Low-Rank Matrix Recovery

Yuxin Chen

Princeton University, Fall 2020

Outline

- Motivation
- Problem setup
- Nuclear norm minimization
  - RIP and low-rank matrix recovery
  - Phase retrieval / solving random quadratic systems of equations
  - Matrix completion

Motivation

Motivation 1: recommendation systems

- Netflix challenge: Netflix provides highly incomplete ratings from 0.5 million users for & 17,770 movies
- How to predict unseen user ratings for movies?

In general, we cannot infer missing ratings

$$
\begin{bmatrix}
\checkmark & ? & ? & ? & \checkmark & ? \\
? & ? & \checkmark & \checkmark & ? & ? \\
\checkmark & ? & ? & \checkmark & ? & ? \\
? & ? & \checkmark & ? & ? & \checkmark \\
\checkmark & ? & ? & ? & ? & ? \\
? & \checkmark & ? & ? & \checkmark & ? \\
? & ? & \checkmark & \checkmark & ? & ?
\end{bmatrix}
$$

Underdetermined system (more unknowns than observations)

... unless rating matrix has other structure

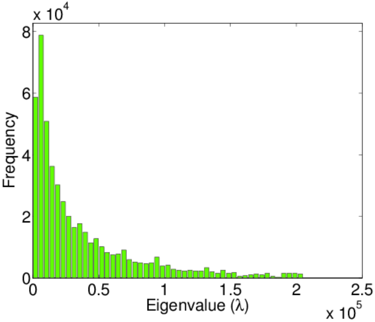

A few factors explain most of the data

... unless rating matrix has other structure

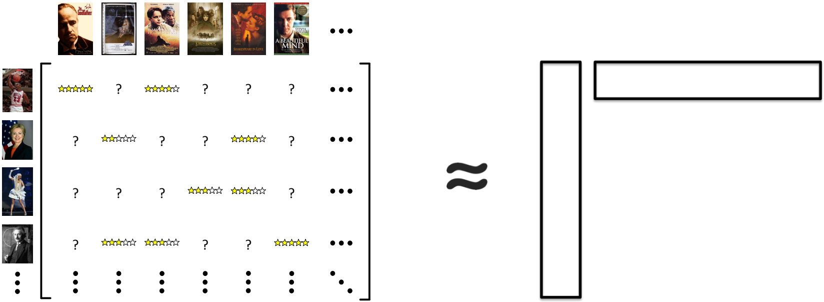

A few factors explain most of the data $\longrightarrow$ low-rank approximation

How to exploit (approx.) low-rank structure in prediction?

Motivation 2: sensor localization

- n sensors / points $\boldsymbol{x}_j \in \mathbb{R}^3$, $j = 1, \cdots, n$
- Observe partial information about pairwise distances

$$D_{i,j} = \|\boldsymbol{x}_i - \boldsymbol{x}_j\|_2^2 = \|\boldsymbol{x}_i\|_2^2 + \|\boldsymbol{x}_j\|_2^2 - 2\boldsymbol{x}_i^\top \boldsymbol{x}_j$$

- Goal: infer distance between every pair of nodes

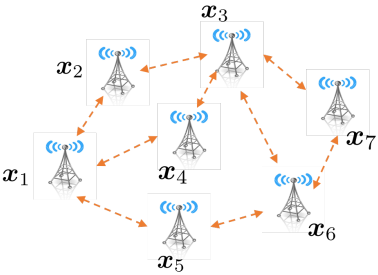

Motivation 2: sensor localization

Introduce

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_1^\top \\ \boldsymbol{x}_2^\top \\ \vdots \\ \boldsymbol{x}_n^\top \end{bmatrix} \in \mathbb{R}^{n \times 3}$$

then distance matrix $\boldsymbol{D} = [D_{i,j}]_{1 \le i,j \le n}$ can be written as

$$\boldsymbol{D} = \underbrace{\underbrace{\begin{bmatrix} \|\boldsymbol{x}_1\|_2^2 \\ \vdots \\ \|\boldsymbol{x}_n\|_2^2 \end{bmatrix} \boldsymbol{1}^\top}_{\text{rank } 1} + \underbrace{\boldsymbol{1} \cdot \left[\|\boldsymbol{x}_1\|_2^2, \cdots, \|\boldsymbol{x}_n\|_2^2\right]}_{\text{rank } 1} - \underbrace{2\boldsymbol{X}\boldsymbol{X}^\top}_{\text{rank } 3}}_{\text{low rank}}$$

$\text{rank}(\boldsymbol{D}) \ll n \quad \longrightarrow \quad$ low-rank matrix completion

Motivation 3: structure from motion

Given multiple images and a few correspondences between image features, how to estimate the locations of 3D points?

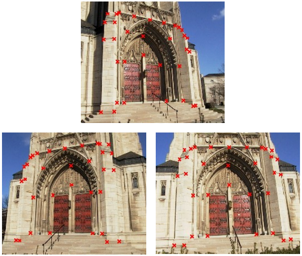

Snavely, Seitz, & Szeliski

Structure from motion: reconstruct 3D scene geometry (structure) and camera poses (motion) from multiple images

Motivation 3: structure from motion

Tomasi and Kanade's factorization:

- Consider n 3D points $\{\boldsymbol{p}_j \in \mathbb{R}^3\}_{1 \le j \le n}$ in m different 2D frames
- $\boldsymbol{x}_{i,j} \in \mathbb{R}^{2\times 1}$: locations of the j^{th} point in the i^{th} frame

$$\boldsymbol{x}_{i,j} = \underbrace{\boldsymbol{M}_i}_{\text{projection matrix } \in \mathbb{R}^{2\times 3}} \underbrace{\boldsymbol{p}_j}_{\text{3D position } \in \mathbb{R}^3}$$

Motivation 3: structure from motion

Tomasi and Kanade's factorization:

- Matrix of all 2D locations

$$\boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}_{1,1} & \cdots & \boldsymbol{x}_{1,n} \\ \vdots & \ddots & \vdots \\ \boldsymbol{x}_{m,1} & \cdots & \boldsymbol{x}_{m,n} \end{bmatrix} = \underbrace{\begin{bmatrix} \boldsymbol{M}_1 \\ \vdots \\ \boldsymbol{M}_m \end{bmatrix} \begin{bmatrix} \boldsymbol{p}_1 & \cdots & \boldsymbol{p}_n \end{bmatrix}}_{\text{low-rank factorization}} \in \mathbb{R}^{2m \times n}$$

Goal: fill in missing entries of $\boldsymbol{X}$ given a small number of entries

Motivation 4: missing phase problem

Detectors record intensities of diffracted rays

- electric field $x(t_1, t_2) \longrightarrow$ Fourier transform $\hat{x}(f_1, f_2)$

Fig credit: Stanford SLAC

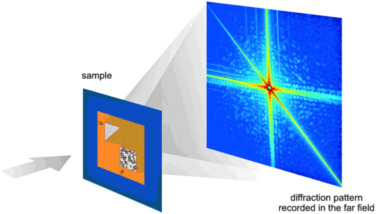

intensity of electrical field: $\left|\hat{x}(f_1, f_2)\right|^2 = \left|\int x(t_1, t_2) e^{-i2\pi(f_1 t_1 + f_2 t_2)} \mathrm{d}t_1 \mathrm{d}t_2\right|^2$

Phase retrieval: recover signal $x(t_1, t_2)$ from intensity $\left|\hat{x}(f_1, f_2)\right|^2$

A discrete-time model: solving quadratic systems

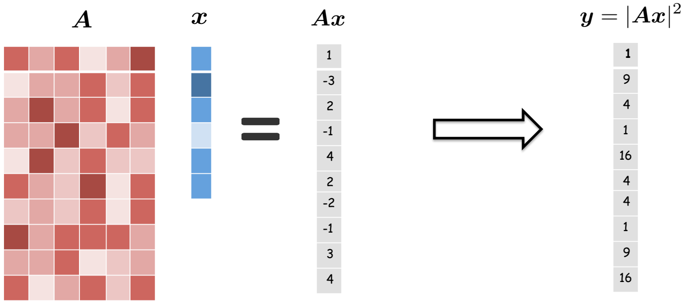

Solve for $\boldsymbol{x} \in \mathbb{R}^n$ in m quadratic equations

$$
\begin{array}{llll}
 & y_k & = & |\boldsymbol{a}_k^\top \boldsymbol{x}|^2, \qquad k = 1, \ldots, m \\
\text{or} & \boldsymbol{y} & = & |\boldsymbol{A}\boldsymbol{x}|^2 \qquad \text{where } |\boldsymbol{z}|^2 := \{|z_1|^2, \cdots, |z_m|^2\}
\end{array}
$$

