Best* Case Approximability of Sparse PCA

refusing to graduate :-)

Aviad Rubinstein (UC Berkeley)

Joint work with Siu-On Chan and Dimitris Papailliopoulos

is on the job market!

Sparse Principal Component Analysis

$$\max x^\top A x$$

$$\text{s.t. } \|x\|_2 = 1 \text{ and } \|x\|_0 \leq k$$

(and A is PSD)

$$\begin{pmatrix} 7.1 & 1.3 & \cdots & 4.5 \\ -2.6 & -3.4 & \cdots & 6.2 \\ \vdots & \vdots & \ddots & \vdots \\ 3.1 & 9.2 & \cdots & -4.8 \end{pmatrix}$$

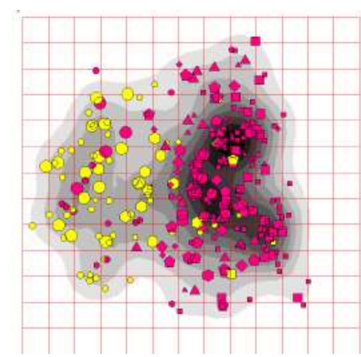

"Yes we SPCA!"

-Obama, 2008

Sparse Spiked Covariance Model ("average case")

$$A = I_n + \begin{pmatrix} \begin{bmatrix} \text{rank } 1 \\ k \times k \text{ block} \end{bmatrix} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 0 \end{pmatrix} \text{ +noise}$$

- Cool sample-complexity / computational-complexity tradeoff [Berthet & Rigollet '13, Wang et al '14, Gao et al '15, Kraugthgamer et al '15, Ma & Wigderson '15]

- Good news: trivial algorithm gives $(1 - o(1))$-approximation
- Bad news: this is not what *your data* looks like!

Our results: "best-case" analysis

Computationally intractable even when given the exact covariance matrix:

- NP-hard to approximate to within $(1-\epsilon)$
- SSE-hard to approximate to within any c
- Quasi-quasi-poly ($e^{e^{\sqrt{\ln \ln n}}}$) integrality gap
- $n^{-1/3}$-approximation algorithm

"Vive la SPCA!"

-Napoleon, 1808