Computational lower bounds for multi-frequency group synchronization

Anastasia Kireeva [1] Afonso S. Bandeira [1] Dmitriy Kunisky [2]

[1]ETH Zurich [2]Johns Hopkins University

General synchronization model

Let G be a compact group.

Goal: recover $g = (g_1, \dots, g_n) \in G^n$ from **pairwise** measurements Y_{kj}, $k, j \in [n]$:

$$Y_{kj} = f(g_k g_j^{-1}) + W_{kj}$$

for i.i.d. Gaussian random variables W_{kj}.

By considering the Peter-Weyl decomposition of $f(g_k g_j^{-1})$, we can factor the observation into different frequencies (irreducible representations).

Multi-frequency synchronization

Draw a vector $g \in G^n$ by sampling independently each coordinate from Haar (uniform) measure on G.

For each irreducible respresentation ρ we get

$$Y^{\rho}_{kj} = \frac{\lambda}{n}\rho(g_k g_j^{-1}) + \frac{1}{\sqrt{n d_\rho}} W^{\rho}_{kj}, \qquad (1)$$

where W^{ρ} are independent Gaussian ensembles (GOE/GUE/GSE depending on the type of ρ) and d_ρ is dimension of ρ.

Example: Angular synchronization

Let $x_j \overset{\text{i.i.d.}}{\sim} \mathrm{Unif}(\{e^{i\varphi}, \varphi \in [0, 2\pi)\})$.

$$\begin{cases} Y_1 = \frac{\lambda}{n} x x^* + \frac{1}{\sqrt{n}} W_1, \\ Y_2 = \frac{\lambda}{n} x^{(2)} (x^{(2)})^* + \frac{1}{\sqrt{n}} W_2, \\ \vdots \\ Y_L = \frac{\lambda}{n} x^{(L)} (x^{(L)})^* + \frac{1}{\sqrt{n}} W_L. \end{cases}$$

Here $x^{(k)}$ denotes the entrywise k-th power and $W_1, \dots, W_L$ are independent Gaussian unitary ensembles.

Prior work: Single frequency

Spectral threshold

In the single frequency case, the model can be seen as a *Wigner spiked matrix model*:

$$Y = \frac{\lambda}{n} x x^* + \frac{1}{\sqrt{n}} W.$$

BBP transition: change of the matrix spectrum at **the spectral threshold** $\lambda = 1$

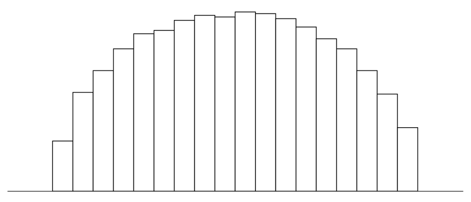

For $\lambda < 1$, PCA does not provide strong detection

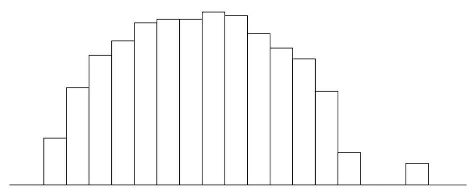

For $\lambda > 1$, detection is possible based on the top eigenvalue of Y

For many dense priors (e.g., angular synchronization, $\mathbb{Z}_2$), no algorithm can surpass the spectral threshold.

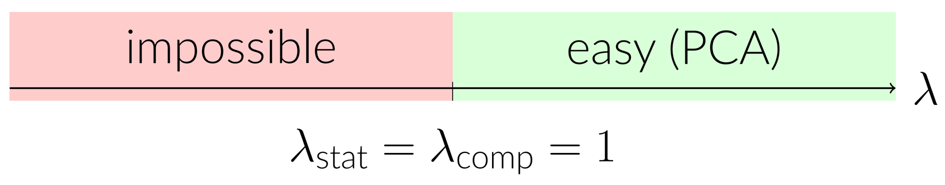

Multi-frequency model: Naïve bound

Given L independent draws of a single frequency,

$$Y_j = \frac{\lambda}{n} x x^* + \frac{1}{\sqrt{n}} W_j, \quad j \in [L],$$

PCA would indeed detect the signal once $\lambda > 1/\sqrt{L}$.

