

# Clustering a mixture of Gaussians with unknown covariance\*

**Mateo Díaz**

Computing and Mathematical Sciences




**Robust and Resilience Workshop - Erice**

May 2022

\*Joint work with Damek Davis and Kaizheng Wang.

# Clustering

**Clustering** is the process of partitioning a heterogeneous, unlabeled dataset into **groups of similar samples**.




Ubiquitous task in ML and statistics with **applications** to computer vision, data analysis, network analysis, genomics, among others.

# The problem today

## Model

We consider a **Gaussian mixture** of the form:

$$\mathbf{X}_i \sim \frac{1}{2}N(-\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*) + \frac{1}{2}N(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$$



Equivalently,

$$\mathbf{X}_i = y_i \boldsymbol{\mu} + \mathbf{z}_i \quad \text{with} \quad \begin{cases} \mathbf{z} \sim N(\mathbf{0}, \boldsymbol{\Sigma}), \\ \mathbb{P}(y_i = 1) = \mathbb{P}(y_i = -1) = 1/2. \end{cases}$$


## Goal

Given a sample  $\{\mathbf{X}_i\}_{i=1}^n$ , we want to recover the labels  $\{y_i\}_{i=1}^n$ .

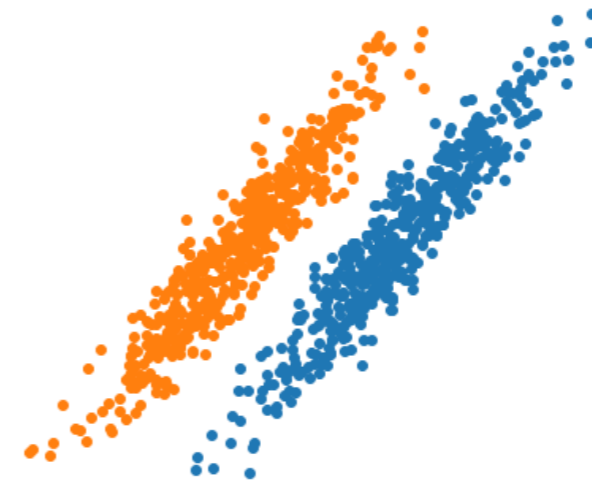
# Challenges

## Stretched mixtures

We are interested in the case when the covariance is **ill-conditioned**.



**Spherical**



**Ill-conditioned**

The de facto solutions (PCA and k-means) struggle in this setting.

## Efficiency

Make **efficient use** of statistical and computational resources.

