Nonconvex optimization landscapes in statistics: benignness, relaxation, and tightness

Andrew D. McRae

CERMICS
École nationale des ponts et chaussées

November 13, 2025

A problem in signal processing/statistics: phase retrieval

Generalized linear model: for unknown $x_* \in \mathbf{C}^d$, suppose we observe

$$y_i \approx |\langle a_i, x_* \rangle|^2, \quad i = 1, \ldots, n,$$

where $a_1, \ldots, a_n \in \mathbf{C}^d$ are known measurement vectors.

Recovery problem: estimate x_*

Motivation: optical imaging

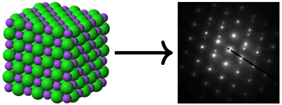

- Electromagnetic field (complex amplitude) is often linear...
- However, measured light **intensity** is the (squared) magnitude

Least-squares estimation

We observe

$$y_i \approx |\langle a_i, x_* \rangle|^2, \qquad a_1, \ldots, a_n \in \mathbf{C}^n \text{ known}, \qquad x_* \in \mathbf{C}^n \text{ unknown}$$

How do we efficiently **compute** as estimate of x_*?

- ($\exists$ vast literature)

Least-squares estimator of x_*:

$$\min_{x \in \mathbf{C}^d} \sum_{i=1}^{n} (|\langle a_i, x \rangle|^2 - y_i)^2$$

Nonconvex: could have bad local minima

- How can we overcome this?

Challenges of nonconvexity

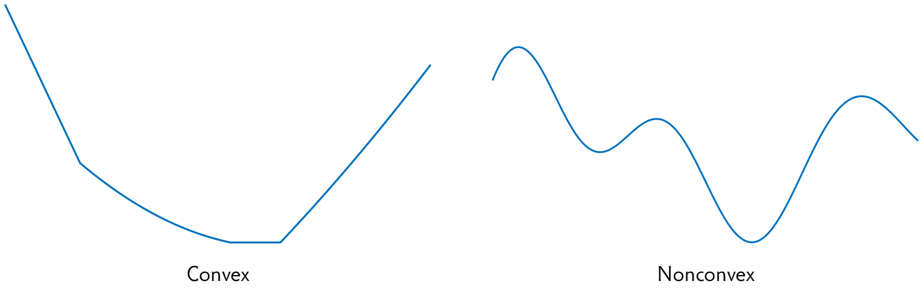

In general, nonconvex optimization problems can have (many) spurious local minima

- May be impossible to solve without exhaustive search
  - Intractable in high dimensions
- Can we do better in some cases?

What is the phase retrieval "landscape"?

Real case, $d = 2$: if $x_* = (1, 0)$, and $a \sim \mathcal{N}(0, I_2)$,

$$\begin{aligned}\mathbf{E}(|\langle a, x\rangle|^2 - |\langle a, x_*\rangle|^2)^2 &= \mathbf{E}(a_1^2 - (x_1 a_1 + x_2 a_2)^2)^2 \\ &= 3(x_1^2 - 1)^2 + 6x_1^2 x_2^2 + 3x_2^4 - 2x_2^2\end{aligned}$$

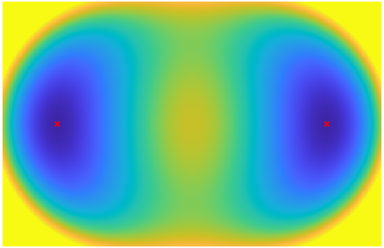

- Only local minima are $\pm x_*$
- Landscape also "benign" in higher dimensions/complex case.

Theorem (Cai et al., 2023)

If the measurement vectors $a_1, \ldots, a_n$ are i.i.d. Gaussian in $\mathbf{C}^d$, and $n \gtrsim d \log d$, then, with high probability, every second-order critical point x of

$$\min_{x \in \mathbf{C}^d} \sum_{i=1}^{n} (|\langle a_i, x\rangle|^2 - |\langle a_i, x_*\rangle|^2)^2$$

satisfies $x = sx_*$ for some $s \in \mathbf{C}, |s| = 1$.

