statistical implications of information constraints internal to the algorithm used are relatively little explored. In this talk, I'll discuss these issues, and describe how a single set of results sheds some light on them, simultaneously for different types of information constraints and for different learning settings.