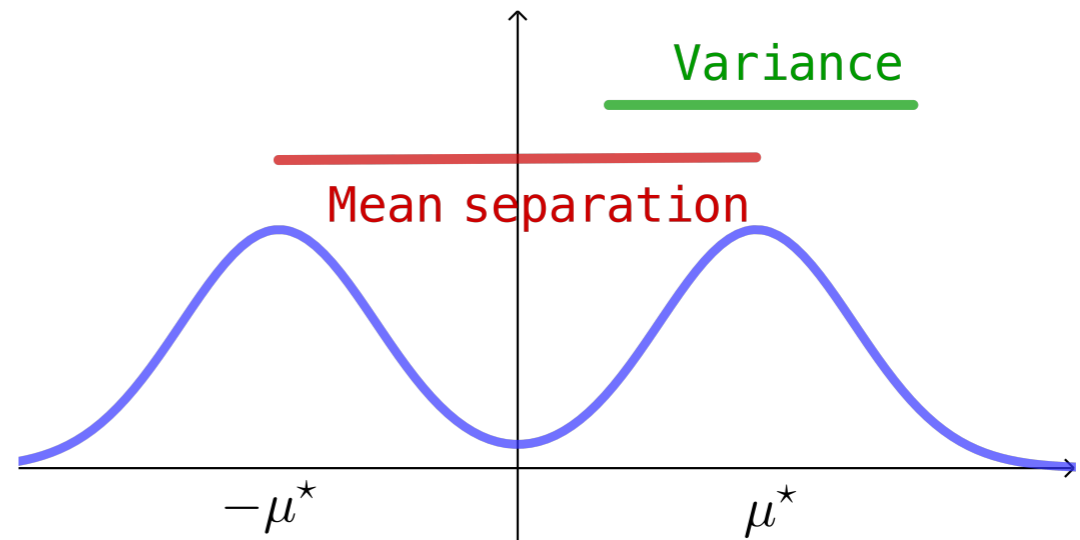


# How to measure separation?

## Baby steps

In **one dimension** a natural way to measure the signal-strength is

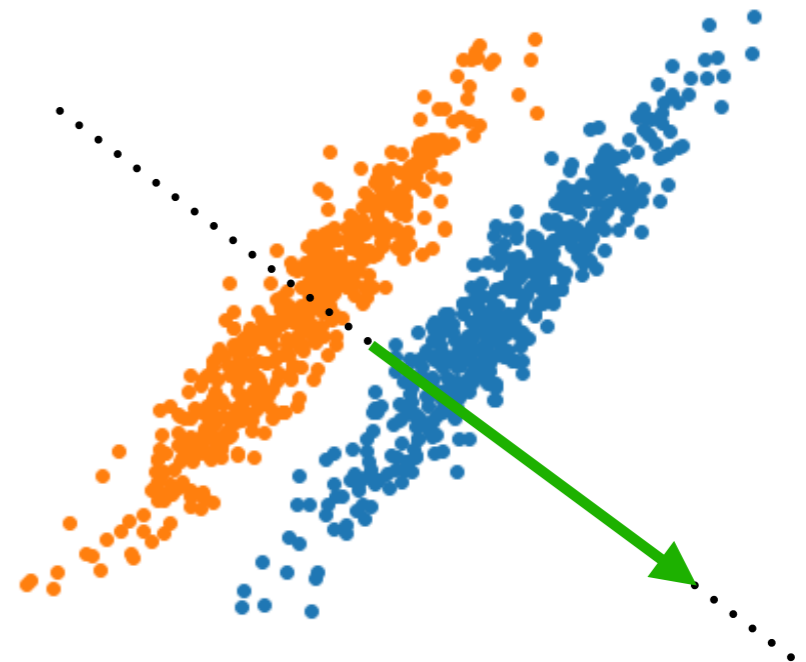
$$\text{SNR} = \left( \frac{\mu^*}{\sigma^*} \right)^2$$



## Signal-to-noise ratio

Generalizing this to higher dimensions:

$$\text{SNR} = \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}$$



# Statistical metrics

## Misclassification error

$$\mathcal{R}(\hat{\mathbf{y}}, \mathbf{y}^*) = n^{-1} \min_{s=\pm 1} |\{i \in [n] : s\hat{y}_i \neq y_i^*\}|.$$

## Baseline

The **Bayes-optimal error**

$$\hat{y}_i = \text{sign}(\langle \boldsymbol{\Sigma}^{*-1} \boldsymbol{\mu}^*, \mathbf{X}_i \rangle)$$

$$\mathbb{E}(\mathcal{R}(\hat{\mathbf{y}}, \mathbf{y}^*)) = \exp(-\Omega(\text{SNR}))$$

## Sample complexity

**Minimum number of samples** necessary to achieve **Bayes-optimal error**.

## Baseline

If the labels are given the **sample complexity** is  $n = \Omega(d)$ .

# Questions

## Statistical question

*When the labels, mean, and covariance are unknown, is it possible to achieve the **Bayes-optimal rate** with (near) **linear sample complexity**?*

**Yes**, using an optimization problem over the discrete hypercube  $\{\pm 1\}^n$ .

## Computational question

*If so, is there a **computationally efficient estimator**?*

**Doesn't seem possible** unless more samples are provided.

# Related work

**Previous works** do not tackle our questions because they either:

## 1. Need to **know the covariance matrix**.

Anandkumar, Ge, Hsu, Kakade, Telgarsky (2014),  
Balakrishnan, Wainwright, Yu (2017),  
Chen, Yang (2021),  
Daskalakis, Tzamos, Zampetakis (2017),  
Dwivedi, Ho, Khamaru, Wainwright, Jordan, Yu (2020),

...

Jin, Ke, Wang (2017),  
Kwon, Caramanis (2020),  
Löffler, Zhang, Zhou (2019),  
Ndaoud (2018)  
Vempala, Wang (2004),

## 2. Have **high sample complexity**.

Bakshi, Diakonikolas, Jia, Kane, Kothari (2020),  
Bakshi, Kothari (2020),  
Belkin, Sinha (2010),  
Brubaker, Vempala (2008),

...

Cai, Ma, Zhang (2019)  
Ge, Huang, Kakade (2015)  
Moitra, Valiant (2010)  
Tan, Vershynin (2018),

## 3. Exhibit **suboptimal dependency on the signal-to-noise ratio**.

Abbe, Fan, Wang (2020),  
Chen, Yang (2021),  
Fei, Chen (2018),  
Giraud, Verzelen (2019),


Lu, Zhou (2016),  
Mixon, Villar, Ward (2017),  
Royer (2017),

...


# Numerical illustration: FashionMNIST

We sampled 1000 T-shirts/tops and Pullovers from **FashionMNIST** and classified them with **k-means** and our **new algorithm**.