Limitation: strong assumptions

Theorem (Cai et al., 2023)

If the measurement vectors $a_1, \ldots, a_n$ are i.i.d. ***Gaussian****, and* $n \gtrsim \boldsymbol{d} \log \boldsymbol{d}$***...***

- Gaussian measurements are unrealistic for applications
- Requirement $n \gtrsim d \log d$ is statistically suboptimal

For "harder" problem instances, nonconvex landscape is **not benign** in general!

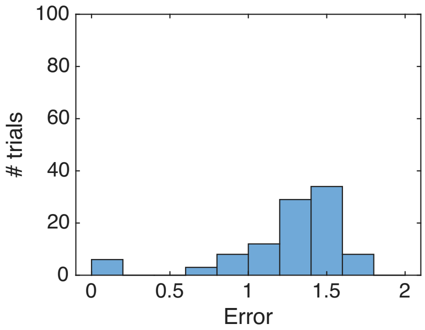

Matrix sensing approach

We observe

$$y_i \approx |\langle a_i, x_* \rangle|^2 \underbrace{= \langle a_i a_i^*, x_* x_*^* \rangle}_{\text{linear in } x_* x_*^*}, \qquad (\text{"lifting" trick})$$

We can then use the techniques of (linear) **low-rank matrix sensing**

- $x_* x_*^*$ is a rank-1 positive semidefinite matrix

"Lifted" matrix estimator ($A_i = a_i a_i^*$):

$$\min_{Z \succeq 0} \sum_{i=1}^{n} (\langle A_i, Z \rangle - y_i)^2 \text{ s.t. } \operatorname{rank}(Z) = 1$$

One approach: drop rank constraint to get **convex** semidefinite program ("PhaseLift")

- State-of-the-art for algorithms with theoretical guarantees
- Computationally expensive ($\approx d^2$ variables)

Relaxation

To try to improve the landscape, we relax the **rank constraint**

$$\min_{Z \succeq 0} \sum_{i=1}^{n} (\langle A_i, Z \rangle - y_i)^2 \text{ s.t. } \operatorname{rank}(Z) \leq \boldsymbol{p} \quad \Longleftrightarrow \quad \min_{X \in \mathbf{C}^{d \times \boldsymbol{p}}} \sum_{i=1}^{n} (\langle A_i, XX^* \rangle - y_i)^2$$

- Motivated by existing work in matrix sensing and synchronization
- $p = d \longleftrightarrow$ SDP

Theoretically, not obvious this helps!

- In general, such "overparametrization" can *introduce* spurious local optima!

Empirically, seems promising:

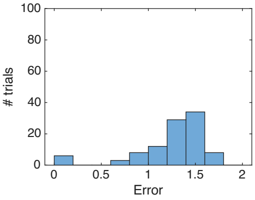

$p = 1$

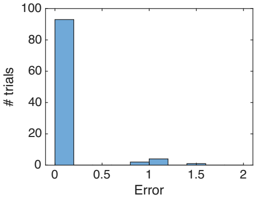

$p = 2$

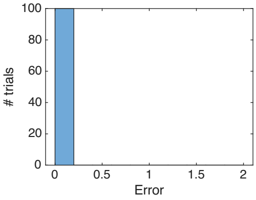

$p = 3$

Some theoretical guarantees

Relaxed nonconvex estimator ($A_i = a_i a_i^*$):

$$\text{For } y_i = \langle A_i, x_* x_*^* \rangle + \xi_i, \qquad \text{solve} \qquad \min_{X \in \mathbf{C}^{d \times p}} \sum_{i=1}^{n} (\langle A_i, XX^* \rangle - y_i)^2 \qquad \text{(BM-}p\text{)}$$

Theorem (McRae, 2025b, representative)

If $a_1, \dots, a_n$ are i.i.d. Gaussian random vectors,[1] *as long as $n \gtrsim \boldsymbol{d}$ and $p \gtrsim 1 + \frac{d \log d}{n}$, with high probability, every second-order critical point of* (BM-p) *satisfies*

$$\|XX^* - x_* x_*^*\|_* \lesssim \left\| \frac{1}{n} \sum_{i=1}^{n} \xi_i A_i \right\|_{\ell_2}$$

- p only needs to be at most $\approx \log d \ll d$
- Without noise, $X = x_* v^*$ for some $v \in \mathbf{C}^p$, $\|v\| = 1$
- In some cases, first statistically optimal result without needing to solve an SDP

[1] Can be extended to isotropic sub-Gaussian with some additional assumptions for identifiability.

