# Computational Barriers in Statistical Estimation 

Andrea Montanari

Stanford University

$$
\text { July 9, } 2017
$$

## Statistical estimation/Statistical learning

Class of models $\left(\Theta \subseteq \mathbb{R}^{d}\right)$

$$
\mathcal{C}_{\Theta} \equiv\left\{\mathbb{P}_{\boldsymbol{\theta}}: \quad \boldsymbol{\theta} \in \Theta\right\}
$$

Data

$$
x_{1}, x_{2}, \ldots, x_{n} \sim_{i i d} \mathbb{P}_{\boldsymbol{\theta}_{0}}(\cdot)
$$

Estimate $\boldsymbol{\theta}_{0}$ from data $\boldsymbol{x}_{1}^{n}=\left(x_{1}, \ldots, x_{n}\right)$

## Statistical estimation/Statistical learning

Class of models $\left(\Theta \subseteq \mathbb{R}^{d}\right)$

$$
\mathcal{C}_{\Theta} \equiv\left\{\mathbb{P}_{\boldsymbol{\theta}}: \quad \boldsymbol{\theta} \in \Theta\right\}
$$

Data

$$
x_{1}, x_{2}, \ldots, x_{n} \sim_{i i d} \mathbb{P}_{\boldsymbol{\theta}_{0}}(\cdot)
$$

Estimate $\boldsymbol{\theta}_{0}$ from data $\boldsymbol{x}_{1}^{n}=\left(\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{n}\right)$

## Minimax theory

## Loss function

$$
\begin{aligned}
& L: \mathbb{R}^{p} \times \mathbb{R}^{p} \rightarrow \mathbb{R} \\
& \quad\left(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_{0}\right) \mapsto L\left(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_{0}\right)
\end{aligned}
$$

Minimax risk

$$
R_{n}^{*}(\Theta)=\inf _{\hat{\boldsymbol{\theta}}(\cdot)} \sup _{\boldsymbol{\theta}_{0} \in \Theta} \mathbb{E}_{\boldsymbol{\theta}_{0}} L\left(\hat{\boldsymbol{\theta}}\left(\boldsymbol{x}_{1}^{n}\right), \boldsymbol{\theta}_{0}\right)
$$

[Wald, 1950]

## Kindergarten example

$$
\begin{aligned}
& \mathbb{P}_{\boldsymbol{\theta}}(\cdot)=\mathrm{N}\left(\boldsymbol{\theta}, \mathrm{I}_{d}\right) \\
& L\left(\boldsymbol{\theta}_{0}, \hat{\boldsymbol{\theta}}\right)=\left\|\boldsymbol{\theta}_{0}-\hat{\boldsymbol{\theta}}\right\|_{2}^{2} \\
& \Theta=\mathbb{R}^{d} .
\end{aligned}
$$



- Foundation of the least squares, maximum likelihood,


## Kindergarten example

$$
\begin{aligned}
& \mathbb{P}_{\boldsymbol{\theta}}(\cdot)=\mathrm{N}\left(\boldsymbol{\theta}, \mathrm{I}_{d}\right) \\
& L\left(\boldsymbol{\theta}_{0}, \hat{\boldsymbol{\theta}}\right)=\left\|\boldsymbol{\theta}_{0}-\hat{\boldsymbol{\theta}}\right\|_{2}^{2} \\
& \Theta=\mathbb{R}^{d} .
\end{aligned}
$$

Theorem

$$
R_{n}^{*}(\Theta)=\frac{d}{n}
$$

- Foundation of the least squares, maximum likelihood,


## Kindergarten example

$$
\begin{aligned}
& \mathbb{P}_{\boldsymbol{\theta}}(\cdot)=\mathrm{N}\left(\boldsymbol{\theta}, \mathrm{I}_{d}\right) \\
& L\left(\boldsymbol{\theta}_{0}, \hat{\boldsymbol{\theta}}\right)=\left\|\boldsymbol{\theta}_{0}-\hat{\boldsymbol{\theta}}\right\|_{2}^{2} \\
& \Theta=\mathbb{R}^{d} .
\end{aligned}
$$

Theorem

$$
R_{n}^{*}(\Theta)=\frac{d}{n}
$$

- Foundation of the least squares, maximum likelihood, ...


## A more sophisticated example

$$
\begin{aligned}
& \mathbb{P}_{\boldsymbol{\theta}}(\cdot)=\mathrm{N}\left(\boldsymbol{\theta}, \mathrm{I}_{d}\right) \\
& L\left(\boldsymbol{\theta}_{0}, \hat{\boldsymbol{\theta}}\right)=\left\|\boldsymbol{\boldsymbol { \theta } _ { 0 }}-\hat{\boldsymbol{\theta}}\right\|_{2}^{2} \\
& \Theta=\left\{\boldsymbol{\theta} \in \mathbb{R}^{d}:\|\boldsymbol{\theta}\|_{0} \leq s_{0}\right\} .
\end{aligned}
$$

Theorem (Donoho, Johnstone 1990s) If $s_{0} / d \rightarrow 0$, then

$$
R_{n}^{*}(\Theta)=\frac{2 s_{0}}{n} \log \left(d / s_{0}\right) \cdot(1+o(1)) .
$$

## A more sophisticated example

$$
\begin{aligned}
& \mathbb{P}_{\boldsymbol{\theta}}(\cdot)=N\left(\boldsymbol{\theta}, \mathrm{I}_{d}\right) \\
& L\left(\boldsymbol{\theta}_{0}, \hat{\boldsymbol{\theta}}\right)=\left\|\boldsymbol{\theta}_{0}-\hat{\boldsymbol{\theta}}\right\|_{2}^{2} \\
& \Theta=\left\{\boldsymbol{\theta} \in \mathbb{R}^{d}:\|\boldsymbol{\theta}\|_{0} \leq s_{0}\right\} .
\end{aligned}
$$

Theorem (Donoho, Johnstone 1990s)
If $s_{0} / d \rightarrow 0$, then

$$
R_{n}^{*}(\Theta)=\frac{2 s_{0}}{n} \log \left(d / s_{0}\right) \cdot(1+o(1)) .
$$

- Key role in compressed sensing, sparse learning,


## A more sophisticated example

$$
\begin{aligned}
& \mathbb{P}_{\boldsymbol{\theta}}(\cdot)=\mathrm{N}\left(\boldsymbol{\theta}, \mathrm{I}_{d}\right) \\
& L\left(\boldsymbol{\theta}_{0}, \hat{\boldsymbol{\theta}}\right)=\left\|\boldsymbol{\theta}_{0}-\hat{\boldsymbol{\theta}}\right\|_{2}^{2} \\
& \Theta=\left\{\boldsymbol{\theta} \in \mathbb{R}^{d}:\|\boldsymbol{\theta}\|_{0} \leq s_{0}\right\} .
\end{aligned}
$$

Theorem (Donoho, Johnstone 1990s)
If $s_{0} / d \rightarrow 0$, then

$$
R_{n}^{*}(\Theta)=\frac{2 s_{0}}{n} \log \left(d / s_{0}\right) \cdot(1+o(1))
$$

- Key role in compressed sensing, sparse learning, ...