**True labels**



**k-means**



**New algorithm**



**Misclassification rate: 44.7% and 7.1%.**

# Agenda

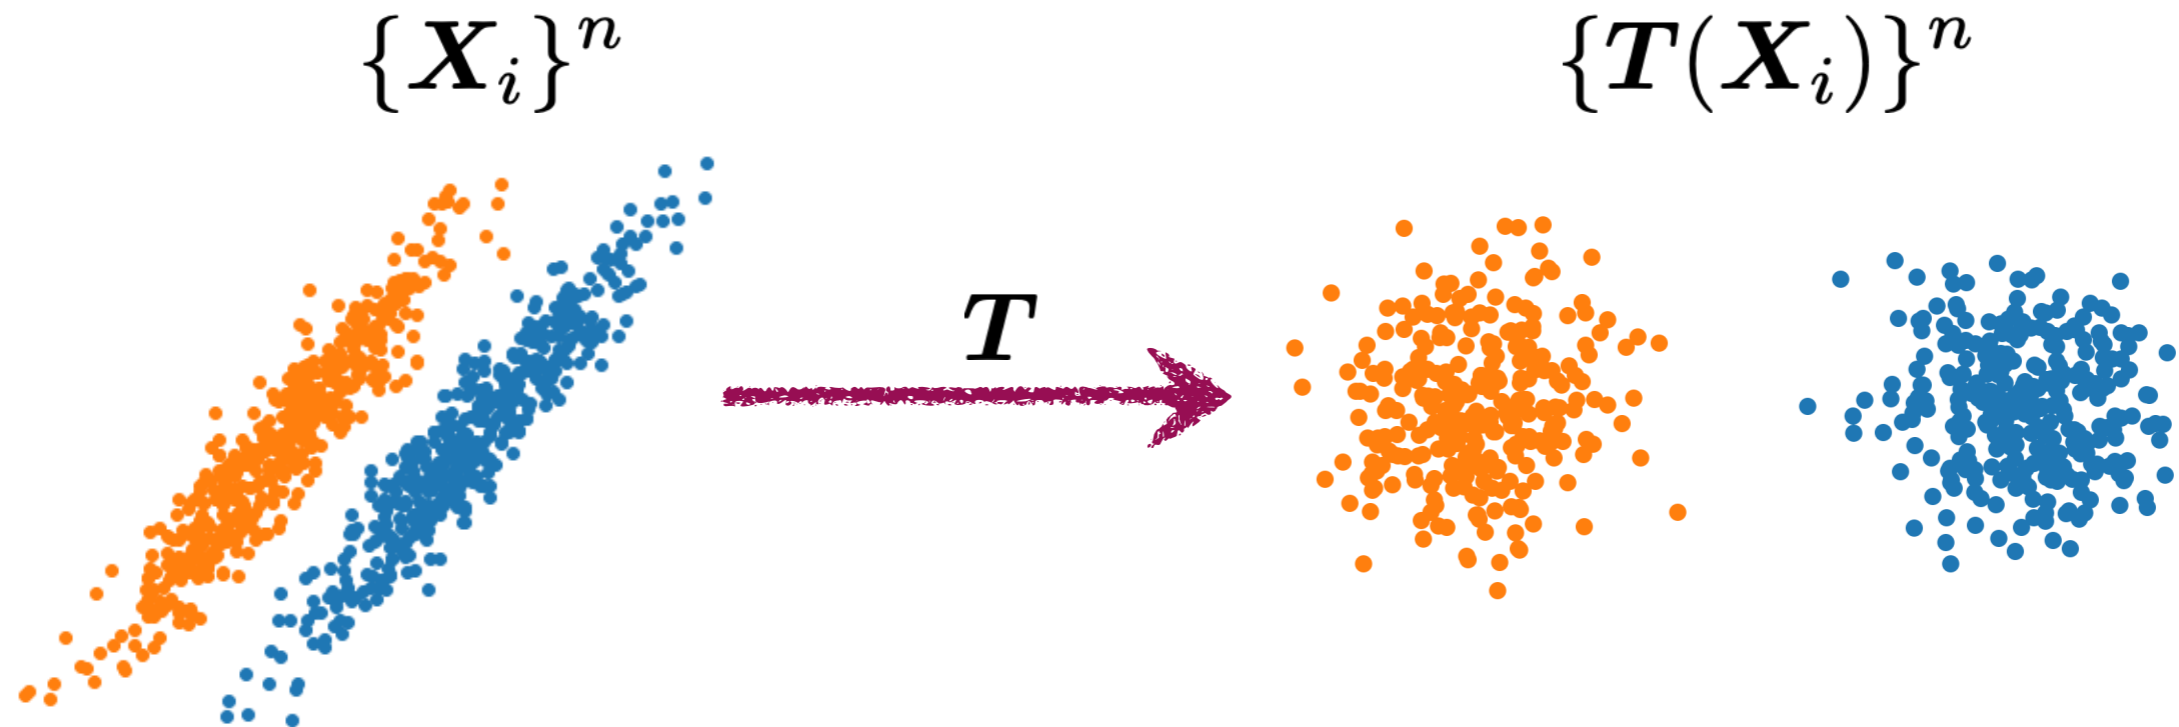
- ▶ Introduction
- ▶ MAX-CUT formulation
- ▶ Efficient spectral algorithm
- ▶ Potential statistical-computational gap



# Max-Cut formulation

# Insight: Invariance

Let  $\mathbf{T}: \mathbb{R}^d \rightarrow \mathbb{R}^d$  be an arbitrary nonsingular linear transformation