One might expect that it is possible to combine information from multiple frequencies in a similar way; however this prediction is too optimistic for the multi-frequency model due to additional algebraic structure. In particular, it is provably impossible to detect the signal below $\Theta(\sqrt{\log L / L})$ for synchronization over $\mathbb{Z}_L$.

At what **signal-to-noise ratio** λ does detection by an efficient algorithm become possible?

Main result: λ_{comp} equals the spectral threshold

Two settings:

- angular synchronization with L frequencies
- synchronization over finite group G of size L

Theorem: Let $L = O(1)$. Assuming the Low-Degree Conjecture holds, if $\lambda \le 1$, any algorithm for strong detection requires runtime at least $\exp(\tilde{\Omega}(n^{1/3}))$.

Example: Statistical-to-Computational gap for $\mathbb{Z}_L$ synchronization $(L \ge 11)$

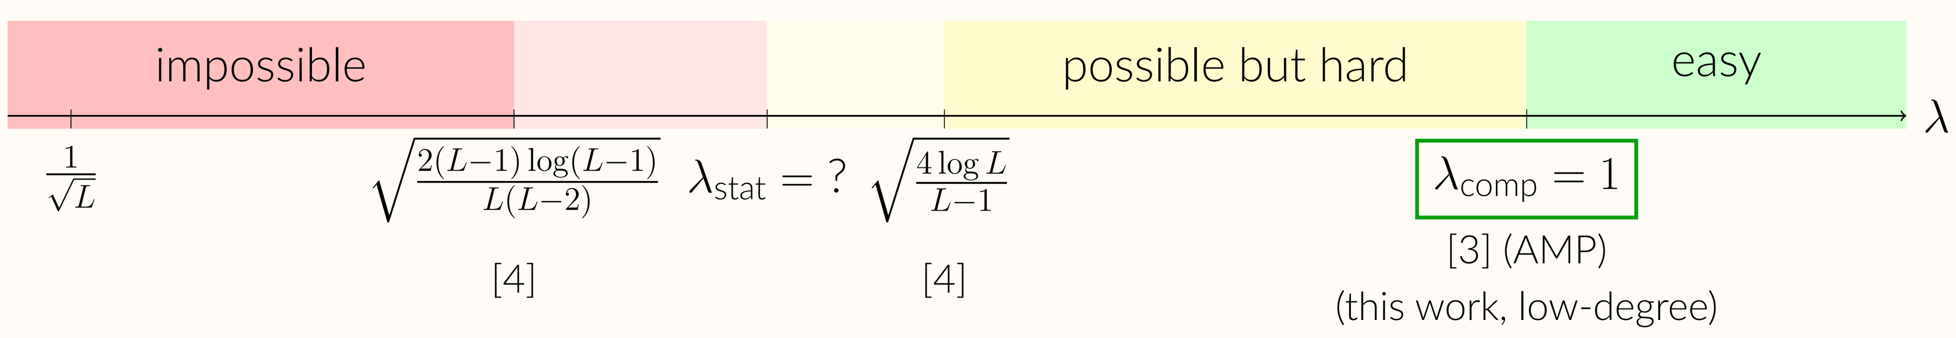

The *"possible but hard"* regime corresponds to the statistical-to-computational gap in the low-degree sense. Strong detection is information-theoretically possible; however, conjecturally, there are no efficient algorithms achieving it.

Takeaway: Adding more frequencies does not give a computational advantage.

Proof idea: Finite groups

Rewrite the low-degree likelihood ratio in terms of integer random variables $n_h := \{j \in [n] : g_j = h\}$, $h \in G$.

$$\|L_n^{\le D}\|^2 = \sum_{d=0}^{D} \frac{1}{d!}\frac{\lambda^{2d}}{n^d} \mathbb{E}\Big(\frac{L-1}{2}\sum_{h \in G} n_h^2 - \frac{1}{2}\sum_{\substack{g,h \in G \\ g \ne f}} n_g n_h\Big)^d.$$

Eliminate each n_h iteratively by taking conditional expectation.