What about tightness?

We were solving

$$\min_{Z \succeq 0} \sum_{i=1}^{n} (\langle A_i, Z \rangle - y_i)^2 \text{ s.t. } \operatorname{rank}(Z) \leq p.$$

But we want a rank-1 $Z = xx^*$ for some $x \in \mathbf{C}^d$.

- Without noise, $Z = XX^* = x_* x_*^*$, so tight
- With noise, in general $\operatorname{rank}(X) > 1$ so not tight...
- However, the error bound on $\|XX^* - x_* x_*^*\|_*$ combined with eigenvector perturbation bounds ensures that the **best rank-1 approximation to X still gives a good statistical estimator**

Open problem—nonconvex estimator with sparsity

Old paper: Andrew D. McRae, Justin Romberg, and Mark A. Davenport (2023). "Optimal convex lifted sparse phase retrieval and PCA with an atomic matrix norm regularizer". In: *IEEE Trans. Inf. Theory* 69.3, pp. 1866–1882

- Promising empirical results with estimator of the form

$$\min_{X \in \mathbf{C}^{d \times p}} \sum_{i=1}^{n} (\langle A_i, XX^* \rangle - y_i)^2 + \theta(X) \quad \longleftarrow \quad \text{penalty based on } \ell_1 \text{ norm}$$

- Difficulty: every version of this I can think of with an ℓ_1 norm has spurious local optima due to nonsmoothness
- **Questions:**
  - Why does it work so well empirically?
  - Is there a formulation more amenable to theory?

Example problem 2: graph clustering

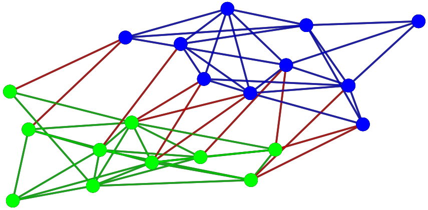

- Graph $G = (V, E)$, $V = \{1, \dots, n\}$
- We want to **label** the vertices in a way that corresponds to the edge information

Signed graph clustering

For simplicity, consider **signed** graph clustering

- (Unsigned clustering also works with some tweaks)

There is some (unknown) ground-truth labeling $z_1, \dots, z_n \in \{\pm 1\}$, and for each edge, we observe (approximately) the relative sign of its vertices:

$$R_{ij} \approx z_i z_j \qquad \text{for} \qquad (i, j) \in E$$

Estimate of clusters is

$$\arg\max_{x \in \{\pm 1\}^n} \underbrace{\sum_{(i,j) \in E} R_{ij} x_i x_j}_{= \langle C, xx^T \rangle}$$

A discrete problem

We end up with a combinatorial optimization problem

$$\max_{x\in\{\pm 1\}^n} \langle C, xx^T \rangle$$

Has similar structure to **NP-hard** max-cut problem

- But graph clustering is not (always) NP-hard, so maybe we can do better
- **Q:** What is a good algorithm (other than brute-force search)?

Continuous relaxation

We are maximizing

$$\langle C, xx^T \rangle = \sum_{i,j} C_{ij} x_i x_j.$$

We can make this continuous and smooth by **relaxing**

$$x_i x_j, \quad x_1, \ldots, x_n \in \{\pm 1\}$$

$$\downarrow$$

$$\langle x_i, x_j \rangle, \quad x_1, \ldots, x_n \in \mathbf{R}^p, \|x_i\| = 1, p \geq 2$$

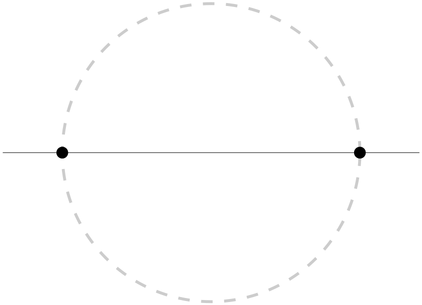

In matrix notation

$$\max_{X \in \mathbf{R}^{n \times p}} \langle C, XX^T \rangle \text{ s.t. } \operatorname{diag}(XX^T) = \mathbf{1}$$