An equivalent view: low-rank factorization

Lifting: introduce $\boldsymbol{X} = \boldsymbol{x}\boldsymbol{x}^*$ to linearize constraints

$$y_k \;=\; |\boldsymbol{a}_k^* \boldsymbol{x}|^2 = \boldsymbol{a}_k^*(\boldsymbol{x}\boldsymbol{x}^*)\boldsymbol{a}_k \quad \Longrightarrow \quad y_k = \boldsymbol{a}_k^* \boldsymbol{X} \boldsymbol{a}_k = \langle \boldsymbol{a}_k \boldsymbol{a}_k^*, \boldsymbol{X} \rangle \tag{13.1}$$

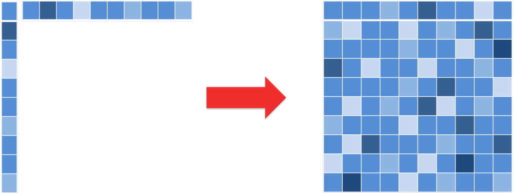

$$\begin{aligned} \text{find} \quad & \boldsymbol{X} \succeq \boldsymbol{0} \\ \text{s.t.} \quad & y_k \;=\; \langle \boldsymbol{a}_k \boldsymbol{a}_k^*, \boldsymbol{X} \rangle, \qquad k = 1, \cdots, m \\ & \text{rank}(\boldsymbol{X}) = 1 \end{aligned}$$

Problem setup

Setup

- Consider $\boldsymbol{M} \in \mathbb{R}^{n \times n}$
- $\text{rank}(\boldsymbol{M}) = r \ll n$
- Singular value decomposition (SVD) of $\boldsymbol{M}$:

$$\boldsymbol{M} = \underbrace{\boldsymbol{U\Sigma V}^\top}_{(2n-r)r \text{ degrees of freedom}} = \sum_{i=1}^{r} \sigma_i \boldsymbol{u}_i \boldsymbol{v}_i^\top$$

where $\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix}$ contains all singular values $\{\sigma_i\}$;

$\boldsymbol{U} := [\boldsymbol{u}_1, \cdots, \boldsymbol{u}_r]$, $\boldsymbol{V} := [\boldsymbol{v}_1, \cdots, \boldsymbol{v}_r]$ consist of singular vectors

Low-rank matrix completion

Observed entries

$$M_{i,j}, \qquad (i,j) \in \underbrace{\Omega}_{\text{sampling set}}$$

Completion via rank minimization

$$\text{minimize}_{\boldsymbol{X}} \quad \text{rank}(\boldsymbol{X}) \qquad \text{s.t.} \quad X_{i,j} = M_{i,j}, \quad (i,j) \in \Omega$$

Low-rank matrix completion

Observed entries

$$M_{i,j}, \qquad (i,j) \in \underbrace{\Omega}_{\text{sampling set}}$$

- An operator $\mathcal{P}_\Omega$: orthogonal projection onto the subspace of matrices supported on Ω

Completion via rank minimization

$$\text{minimize}_{\boldsymbol{X}} \quad \text{rank}(\boldsymbol{X}) \qquad \text{s.t.} \quad \mathcal{P}_\Omega(\boldsymbol{X}) = \mathcal{P}_\Omega(\boldsymbol{M})$$

More general: low-rank matrix recovery

Linear measurements

$$y_i = \langle \boldsymbol{A}_i, \boldsymbol{M} \rangle = \mathsf{Tr}(\boldsymbol{A}_i^\top \boldsymbol{M}), \qquad i = 1, \ldots m$$

- An operator form

$$\boldsymbol{y} = \mathcal{A}(\boldsymbol{M}) := \begin{bmatrix} \langle \boldsymbol{A}_1, \boldsymbol{M} \rangle \\ \vdots \\ \langle \boldsymbol{A}_m, \boldsymbol{M} \rangle \end{bmatrix} \in \mathbb{R}^m$$

Recovery via rank minimization

$$\text{minimize}_{\boldsymbol{X}} \quad \text{rank}(\boldsymbol{X}) \qquad \text{s.t.} \quad \boldsymbol{y} = \mathcal{A}(\boldsymbol{X})$$

Nuclear norm minimization

Convex relaxation

$$\begin{aligned} \text{minimize}_{\boldsymbol{X} \in \mathbb{R}^{n \times n}} \quad & \underbrace{\text{rank}(\boldsymbol{X})}_{\textcolor{blue}{\text{nonconvex}}} \\ \text{s.t.} \quad & \mathcal{P}_{\Omega}(\boldsymbol{X}) = \mathcal{P}_{\Omega}(\boldsymbol{M}) \end{aligned}$$

$$\begin{aligned} \text{minimize}_{\boldsymbol{X} \in \mathbb{R}^{n \times n}} \quad & \underbrace{\text{rank}(\boldsymbol{X})}_{\textcolor{blue}{\text{nonconvex}}} \\ \text{s.t.} \quad & \mathcal{A}(\boldsymbol{X}) = \mathcal{A}(\boldsymbol{M}) \end{aligned}$$

Question: what is the convex surrogate for $\text{rank}(\cdot)$?

Nuclear norm

Definition 13.1

The nuclear norm of $\boldsymbol{X}$ is

$$\|\boldsymbol{X}\|_* := \sum_{i=1}^{n} \underbrace{\sigma_i(\boldsymbol{X})}_{i^{\text{th}} \text{ largest singular value}}$$

- Nuclear norm is a counterpart of ℓ_1 norm for rank function
- Relations among different norms

$$\|\boldsymbol{X}\| \leq \|\boldsymbol{X}\|_{\mathsf{F}} \leq \|\boldsymbol{X}\|_* \leq \sqrt{r}\|\boldsymbol{X}\|_{\mathsf{F}} \leq r\|\boldsymbol{X}\|$$

- **(Tightness)** $\{\boldsymbol{X} : \|\boldsymbol{X}\|_* \leq 1\}$ is the convex hull of rank-1 matrices obeying $\|\boldsymbol{u}\boldsymbol{v}^\top\| \leq 1$ (Fazel '02)

Additivity of nuclear norm

Fact 13.2

Let $\boldsymbol{A}$ *and* $\boldsymbol{B}$ *be matrices of the same dimensions. If* $\boldsymbol{A}\boldsymbol{B}^\top = \boldsymbol{0}$ *and* $\boldsymbol{A}^\top \boldsymbol{B} = \boldsymbol{0}$, *then* $\|\boldsymbol{A}+\boldsymbol{B}\|_* = \|\boldsymbol{A}\|_* + \|\boldsymbol{B}\|_*$.

- If row (resp. column) spaces of $\boldsymbol{A}$ and $\boldsymbol{B}$ are orthogonal, then $\|\boldsymbol{A}+\boldsymbol{B}\|_* = \|\boldsymbol{A}\|_* + \|\boldsymbol{B}\|_*$

- Similar to ℓ_1 norm: when $\boldsymbol{x}$ and $\boldsymbol{y}$ have disjoint support,

$$\|\boldsymbol{x}+\boldsymbol{y}\|_1 = \|\boldsymbol{x}\|_1 + \|\boldsymbol{y}\|_1 \quad \text{(a key to study } \ell_1\text{-min under RIP)}$$

Proof of Fact 13.2

Suppose $\boldsymbol{A} = \boldsymbol{U}_A \boldsymbol{\Sigma}_A \boldsymbol{V}_A^\top$ and $\boldsymbol{B} = \boldsymbol{U}_B \boldsymbol{\Sigma}_B \boldsymbol{V}_B^\top$, which gives

$$\begin{aligned} \boldsymbol{A}\boldsymbol{B}^\top &= \boldsymbol{0} \\ \boldsymbol{A}^\top \boldsymbol{B} &= \boldsymbol{0} \end{aligned} \quad \Longleftrightarrow \quad \begin{aligned} \boldsymbol{V}_A^\top \boldsymbol{V}_B &= \boldsymbol{0} \\ \boldsymbol{U}_A^\top \boldsymbol{U}_B &= \boldsymbol{0} \end{aligned}$$

Thus, one can write

$$\begin{aligned} \boldsymbol{A} &= [\boldsymbol{U}_A, \boldsymbol{U}_B, \boldsymbol{U}_C] \begin{bmatrix} \boldsymbol{\Sigma}_A & & \\ & \boldsymbol{0} & \\ & & \boldsymbol{0} \end{bmatrix} [\boldsymbol{V}_A, \boldsymbol{V}_B, \boldsymbol{V}_C]^\top \\ \boldsymbol{B} &= [\boldsymbol{U}_A, \boldsymbol{U}_B, \boldsymbol{U}_C] \begin{bmatrix} \boldsymbol{0} & & \\ & \boldsymbol{\Sigma}_B & \\ & & \boldsymbol{0} \end{bmatrix} [\boldsymbol{V}_A, \boldsymbol{V}_B, \boldsymbol{V}_C]^\top \end{aligned}$$

and hence

$$\|\boldsymbol{A} + \boldsymbol{B}\|_* = \left\| [\boldsymbol{U}_A, \boldsymbol{U}_B] \begin{bmatrix} \boldsymbol{\Sigma}_A & \\ & \boldsymbol{\Sigma}_B \end{bmatrix} [\boldsymbol{V}_A, \boldsymbol{V}_B]^\top \right\|_* = \|\boldsymbol{A}\|_* + \|\boldsymbol{B}\|_*$$