## What is this talk about?

$$
R_{n}^{\text {Poly }}(\Theta)=\inf _{\hat{\boldsymbol{\theta}}(\cdot) \in \text { Poly }} \sup _{\boldsymbol{\theta}_{0} \in \Theta} \mathbb{E}_{\boldsymbol{\theta}_{0}} L\left(\hat{\boldsymbol{\theta}}\left(\boldsymbol{x}_{1}^{n}\right), \boldsymbol{\theta}_{0}\right)
$$

Developments

- Often we expect $R_{n}^{\text {Poly }}(\Theta) \geq R_{n}(\Theta)$
- Accurate predictions
- Convergence of fields (Statistics, CS Theory, Physics)
- New algorithms?


## What is this talk about?

$$
R_{n}^{\text {Poly }}(\Theta)=\inf _{\hat{\boldsymbol{\theta}}(\cdot) \in \text { Poly }} \sup _{\boldsymbol{\theta}_{0} \in \Theta} \mathbb{E}_{\boldsymbol{\theta}_{0}} L\left(\hat{\boldsymbol{\theta}}\left(\boldsymbol{x}_{1}^{n}\right), \boldsymbol{\theta}_{0}\right)
$$

## Developments

- Often we expect $R_{n}^{\text {Poly }}(\Theta) \geqslant R_{n}(\Theta)$
- Accurate predictions
- Convergence of fields (Statistics, CS Theory, Physics)
- New algorithms?


## Outline

(1) Local algorithms
(2) Landscapes
(3) SDP relaxations
(4) Conclusion

Local and message passing algorithms

## Example \#1: Sparse low-rank matrix

## Unknowns:

$$
\Theta\left(s_{0}, d\right)=\left\{\boldsymbol{\theta} \in\{0,1\}^{d}:\|\boldsymbol{\theta}\|_{0}=s_{0}\right\} .
$$

## Loss:

$$
L(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta})=\frac{1}{d} \operatorname{Hamming}(\hat{\boldsymbol{\theta}}, \boldsymbol{\theta})
$$

## Example \#1: Sparse low-rank matrix

Data: $\left(x_{1}, \ldots, x_{n}\right)$

$$
\begin{aligned}
& x_{\ell}=\left(i_{\ell}, j_{\ell}, Y_{i_{\ell}, j_{\ell}}\right) \in[d] \times[d] \times \mathbb{R}, \\
& i_{\ell,}, j_{\ell} \sim_{i i d} \operatorname{Unif}([d]), \\
& \left.Y_{i_{\ell}, j_{\ell}}\right|_{i_{\ell}, j_{\ell}} \sim N\left(\theta_{0, i_{\ell}} \theta_{0, j_{\ell}}, \sigma^{2}\right)
\end{aligned}
$$

## Pictorially



## Example \#1: A different formulation

Data $\boldsymbol{Y} \in \mathbb{R}^{n \times n}$

$$
\boldsymbol{Y}=\mathcal{P}_{E}\left(\boldsymbol{\theta}_{0} \boldsymbol{\theta}_{0}^{\top}+\boldsymbol{W}\right)
$$

- $E \subseteq\binom{[d]}{2}$ uniformly random s.t. $|E|=n$
- $\mathcal{P}_{E}: \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d \times d}:$ projector that zeroes entries not in $E$
- $\left(W_{i j}\right)_{i<j} \sim_{i i d} \mathrm{~N}\left(0, \sigma^{2}\right), \boldsymbol{W}=\boldsymbol{W}^{\top}$
$\boldsymbol{\theta}_{0} \boldsymbol{\theta}_{0}^{\top}$



## Example \#1: Yet another formulation

- $G=(V, E) \sim \mathcal{G}(d, n)$
(uniform random with $d$ vertices, $n$ edges)
- For each $(i, j) \in E$

$$
Y_{i j}=\theta_{i} \theta_{j}+W_{i j}, \quad W_{i j} \sim \mathrm{~N}\left(0, \sigma^{2}\right)
$$

## Problem parameters

$$
\begin{array}{ll}
\delta=\frac{n}{d} & \text { (half) average graph degree } \\
\varepsilon=\frac{s_{0}}{d} & \text { sparsity }
\end{array}
$$

## Asymptotics

- Dense graph: $n \asymp d^{2}, s_{0} \asymp \sqrt{d} \quad(\sim$ Planted clique problem $)$
- Sparse graph: $n \asymp d, s_{0} \asymp d$ ( $\delta, \varepsilon$ fixed)


## Problem parameters

$$
\begin{array}{ll}
\delta=\frac{n}{d} & \text { (half) average graph degree } \\
\varepsilon=\frac{s_{0}}{d} & \text { sparsity }
\end{array}
$$

Asymptotics

- Dense graph: $n \asymp d^{2}, s_{0} \asymp \sqrt{d}$
- Sparse graph: $n \asymp d, s_{0} \asymp d$
( $\sim$ Planted clique problem) ( $\delta, \varepsilon$ fixed)

$$
R^{\text {Poly }}(\delta, \varepsilon ; d) \equiv R_{n=d \delta}^{\text {Poly }}\left(\Theta\left(s_{0}=d \varepsilon, d\right)\right)
$$

$$
\lim _{d \rightarrow \infty} R^{\text {Poly }}(\delta, \varepsilon ; d)=?
$$

## We do not know, but

$$
R^{\text {Poly }}(\delta, \varepsilon ; d) \equiv R_{n=d \delta}^{\text {Poly }}\left(\Theta\left(s_{0}=d \varepsilon, d\right)\right)
$$

$$
\lim _{d \rightarrow \infty} R^{\text {Poly }}(\delta, \varepsilon ; d)=?
$$

We do not know, but ...

## First simplifications

- Worst case prior

$$
\boldsymbol{\theta} \sim \operatorname{Unif}\left(\Theta\left(s_{0}, d\right)\right)
$$

- Roughly $\left(\varepsilon=s_{0} / d\right)$

$$
\left(\theta_{i}\right)_{i \leq d} \sim_{i i d} \operatorname{Bern}(\varepsilon)
$$

## First simplifications

- Worst case prior

$$
\boldsymbol{\theta} \sim \operatorname{Unif}\left(\Theta\left(s_{0}, d\right)\right)
$$

- Roughly $\left(\varepsilon=s_{0} / d\right)$

$$
\left(\theta_{i}\right)_{i \leq d} \sim_{i i d} \operatorname{Bern}(\varepsilon)
$$

## Local weak limit: $d \rightarrow \infty$



$$
G \stackrel{l w c}{\Rightarrow} \mathrm{GW}(\operatorname{Pois}(2 \delta))
$$

## Local weak limit: $d \rightarrow \infty$



$$
(G, \theta, Y) \stackrel{\text { lwc } c}{\Rightarrow}(\operatorname{GW}(\operatorname{Pois}(2 \delta)), \theta, Y)
$$

## First guess



$$
\lim _{d \rightarrow \infty} R^{\text {Poly }}(\delta, \varepsilon, d) \stackrel{?}{=} \inf _{\hat{\boldsymbol{\theta}}} \mathbb{P}_{\text {Tree }}\left(\hat{\theta}_{\varnothing}(\boldsymbol{Y}) \neq \theta_{0, \varnothing}\right)
$$

## It gets interesting

$$
\lim _{d \rightarrow \infty} R^{\text {Poly }}(\delta, \varepsilon, d) \stackrel{?}{\stackrel{i n f}{\hat{\theta}}} \mathbb{P}_{\text {Tree }}\left(\hat{\theta}_{\phi}(Y) \neq \theta_{0, \phi}\right)
$$

How do you define the r.h.s.?