Smooth but **nonconvex**

More general problem: orthogonal group synchronization on graph

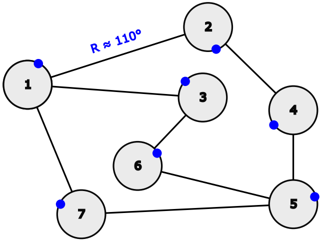

- Graph $G = (V, E)$ with vertices $V = \{1, \ldots, n\}$
- Each node i has associated $r \times r$ orthogonal matrix Z_i ($Z_i Z_i^T = I_r$)
- Observed data: $R_{ij} \approx Z_i Z_j^T$ for $(i, j) \in E$
- Goal: estimate $Z_1, \ldots, Z_n$
- Many applications in robotics, computer vision, signal processing...

General optimization problem

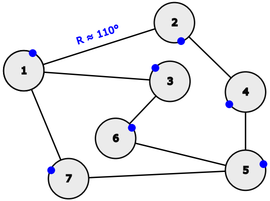

Setup:

- Graph $G = (V, E)$ with vertices $V = \{1, \dots, n\}$
- Want to estimate $r \times r$ orthogonal matrices $Z_1, \dots, Z_n$
- Observed data: $R_{ij} \approx Z_i Z_j^T$ for $(i, j) \in E$

Estimator:

$$\max_{Y_i \in \mathbf{R}^{r \times r}} \sum_{(i,j) \in E} \langle R_{ij}, Y_i Y_j^T \rangle \qquad \text{s.t. } Y_i Y_i^T = I_r, i = 1, \dots, n, \qquad \text{or}$$

$$\max_{Y \in \mathbf{R}^{rn \times r}} \langle C, YY^T \rangle \qquad \text{s.t. } \underbrace{\text{blkdiag}}_{n \text{ diag. } r \times r \text{ blks}}(YY^T) = I_{rn}$$

Relaxed version (orthogonal group → Stiefel manifold): for $p > r$,

$$\max_{Y \in \mathbf{R}^{rn \times p}} \langle C, YY^T \rangle \text{ s.t. } \text{blkdiag}(YY^T) = I_{rn}$$

A landscape guarantee (noiseless case)

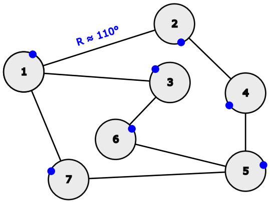

Theorem (McRae and Boumal, 2024)

Suppose

- *Measurement graph G is connected*
- *We observe exactly* $R_{ij} = Z_i Z_j^T \in \mathbf{R}^{r \times r}, (i, j) \in E$
- $\boldsymbol{p \geq r + 2}$*, and we solve*

Then every second-order critical point Y of

$$\max_{Y \in \mathbf{R}^{rn \times p}} \langle C, YY^T \rangle \text{ s.t. } \operatorname{blkdiag}(YY^T) = I_{rn}, \qquad C_{ij} = R_{ij} \text{ for } (i, j) \in E$$

satisfies $YY^T = ZZ^T$, *where* $Z = \begin{bmatrix} Z_1 \\ \vdots \\ Z_n \end{bmatrix}.$

- No dependence on the graph other than that it is connected.
- Condition on p is optimal (Markdahl, 2021)

Why $p \geq r + 2$?