Dual norm

Definition 13.3 (Dual norm)

For a given norm $\|\cdot\|_{\mathcal{A}}$, the dual norm is defined as

$$\|\boldsymbol{X}\|_{\mathcal{A}}^{\star} := \max\{\langle \boldsymbol{X}, \boldsymbol{Y} \rangle : \|\boldsymbol{Y}\|_{\mathcal{A}} \leq 1\}$$

- ℓ_1 norm $\overset{\text{dual}}{\longleftrightarrow}$ ℓ_∞ norm
- nuclear norm $\overset{\text{dual}}{\longleftrightarrow}$ spectral norm
- ℓ_2 norm $\overset{\text{dual}}{\longleftrightarrow}$ ℓ_2 norm
- Frobenius norm $\overset{\text{dual}}{\longleftrightarrow}$ Frobenius norm

Representing nuclear norm via SDP

Since the spectral norm is the dual norm of the nuclear norm,

$$\|\boldsymbol{X}\|_* = \max\{\langle \boldsymbol{X}, \boldsymbol{Y}\rangle : \|\boldsymbol{Y}\| \le 1\}$$

The constraint is equivalent to

$$\|\boldsymbol{Y}\| \le 1 \quad \Longleftrightarrow \quad \boldsymbol{Y}\boldsymbol{Y}^\top \preceq \boldsymbol{I} \quad \overset{\text{Schur complement}}{\Longleftrightarrow} \quad \begin{bmatrix} \boldsymbol{I} & \boldsymbol{Y} \\ \boldsymbol{Y}^\top & \boldsymbol{I} \end{bmatrix} \succeq \boldsymbol{0}$$

Fact 13.4

$$\|\boldsymbol{X}\|_* = \max_{\boldsymbol{Y}} \left\{ \langle \boldsymbol{X}, \boldsymbol{Y}\rangle \;\middle|\; \begin{bmatrix} \boldsymbol{I} & \boldsymbol{Y} \\ \boldsymbol{Y}^\top & \boldsymbol{I} \end{bmatrix} \succeq \boldsymbol{0} \right\}$$

Representing nuclear norm via SDP

Since the spectral norm is the dual norm of the nuclear norm,

$$\|\boldsymbol{X}\|_* = \max\{\langle \boldsymbol{X}, \boldsymbol{Y} \rangle : \|\boldsymbol{Y}\| \le 1\}$$

The constraint is equivalent to

$$\|\boldsymbol{Y}\| \le 1 \quad \Longleftrightarrow \quad \boldsymbol{Y}\boldsymbol{Y}^\top \preceq \boldsymbol{I} \quad \overset{\text{Schur complement}}{\Longleftrightarrow} \quad \begin{bmatrix} \boldsymbol{I} & \boldsymbol{Y} \\ \boldsymbol{Y}^\top & \boldsymbol{I} \end{bmatrix} \succeq \boldsymbol{0}$$

Fact 13.5 (Dual characterization)

$$\|\boldsymbol{X}\|_* = \min_{\boldsymbol{W}_1, \boldsymbol{W}_2} \left\{ \frac{1}{2}\mathsf{Tr}(\boldsymbol{W}_1) + \frac{1}{2}\mathsf{Tr}(\boldsymbol{W}_2) \;\middle|\; \begin{bmatrix} \boldsymbol{W}_1 & \boldsymbol{X} \\ \boldsymbol{X}^\top & \boldsymbol{W}_2 \end{bmatrix} \succeq \boldsymbol{0} \right\}$$

- Optimal point: $\boldsymbol{W}_1 = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{U}^\top$, $\boldsymbol{W}_2 = \boldsymbol{V}\boldsymbol{\Sigma}\boldsymbol{V}^\top$ (where $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\top$)

Aside: dual of semidefinite program

$$\begin{aligned}
\text{(primal)} \quad \text{minimize}_{\boldsymbol{X}} \quad & \langle \boldsymbol{C}, \boldsymbol{X} \rangle \\
\text{s.t.} \quad & \langle \boldsymbol{A}_i, \boldsymbol{X} \rangle = b_i, \quad 1 \le i \le m \\
& \boldsymbol{X} \succeq \boldsymbol{0}
\end{aligned}$$

$$\Updownarrow$$

$$\begin{aligned}
\text{(dual)} \quad \text{maximize}_{\boldsymbol{y}} \quad & \boldsymbol{b}^\top \boldsymbol{y} \\
\text{s.t.} \quad & \sum_{i=1}^{m} y_i \boldsymbol{A}_i + \boldsymbol{S} = \boldsymbol{C} \\
& \boldsymbol{S} \succeq \boldsymbol{0}
\end{aligned}$$

Exercise: use this to verify Fact 13.5

Nuclear norm minimization via SDP

Convex relaxation of rank minimization

$$\hat{\boldsymbol{M}} = \text{argmin}_{\boldsymbol{X}} \|\boldsymbol{X}\|_* \quad \text{s.t.} \quad \boldsymbol{y} = \mathcal{A}(\boldsymbol{X})$$

This is solvable via SDP

$$\begin{aligned} \text{minimize}_{\boldsymbol{X},\boldsymbol{W}_1,\boldsymbol{W}_2} \quad & \frac{1}{2}\text{Tr}(\boldsymbol{W}_1) + \frac{1}{2}\text{Tr}(\boldsymbol{W}_2) \\ \text{s.t.} \quad & \boldsymbol{y} = \mathcal{A}(\boldsymbol{X}), \quad \begin{bmatrix} \boldsymbol{W}_1 & \boldsymbol{X} \\ \boldsymbol{X}^\top & \boldsymbol{W}_2 \end{bmatrix} \succeq \boldsymbol{0} \end{aligned}$$

RIP and low-rank matrix recovery

RIP for low-rank matrices

Almost parallel results to compressed sensing …[1]

Definition 13.6

The r-restricted isometry constants $\delta_r^{\mathrm{ub}}(\mathcal{A})$ and $\delta_r^{\mathrm{lb}}(\mathcal{A})$ are the smallest quantities s.t.

$$(1-\delta_r^{\mathrm{lb}})\|\boldsymbol{X}\|_{\mathsf{F}} \le \|\mathcal{A}(\boldsymbol{X})\|_{\mathsf{F}} \le (1+\delta_r^{\mathrm{ub}})\|\boldsymbol{X}\|_{\mathsf{F}}, \qquad \forall \boldsymbol{X} : \mathrm{rank}(\boldsymbol{X}) \le r$$

[1]One can also define RIP w.r.t. $\|\cdot\|_{\mathrm{F}}^2$ rather than $\|\cdot\|_{\mathrm{F}}$.

RIP and low-rank matrix recovery

Theorem 13.7 (Recht, Fazel, Parrilo '10, Candes, Plan '11)

Suppose $\mathrm{rank}(\boldsymbol{M}) = r$. *For any fixed integer* $K > 0$, *if* $\frac{1+\delta_{Kr}^{\mathrm{ub}}}{1-\delta_{(2+K)r}^{\mathrm{lb}}} < \sqrt{\frac{K}{2}}$, *then nuclear norm minimization is exact*

- It allows $\delta_{Kr}^{\mathrm{ub}}$ to be larger than 1
- Can be easily extended to account for noisy case and approximately low-rank matrices

Geometry of nuclear norm ball

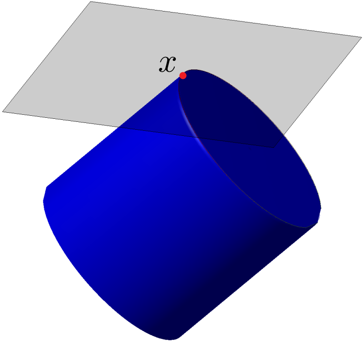

Level set of nuclear norm ball: $\left\| \begin{bmatrix} x & y \\ y & z \end{bmatrix} \right\|_* \leq 1$

Fig. credit: Candes '14

Some notation

Recall $\boldsymbol{M} = \boldsymbol{U\Sigma V}^\top$

- Let T be the span of matrices of the form (called *tangent space*)

$$T = \{\boldsymbol{U X}^\top + \boldsymbol{Y V}^\top : \boldsymbol{X}, \boldsymbol{Y} \in \mathbb{R}^{n\times r}\}$$

- Let $\mathcal{P}_T$ be the orthogonal projection onto T:

$$\mathcal{P}_T(\boldsymbol{X}) = \boldsymbol{U U}^\top \boldsymbol{X} + \boldsymbol{X V V}^\top - \boldsymbol{U U}^\top \boldsymbol{X V V}^\top$$

- Its complement $\mathcal{P}_{T^\perp} = \mathcal{I} - \mathcal{P}_T$:

$$\mathcal{P}_{T^\perp}(\boldsymbol{X}) = (\boldsymbol{I} - \boldsymbol{U U}^\top)\boldsymbol{X}(\boldsymbol{I} - \boldsymbol{V V}^\top)$$

  - $\boldsymbol{M}\mathcal{P}_{T^\perp}^\top(\boldsymbol{X}) = \boldsymbol{0}$ and $\boldsymbol{M}^\top\mathcal{P}_{T^\perp}(\boldsymbol{X}) = \boldsymbol{0}$

Proof of Theorem 13.7

Suppose $\boldsymbol{X} = \boldsymbol{M} + \boldsymbol{H}$ is feasible and obeys $\|\boldsymbol{M} + \boldsymbol{H}\|_* \leq \|\boldsymbol{M}\|_*$. The goal is to show that $\boldsymbol{H} = \boldsymbol{0}$ under RIP.