## A natural idea:

$$
\begin{aligned}
Y_{\ell}^{0} & =\left(Y_{i j}: d(\phi, i) \leq \ell\right), \\
\hat{\theta}_{\phi}\left(Y_{\ell}^{0}\right) & =\underset{\sigma \in\{0,1\}}{\arg \max ^{1}\left(\theta_{\phi}=\sigma \mid Y_{\ell}^{0}\right) .} .
\end{aligned}
$$

## It gets interesting

$$
\lim _{d \rightarrow \infty} R^{\text {Poly }}(\delta, \varepsilon, d) \stackrel{?}{=} \inf _{\hat{\theta}} \mathbb{P}_{\text {Tree }}\left(\hat{\theta}_{\phi}(Y) \neq \theta_{0, \phi}\right)
$$

How do you define the r.h.s.?

A natural idea:

$$
\begin{aligned}
Y_{\ell}^{0} & =\left(Y_{i j}: d(\varnothing, i) \leq \ell\right), \\
\hat{\theta}_{\varnothing}\left(Y_{\ell}^{0}\right) & =\underset{\sigma \in\{0,1\}}{\arg \max ^{2}} \mathbb{P}\left(\theta_{\varnothing}=\sigma \mid Y_{\ell}^{0}\right) .
\end{aligned}
$$

## Risk of local algorithms

$$
R^{\text {loc }}(\delta, \varepsilon)=\lim _{\ell \rightarrow \infty} \mathbb{P}_{\text {Tree }}\left(\hat{\theta}_{\varnothing}\left(\boldsymbol{Y}_{\ell}^{0}\right) \neq \theta_{0, \varnothing}\right)
$$

## Remark

Belief nropagation achieves $R^{100}(\delta, \varepsilon)$

## Risk of local algorithms

$$
R^{\mathrm{loc}}(\delta, \varepsilon)=\lim _{\ell \rightarrow \infty} \mathbb{P}_{\text {Tree }}\left(\hat{\theta}_{\varnothing}\left(\boldsymbol{Y}_{\ell}^{0}\right) \neq \theta_{0, \varnothing}\right)
$$

## Remark <br> Belief propagation achieves $R^{\text {loc }}(\delta, \varepsilon)$

## Can be computed

$$
(\varepsilon=0.02, \sigma=1.5)
$$



$$
R^{\text {loc }}(\delta, \varepsilon)=\lim _{\ell \rightarrow \infty} \mathbb{P}_{\text {Tree }}\left(\hat{\theta}_{\varnothing}\left(\boldsymbol{Y}_{\ell}^{0}\right) \neq \theta_{0, \varnothing}\right)
$$

## Phase transition for local algorithms

Theorem (~ Deshpande, Montanari, 2015)
As $\varepsilon \rightarrow 0$ (with $\delta \rightarrow \infty, \sigma$ fixed)

$$
\frac{1}{\varepsilon} R^{\mathrm{loc}}(\delta, \varepsilon) \rightarrow \begin{cases}1 & \text { if } \delta \leq\left(1-o_{\varepsilon}(1)\right) \cdot \frac{\sigma^{2}}{2 e \varepsilon^{2}} \\ 0 & \text { if } \delta \geq\left(1+o_{\varepsilon}(1)\right) \cdot \frac{\sigma^{2}}{2 e \varepsilon^{2}}\end{cases}
$$

## A different definition of the limit

$$
\begin{aligned}
\boldsymbol{Y}_{\ell}^{+} & =\left\{\left(Y_{i j}: d(\varnothing, i) \leq \ell\right) ; \quad\left(\theta_{i}: d(\varnothing, i>\ell)\right\}\right. \\
\hat{\theta}_{\varnothing}\left(\boldsymbol{Y}_{\ell}^{+}\right) & =\arg \max _{\sigma \in\{0,1\}} \mathbb{P}\left(\theta_{\varnothing}=\sigma \mid \boldsymbol{Y}_{\ell}^{+}\right)
\end{aligned}
$$

$$
R^{\text {Ora }}(\delta, \varepsilon)=\lim _{\ell \rightarrow \infty} \mathbb{P}_{\text {Tree }}\left(\hat{\theta}_{\phi}\left(\boldsymbol{Y}_{\ell}^{0}\right) \neq \theta_{0, \varnothing}\right)
$$

## It also can be computed



$$
R^{\text {Ora }}(\delta, \varepsilon) \leq R^{*}(\delta, \varepsilon) \leq R^{\text {Poly }}(\delta, \varepsilon) \leq R^{\text {loc }}(\delta, \varepsilon)
$$

## It also can be computed



$$
R^{\mathrm{Ora}}(\delta, \varepsilon) \leq R^{*}(\delta, \varepsilon) \leq R^{\mathrm{Poly}}(\delta, \varepsilon) \leq R^{\mathrm{loc}}(\delta, \varepsilon)
$$

## It also can be computed



$$
R^{\mathrm{Ora}}(\delta, \varepsilon) \leq R^{*}(\delta, \varepsilon) \leq R^{\mathrm{Poly}}(\delta, \varepsilon) \leq R^{\mathrm{loc}}(\delta, \varepsilon)
$$

## Minimax risk

Theorem ( $\sim$ Montanari, 2015)
As $\varepsilon \rightarrow 0$ (with $\delta \rightarrow \infty, \sigma$ fixed)

$$
\frac{1}{\varepsilon} R^{*}(\delta, \varepsilon) \rightarrow \begin{cases}1 & \text { if } \delta \leq\left(1-o_{\varepsilon}(1)\right) \cdot \frac{\sigma^{2}}{2 \varepsilon} \log (1 / \varepsilon) \\ 0 & \text { if } \delta \geq\left(1+o_{\varepsilon}(1)\right) \cdot \frac{\sigma^{2}}{2 \varepsilon} \log (1 / \varepsilon)\end{cases}
$$



## Minimax risk

Theorem ( $\sim$ Montanari, 2015)
As $\varepsilon \rightarrow 0$ (with $\delta \rightarrow \infty, \sigma$ fixed)

$$
\frac{1}{\varepsilon} R^{*}(\delta, \varepsilon) \rightarrow \begin{cases}1 & \text { if } \delta \leq\left(1-o_{\varepsilon}(1)\right) \cdot \frac{\sigma^{2}}{2 \varepsilon} \log (1 / \varepsilon) \\ 0 & \text { if } \delta \geq\left(1+o_{\varepsilon}(1)\right) \cdot \frac{\sigma^{2}}{2 \varepsilon} \log (1 / \varepsilon)\end{cases}
$$

- Sharp threshold (dense): Lelarge, Miolane 2017; Barbier et al. 2017


## Minimax risk

Theorem ( $\sim$ Montanari, 2015)
As $\varepsilon \rightarrow 0$ (with $\delta \rightarrow \infty, \sigma$ fixed)

$$
\frac{1}{\varepsilon} R^{*}(\delta, \varepsilon) \rightarrow \begin{cases}1 & \text { if } \delta \leq\left(1-o_{\varepsilon}(1)\right) \cdot \frac{\sigma^{2}}{2 \varepsilon} \log (1 / \varepsilon) \\ 0 & \text { if } \delta \geq\left(1+o_{\varepsilon}(1)\right) \cdot \frac{\sigma^{2}}{2 \varepsilon} \log (1 / \varepsilon)\end{cases}
$$

- Sharp threshold (dense): Lelarge, Miolane 2017; Barbier et al. 2017
- Do not know of any polytime algorithm working for

$$
1.01 \cdot \frac{\sigma^{2}}{2 \varepsilon} \log (1 / \varepsilon)<\delta<0.99 \cdot \frac{\sigma^{2}}{2 e \varepsilon^{2}}
$$

## Open problems

- Can we beat $R^{\text {loc }}(\delta, \varepsilon)$ by Gibbs sampling?
- Can we beat $R^{\text {loc }}(\delta, \varepsilon)$ by convex optimization?
- Is $R^{\text {Poly }}(\delta, \varepsilon)=R^{\text {loc }}(\delta, \varepsilon)$ ?


