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, Jan 08, 2020
at 11:15

Yair Carmon
Stanford

Fundamental limits of modern machine learning and how to get around them

Abstract:

This talk presents new computational and statistical barriers in machine learning, along with the algorithmic developments that they inspire.

The computational barriers arise in nonconvex optimization: we prove lower bounds on the (oracle) complexity of finding stationary points using (stochastic) gradient methods, showing that gradient descent is unimprovable for a natural class of problems. We bypass this barrier by designing an algorithm that outperforms gradient descent for a large subclass of problems with high-order smoothness. Our algorithm leverages classical momentum techniques from convex optimization using a "convex until proven guilty" principle that we develop.

The statistical barrier is the large amount of data required for adversarially robust learning. In a Gaussian model, we prove that unlabeled data allows us to circumvent an information theoretic gap between robust and standard classification. Our analysis directly leads to a general robust self-training procedure; we use it to significantly improve state-of-the-art performance on the challenging and extensively studied CIFAR-10 adversarial robustness benchmark.