The key is to decompose $\boldsymbol{H}$ into $\boldsymbol{H}_0 + \underbrace{\boldsymbol{H}_1 + \boldsymbol{H}_2 + \ldots}_{\boldsymbol{H}_{\mathrm{c}}}$

- $\boldsymbol{H}_0 = \mathcal{P}_T(\boldsymbol{H})$ (rank $2r$)
- $\boldsymbol{H}_{\mathrm{c}} = \mathcal{P}_T^{\perp}(\boldsymbol{H})$ (obeying $\boldsymbol{M}\boldsymbol{H}_{\mathrm{c}}^{\top} = \boldsymbol{0}$ and $\boldsymbol{M}^{\top}\boldsymbol{H}_{\mathrm{c}} = \boldsymbol{0}$)
- $\boldsymbol{H}_1$: the best rank-(Kr) approximation of $\boldsymbol{H}_{\mathrm{c}}$ (K is const)
- $\boldsymbol{H}_2$: the best rank-(Kr) approximation of $\boldsymbol{H}_{\mathrm{c}} - \boldsymbol{H}_1$
- ...

Proof of Theorem 13.7

Informally, the proof proceeds by showing that

1. $\boldsymbol{H}_0$ "dominates" $\sum_{i\geq 2} \boldsymbol{H}_i$ (by objective function)
   — *see Step 1*
2. (converse) $\sum_{i\geq 2} \boldsymbol{H}_i$ "dominates" $\boldsymbol{H}_0 + \boldsymbol{H}_1$ (by RIP + feasibility)
   — *see Step 2*

These cannot happen simultaneously unless $\boldsymbol{H} = \boldsymbol{0}$

Proof of Theorem 13.7

Step 1 (which does not rely on RIP). Show that

$$\sum_{j\geq 2} \|\boldsymbol{H}_j\|_{\mathrm{F}} \leq \|\boldsymbol{H}_0\|_* / \sqrt{Kr}. \tag{13.2}$$

This follows immediately by combining the following 2 observations:

(i) Since $\boldsymbol{M}+\boldsymbol{H}$ is assumed to be a better estimate:

$$\begin{aligned} \|\boldsymbol{M}\|_* \geq \|\boldsymbol{M}+\boldsymbol{H}\|_* &\geq \|\boldsymbol{M}+\boldsymbol{H}_{\mathrm{c}}\|_* - \|\boldsymbol{H}_0\|_* \qquad (13.3) \\ &\geq \underbrace{\|\boldsymbol{M}\|_* + \|\boldsymbol{H}_{\mathrm{c}}\|_*}_{\text{Fact 13.2 } (\boldsymbol{M}\boldsymbol{H}_{\mathrm{c}}^\top = \boldsymbol{0} \text{ and } \boldsymbol{M}^\top \boldsymbol{H}_{\mathrm{c}} = \boldsymbol{0})} - \|\boldsymbol{H}_0\|_* \end{aligned}$$

$$\Longrightarrow \quad \|\boldsymbol{H}_{\mathrm{c}}\|_* \leq \|\boldsymbol{H}_0\|_* \tag{13.4}$$

(ii) Since nonzero singular values of $\boldsymbol{H}_{j-1}$ dominate those of $\boldsymbol{H}_j$ $(j \geq 2)$:

$$\|\boldsymbol{H}_j\|_{\mathrm{F}} \leq \sqrt{Kr}\|\boldsymbol{H}_j\| \leq \sqrt{Kr}\left[\|\boldsymbol{H}_{j-1}\|_*/(Kr)\right] \leq \|\boldsymbol{H}_{j-1}\|_*/\sqrt{Kr}$$

$$\Longrightarrow \quad \sum_{j\geq 2} \|\boldsymbol{H}_j\|_{\mathrm{F}} \leq \frac{1}{\sqrt{Kr}} \sum_{j\geq 2} \|\boldsymbol{H}_{j-1}\|_* \leq \frac{1}{\sqrt{Kr}} \|\boldsymbol{H}_{\mathrm{c}}\|_* \tag{13.5}$$

Proof of Theorem 13.7

Step 2 (using feasibility + RIP). Show that $\exists \rho < \sqrt{K/2}$ s.t.

$$\|\boldsymbol{H}_0 + \boldsymbol{H}_1\|_{\mathrm{F}} \le \rho \sum\nolimits_{j\ge 2} \|\boldsymbol{H}_j\|_{\mathrm{F}} \tag{13.6}$$

If this claim holds, then

$$\begin{aligned}\|\boldsymbol{H}_0 + \boldsymbol{H}_1\|_{\mathrm{F}} &\le \rho \sum\nolimits_{j\ge 2} \|\boldsymbol{H}_j\|_{\mathrm{F}} \overset{(13.2)}{\le} \rho \frac{1}{\sqrt{Kr}} \|\boldsymbol{H}_0\|_* \\ &\le \rho \frac{1}{\sqrt{Kr}} \left(\sqrt{2r}\|\boldsymbol{H}_0\|_{\mathrm{F}}\right) = \rho \sqrt{\frac{2}{K}} \|\boldsymbol{H}_0\|_{\mathrm{F}} \\ &\le \rho \sqrt{\frac{2}{K}} \|\boldsymbol{H}_0 + \boldsymbol{H}_1\|_{\mathrm{F}} \end{aligned} \tag{13.7}$$

where the last line holds since, by construction, $\boldsymbol{H}_0$ and $\boldsymbol{H}_1$ lie in orthogonal subspaces.

This bound (13.7) cannot hold with $\rho < \sqrt{K/2}$ unless $\underbrace{\boldsymbol{H}_0 + \boldsymbol{H}_1 = \boldsymbol{0}}_{\text{equivalently, } \boldsymbol{H}_0 = \boldsymbol{H}_1 = \boldsymbol{0}}$

Proof of Theorem 13.7

We now prove (13.6). To connect $\boldsymbol{H}_0 + \boldsymbol{H}_1$ with $\sum_{j\geq 2} \boldsymbol{H}_j$, we use feasibility:

$$\mathcal{A}(\boldsymbol{H}) = \boldsymbol{0} \quad \Longleftrightarrow \quad \mathcal{A}(\boldsymbol{H}_0 + \boldsymbol{H}_1) = -\sum_{j\geq 2} \mathcal{A}(\boldsymbol{H}_j),$$

which taken collectively with RIP yields

$$\begin{aligned}
(1-\delta^{\text{lb}}_{(2+K)r})\|\boldsymbol{H}_0 + \boldsymbol{H}_1\|_{\text{F}} &\leq \big\|\mathcal{A}(\boldsymbol{H}_0 + \boldsymbol{H}_1)\big\|_{\text{F}} = \Big\|\sum_{j\geq 2} \mathcal{A}(\boldsymbol{H}_j)\Big\|_{\text{F}} \\
&\leq \sum_{j\geq 2} \big\|\mathcal{A}(\boldsymbol{H}_j)\big\|_{\text{F}} \\
&\leq \sum_{j\geq 2} (1+\delta^{\text{ub}}_{Kr})\|\boldsymbol{H}_j\|_{\text{F}}
\end{aligned}$$

This establishes (13.6) as long as $\rho := \frac{1+\delta^{\text{ub}}_{Kr}}{1-\delta^{\text{lb}}_{(2+K)r}} < \sqrt{\frac{K}{2}}$.