## Open problems

- Can we beat $R^{\text {loc }}(\delta, \varepsilon)$ by Gibbs sampling?
- Can we beat $R^{\text {loc }}(\delta, \varepsilon)$ by convex optimization?
- Is $R^{\text {Poly }}(\delta, \varepsilon)=R^{\text {loc }}(\delta, \varepsilon)$ ?


## Open problems

- Can we beat $R^{\text {loc }}(\delta, \varepsilon)$ by Gibbs sampling?
- Can we beat $R^{\text {loc }}(\delta, \varepsilon)$ by convex optimization?
- Is $R^{\text {Poly }}(\delta, \varepsilon)=R^{\text {loc }}(\delta, \varepsilon)$ ?


## Ubiquitous

Sparse principal component analysis

- Berthet, Rigollet, 2013
- Deshpande, Montanari, 2014
- Barbier et al 2016; Miolane 2017

Hidden clique problem
(2 asymmetric communities)

- Jerrum, 1992
- Feige, Krauthgamer, 200
- Deshpande, Montanari, 2015; Montanari 2015

Community detection ( $k \geq 5$ symmetric communities)

- Decelle, Krzakala, Moore, Zdeborova 2011
- Bordenave, Lelarge, Massoulie 2015
- Abbe, Sandon 2015


## Tensor PCA

- See below


## Landscapes

## Empirical risk minimization/M-estimation

Class of models $\left(\Theta \subseteq \mathbb{R}^{d}\right)$

$$
C_{\Theta} \equiv\left\{\mathbb{P}_{\boldsymbol{\theta}}: \quad \theta \in \Theta\right\}
$$

## Data

$$
x_{1}, x_{2}, \ldots, x_{n} \sim_{i i d} \mathbb{P}_{\boldsymbol{\theta}_{0}}(\cdot)
$$

## Empirical risk minimization/M-estimation

Class of models $\left(\Theta \subseteq \mathbb{R}^{d}\right)$

$$
C_{\Theta} \equiv\left\{\mathbb{P}_{\boldsymbol{\theta}}: \quad \theta \in \Theta\right\}
$$

## Data

$$
x_{1}, x_{2}, \ldots, x_{n} \sim_{i i d} \mathbb{P}_{\boldsymbol{\theta}_{0}}(\cdot)
$$

$$
\operatorname{minimize} \quad \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta}) \equiv \frac{1}{n} \sum_{i=1}^{n} \ell\left(\boldsymbol{\theta} ; \boldsymbol{x}_{i}\right)
$$

## Empirical risk minimization/M-estimation

## Rationale

$$
\boldsymbol{\theta}_{0}=\arg \min _{\boldsymbol{\theta} \in \mathbb{R}^{d}} \mathcal{L}(\boldsymbol{\theta}), \quad \mathcal{L}(\boldsymbol{\theta})=\mathbb{E} \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})
$$

- What can we say generically?
- How does complexity show up?
- What can we say generically?
- How does complexity show up?


## Uniform convergence

Theorem (Vapnik, Chervonenkis, 1968; ...)
Under conditions [omitted], with high probability

$$
\begin{aligned}
& \sup _{\boldsymbol{\theta} \in \Theta}\left|\widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})-\mathcal{L}(\boldsymbol{\theta})\right| \leq C \sqrt{\frac{d_{*}}{n}} \\
&\left(d_{*}=\text { VC dimension; } \ldots\right)
\end{aligned}
$$

## Uniform convergence



Population risk


Empirical risk

## Will optimization algorithms get stuck in local minima? Landscape analysis

## Uniform convergence



Population risk


Empirical risk

Will optimization algorithms get stuck in local minima?

## Uniform convergence



Population risk


Empirical risk

Will optimization algorithms get stuck in local minima? Landscape analysis

## Assumptions

$\boldsymbol{\theta} \in \mathrm{B}^{p}(r)=$ Ball of radius $r$ in $\mathbb{R}^{p}$
Data: $\boldsymbol{Z}_{1}, Z_{2}, \ldots, Z_{n}$ iid

> A1 $\nabla_{\theta} \ell(\theta ; Z)$ is $\tau^{2}$-sub-Gaussian
> A2 For any $\lambda \in \mathrm{B}^{p}(1), \mathcal{Z}_{\lambda} \equiv\left\langle\lambda, \nabla^{2} \ell(\theta ; Z) \lambda\right\rangle$ is $\tau^{2}$-sub-Exponential.
> A3 The Hessian of the population risk at 0 is bounded by a polynomial

$$
\left\|\nabla^{2} \mathcal{L}(0)\right\|_{\mathrm{op}} \leq \tau^{2} d^{C} .
$$

A4 The Hessian of the loss is Lipschitz continuous with integrable constant

$$
\mathbb{E}\left\{\left\|\nabla^{2} \ell(\cdot ; Z)\right\|_{\text {Lip }}\right\} \leq \tau^{3} d^{C} .
$$

## Assumptions

$\boldsymbol{\theta} \in \mathrm{B}^{p}(r)=$ Ball of radius $r$ in $\mathbb{R}^{p}$
Data: $Z_{1}, Z_{2}, \ldots, Z_{n}$ iid
A1 $\nabla_{\theta} \ell(\theta ; Z)$ is $\tau^{2}$-sub-Gaussian
A2 For any $\lambda \in B^{p}(1), Z_{\lambda} \equiv\left\langle\lambda, \nabla^{2} \ell(\theta ; Z) \lambda\right\rangle$ is $\tau^{2}$-sub-Exponential.
A3 The Hessian of the population risk at 0 is bounded by a polynomial


A4 The Hessian of the loss is Lipschitz continuous with integrable constant


## Assumptions

$\boldsymbol{\theta} \in \mathrm{B}^{p}(r)=$ Ball of radius $r$ in $\mathbb{R}^{p}$
Data: $\boldsymbol{Z}_{1}, \boldsymbol{Z}_{2}, \ldots, \boldsymbol{Z}_{n}$ iid
A1 $\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta} ; \boldsymbol{Z})$ is $\tau^{2}$-sub-Gaussian
A2 For any $\boldsymbol{\lambda} \in \mathrm{B}^{p}(1), \mathcal{Z}_{\boldsymbol{\lambda}} \equiv\left\langle\boldsymbol{\lambda}, \nabla^{2} \ell(\boldsymbol{\theta} ; \boldsymbol{Z}) \boldsymbol{\lambda}\right\rangle$ is $\tau^{2}$-sub-Exponential. polynomial