Then, we have

$$\text{SNR}(\{\mathbf{X}_i\}^n) = \text{SNR}(\{\mathbf{T}\mathbf{X}_i\}^n)$$


## Key insight

Any estimator that depends on SNR **has to be invariant.**

# Canonical form

## Lemma (Canonical form)

There is a map  $\mathbf{T}: \mathbb{R}^d \rightarrow \mathbb{R}^d$  such that:

$$\begin{pmatrix} \text{---} & (\mathbf{T}\mathbf{X}_1)^\top & \text{---} \\ & \vdots & \\ \text{---} & (\mathbf{T}\mathbf{X}_n)^\top & \text{---} \end{pmatrix} = \begin{pmatrix} \sqrt{1 - \sigma^2} \mathbf{y}^* + \sigma \mathbf{g}_1 & \mathbf{g}_2 & \dots & \mathbf{g}_d \end{pmatrix} \cdot$$

$$\sigma = \frac{1}{\sqrt{\text{SNR} + 1}}$$

where  $\mathbf{g}_i \sim N(0, \mathbf{I})$  are iid and independent from labels.

## Intuition

Every datapoint has its **label in the first component** and noise everywhere else.


# Maximum likelihood estimator

Define  $\mathbf{H} \in \mathbb{R}^{n \times n}$  to be the projection onto the range of  $(\mathbf{X}_1, \dots, \mathbf{X}_n)^\top$ .

## Lemma (From MLE to Max-Cut)

The **MLE** is given by the following problem over the hypercube:


$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y} \in \{\pm 1\}^n} \mathbf{y}^\top \mathbf{H} \mathbf{y} = \arg \min_{\mathbf{y} \in \{\pm 1\}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{H} \mathbf{y}\|^2$$



## Intuition

Minimize the distance to the subspace generated by

$$\begin{pmatrix} \sqrt{1 - \sigma^2} \mathbf{y}^* + \sigma \mathbf{g}_1 & \mathbf{g}_2 & \dots & \mathbf{g}_d \end{pmatrix}$$



# Optimality of Max-Cut

## Theorem (Max-Cut)

Assume that  $\text{SNR} \rightarrow \infty$  and  $n/(d \log n) \rightarrow \infty$ . The solution of the Max-Cut problem satisfies:

1. (**Expected error**)

$$\text{SNR} \leq C \log n \implies \mathbb{E} \mathcal{R}(\hat{\mathbf{y}}, \mathbf{y}^*) \leq e^{-\text{SNR}/[2+o(1)]}.$$

2. (**Probability of zero error**)

$$\text{SNR} \geq (2 + \varepsilon) \log n \implies \mathbb{P}[\mathcal{R}(\hat{\mathbf{y}}, \mathbf{y}^*) = 0] = 1 - o(1).$$

## Remarks.

1. Expected error matches Bayes-optimal error (up to the little-o term).
2. Separation holds with high probability if and only if  $\text{SNR} > 2 \log n$ .
3. Best known algorithm to solve Max-Cut takes exponential time.

# Why is this Max-Cut instance hard?

Natural ideas to solve the Max-Cut:

## Spectral relaxation

Relax the constrained set to a sphere:

$$\arg \max_{\|\mathbf{y}\|^2 = n} \mathbf{y}^\top \mathbf{H} \mathbf{y}.$$

**Issue:** the leading eigenvector is not unique.

## Semidefinite relaxation

Relax to a spectrahedron a la **Goemans-Williamson** (1995):

$$\max_{\mathbf{Y} \in \mathbb{R}^{n \times n}} \langle \mathbf{H}, \mathbf{Y} \rangle \quad \text{s.t.} \quad \mathbf{Y} \succeq 0, \quad \text{diag}(\mathbf{Y}) = \mathbf{1}.$$