Gaussian sampling operators satisfy RIP

If the entries of $\{\boldsymbol{A}_i\}_{i=1}^m$ are i.i.d. $\mathcal{N}(0, 1/m)$, then

$$\delta_{5r}(\mathcal{A}) < \frac{\sqrt{3}-\sqrt{2}}{\sqrt{3}+\sqrt{2}}$$

with high prob., provided that

$$m \;\gtrsim\; nr \quad \text{(near-optimal sample size)}$$

This satisfies the assumption of Theorem 13.7 with $K = 3$

Precise phase transition

Using the statistical dimension machienry, we can locate precise phase transition (Amelunxen, Lotz, McCoy & Tropp '13)

$$\text{nuclear norm min} \begin{cases} \text{works if} & m > \text{stat-dim}\big(\mathcal{D}\left(\|\cdot\|_*, \boldsymbol{X}\right)\big) \\ \text{fails if} & m < \text{stat-dim}\big(\mathcal{D}\left(\|\cdot\|_*, \boldsymbol{X}\right)\big) \end{cases}$$

where

$$\text{stat-dim}\big(\mathcal{D}\left(\|\cdot\|_*, \boldsymbol{X}\right)\big) \;\approx\; n^2 \psi\left(\frac{r}{n}\right)$$

and

$$\psi\left(\rho\right) = \inf_{\tau \geq 0} \left\{ \rho + (1-\rho)\left[\rho(1+\tau^2) + (1-\rho)\int_{\tau}^{2} \left(u-\tau\right)^2 \frac{\sqrt{4-u^2}}{\pi} \mathrm{d}u \right] \right\}$$

Aside: subgradient of nuclear norm

Subdifferential (set of subgradients) of $\|\cdot\|_*$ at $\boldsymbol{M}$ is

$$\partial\|\boldsymbol{M}\|_* = \Big\{\boldsymbol{U}\boldsymbol{V}^\top + \boldsymbol{W}: \quad \mathcal{P}_T(\boldsymbol{W}) = 0,\ \|\boldsymbol{W}\| \le 1\Big\}$$

- Does not depend on the singular values of $\boldsymbol{M}$
- $\boldsymbol{Z} \in \partial\|\boldsymbol{M}\|_*$ iff

$$\mathcal{P}_T(\boldsymbol{Z}) = \boldsymbol{U}\boldsymbol{V}^\top, \quad \|\mathcal{P}_{T^\perp}(\boldsymbol{Z})\| \le 1.$$

Derivation of the statistical dimension

WLOG, suppose $\boldsymbol{X} = \begin{bmatrix} \boldsymbol{I}_r & \\ & \boldsymbol{0} \end{bmatrix}$, then $\partial \|\boldsymbol{X}\|_* = \left\{ \begin{bmatrix} \boldsymbol{I}_r & \\ & \boldsymbol{W} \end{bmatrix} \mid \|\boldsymbol{W}\| \leq 1 \right\}$.
Let $\boldsymbol{G} = \begin{bmatrix} \boldsymbol{G}_{11} & \boldsymbol{G}_{12} \\ \boldsymbol{G}_{21} & \boldsymbol{G}_{22} \end{bmatrix}$ be i.i.d. standard Gaussian.

From the convex geometry lecture, we know that

$$\begin{aligned} \text{stat-dim}\big(\mathcal{D}(\|\cdot\|_*, \boldsymbol{X})\big) \;&\approx\; \inf_{\tau \geq 0} \mathbb{E}\left[\inf_{\boldsymbol{Z} \in \partial \|\boldsymbol{X}\|_*} \|\boldsymbol{G} - \tau \boldsymbol{Z}\|_{\mathrm{F}}^2 \right] \\ &= \inf_{\tau \geq 0} \mathbb{E}\left[\inf_{\boldsymbol{W} : \|\boldsymbol{W}\| \leq 1} \left\| \begin{bmatrix} \boldsymbol{G}_{11} & \boldsymbol{G}_{12} \\ \boldsymbol{G}_{21} & \boldsymbol{G}_{22} \end{bmatrix} - \tau \begin{bmatrix} \boldsymbol{I}_r & \\ & \boldsymbol{W} \end{bmatrix} \right\|_{\mathrm{F}}^2 \right] \end{aligned}$$

Derivation of statistical dimension

Observe that

$$\mathbb{E}\left[\inf_{\boldsymbol{W}:\|\boldsymbol{W}\|\leq 1}\left\|\left[\begin{array}{cc}\boldsymbol{G}_{11} & \boldsymbol{G}_{12}\\ \boldsymbol{G}_{21} & \boldsymbol{G}_{22}\end{array}\right]-\tau\left[\begin{array}{cc}\boldsymbol{I}_r & \\ & \boldsymbol{W}\end{array}\right]\right\|_{\mathrm{F}}^2\right]$$

$$=\mathbb{E}\left[\|\boldsymbol{G}_{11}-\tau\boldsymbol{I}_r\|_{\mathrm{F}}^2+\|\boldsymbol{G}_{21}\|_{\mathrm{F}}^2+\|\boldsymbol{G}_{12}\|_{\mathrm{F}}^2+\inf_{\|\boldsymbol{W}\|\leq 1}\|\boldsymbol{G}_{22}-\tau\boldsymbol{W}\|_{\mathrm{F}}^2\right]$$

$$=r\left(2n-r+\tau^2\right)+\mathbb{E}\left[\sum_{i=1}^{n-r}\left(\sigma_i\left(\boldsymbol{G}_{22}\right)-\tau\right)_+^2\right].$$

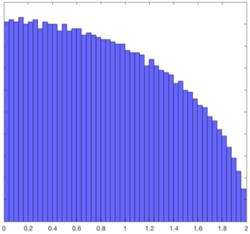

empirical distributions of $\{\sigma_i(\boldsymbol{G}_{22})/\sqrt{n-r}\}$

Derivation of statistical dimension

Observe that

$$\begin{aligned}
&\mathbb{E}\left[\inf_{\boldsymbol{W}:\|\boldsymbol{W}\|\leq 1}\left\|\begin{bmatrix}\boldsymbol{G}_{11} & \boldsymbol{G}_{12}\\ \boldsymbol{G}_{21} & \boldsymbol{G}_{22}\end{bmatrix} - \tau\begin{bmatrix}\boldsymbol{I}_r & \\ & \boldsymbol{W}\end{bmatrix}\right\|_{\mathrm{F}}^2\right]\\
&= \mathbb{E}\left[\|\boldsymbol{G}_{11} - \tau\boldsymbol{I}_r\|_{\mathrm{F}}^2 + \|\boldsymbol{G}_{21}\|_{\mathrm{F}}^2 + \|\boldsymbol{G}_{12}\|_{\mathrm{F}}^2 + \inf_{\|\boldsymbol{W}\|\leq 1}\|\boldsymbol{G}_{22} - \tau\boldsymbol{W}\|_{\mathrm{F}}^2\right]\\
&= r\left(2n - r + \tau^2\right) + \mathbb{E}\left[\sum\nolimits_{i=1}^{n-r}\left(\sigma_i\left(\boldsymbol{G}_{22}\right) - \tau\right)_+^2\right].
\end{aligned}$$

Recall from random matrix theory (Marchenko-Pastur law)

$$\frac{1}{n-r}\mathbb{E}\left[\sum_{i=1}^{n-r}\left(\sigma_i\left(\tilde{\boldsymbol{G}}_{22}\right) - \tau\right)_+^2\right] \to \int_0^2 (u-\tau)_+^2\frac{\sqrt{4-u^2}}{\pi}\mathrm{d}u,$$

where $\tilde{\boldsymbol{G}}_{22} \sim \mathcal{N}\left(\mathbf{0}, \frac{1}{n-r}\boldsymbol{I}\right)$. Taking $\rho = r/n$ and minimizing over τ lead to closed-form expression for phase transition boundary.

Numerical phase transition ($n = 30$)

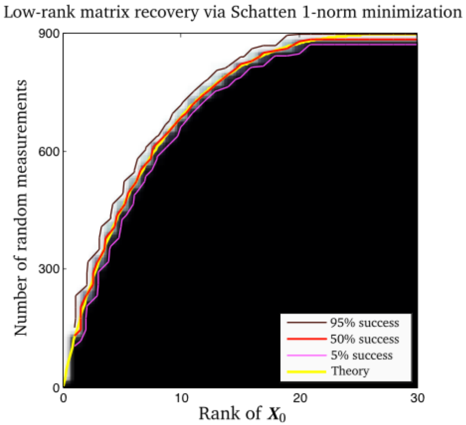

Figure credit: Amelunxen, Lotz, McCoy, & Tropp '13

Sampling operators that do NOT satisfy RIP

Unfortunately, many sampling operators fail to satisfy RIP
(e.g. none of the 4 motivating examples in this lecture satisfies RIP)

Two important examples:

- Phase retrieval / solving random quadratic systems of equations
- Matrix completion

