Computer Science Theory Seminar

Understanding statistical-vs-computational tradeoffs via the low-degree likelihood ratio

Alex Wein (Courant Institute)

Abstract: High-dimensional inference problems such as sparse PCA and planted clique often exhibit statistical-vs-computational tradeoffs whereby there is no known polynomial-time algorithm matching the performance of the optimal estimator. I will discuss an emerging framework – based on the so-called low-degree likelihood ratio – for precisely predicting these tradeoffs and giving rigorous evidence for computational hardness in the conjectured hard regime. This method was originally proposed in a sequence of works on the sum-of-squares hierarchy, and the key idea is to study whether or not there exists a low-degree polynomial that succeeds at a given statistical task.

In the second part of the talk, I will show how to use the above framework to give new results for the sparse PCA problem. Here the goal is to recover a rank-1 signal $xx^\top$ planted in a random matrix (either Wigner or Wishart), where x is rho-sparse (rho fraction of entries nonzero). Polynomial-time algorithms are known when rho << 1/sqrt(n) and naive exhaustive search succeeds when rho << 1; however, no efficient algorithm is known when rho >> 1/sqrt(n). We explore precisely how hard the "hard" regime is by showing that for any rho >> 1/sqrt(n) there is a subexponential-time algorithm of runtime exp(rho^2 n), and the low-degree likelihood ratio suggests that this is optimal. In contrast, naive exhaustive search has runtime exp(rho n).

Wednesday, November 20 at 4:15 PM in 1325 SEO

Based on joint work with Afonso Bandeira, Yunzi Ding, and Tim Kunisky (https://arxiv.org/abs/1907.11635 and https://arxiv.org/

Wednesday, November 20 at 4:15 PM in 1325 SEO