Simple problem instance:[2] $r = 1,\ z_1 = \cdots = z_n = 1$

$$\max_{Y \in \mathbf{R}^{n \times p}} \sum_{i,j} A_{ij} \langle Y_i, Y_j \rangle \text{ s.t. } \|Y_i\| = 1, i = 1, \ldots, n, \qquad A \text{ is a graph adjacency matrix}$$

Global optima are $Y_1 = \cdots = Y_n$ ("synchronized states")

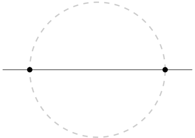

$p = 1$
discrete

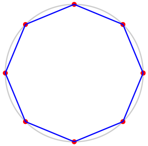

$p = 2$
spurious local optimum

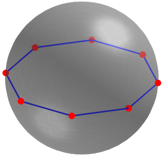

$p = 3$
no spurious optima

This is related to the topological notion of simple connectedness (Markdahl, 2021)

- Stiefel manifold $\mathrm{St}(p, r)$ simply connected $\iff p \geq r + 2$

[2] This is also known as Kuramoto oscillator synchronization, an important problem in dynamical systems.

Tightness

When can we ensure that a solution/local optimum to

$$\max_{Y \in \mathbf{R}^{rn \times p}} \langle C, YY^T \rangle \text{ s.t. } \operatorname{blkdiag}(YY^T) = I_{rn}$$

has rank exactly r?

- As with phase retrieval, noiseless → exact recovery → tight
- Extension of result gives tightness with (some) noise if $p > r + 2$ (strict inequality)
- Stronger results possible with more assumptions (Ling, 2025)
  - Especially in the simplest case $r = 1$ (McRae, Abdalla, et al., 2025; Rakoto Endor and Waldspurger, 2024; McRae, 2025a)

Tightness with noise possible when the SDP relaxation is tight and has **strict complementarity**

- Dual certificate matrix has rank exactly $(n-1)r$
- Quantitatively, depends on spectral properties of the graph

Conclusions

- Nonconvex optimization problems arising in practical applications can be **surprisingly tractable** to solve
- My recent work: theoretical guarantees of **benign landscape**
- Mild relaxation (adding variables) can help

Examples presented

- Phase retrieval
- Clustering/synchronization problems

What's next?

- Many questions on these topics and more general low-rank matrix optimization
- What interesting statistics/optimization problems are **you** working on?

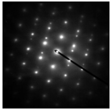

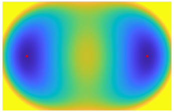

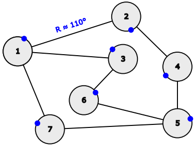

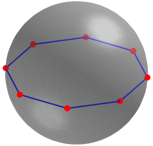

Thanks!

References I

- Cai, Jian-Feng, Meng Huang, Dong Li, and Yang Wang (2023). "Nearly optimal bounds for the global geometric landscape of phase retrieval". In: *Inverse Probl.* 39.7.
- Ling, Shuyang (2025). "Local Geometry Determines Global Landscape in Low-Rank Factorization for Synchronization". In: *Found. Comput. Math.*
- Markdahl, Johan (2021). "Synchronization on Riemannian Manifolds: Multiply Connected Implies Multistable". In: *IEEE Trans. Autom. Control* 66.9, pp. 4311–4318.
- McRae, Andrew D. (2025a). "Benign landscapes for synchronization on spheres via normalized Laplacian matrices". In: arXiv: `2503.18801 [math.OC]`.
- — (2025b). "Phase retrieval and matrix sensing via benign and overparametrized nonconvex optimization". In: arXiv: `2505.02636 [math.OC]`.
- McRae, Andrew D., Pedro Abdalla, Afonso S. Bandeira, and Nicolas Boumal (2025). "Nonconvex landscapes for $\mathbf{Z}_2$ synchronization and graph clustering are benign near exact recovery thresholds". In: *Inform. Inference.* 14.2.

References II

McRae, Andrew D. and Nicolas Boumal (2024). "Benign Landscapes of Low-Dimensional Relaxations for Orthogonal Synchronization on General Graphs". In: *SIAM J. Optim.* 34.2, pp. 1427–1454.

McRae, Andrew D., Justin Romberg, and Mark A. Davenport (2023). "Optimal convex lifted sparse phase retrieval and PCA with an atomic matrix norm regularizer". In: *IEEE Trans. Inf. Theory* 69.3, pp. 1866–1882.

Rakoto Endor, Faniriana and Irène Waldspurger (2024). "Benign landscape for Burer-Monteiro factorizations of MaxCut-type semidefinite programs". In: arXiv: `2411.03103 [math.OC]`.