Phase retrieval / solving random quadratic systems of equations

Rank-one measurements

Measurements: see (13.1)

$$y_i = \boldsymbol{a}_i^\top \underbrace{\boldsymbol{x}\boldsymbol{x}^\top}_{:=\boldsymbol{M}} \boldsymbol{a}_i = \langle \underbrace{\boldsymbol{a}_i\boldsymbol{a}_i^\top}_{:=\boldsymbol{A}_i}, \boldsymbol{M} \rangle, \qquad 1 \le i \le m$$

$$\mathcal{A}(\boldsymbol{X}) = \begin{bmatrix} \langle \boldsymbol{A}_1, \boldsymbol{X} \rangle \\ \langle \boldsymbol{A}_2, \boldsymbol{X} \rangle \\ \vdots \\ \langle \boldsymbol{A}_m, \boldsymbol{X} \rangle \end{bmatrix} = \begin{bmatrix} \langle \boldsymbol{a}_1\boldsymbol{a}_1^\top, \boldsymbol{X} \rangle \\ \langle \boldsymbol{a}_2\boldsymbol{a}_2^\top, \boldsymbol{X} \rangle \\ \vdots \\ \langle \boldsymbol{a}_m\boldsymbol{a}_m^\top, \boldsymbol{X} \rangle \end{bmatrix}$$

Rank-one measurements

Suppose $\boldsymbol{a}_i \overset{\text{ind.}}{\sim} \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_n)$

- If $\boldsymbol{x}$ is independent of $\{\boldsymbol{a}_i\}$, then

$$\langle \boldsymbol{a}_i \boldsymbol{a}_i^\top, \boldsymbol{x}\boldsymbol{x}^\top \rangle = \left| \boldsymbol{a}_i^\top \boldsymbol{x} \right|^2 \asymp \|\boldsymbol{x}\|_2^2 \ \Rightarrow \ \left\| \mathcal{A}(\boldsymbol{x}\boldsymbol{x}^\top) \right\|_{\mathrm{F}} \asymp \sqrt{m} \|\boldsymbol{x}\boldsymbol{x}^\top\|_{\mathrm{F}}$$

- Consider $\boldsymbol{A}_i = \boldsymbol{a}_i \boldsymbol{a}_i^\top$: with high prob.,

$$\langle \boldsymbol{a}_i \boldsymbol{a}_i^\top, \boldsymbol{A}_i \rangle = \|\boldsymbol{a}_i\|_2^4 \approx n \|\boldsymbol{a}_i \boldsymbol{a}_i^\top\|_{\mathrm{F}}$$

$$\Longrightarrow \quad \|\mathcal{A}(\boldsymbol{A}_i)\|_{\mathrm{F}} \geq |\langle \boldsymbol{a}_i \boldsymbol{a}_i^\top, \boldsymbol{A}_i \rangle| \approx n \|\boldsymbol{A}_i\|_{\mathrm{F}}$$

Rank-one measurements

Suppose $\boldsymbol{a}_i \overset{\text{ind.}}{\sim} \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}_n)$

- If the sample size $m \asymp n$ (information limit) and $K \asymp 1$, then

$$\frac{\max_{\boldsymbol{X}:\, \mathsf{rank}(\boldsymbol{X})=1} \frac{\|\mathcal{A}(\boldsymbol{X})\|_{\mathrm{F}}}{\|\boldsymbol{X}\|_{\mathrm{F}}}}{\min_{\boldsymbol{X}:\, \mathsf{rank}(\boldsymbol{X})=1} \frac{\|\mathcal{A}(\boldsymbol{X})\|_{\mathrm{F}}}{\|\boldsymbol{X}\|_{\mathrm{F}}}} \gtrsim \frac{n}{\sqrt{m}} \gtrsim \sqrt{n}$$

$$\Longrightarrow \quad \frac{1+\delta_K^{\mathrm{ub}}}{1-\delta_{2+K}^{\mathrm{lb}}} \geq \frac{\max_{\boldsymbol{X}:\, \mathsf{rank}(\boldsymbol{X})=1} \frac{\|\mathcal{A}(\boldsymbol{X})\|_{\mathrm{F}}}{\|\boldsymbol{X}\|_{\mathrm{F}}}}{\min_{\boldsymbol{X}:\, \mathsf{rank}(\boldsymbol{X})=1} \frac{\|\mathcal{A}(\boldsymbol{X})\|_{\mathrm{F}}}{\|\boldsymbol{X}\|_{\mathrm{F}}}} \gtrsim \sqrt{n} \gg \sqrt{K}$$

  - Violate RIP condition in Theorem 13.7 unless K is exceeding large

Why do we lose RIP?

Problems:

- Some low-rank matrices $\boldsymbol{X}$ (e.g. $\boldsymbol{a}_i\boldsymbol{a}_i^\top$) might be too aligned with some (rank-1) measurement matrices
  - loss of "incoherence" in some measurements

- Some measurements $\langle \boldsymbol{A}_i, \boldsymbol{X} \rangle$ might have too high of a leverage on $\mathcal{A}(\boldsymbol{X})$ when measured in $\|\cdot\|_\mathrm{F}$
  - Solution: replace $\|\cdot\|_\mathrm{F}$ by other norms!

Mixed-norm RIP

Solution: modify RIP appropriately ...

Definition 13.8 (RIP-ℓ_2/ℓ_1)

Let $\xi_r^{\mathrm{ub}}(\mathcal{A})$ and $\xi_r^{\mathrm{lb}}(\mathcal{A})$ be the smallest quantities s.t.

$$(1-\xi_r^{\mathrm{lb}})\|\boldsymbol{X}\|_{\mathsf{F}} \leq \|\mathcal{A}(\boldsymbol{X})\|_1 \leq (1+\xi_r^{\mathrm{ub}})\|\boldsymbol{X}\|_{\mathsf{F}}, \qquad \forall \boldsymbol{X} : \mathsf{rank}(\boldsymbol{X}) \leq r$$

- More generally, it only requires $\mathcal{A}$ to satisfy

$$\frac{\sup_{\boldsymbol{X}:\mathsf{rank}(\boldsymbol{X})\leq r} \frac{\|\mathcal{A}(\boldsymbol{X})\|_1}{\|\boldsymbol{X}\|_{\mathrm{F}}}}{\inf_{\boldsymbol{X}:\mathsf{rank}(\boldsymbol{X})\leq r} \frac{\|\mathcal{A}(\boldsymbol{X})\|_1}{\|\boldsymbol{X}\|_{\mathrm{F}}}} \leq \frac{1+\xi_r^{\mathrm{ub}}}{1-\xi_r^{\mathrm{lb}}} \tag{13.8}$$

Analyzing phase retrieval via RIP-ℓ_2/ℓ_1

Theorem 13.9 (Chen, Chi, Goldsmith '15)

Theorem 13.7 continues to hold if we replace δ_r^{ub} and δ_r^{lb} with ξ_r^{ub} and ξ_r^{lb} (defined in (13.8)), respectively

- Follows the same proof as for Theorem 13.7, except that $\|\cdot\|_{\mathrm{F}}$ (highlighted in red) is replaced by $\|\cdot\|_1$ in Slide 13-36

Analyzing phase retrieval via RIP-ℓ_2/ℓ_1

Theorem 13.9 (Chen, Chi, Goldsmith '15)

Theorem 13.7 continues to hold if we replace δ_r^{ub} *and* δ_r^{lb} *with* ξ_r^{ub} *and* ξ_r^{lb} *(defined in* (13.8)*), respectively*

- Back to the example in Slide 13-46:
  - If $\boldsymbol{x}$ is independent of $\{\boldsymbol{a}_i\}$, then

$$\left\langle \boldsymbol{a}_i\boldsymbol{a}_i^\top, \boldsymbol{x}\boldsymbol{x}^\top \right\rangle = \left|\boldsymbol{a}_i^\top \boldsymbol{x}\right|^2 \asymp \|\boldsymbol{x}\|_2^2 \quad \Rightarrow \quad \left\|\mathcal{A}\big(\boldsymbol{x}\boldsymbol{x}^\top\big)\right\|_1 \asymp m\|\boldsymbol{x}\boldsymbol{x}^\top\|_{\mathrm{F}}$$

  - $\|\mathcal{A}(\boldsymbol{A}_i)\|_1 = |\langle \boldsymbol{a}_i\boldsymbol{a}_i^\top, \boldsymbol{A}_i\rangle| + \sum_{j:j\neq i} |\langle \boldsymbol{a}_i\boldsymbol{a}_i^\top, \boldsymbol{A}_j\rangle| \approx (n+m)\|\boldsymbol{A}_i\|_{\mathrm{F}}$
  - For both cases, $\frac{\|\mathcal{A}(\boldsymbol{X})\|_1}{\|\boldsymbol{X}\|_{\mathrm{F}}}$ are of the same order if $m \gg n$

Analyzing phase retrieval via RIP-ℓ_2/ℓ_1

Informally, a debiased operator satisfies RIP condition of Theorem 13.9 when $m \gtrsim nr$ (Chen, Chi, Goldsmith '15)

$$\mathcal{B}(\boldsymbol{X}) := \begin{bmatrix} \langle \boldsymbol{A}_1 - \boldsymbol{A}_2, \boldsymbol{X} \rangle \\ \langle \boldsymbol{A}_3 - \boldsymbol{A}_4, \boldsymbol{X} \rangle \\ \vdots \end{bmatrix} \in \mathbb{R}^{m/2}$$

