מכון ויצמן למדע

THE WEIZMANN INSTITUTE OF SCIENCE

FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Machine Learning and Statistics Seminar

Room 1 ,Ziskind Building
on Wednesday, Dec 16, 2015
at 11:15

Tomer Koren
Technion

The Computational Power of Optimization in Online Learning

Abstract:

We consider the fundamental problem of prediction with expert advice where the experts are "optimizable": there is a black-box optimization oracle that can be used to compute, in constant time, the leading expert in retrospect at any point in time. In this setting, we give a novel online algorithm that attains vanishing regret with respect to N experts in total $\sqrt{N}$ computation time. We also give a lower bound showing that this running time cannot be improved (up to log factors) in the oracle model, thereby exhibiting a quadratic speedup as compared to the standard, oracle-free setting where the required time for vanishing regret is linear in N. These results demonstrate an exponential gap between the power of optimization in online learning and its power in statistical learning: in the latter, an optimization oracle---i.e., an efficient empirical risk minimizer---allows to learn a finite hypothesis class of size N in time $\log{N}$.

We also study the implications of our results to learning in repeated zero-sum games, in a setting where the players have access to oracles that compute, in constant time, their best-response to any mixed strategy of their opponent. We show that the runtime required for approximating the minimax value of the game in this setting is $\sqrt{N}$, yielding again a quadratic improvement upon the oracle-free setting, where linear time in N is known to be tight.