**Issue:** Guarantees do not apply since  $\mathbf{H}$  might have **negative entries**.



**Efficient algorithm**

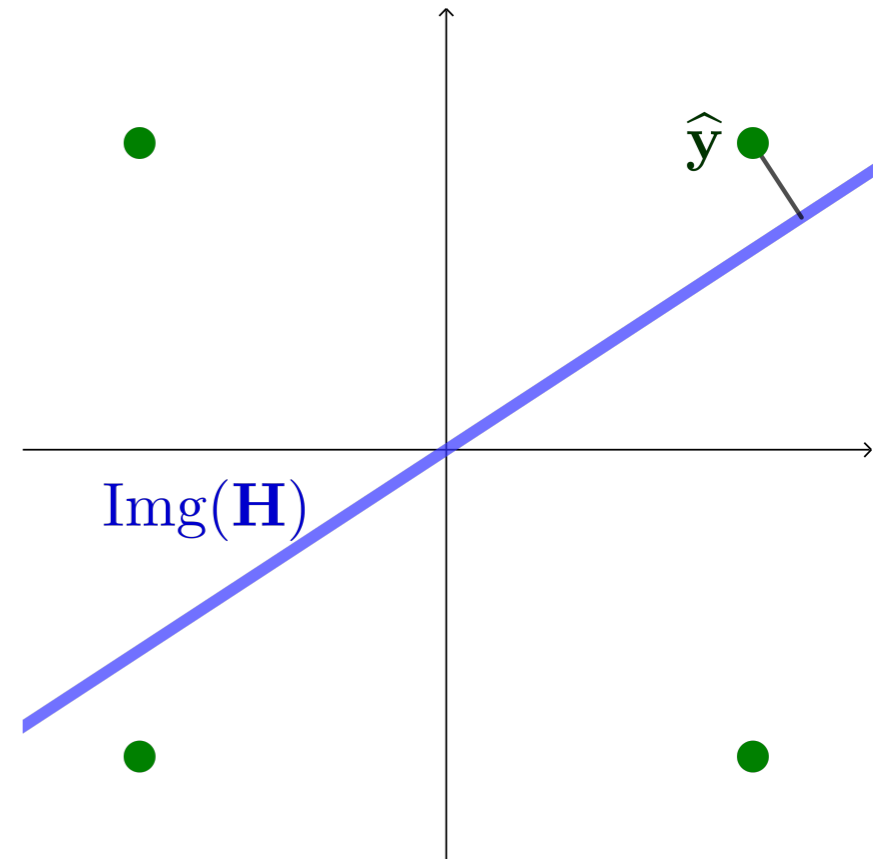
# Two stage algorithm

We want to solve the **nonconvex problem**:

$$\arg \min_{\mathbf{y} \in \{\pm 1\}^n} \frac{1}{2} \|\mathbf{y} - \mathbf{H}\mathbf{y}\|^2$$

**Strategy** split the algorithm into two:

- **Stage 1: Initialization** (spectral method)  
Finds a point **close to the solution**.
- **Stage 2: Local refinement** (optimization)  
**Iterative algorithm** that solves the optimization problem.



# Projected power iteration

---

## Algorithm 1 Projected power iteration

---

Input data matrix  $\mathbf{X} \in \mathbb{R}^{n \times d}$ , initial guess  $\mathbf{y}^0 \in \{\pm 1\}^n$ .

Compute  $\mathbf{H} = \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$  and set  $T = 4\lceil \log_2 n \rceil + 4$ .

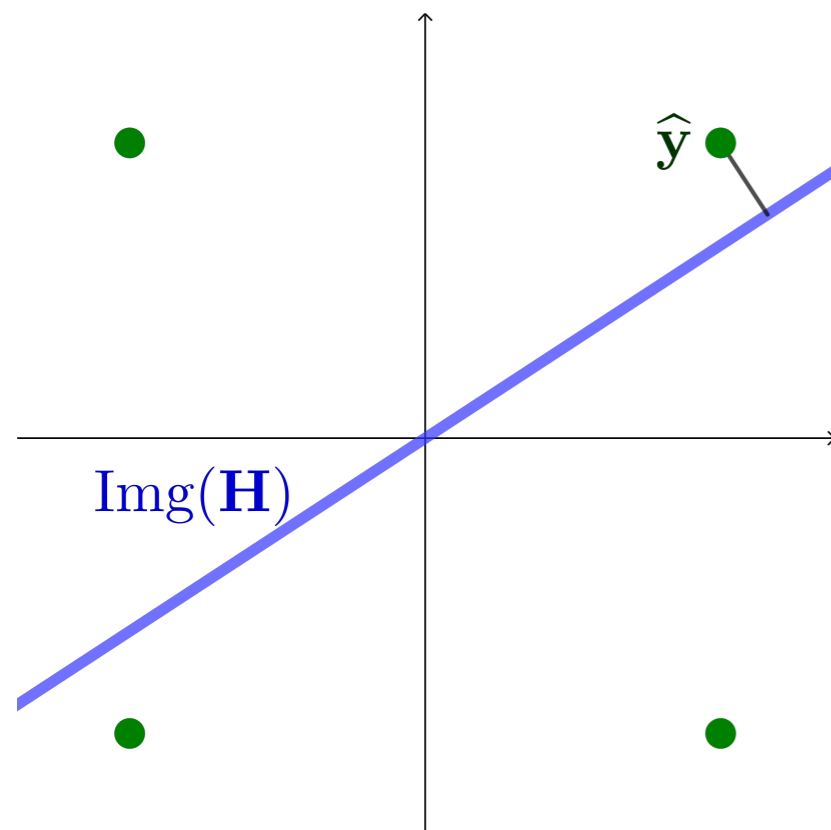
For  $t = 0, 1, \dots, T - 1$

$\mathbf{y}^{t+1} = \text{sgn}(\mathbf{H}\mathbf{y}^t)$  // applied in an entry-wise manner

Return  $\hat{\mathbf{y}}^{\text{PPI}} = \mathbf{y}^T$ .

---

This algorithm is **simply alternating projections!**



At each iteration we **project onto a subspace** and then **onto the discrete hypercube**.