- Debiasing is crucial when $r \gg 1$
- A consequence of the Hanson-Wright inequality for quadratic form (Hanson & Wright '71, Rudelson & Vershynin '03)

Theoretical guarantee for phase retrieval

$$(\textbf{PhaseLift}) \qquad \underset{\boldsymbol{X}\in\mathbb{R}^{n\times n}}{\text{minimize}} \qquad \underbrace{\operatorname{tr}\boldsymbol{X}}_{\|\cdot\|_* \text{ for PSD matrices}}$$

$$\text{s.t.} \qquad y_i = \boldsymbol{a}_i^\top \boldsymbol{X} \boldsymbol{a}_i, \quad 1 \le i \le m$$

$$\boldsymbol{X} \succeq \boldsymbol{0} \quad (\text{since } \boldsymbol{X} = \boldsymbol{x}\boldsymbol{x}^\top)$$

Theorem 13.10 (Candès, Strohmer, Voroninski '13, Candès, Li '14)

Suppose $\boldsymbol{a}_i \overset{\text{ind.}}{\sim} \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$. *With high prob., PhaseLift recovers* $\boldsymbol{x}\boldsymbol{x}^\top$ *exactly as soon as* $m \gtrsim n$

Extension of phase retrieval

$$
\text{(PhaseLift)} \qquad \underset{\boldsymbol{X} \in \mathbb{R}^{n \times n}}{\text{minimize}} \quad \underbrace{\operatorname{tr} \boldsymbol{X}}_{\|\cdot\|_* \text{ for PSD matrices}}
$$

$$
\text{s.t.} \quad \boldsymbol{a}_i^\top \boldsymbol{X} \boldsymbol{a}_i = \boldsymbol{a}_i^\top \boldsymbol{M} \boldsymbol{a}_i, \quad 1 \leq i \leq m
$$

$$
\boldsymbol{X} \succeq \boldsymbol{0}
$$

Theorem 13.11 (Chen, Chi, Goldsmith '15, Cai, Zhang '15)

Suppose $\boldsymbol{M} \succeq \boldsymbol{0}$, rank$(\boldsymbol{M}) = r$, and $\boldsymbol{a}_i \overset{\text{ind.}}{\sim} \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$. With high prob., PhaseLift recovers $\boldsymbol{M}$ exactly as soon as $m \gtrsim nr$

Matrix completion

Sampling operators for matrix completion

Observation operator (projection onto matrices supported on Ω)

$$\boldsymbol{Y} = \mathcal{P}_\Omega(\boldsymbol{M})$$

where $(i,j) \in \Omega$ with prob. p (random sampling)

- $\mathcal{P}_\Omega$ does NOT satisfy RIP when $p \ll 1$!
- For example,

$$\underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}}_{\boldsymbol{M}} \qquad \underbrace{\begin{bmatrix} ? & \checkmark & ? & \checkmark & \checkmark \\ \checkmark & ? & \checkmark & ? & \checkmark \\ ? & \checkmark & \checkmark & ? & ? \\ \checkmark & ? & ? & \checkmark & ? \\ \checkmark & ? & \checkmark & ? & \checkmark \end{bmatrix}}_{\Omega}$$

$\|\mathcal{P}_\Omega(\boldsymbol{M})\|_{\mathrm{F}} = 0$, or equivalently, $\frac{1+\delta_K^{\mathrm{ub}}}{1-\delta_{2+K}^{\mathrm{lb}}} = \infty$

Which sampling pattern?

Consider the following sampling pattern

$$\begin{bmatrix} \checkmark & \checkmark & \checkmark & \checkmark & \checkmark \\ ? & ? & ? & ? & ? \\ \checkmark & \checkmark & \checkmark & \checkmark & \checkmark \\ \checkmark & \checkmark & \checkmark & \checkmark & \checkmark \\ \checkmark & \checkmark & \checkmark & \checkmark & \checkmark \end{bmatrix}$$

- If some rows / columns are not sampled, recovery is impossible

Which low-rank matrices can we recover?

Compare the following rank-1 matrices:

$$\begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \qquad \longleftarrow \qquad \begin{bmatrix} ? & 0 & ? & \cdots & 0 \\ 0 & ? & 0 & \cdots & ? \\ \vdots & \vdots & \vdots & & \\ ? & 0 & ? & \cdots & 0 \end{bmatrix}$$

if we miss the top-left entry, then we cannot hope to recover the matrix

Which low-rank matrices can we recover?

Compare the following rank-1 matrices:

$$\begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \vdots & & \\ 1 & 1 & 1 & \cdots & 1 \end{bmatrix} \qquad \longleftarrow \qquad \begin{bmatrix} {\color{red}?} & 1 & {\color{red}?} & \cdots & 1 \\ 1 & {\color{red}?} & 1 & \cdots & {\color{red}?} \\ \vdots & \vdots & \vdots & & \\ {\color{red}?} & 1 & {\color{red}?} & \cdots & 1 \end{bmatrix}$$

it is possible to fill in all missing entries by exploiting the rank-1 structure

Which low-rank matrices can we recover?

Compare the following rank-1 matrices:

$$\underbrace{\begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix}}_{\text{hard}} \quad \text{vs.} \quad \underbrace{\begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \vdots & & \\ 1 & 1 & 1 & \cdots & 1 \end{bmatrix}}_{\text{easy}}$$

Column / row spaces cannot be aligned with canonical basis vectors

Coherence

Definition 13.12

Coherence parameter μ of $\boldsymbol{M} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\top$ is the smallest quantity s.t.

$$\max_i \|\boldsymbol{U}^\top \boldsymbol{e}_i\|_2^2 \leq \frac{\mu r}{n} \quad \text{and} \quad \max_i \|\boldsymbol{V}^\top \boldsymbol{e}_i\|_2^2 \leq \frac{\mu r}{n}$$

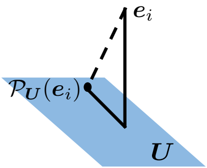

- $\mu \geq 1$ (since $\sum_{i=1}^n \|\boldsymbol{U}^\top \boldsymbol{e}_i\|_2^2 = \|\boldsymbol{U}\|_\mathrm{F}^2 = r$)
- $\mu = 1$ if $\frac{1}{\sqrt{n}}\mathbf{1} = \boldsymbol{U} = \boldsymbol{V}$ (most incoherent)
- $\mu = \frac{n}{r}$ if $\boldsymbol{e}_i \in \boldsymbol{U}$ (most coherent)

Performance guarantee

Theorem 13.13 (Candes & Recht '09, Candes & Tao '10, Gross '11, ...)

Nuclear norm minimization is exact and unique with high probability, provided that

$$m \gtrsim \mu n r \log^2 n$$

- This result is optimal up to a logarithmic factor
- Established via a RIPless theory

Numerical performance of nuclear-norm minimization

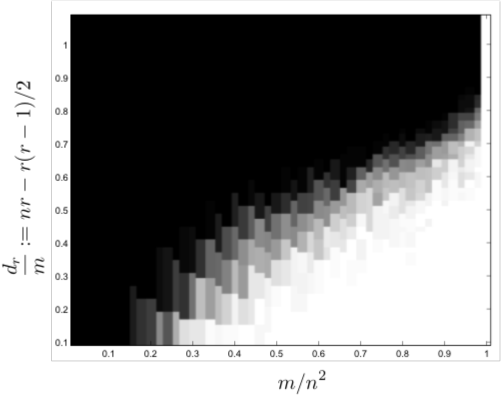

$$n = 50$$

Fig. credit: Candes, Recht '09

KKT condition

Lagrangian:

$$\mathcal{L}(\boldsymbol{X}, \boldsymbol{\Lambda}) = \|\boldsymbol{X}\|_* + \langle \boldsymbol{\Lambda}, \mathcal{P}_\Omega(\boldsymbol{X}) - \mathcal{P}_\Omega(\boldsymbol{M}) \rangle = \|\boldsymbol{X}\|_* + \langle \mathcal{P}_\Omega(\boldsymbol{\Lambda}), \boldsymbol{X} - \boldsymbol{M} \rangle$$

When $\boldsymbol{M}$ is the minimizer, the KKT condition reads

$$\boldsymbol{0} \in \partial_{\boldsymbol{X}} \mathcal{L}(\boldsymbol{X}, \boldsymbol{\Lambda}) \Big|_{\boldsymbol{X}=\boldsymbol{M}} \Longleftrightarrow \quad \exists \boldsymbol{\Lambda} \text{ s.t. } -\mathcal{P}_\Omega(\boldsymbol{\Lambda}) \in \partial \|\boldsymbol{M}\|_*$$

$$\Longleftrightarrow \quad \exists \boldsymbol{W} \text{ s.t.} \qquad \begin{array}{l} \boldsymbol{U}\boldsymbol{V}^\top + \boldsymbol{W} \text{ is supported on } \Omega, \\ \mathcal{P}_T(\boldsymbol{W}) = \boldsymbol{0}, \text{ and } \|\boldsymbol{W}\| \le 1 \end{array}$$

