

The hidden submatrix problem

Andrea Montanari

Stanford University - US

Abstract:

Given a large random matrix with independent entries, the hidden submatrix problem requires to find a small subset of its rows and columns, such that the corresponding entries have mean larger than the others.

This problem arises in a variety of applications ranging from genomics to computer vision and signal processing.

Two natural questions arise in this context. First: under which conditions on the entries distributions the submatrix can be identified reliably? Second: is there an efficient (polynomial) algorithm that allows to identify it? It turns out that this problem presents a rich interplay between statistical limits, and computational barriers.

I will present recent progress, and connections with related questions.

[Based on joint work with Yash Deshpande]