# Spectral algorithm

---

## Algorithm 2 Spectral initialization

---

**Input** Data matrix  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top \in \mathbb{R}^{n \times d}$ .

**Step 1.** Compute  $\mathbf{W} = \sqrt{n} \mathbf{X} (\mathbf{X}^\top \mathbf{X})^{-1/2}$  and let  $\mathbf{w}_i$  be the  $i$ -th column of  $\mathbf{W}^\top$ .

**Step 2.** Compute the weighted sample covariance matrix

$$\mathbf{S} = \frac{1}{n} \sum_{i=1}^n (\|\mathbf{w}_i\|_2^2 - d) \mathbf{w}_i \mathbf{w}_i^\top.$$

**Step 3.** Compute the eigenvector  $\mathbf{v} \in \mathbb{S}^{d-1}$  of  $\mathbf{S}$  associated with its smallest eigenvalue.

**Output**  $\hat{\mathbf{y}}^{\text{spec}} = \text{sgn}(\mathbf{W} \mathbf{v})$ .

---

### Key insight

Once  $n = \tilde{\Omega}(d^2)$ , the matrix concentrates around:

$$\mathbf{S} \approx 2 \left( \mathbf{I} - c \frac{\mathbf{v}^* \mathbf{v}^{*\top}}{\|\mathbf{v}^*\|^2} \right)$$

where the vector  $\mathbf{v}^*$  gives the optimal classifier for  $\mathbf{w}_1, \dots, \mathbf{w}_n$ .

# Global convergence guarantee

Let  $\hat{\mathbf{y}}$  be the output of combining the spectral method and alternating projections algorithm.

## Theorem (Global convergence)

Assume that  $\text{SNR} \rightarrow \infty$ ,  $n/(d^2 \log^3 n) \rightarrow \infty$ . Then:

1. (**Expected error**)

$$\text{SNR} \leq C \log n \implies \mathbb{E} \mathcal{R}(\hat{\mathbf{y}}, \mathbf{y}^*) \leq e^{-\text{SNR}/[2+o(1)]}.$$

2. (**Probability of zero error**)

$$\text{SNR} \geq (2 + \varepsilon) \log n \implies \mathbb{P}[\mathcal{R}(\hat{\mathbf{y}}, \mathbf{y}^*) = 0] = 1 - o(1).$$

## Remark


Exactly the same guarantees as Max-Cut with **quadratic sample complexity**.

**A potential  
statistical-computational gap**



# A statistical-computational gap?

So far we have established that when  $S \lesssim 1 \ll \text{SNR}$ , then




We conjecture that **no polynomial time algorithm performs better than a random guess** in the regime:

$$d \lesssim n \ll d^2.$$

Note that there is no statistical-computational gap when  $S \gg 1$ .

# Caveat: Conjecture requires noise!

Recent papers by **Zadik, Song, Wein, and Bruna (2021)** and **Diakonikolas and Kane (2022)** disproved the conjecture when there is **no noise**, e.g., infinite SNR.



Their result is based on **Lenstra–Lenstra–Lovász lattice basis reduction** and only require  $d + 1$  samples.

# A statistical-computational gap?

We collect three pieces of evidence:

## Numerical evidence

Popular polynomial time methods seem to need  $n = \Omega(d^2)$ .

## Reduction from a hard testing problem

We reduce the problem from a **hypothesis testing problem** believed to be hard in the regime  $n \ll d^2$ .

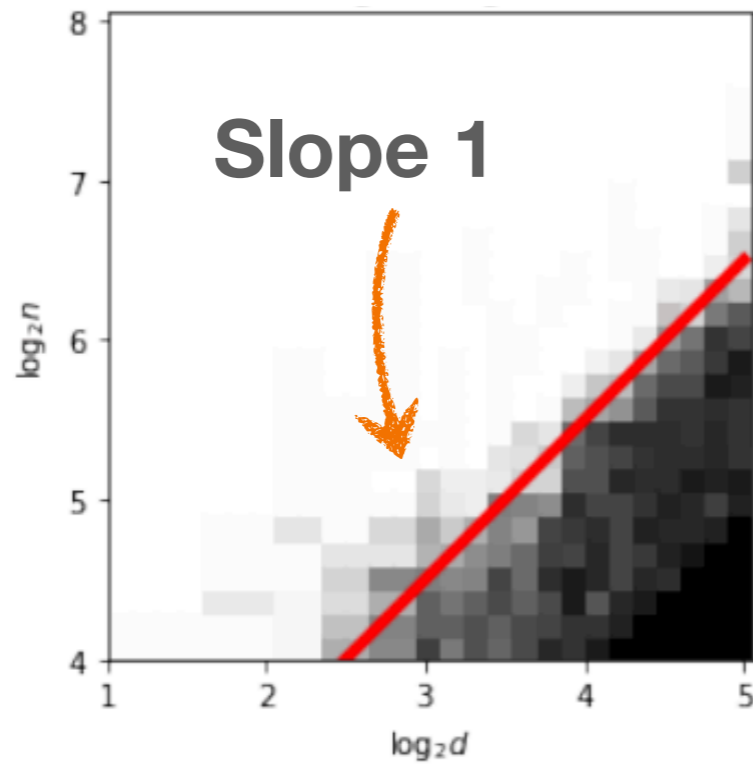
## Lower bound for sum-of-squares relaxations

We prove that **SoS relaxations** of the Max-Cut formulations **fail** when  $n \ll d^{3/2}$ .