Proof idea: Angular model

Detection by low-degree polynomials in the angular synchronization model with L frequencies is **at least as hard** as detection in model over $\mathbb{Z}_L$:

$$\|L^{\le D}_{n,\mathbb{S}}\|^2 \le \|L^{\le D}_{n,\mathbb{Z}_L}\|^2,$$

where $L^{\le D}_{n,\mathbb{S}}$ is the low-degree likelihood ratio for detection in the angular model, and $L^{\le D}_{n,\mathbb{Z}_L}$ for $\mathbb{Z}_L$ model.

Motivation

We can model receiving pairwise information as receiving a "score" $z_{kj}(h)$ for each possible group element $h \in G$ measuring how likely it is that $g_k g_j^{-1} = h$.

Consider the score function as a noisy indicator of a form

$$z_{kj}(h) = \begin{cases} \gamma + w_{kj}(h) & \text{if } h = g_k g_j^{-1}, \\ w_{kj}(h) & \text{otherwise,} \end{cases}$$

where $\gamma > 0$ and $w_{kj}(h) \sim \mathcal{N}(0, 1)$.

This model is equivalent to the multi-frequency synchronization (1) with $\gamma = \lambda\sqrt{2L/n}$.

Low-degree polynomials

Statistical distinguishability

Definition A sequence of functions $f_n : \mathcal{S} \to \{\mathrm{p}, \mathrm{q}\}$ achieves *strong detection* between $\mathbb{P}_n$ and $\mathbb{Q}_n$ if

$$\begin{aligned} &\text{if } Y \sim \mathbb{P}_n \text{ then } f_n(Y) = \mathrm{p} \quad \text{with probability } 1 - o(1); \\ &\text{if } Y \sim \mathbb{Q}_n \text{ then } f_n(Y) = \mathrm{q} \quad \text{with probability } 1 - o(1). \end{aligned}$$

Low-degree likelihood ratio

The low-degree polynomials framework provides a criterion for analyzing the hardness of statistical inference problems.

The *degree-D likelihood ratio* $L_n^{\le D}$ is defined as the projection of L_n to the linear subspace of polynomials of degree at most D, i.e.,

$$L_n^{\le D} := \mathcal{P}^{\le D} L_n = \mathcal{P}^{\le D}\left(\frac{d\mathbb{P}_n}{d\mathbb{Q}_n}(Y)\right),$$

where $\mathcal{P}^{\le D}$ is an orthogonal projection operator with respect to the inner product $\langle p, q \rangle = \mathbb{E}_{Y \sim \mathbb{Q}_n} p(Y) q(Y)$.

Low-degree conjecture

Conjecture: For "sufficiently natural" sequences of distributions $\mathbb{P}_n$, $\mathbb{Q}_n$, if $\|L_n^{\le D}\|^2$ remains bounded as $n \to \infty$, then strong detection requires runtime at least $\exp(\tilde{\Omega}(D))$.

Open problems

Statistical thresholds

The statistical thresholds for multi-frequency models are unknown, however, can be analyzed through studying the landscape of the replica potential [5].

Synchronization over infinite groups

Extend results to $SO(d)$, $d \ge 3$, for applications like Cryo-EM.

Non-constant number of frequencies

Numerical simulations in [1] suggest the possibility of surpassing the spectral threshold using an efficient algorithm when $L = \Omega(1)$. The computational threshold for low-degree polynomials in this setting is unknown.

References

[1] Tingran Gao and Zhizhen Zhao. Multi-frequency phase synchronization. In *International Conference on Machine Learning*, 2019.

[2] Anastasia Kireeva, Afonso S Bandeira, and Dmitriy Kunisky. Computational lower bounds for multi-frequency group synchronization. *arXiv preprint arXiv:2406.03424*, 2024.

[3] Amelia Perry, Alexander Wein, Afonso Bandeira, and Ankur Moitra. Message-passing algorithms for synchronization problems over compact groups. *Communications on Pure and Applied Mathematics*, 71, 2016.

[4] Amelia Perry, Alexander S Wein, Afonso S Bandeira, and Ankur Moitra. Optimality and sub-optimality of pca for spiked random matrices and synchronization. *arXiv preprint arXiv:1609.05573*, 2016.

[5] Kaylee Y Yang, Timothy LH Wee, and Zhou Fan. Asymptotic mutual information in quadratic estimation problems over compact groups. *arXiv preprint arXiv:2404.10169*, 2024.