The Hessian of the loss is Lipschitz continuous with integrable constant


## Assumptions

$\boldsymbol{\theta} \in \mathrm{B}^{p}(r)=$ Ball of radius $r$ in $\mathbb{R}^{p}$
Data: $\boldsymbol{Z}_{1}, \boldsymbol{Z}_{2}, \ldots, \boldsymbol{Z}_{n}$ iid
A1 $\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta} ; \boldsymbol{Z})$ is $\tau^{2}$-sub-Gaussian
A2 For any $\boldsymbol{\lambda} \in \mathrm{B}^{p}(1), \mathcal{Z}_{\boldsymbol{\lambda}} \equiv\left\langle\boldsymbol{\lambda}, \nabla^{2} \ell(\boldsymbol{\theta} ; \boldsymbol{Z}) \boldsymbol{\lambda}\right\rangle$ is $\boldsymbol{\tau}^{2}$-sub-Exponential.
A3 The Hessian of the population risk at 0 is bounded by a polynomial

$$
\left\|\nabla^{2} \mathcal{L}(0)\right\|_{\mathrm{op}} \leq \tau^{2} d^{C}
$$

A4 The Hessian of the loss is Lipschitz continuous with integrable constant


## Assumptions

$\boldsymbol{\theta} \in \mathrm{B}^{p}(r)=$ Ball of radius $r$ in $\mathbb{R}^{p}$
Data: $\boldsymbol{Z}_{1}, \boldsymbol{Z}_{2}, \ldots, \boldsymbol{Z}_{n}$ iid
A1 $\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta} ; \boldsymbol{Z})$ is $\tau^{2}$-sub-Gaussian
A2 For any $\boldsymbol{\lambda} \in \mathrm{B}^{p}(1), \mathcal{Z}_{\boldsymbol{\lambda}} \equiv\left\langle\boldsymbol{\lambda}, \nabla^{2} \ell(\boldsymbol{\theta} ; \boldsymbol{Z}) \boldsymbol{\lambda}\right\rangle$ is $\boldsymbol{\tau}^{2}$-sub-Exponential.
A3 The Hessian of the population risk at 0 is bounded by a polynomial

$$
\left\|\nabla^{2} \mathcal{L}(0)\right\|_{\mathrm{op}} \leq \tau^{2} d^{C}
$$

A4 The Hessian of the loss is Lipschitz continuous with integrable constant

$$
\mathbb{E}\left\{\left\|\nabla^{2} \ell(\cdot ; \boldsymbol{Z})\right\|_{\text {Lip }}\right\} \leq \tau^{3} d^{C}
$$

## Assumptions

$\boldsymbol{\theta} \in \mathrm{B}^{p}(r)=$ Ball of radius $r$ in $\mathbb{R}^{p}$
Data: $\boldsymbol{Z}_{1}, \boldsymbol{Z}_{2}, \ldots, \boldsymbol{Z}_{n}$ iid
A1 $\nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta} ; \boldsymbol{Z})$ is $\tau^{2}$-sub-Gaussian
A2 For any $\boldsymbol{\lambda} \in \mathrm{B}^{p}(1), \mathcal{Z}_{\boldsymbol{\lambda}} \equiv\left\langle\boldsymbol{\lambda}, \nabla^{2} \ell(\boldsymbol{\theta} ; \boldsymbol{Z}) \boldsymbol{\lambda}\right\rangle$ is $\boldsymbol{\tau}^{2}$-sub-Exponential.
A3 The Hessian of the population risk at 0 is bounded by a polynomial

$$
\left\|\nabla^{2} \mathcal{L}(0)\right\|_{\mathrm{op}} \leq \tau^{2} d^{C}
$$

A4 The Hessian of the loss is Lipschitz continuous with integrable constant

$$
\mathbb{E}\left\{\left\|\nabla^{2} \ell(\cdot ; \boldsymbol{Z})\right\|_{\text {Lip }}\right\} \leq \tau^{3} d^{C}
$$

'Under mild assumptions'

## Lemma

Under assumptions A1, A2, A3, A4, if $n \geq C p \log p$, then with probability at least $1-\delta$, the following hold:

$$
\begin{gathered}
\sup _{\boldsymbol{\theta} \in \mathrm{B}^{p}(r)}\left\|\nabla \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})-\nabla \mathcal{L}(\boldsymbol{\theta})\right\|_{2} \leq \tau \sqrt{\frac{C d \log n}{n}} \\
\sup _{\boldsymbol{\theta} \in \mathrm{B}^{p}(r)}\left\|\nabla^{2} \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})-\nabla^{2} \mathcal{L}(\boldsymbol{\theta})\right\|_{\mathrm{op}} \leq \tau^{2} \sqrt{\frac{C d \log n}{n}} .
\end{gathered}
$$

## This cannot happen!



Population risk


Empirical risk

## This can happen!



## Nice population risk $\Rightarrow$ Nice empirical risk

## This can happen!



Population risk


Empirical risk

Nice population risk $\Rightarrow$ Nice empirical risk

## Example: Binary classification


$\boldsymbol{z}_{i}=\left(y_{i}, \boldsymbol{x}_{i}\right), y_{i} \in\{0,1\}, \boldsymbol{x}_{i} \in \mathbb{R}^{d}$

$$
\begin{aligned}
\mathbb{P}\left(y_{i}=1 \mid x_{i}\right) & =\sigma\left(\left\langle\boldsymbol{\theta}_{0}, \boldsymbol{x}_{i}\right\rangle\right) \\
\widehat{\mathcal{L}}_{n}(\boldsymbol{\theta}) & =\frac{1}{n} \sum_{i=1}^{n}\left(y_{i}-\sigma\left(\left\langle\boldsymbol{\theta}, \boldsymbol{x}_{i}\right\rangle\right)\right)^{2} .
\end{aligned}
$$

- Rosenblatt 1958 (Perceptron); . . . many extensions
- More robust than logistic regression


## Sample application: Binary classification

Theorem (Mei, Bai, Montanari 2017)
Assume $\boldsymbol{X}_{i}$ to be centered, sub-Gaussian, with $\mathbb{E}\left\{\boldsymbol{X} \boldsymbol{X}^{\top}\right\} \succeq \delta I_{d}$. For nice ${ }^{a}$ functions $\sigma$, whp:

1. The population risk has a unique critical point $\hat{\theta}_{n}$.
2. Gradient descent converges exponentially fast to $\hat{\boldsymbol{\theta}}_{n}$.
3. The estimation error is $\left\|\hat{\theta}_{n}-0_{0}\right\|_{2} \leq C \sqrt{(d \log n) / n}$.

$$
{ }^{a} \sigma^{\prime}(x)>0,\left\|\sigma^{\prime}\right\|_{\infty},\left\|\sigma^{\prime \prime}\right\|_{\infty},\left\|\sigma^{\prime \prime \prime}\right\|_{\infty} \leq C .
$$

Similar results for robust regression, one-bit compressed sensing, [see pape]

## Sample application: Binary classification

Theorem (Mei, Bai, Montanari 2017)
Assume $\boldsymbol{X}_{i}$ to be centered, sub-Gaussian, with $\mathbb{E}\left\{\boldsymbol{X} \boldsymbol{X}^{\top}\right\} \succeq \delta I_{d}$. For nice ${ }^{a}$ functions $\sigma$, whp:

1. The population risk has a unique critical point $\hat{\boldsymbol{\theta}}_{n}$.
2. Gradient descent converges exponentially fast to $\hat{\boldsymbol{\theta}}_{n}$.
3. The estimation error is $\left\|\hat{\boldsymbol{\theta}}_{n}-\theta_{0}\right\|_{2} \leq C \sqrt{(d \log n) / n}$.

$$
{ }^{a} \sigma^{\prime}(x)>0,\left\|\sigma^{\prime}\right\|_{\infty},\left\|\sigma^{\prime \prime}\right\|_{\infty},\left\|\sigma^{\prime \prime \prime}\right\|_{\infty} \leq C .
$$

Similar results for robust regression, one-bit compressed sensing,

## Sample application: Binary classification

## Theorem (Mei, Bai, Montanari 2017)

Assume $X_{i}$ to be centered, sub-Gaussian, with $\mathbb{E}\left\{X X^{\top}\right\} \succeq \delta I_{d}$. For nice ${ }^{a}$ functions $\sigma$, whp:

1. The population risk has a unique critical point $\hat{\boldsymbol{\theta}}_{n}$.
2. Gradient descent converges exponentially fast to $\hat{\boldsymbol{\theta}}_{n}$.
3. The estimation error is $\left\|\hat{\theta}_{n}-\theta_{0}\right\|_{2} \leq C \sqrt{(d \log n) / n}$.

$$
{ }^{a} \sigma^{\prime}(x)>0,\left\|\sigma^{\prime}\right\|_{\infty},\left\|\sigma^{\prime \prime}\right\|_{\infty},\left\|\sigma^{\prime \prime \prime}\right\|_{\infty} \leq C .
$$

Similar results for robust regression, one-bit compressed sensing,

## Sample application: Binary classification

## Theorem (Mei, Bai, Montanari 2017)

Assume $X_{i}$ to be centered, sub-Gaussian, with $\mathbb{E}\left\{X X^{\top}\right\} \succeq \delta I_{d}$. For nice ${ }^{a}$ functions $\sigma$, whp:

1. The population risk has a unique critical point $\hat{\boldsymbol{\theta}}_{n}$.
2. Gradient descent converges exponentially fast to $\hat{\boldsymbol{\theta}}_{n}$.
3. The estimation error is $\left\|\hat{\boldsymbol{\theta}}_{n}-\boldsymbol{\theta}_{0}\right\|_{2} \leq C \sqrt{(d \log n) / n}$.

$$
{ }^{2} \sigma^{\prime}(x)>0,\left\|\sigma^{\prime}\right\|_{\infty},\left\|\sigma^{\prime \prime}\right\|_{\infty},\left\|\sigma^{\prime \prime \prime}\right\|_{\infty} \leq C .
$$

Similar results for robust regression, one-bit compressed sensing,

## Sample application: Binary classification

## Theorem (Mei, Bai, Montanari 2017)

Assume $X_{i}$ to be centered, sub-Gaussian, with $\mathbb{E}\left\{X^{\top} X^{\top}\right\} \succeq \delta I_{d}$. For nice ${ }^{a}$ functions $\sigma$, whp:

1. The population risk has a unique critical point $\hat{\boldsymbol{\theta}}_{n}$.
2. Gradient descent converges exponentially fast to $\hat{\boldsymbol{\theta}}_{n}$.
3. The estimation error is $\left\|\hat{\boldsymbol{\theta}}_{n}-\boldsymbol{\theta}_{0}\right\|_{2} \leq C \sqrt{(d \log n) / n}$.

$$
{ }^{2} \sigma^{\prime}(x)>0,\left\|\sigma^{\prime}\right\|_{\infty},\left\|\sigma^{\prime \prime}\right\|_{\infty},\left\|\sigma^{\prime \prime \prime}\right\|_{\infty} \leq C .
$$

Similar results for robust regression, one-bit compressed sensing, ...
[see paper]

## Non-convex literature

Convergence to 'statistical neighborhood'

- Loh, Wainwright, 2012
- Loh Wainwright, 2013
- Yang, Wang, Liu, Eldar, Zhang, 2015

Smart initialization

- TKeshavan, Montanari, Oh, 2009
- Chen, Candés, 2015
- Anandkumar Ge Jenramin, 2015


## Unique local minimum

- Ioh, 2015
- Sun, Qu, Wright, 2016
[High-dim. regression, $n \gtrsim s_{0}^{2}$ ]
[Phase retrievall
- Ge, Lee, Ma, 2016


## Non-convex literature

Convergence to 'statistical neighborhood'

- Loh, Wainwright, 2012
- Loh Wainwright, 2013
- Yang, Wang, Liu, Eldar, Zhang, 2015

Smart initialization

- Keshavan, Montanari, Oh, 2009
- Chen, Candés, 2015
- Anandkumar, Ge, Jenzamin, 2015

Unique local minimum

- Loh, 2015
- Sun, Qu, Wright, 2016
- Ge, Lee, Ma, 2016
[High-dim. regression, $n \gtrsim s_{0}^{2}$ ]
[Phase retrieval]
[Matrix completion]
- What can we say generically?
- How does complexity show up?


## Intuition

## Population risk very flat $\Leftrightarrow$ Many local minima

[Close to where local algorithms fail?]

## Simplest example: Spiked tensor model

- Unknown parameter $\boldsymbol{\theta}_{0} \in \mathbb{R}^{n},\left\|\boldsymbol{\theta}_{0}\right\|_{2}=1$
- Data ${ }^{1}$

$$
\boldsymbol{Y}=\lambda \boldsymbol{\theta}_{0}^{\otimes k}+\boldsymbol{W}
$$

- $\boldsymbol{W}=$ symmetric Gaussian noise tensor. $\left(W_{i_{1}, \ldots, i_{k}}\right)_{i_{1}<\cdots<i_{k}} \sim_{i i d} \mathrm{~N}(0,1 / n)$

[^0]
## Spiked tensor model: What do we know?

$$
\boldsymbol{Y}=\lambda \boldsymbol{\theta}_{0}^{\otimes k}+\boldsymbol{W}
$$

Theorem (Montanari, Richard, 2014; Hopkins, Shi, Steurer, 2015)
For any $\varepsilon>0$, there exist constants $\lambda_{\mathrm{IT}}$, $\lambda_{\mathrm{ML}}(\varepsilon), C(\varepsilon)$, such that:

- If $\lambda>\lambda_{\mathrm{ML}}(\varepsilon)$, then $\mathbb{E}\left\{\left|\left\langle\hat{\theta}^{\mathrm{ML}}, \theta_{0}\right\rangle\right|\right\} \geq 1-\varepsilon$.
- No estimator can achieve $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}, \theta_{0}\right\rangle\right|\right\} \geq \varepsilon$ unless $\lambda>\lambda_{\text {IT }}$.
- There exists poly-time estimator achieving
$\mathbb{E}\left\{\mid\left\langle\hat{\theta}^{\text {Dolty }}, \theta_{0}\right|\right\} \geq 1-\varepsilon$, provided $\lambda \geq C(\varepsilon) n^{(i \lambda-2) / 4}$.

