

Title: Elephant in the Room: Non-Smooth Non-Convex Optimization

Authors: Ohad Shamir

Abstract: It is well-known that finding global minima of non-convex optimization problems is computationally hard in general. However, the problem of finding local minima-like points (at least in terms of gradient and Hessian properties) is tractable, and received much attention in the machine learning community in recent years. The resulting literature has been largely motivated by the rising importance of non-convex optimization problems such as deep learning, but in fact, does not quite address them: Nearly all computationally efficient guarantees in this area require the objective function to be smooth, which is seldom satisfied in deep learning problems. This highlights the importance of understanding what we can do efficiently on such non-convex, non-smooth optimization problems. In this talk, I'll describe some results, challenges, and possible approaches to tackle this fundamental question.