# Numerical evidence

Recovery frequency with  $\text{SNR} = 3 \log n$ .

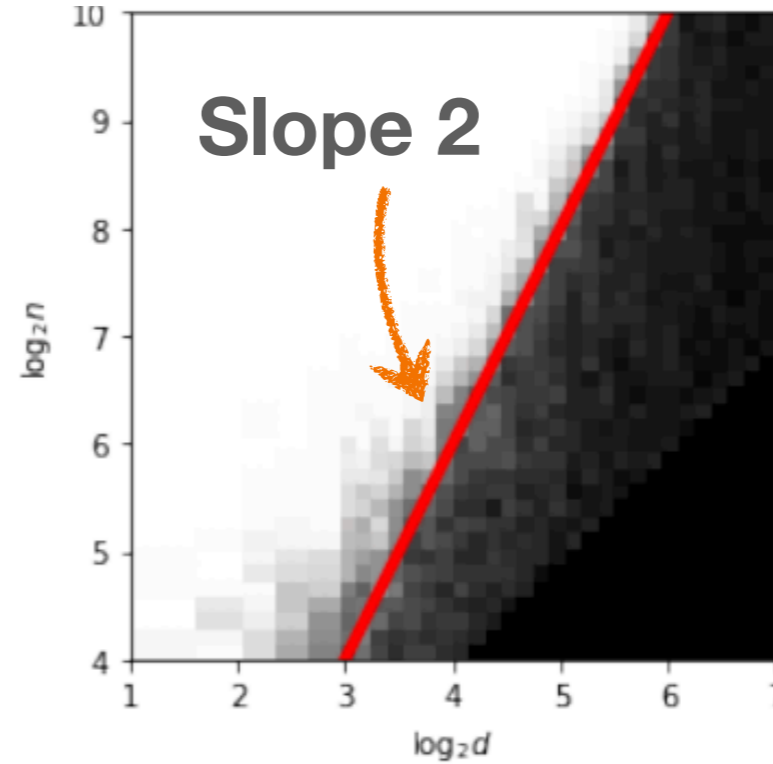
**Max-Cut**



**Spectral +  
alternating  
projections**




**Slope 2**



**Semidefinite  
relaxation**

**Expectation-  
Maximization  
algorithm**



# A hard testing problem

**Strategy:** reduce from a problem that we believe is hard.

## Testing problem

We observe  $\mathbf{X} \in \mathbb{R}^{n \times d}$ . We want a **test**  $\psi : \mathbb{R}^{n \times d} \rightarrow \{H_0, H_1\}$  that can **decide between two hypothesis:**

**Null hypothesis:**

$$\mathbf{X} = (\mathbf{g}_1, \dots, \mathbf{g}_d) \quad \text{with i.i.d. } \mathbf{g}_i \sim N(\mathbf{0}, \mathbf{I})$$

**Alternative hypothesis:**

$$\mathbf{X} = \tilde{\mathbf{X}}\mathbf{Q} \quad \text{with } \begin{cases} \tilde{\mathbf{X}} = (\mathbf{y}^*, \mathbf{g}_2, \dots, \mathbf{g}_d) \\ \mathbf{Q} \text{ unknown rotation.} \end{cases}$$

We want a test with small **Type I** and **Type II** errors.

$$\mathbb{P}\left(\psi(\mathbf{X}) = H_1 \mid H_0\right) + \mathbb{P}\left(\psi(\mathbf{X}) = H_0 \mid H_1\right)$$

# Spectral methods lower bound

## A family of spectral tests

For a fixed degree  $p$ . Let  $M: \mathbb{R}^{n \times d} \rightarrow \mathcal{S}^{n^p}$  with polynomial entries of degree at most  $p$ . Then, consider the test

$$\psi(\mathbf{X}) = \begin{cases} H_0 & \text{if } \|M(\mathbf{X})\| < t, \\ H_1 & \text{otherwise.} \end{cases}$$

## Informal theorem (Mao and Wein, 2021)

Assume that  $n \ll d^2$ . Any test  $\psi: \mathbb{R}^{n \times d} \rightarrow \{H_0, H_1\}$  coming from the above family has to satisfy

$$\mathbb{P}(\psi = H_1 \mid H_0) + \mathbb{P}(\psi = H_0 \mid H_1) \geq n^{-C+o(1)}.$$



# A reduction from testing

## Theorem (Reduction)

Assume that  $n = \Omega(d)$  and  $\text{SNR} = c \log n$ . Then if there is an **estimator for the clustering** problem such that

$$\mathbb{E} \mathcal{R}[\varphi(\mathbf{X}), \mathbf{y}^*] = e^{-\text{SNR}/[2+o(1)]}$$

Then, one can construct a **simple test**  $\psi : \mathbb{R}^{n \times d} \rightarrow \{H_0, H_1\}$  that achieves:

$$\mathbb{P}(\psi = H_1 \mid H_0) + \mathbb{P}(\psi = H_0 \mid H_1) \leq n^{-c+o(1)}.$$

## Conclusion

It is unlikely that there is a polynomial time estimator.