No efficient estimator is known for $1 \ll \lambda \ll n^{(k-2) / 4}$ !

## Spiked tensor model: What do we know?

$$
\boldsymbol{Y}=\lambda \boldsymbol{\theta}_{0}^{\otimes k}+\boldsymbol{W}
$$

Theorem (Montanari, Richard, 2014; Hopkins, Shi, Steurer, 2015)
For any $\varepsilon>0$, there exist constants $\lambda_{\mathrm{IT}}, \lambda_{\mathrm{ML}}(\varepsilon), C(\varepsilon)$, such that:

- If $\lambda>\lambda_{\mathrm{ML}}(\varepsilon)$, then $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}^{\mathrm{ML}}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq 1-\boldsymbol{\varepsilon}$.


No efficient estimator is known for $1 \ll \lambda \ll n^{(k-2) / 4}$ !

## Spiked tensor model: What do we know?

$$
\boldsymbol{Y}=\lambda \boldsymbol{\theta}_{0}^{\otimes k}+\boldsymbol{W}
$$

Theorem (Montanari, Richard, 2014; Hopkins, Shi, Steurer, 2015)
For any $\varepsilon>0$, there exist constants $\lambda_{\text {IT }}$, $\lambda_{\text {МL }}(\varepsilon), C(\varepsilon)$, such that:

- If $\lambda>\lambda_{\mathrm{ML}}(\varepsilon)$, then $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}^{\mathrm{ML}}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq 1-\boldsymbol{\varepsilon}$.
- No estimator can achieve $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq \boldsymbol{\varepsilon}$ unless $\lambda>\lambda_{\mathrm{IT}}$.

No efficient estimator is known for $1 \ll \lambda \ll n^{(k-2) / 4}$ !

## Spiked tensor model: What do we know?

$$
\boldsymbol{Y}=\lambda \boldsymbol{\theta}_{0}^{\otimes k}+\boldsymbol{W}
$$

Theorem (Montanari, Richard, 2014; Hopkins, Shi, Steurer, 2015)
For any $\varepsilon>0$, there exist constants $\lambda_{\mathrm{IT}}, \lambda_{\mathrm{ML}}(\varepsilon), C(\varepsilon)$, such that:

- If $\lambda>\lambda_{\mathrm{ML}}(\varepsilon)$, then $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}^{\mathrm{ML}}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq 1-\boldsymbol{\varepsilon}$.
- No estimator can achieve $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq \varepsilon$ unless $\lambda>\lambda_{\mathrm{IT}}$.
- There exists poly-time estimator achieving $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}^{\text {Poly }}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq 1-\boldsymbol{\varepsilon}$, provided $\lambda \geq C(\varepsilon) n^{(k-2) / 4}$.


## Spiked tensor model: What do we know?

$$
\boldsymbol{Y}=\lambda \boldsymbol{\theta}_{0}^{\otimes k}+\boldsymbol{W}
$$

Theorem (Montanari, Richard, 2014; Hopkins, Shi, Steurer, 2015)
For any $\varepsilon>0$, there exist constants $\lambda_{\mathrm{IT}}, \lambda_{\mathrm{ML}}(\varepsilon), C(\varepsilon)$, such that:

- If $\lambda>\lambda_{\mathrm{ML}}(\varepsilon)$, then $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}^{\mathrm{ML}}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq 1-\boldsymbol{\varepsilon}$.
- No estimator can achieve $\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq \boldsymbol{\varepsilon}$ unless $\lambda>\lambda_{\mathrm{IT}}$.
- There exists poly-time estimator achieving
$\mathbb{E}\left\{\left|\left\langle\hat{\boldsymbol{\theta}}^{\text {Poly }}, \boldsymbol{\theta}_{0}\right\rangle\right|\right\} \geq 1-\boldsymbol{\varepsilon}$, provided $\lambda \geq C(\varepsilon) n^{(k-2) / 4}$.

No efficient estimator is known for $1 \ll \lambda \ll n^{(k-2) / 4}$ !

## More precise results

- Montanari, Reichman, Zeitouni, 2015
- Bandeira, Perry, Wein, 2017
- Krzakala, Lelarge, Miolane, Zdeborova, 2017


## In practice $(k=3)_{n=3 s}$








## What does the landscape look like?

## Maximum likelihood

$$
\begin{aligned}
\text { minimize } & \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})=\left\|\boldsymbol{Y}-\lambda \boldsymbol{\theta}^{\otimes k}\right\|_{F} \\
\text { subject to } & \|\boldsymbol{\theta}\|_{2}=1
\end{aligned}
$$



## Maximum likelihood

$$
\begin{aligned}
\operatorname{minimize} & \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})=\left\|\boldsymbol{Y}-\lambda \boldsymbol{\theta}^{\otimes k}\right\|_{F} \\
\text { subject to } & \|\boldsymbol{\theta}\|_{2}=1
\end{aligned}
$$

$$
\widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})=\text { const. }-2 \lambda\left\langle\boldsymbol{Y}, \boldsymbol{\theta}^{\otimes k}\right\rangle
$$

## Risk



Maximum likelihood

$$
\begin{aligned}
\operatorname{minimize} & \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})=-\left\langle\boldsymbol{Y}, \boldsymbol{\theta}^{\otimes k}\right\rangle \\
\text { subject to } & \|\boldsymbol{\theta}\|_{2}=1
\end{aligned}
$$

'Population' risk

$$
\mathcal{L}(\boldsymbol{\theta})=-\lambda\left\langle\boldsymbol{\theta}_{0}, \boldsymbol{\theta}\right\rangle^{k}
$$

## Back-of-the-envelope

Expected gradient

$$
\nabla \mathcal{L}(\boldsymbol{\theta})=-k \lambda\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle^{k-1} \boldsymbol{\theta}_{0}
$$

Random initialization $\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle=\Theta\left(n^{-1 / 2}\right)$ :

$$
\begin{aligned}
\left\langle\theta_{0}, \nabla \widehat{\mathcal{L}}_{n}(\theta)\right\rangle & =-k \lambda\left\langle\theta_{,} \theta_{0}\right\rangle^{k-1}-k\left\langle W, \theta_{0} \otimes \theta^{\otimes(k-1)}\right\rangle \\
& =-\lambda \Theta\left(n^{-(k-1) / 2}\right)+\Theta\left(n^{-1 / 2}\right)
\end{aligned}
$$

## Back-of-the-envelope

Expected gradient

$$
\nabla \mathcal{L}(\boldsymbol{\theta})=-k \lambda\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle^{k-1} \boldsymbol{\theta}_{0}
$$

Random initialization $\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle=\Theta\left(n^{-1 / 2}\right)$ :

$$
\begin{aligned}
\left\langle\boldsymbol{\theta}_{0}, \nabla \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})\right\rangle & =-k \lambda\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle^{k-1}-k\left\langle\boldsymbol{W}, \boldsymbol{\theta}_{0} \otimes \boldsymbol{\theta}^{\otimes(k-1)}\right\rangle \\
& =-\lambda \Theta\left(n^{-(k-1) / 2}\right)+\Theta\left(n^{-1 / 2}\right)
\end{aligned}
$$

## Back-of-the-envelope

Expected gradient

$$
\nabla \mathcal{L}(\boldsymbol{\theta})=-k \lambda\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle^{k-1} \boldsymbol{\theta}_{0}
$$