Optimality condition via dual certificate

Slightly stronger condition than KKT guarantees uniqueness:

Lemma 13.14

$\boldsymbol{M}$ is the unique minimizer of nuclear norm minimization if

- *the sampling operator $\mathcal{P}_\Omega$ restricted to T is injective, i.e.*

$$\mathcal{P}_\Omega(\boldsymbol{H}) \neq \boldsymbol{0}, \qquad \forall \text{ nonzero } \boldsymbol{H} \in T$$

- *$\exists \boldsymbol{W}$ s.t.*

$$\boldsymbol{U}\boldsymbol{V}^\top + \boldsymbol{W} \text{ is supported on } \Omega,$$
$$\mathcal{P}_T(\boldsymbol{W}) = \boldsymbol{0}, \text{ and } \|\boldsymbol{W}\| < 1$$

Proof of Lemma 13.14

For any $\boldsymbol{W}_0$ obeying $\|\boldsymbol{W}_0\| \leq 1$ and $\mathcal{P}_T(\boldsymbol{W}_0) = \boldsymbol{0}$, one has

$$\begin{aligned}
\|\boldsymbol{M}+\boldsymbol{H}\|_* &\geq \|\boldsymbol{M}\|_* + \left\langle \boldsymbol{U}\boldsymbol{V}^\top + \boldsymbol{W}_0, \boldsymbol{H} \right\rangle \\
&= \|\boldsymbol{M}\|_* + \left\langle \boldsymbol{U}\boldsymbol{V}^\top + \boldsymbol{W}, \boldsymbol{H} \right\rangle + \langle \boldsymbol{W}_0 - \boldsymbol{W}, \boldsymbol{H} \rangle \\
&= \|\boldsymbol{M}\|_* + \left\langle \mathcal{P}_\Omega\left(\boldsymbol{U}\boldsymbol{V}^\top + \boldsymbol{W}\right), \boldsymbol{H} \right\rangle + \langle \mathcal{P}_{T^\perp}(\boldsymbol{W}_0 - \boldsymbol{W}), \boldsymbol{H} \rangle \\
&= \underbrace{\|\boldsymbol{M}\|_* + \left\langle \boldsymbol{U}\boldsymbol{V}^\top + \boldsymbol{W}, \mathcal{P}_\Omega(\boldsymbol{H}) \right\rangle + \langle \boldsymbol{W}_0 - \boldsymbol{W}, \mathcal{P}_{T^\perp}(\boldsymbol{H}) \rangle}_{\text{if we take } \boldsymbol{W}_0 \text{ s.t. } \langle \boldsymbol{W}_0, \mathcal{P}_{T^\perp}(\boldsymbol{H}) \rangle = \|\mathcal{P}_{T^\perp}(\boldsymbol{H})\|_* \text{ — exercise: how to find such an } \boldsymbol{W}_0} \\
&\geq \|\boldsymbol{M}\|_* + \|\mathcal{P}_{T^\perp}(\boldsymbol{H})\|_* - \|\boldsymbol{W}\| \cdot \|\mathcal{P}_{T^\perp}(\boldsymbol{H})\|_* \\
&= \|\boldsymbol{M}\|_* + (1 - \|\boldsymbol{W}\|) \|\mathcal{P}_{T^\perp}(\boldsymbol{H})\|_* > \|\boldsymbol{M}\|_*
\end{aligned}$$

unless $\mathcal{P}_{T^\perp}(\boldsymbol{H}) = \boldsymbol{0}$.

But if $\mathcal{P}_{T^\perp}(\boldsymbol{H}) = \boldsymbol{0}$, then $\boldsymbol{H} = \boldsymbol{0}$ by injectivity. Thus, $\|\boldsymbol{M}+\boldsymbol{H}\|_* > \|\boldsymbol{M}\|_*$ for any $\boldsymbol{H} \neq \boldsymbol{0}$. This concludes the proof.

Constructing dual certificates

Use the "golfing scheme" to produce an approximate dual certificate (Gross '11)

- Think of it as an iterative algorithm (with sample splitting) to find a solution satisfying the KKT condition

(Optional) Proximal algorithm

In the presence of noise, one needs to solve

$$\text{minimize}_{\boldsymbol{X}} \quad \frac{1}{2}\|\boldsymbol{y} - \mathcal{A}(\boldsymbol{X})\|_{\mathrm{F}}^2 + \lambda\|\boldsymbol{X}\|_*$$

which can be solved via proximal methods

Proximal operator:

$$\begin{aligned}\text{prox}_{\lambda\|\cdot\|_*}(\boldsymbol{X}) &= \arg\min_{\boldsymbol{Z}} \left\{\frac{1}{2}\|\boldsymbol{Z} - \boldsymbol{X}\|_{\mathrm{F}}^2 + \lambda\|\boldsymbol{Z}\|_*\right\} \\ &= \boldsymbol{U}\mathcal{T}_\lambda(\boldsymbol{\Sigma})\boldsymbol{V}^\top\end{aligned}$$

where SVD of $\boldsymbol{X}$ is $\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\top$ with $\boldsymbol{\Sigma} = \text{diag}(\{\sigma_i\})$, and

$$\mathcal{T}_\lambda(\boldsymbol{\Sigma}) = \text{diag}(\{(\sigma_i - \lambda)_+\})$$

(Optional) Proximal algorithm

Algorithm 13.1 Proximal gradient methods

for $t = 0, 1, \cdots$**:**

$$\boldsymbol{X}^{t+1} = \mathcal{T}_{\mu_t}\left(\boldsymbol{X}^t - \mu_t \mathcal{A}^* \mathcal{A}(\boldsymbol{X}^t)\right)$$

where μ_t: step size / learning rate

Reference

- "*Lecture notes, Advanced topics in signal processing (ECE 8201)*," Y. Chi, 2015.
- "*Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization*," B. Recht, M. Fazel, P. Parrilo, *SIAM Review*, 2010.
- "*Exact matrix completion via convex optimization*," E. Candes, and B. Recht, *Foundations of Computational Mathematics*, 2009
- "*Matrix completion from a few entries*," R. Keshavan, A. Montanari, S. Oh, *IEEE Transactions on Information Theory*, 2010.
- "*Matrix rank minimization with applications*," M. Fazel, *Ph.D. Thesis*, 2002.
- "*Modeling the world from internet photo collections*," N. Snavely, S. Seitz, and R. Szeliski, *International Journal of Computer Vision*, 2008.

Reference

- "*Computer vision: algorithms and applications*," R. Szeliski, *Springer*, 2011.
- "*Shape and motion from image streams under orthography: a factorization method*," C. Tomasi and T. Kanade, *International Journal of Computer Vision*, 1992.
- "*Topics in random matrix theory*," T. Tao, *American mathematical society*, 2012.
- "*A singular value thresholding algorithm for matrix completion*," J. Cai, E. Candes, Z. Shen, *SIAM Journal on Optimization*, 2010.
- "*Recovering low-rank matrices from few coefficients in any basis*," D. Gross, *IEEE Transactions on Information Theory*, 2011.
- "*Incoherence-optimal matrix completion*," Y. Chen, *IEEE Transactions on Information Theory*, 2015.

Reference

- "*The power of convex relaxation: Near-optimal matrix completion*," E. Candes, and T. Tao, 2010.
- "*Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements*," E. Candes, and Y. Plan, *IEEE Transactions on Information Theory*, 2011.
- "*PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming*," E. Candes, T. Strohmer, and V. Voroninski, *Communications on Pure and Applied Mathematics*, 2013.
- "*Solving quadratic equations via PhaseLift when there are about as many equations as unknowns*," E. Candes, X. Li, *Foundations of Computational Mathematics*, 2014.

Reference

- "*A bound on tail probabilities for quadratic forms in independent random variables*," D. Hanson, F. Wright, *Annals of Mathematical Statistics*, 1971.
- "*Hanson-Wright inequality and sub-Gaussian concentration*," M. Rudelson, R. Vershynin, *Electronic Communications in Probability*, 2013.
- "*Exact and stable covariance estimation from quadratic sampling via convex programming* ," Y. Chen, Y. Chi, and A. Goldsmith, *IEEE Transactions on Information Theory*, 2015.
- "*ROP: Matrix recovery via rank-one projections*," T. Cai, and A. Zhang, *Annals of Statistics*, 2015.
- "*Low rank matrix recovery from rank one measurements*," R. Kueng, H. Rauhut, and U. Terstiege, *Applied and Computational Harmonic Analysis*, 2017.

Reference

- "*Mathematics of sparsity (and a few other things)*," E. Candes, *International Congress of Mathematicians*, 2014.