Is non-convex optimization really hard? A couple of recent stories

Emmanuel J. Candes, Stanford University candes@stanford.edu

In recent years, there has been astounding progress in the theory and practice (algorithms, professionalgrade software development, applications) of convex optimization to the point that it has become a real pillar of modern engineering. On the other hand, the field of non-convex optimization is far less mature and may draw comparisons with 17th century medicine (ad-hoc methods, no performance guarantees, unreliable results, and so on). This is unfortunate because most problems of interest to information scientists are non-convex in nature; e.g. many maximum likelihood estimates are, in fact, solutions to non-convex problems, some of which being notoriously hard. This talk will briefly review a rapidly emerging literature showing that, perhaps surprisingly, some important non-convex problems may not be as hard as they seem. We will discuss some of this exciting research emphasizing applications in signal and image processing such as phase retrieval, and in machine learning such as low-rank factorization.

Joint work with: Yuxin Chen Princeton University (United States)