Random initialization $\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle=\Theta\left(n^{-1 / 2}\right)$ :

$$
\begin{aligned}
\left\langle\boldsymbol{\theta}_{0}, \nabla \widehat{\mathcal{L}}_{n}(\boldsymbol{\theta})\right\rangle & =-k \lambda\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle^{k-1}-k\left\langle\boldsymbol{W}, \boldsymbol{\theta}_{0} \otimes \boldsymbol{\theta}^{\otimes(k-1)}\right\rangle \\
& =-\lambda \Theta\left(n^{-(k-1) / 2}\right)+\Theta\left(n^{-1 / 2}\right)
\end{aligned}
$$

- Convergence: $\lambda \gg n^{(k-2) / 2}$


## Expected number of local minima: $k=3, \lambda=3$





- Exponential in black region

$$
\left(m=\left\langle\boldsymbol{\theta}, \boldsymbol{\theta}_{0}\right\rangle, x=\left\langle\boldsymbol{Y}, \boldsymbol{\theta}^{\otimes k}\right\rangle\right)
$$

- $N=$ number of local minima

$$
\mathbb{E} N(m, x)=e^{n S_{0}(m, x)+o(n)}
$$

[Ben Arous, Mei, Montanari, Nica, 2017]

Complexity of landscape $\stackrel{?}{\leftrightarrow}$ Complexity for local algorithms

## An emerging dichotomy

## In several statistical estimation problems

- Dither local (or message-passing) algorithms work. - ... or SDP hierarchies do not work

Why?

## An emerging dichotomy

In several statistical estimation problems

- Either local (or message-passing) algorithms work...
- . . . or SDP hierarchies do not work


## An emerging dichotomy

In several statistical estimation problems

- Either local (or message-passing) algorithms work...
- ... or SDP hierarchies do not work

Why?

## A possible explanation

Perhaps SDPs on random instances can be solved by local algorithms...

## Simplest example

Centered adjacency matrix of $G=(V, E) \quad$ ( $d=$ average degree)

$$
A_{i j}^{\text {cen }}= \begin{cases}1-\frac{d}{n} & \text { if }(i, j) \in E \\ -\frac{d}{n} & \text { otherwise }\end{cases}
$$

$\operatorname{SDP}\left(\boldsymbol{A}^{\mathrm{cen}}\right):$

$$
\begin{aligned}
\operatorname{maximize} & \left\langle\boldsymbol{A}^{\mathrm{cen}}, \boldsymbol{X}\right\rangle \\
\text { subject to } & \boldsymbol{X} \in \mathbb{R}^{n \times n}, \boldsymbol{X} \succeq 0 \\
& X_{i i}=1
\end{aligned}
$$

- Graph clustering, embedding, testing latent structure,...


## What does it mean?

Input: Graph $G_{n}=\left(V_{n}, E_{n}\right)$

1. Generate $\boldsymbol{z}=(z(i))_{i \in V} \sim_{i i d} \mathrm{~N}(0,1)$
2. Compute, for each $v \in V_{n}, \xi_{v}=F\left(\mathrm{~B}_{\ell}\left(v ; G_{n}\right),\left.z\right|_{\mathrm{B}_{\ell}\left(v ; G_{n}\right)}\right)$
3. Output $\boldsymbol{X}=\mathbb{E}_{\boldsymbol{z}}\left\{\boldsymbol{\xi} \xi^{\top}\right\}+\cdots \in \mathbb{R}^{n \times n}$

Can this achieve $\left\langle\boldsymbol{A}^{\text {cen }}, \boldsymbol{X}\right\rangle \geq\left(1-o_{n}(1)\right) \operatorname{SDP}\left(\boldsymbol{A}^{\text {cen }}\right)$ ?

## Erdős-Renyi random graphs

Theorem (Fan, Montanari, 2016)
Let $G \sim \mathcal{G}(n, d / n)$ and $A^{\text {cen }}=A_{G}^{\mathrm{cen}}$. Then, a.s.,

$$
\begin{aligned}
2 \sqrt{d}\left(1-\frac{1}{d+1}\right) & \leq \lim \inf _{n \rightarrow \infty} \frac{1}{n} \operatorname{SDP}\left(A^{\mathrm{cen}}\right) \leq \\
& \leq \lim \sup _{n \rightarrow \infty} \frac{1}{n} \operatorname{SDP}\left(A^{\mathrm{cen}}\right) \leq 2 \sqrt{d}\left(1-\frac{1}{2 d}\right)
\end{aligned}
$$

Further, the lower bound is achieved by local algorithms.

- A local algorithm achieves $8 / 9$ of $\operatorname{SDP}\left(\boldsymbol{A}^{\mathrm{cen}}\right)$.


## Erdős-Renyi random graphs

Theorem (Fan, Montanari, 2016)
Let $G \sim \mathcal{G}(n, d / n)$ and $A^{\text {cen }}=A_{G}^{\mathrm{cen}}$. Then, a.s.,

$$
\begin{aligned}
2 \sqrt{d}\left(1-\frac{1}{d+1}\right) & \leq \lim \inf _{n \rightarrow \infty} \frac{1}{n} \operatorname{SDP}\left(\boldsymbol{A}^{\mathrm{cen}}\right) \leq \\
& \leq \lim \sup _{n \rightarrow \infty} \frac{1}{n} \operatorname{SDP}\left(\boldsymbol{A}^{\mathrm{cen}}\right) \leq 2 \sqrt{d}\left(1-\frac{1}{2 d}\right)
\end{aligned}
$$

Further, the lower bound is achieved by local algorithms.

## Erdős-Renyi random graphs

## Theorem (Fan, Montanari, 2016)

Let $G \sim \mathcal{G}(n, d / n)$ and $A^{\text {cen }}=A_{G}^{\mathrm{cen}}$. Then, a.s.,

$$
\begin{aligned}
2 \sqrt{d}\left(1-\frac{1}{d+1}\right) & \leq \lim \inf _{n \rightarrow \infty} \frac{1}{n} \operatorname{SDP}\left(A^{\mathrm{cen}}\right) \leq \\
& \leq \lim \sup _{n \rightarrow \infty} \frac{1}{n} \operatorname{SDP}\left(A^{\mathrm{cen}}\right) \leq 2 \sqrt{d}\left(1-\frac{1}{2 d}\right)
\end{aligned}
$$

Further, the lower bound is achieved by local algorithms.

- A local algorithm achieves $8 / 9$ of $\operatorname{SDP}\left(\boldsymbol{A}^{\mathrm{cen}}\right)$.


## Bounds vs numerical simulations


[Related results for planted partition model, regular graphs, ...]

## A better local algorithm (see paper)


[Related results for planted partition model, regular graphs, ...]

## Conclusion

## Conclusion

- Which statistical problems are tractable?
- Multiple points of view:
- Local-message passing algorithms
- Landscape analysis
- SDP hierarchies


## Conclusion

- Which statistical problems are tractable?
- Multiple points of view:
- Local-message passing algorithms
- Landscape analysis
- SDP hierarchies


## Thanks!


[^0]:    ${ }^{1}$ Equivalently, $Y_{i_{1}, \ldots, i_{k}}=\lambda \theta_{0, i_{1}} \cdots \theta_{0, i_{k}}+W_{i_{1}, \ldots, i_{k}}$.