# Max-Cut Semidefinite relaxation

## Observation

We can rewrite the Max-Cut as a linear program:

$$\max_{\mathbf{Y} \in \mathcal{C}} \langle \mathbf{H}, \mathbf{Y} \rangle \quad \text{where} \quad \mathcal{C} = \text{conv}(\{\mathbf{y}\mathbf{y}^T : \mathbf{y} \in \{\pm 1\}^n\}),$$

Cut polytope



and LPs achieve their **maximum at vertices**.

**Obstruction:** Optimizing over the Cut polytope is **NP-hard**.

## Idea (Goemans-Williamson 1995)

Substitute  $\mathcal{C}$  by a set with a tractable SDP representation

$$\mathcal{S}_2 = \{\mathbf{Y} \in \mathbb{R}^{n \times n} : \mathbf{Y} \succeq 0, \text{diag}(\mathbf{Y}) = \mathbf{1}\}$$

Intuitively, we match all the second-moment information we have.

# Sum-of-squares relaxations

## Sum-of-squares (Parrilo 2000, Lasserre 2001)

We could take a hierarchy of tractable sets such that

$$\mathcal{S}_2 \supseteq \mathcal{S}_4 \supseteq \cdots \supseteq \mathcal{S}_m = \mathcal{C}$$

where each set has a **tractable SDP representation** and intuitively

$\mathcal{S}_{2k}$  = “Matches the moment information up to degree  $2k$ .”

For each fixed level we can solve the following in **polynomial time**

$$\max_{Y \in \mathcal{S}_k} \langle \mathbf{H}, \mathbf{Y} \rangle$$

This **strategy has been very successful** for several problems in machine learning, and theoretical computer science.

# Sum-of-squares lower bound

**Theorem (lower-bound SoS)** based on an obstruction by Ghosh et al. (2020)

Assume that  $n \leq d^{3/2-\epsilon}$  and let  $k \leq n^{c\epsilon}$ . Then, with high probability there **exists a solution**

$$\hat{\mathbf{Y}} \in \arg \max_{\mathbf{Y} \in \mathcal{S}_k} \langle \mathbf{H}, \mathbf{Y} \rangle,$$

that is **statistically independent** from the true labels  $\mathbf{y}^*$ .

## Intuition

When  $n = O(d^{3/2})$ , any relaxation of small degree  $k$  has a solution that performs as a **random guess**.

# Summary

## Statistically optimal procedure

The **Max-Cut** formulation

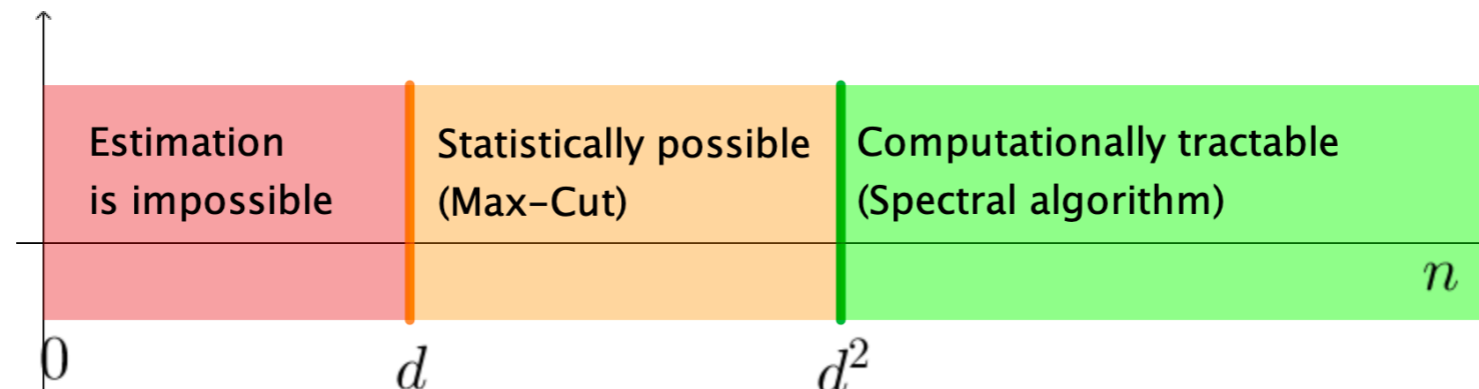
$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y} \in \{\pm 1\}^n} \mathbf{y}^\top \mathbf{H} \mathbf{y}$$

gives **Bayes-optimal error** with (near) **linear sample size**.

## Computationally tractable procedure

**Spectral method + alternating projections** yield **optimal error** provided **quadratic sample size**.

## Conjecture: statistical-computational gap



**Thank you**