From Average Case Complexity to Improper Learning Complexity

Amit Daniely, The Hebrew University

Abstract

It is presently still unknown how to show hardness of learning problems. There are still huge gaps between our upper and lower bounds in the area. The main obstacle is that standard NP-reductions do not yield hardness of improper learning (a.k.a. representation-independent learning). All known lower bounds on improper learning rely on (unproved) cryptographic assumptions.

We introduce a new technique to this area, using reductions from problems that are hard on average. Under Feige's assumption (2002), we show that in learning half-spaces over sparse vectors, more examples reduce the training time. Our newly formulated (and strong) generalization of Feige's assumption yields far reaching implications, such as:

1. Learning DNFs is hard.
2. Agnostically learning half-spaces with a constant approximation ratio is hard.
3. Learning an intersection of any super constant number of half-spaces is hard.

Joint work with Nati Linial and Shai Shalev-Shwartz.