Column Thresholding for Sparse Spiked Wigner Models: Improved Signal Strength Requirements

Jian-Feng Cai[1,2] Zhuozhi Xian[3] Jiaxi Ying[1]

[1]Department of Mathematics
Hong Kong University of Science and Technology

[2]IAS Center for AI for Scientific Discoveries
Hong Kong University of Science and Technology

[3]Department of Electronic and Computer Engineering
Hong Kong University of Science and Technology

Forty-Third International Conference on Machine Learning

Spiked Wigner model

- The sparse spiked Wigner model was introduced by Deshpande and Montanari, 2014.

- In the sparse spiked Wigner model, it aims to recover a vector $\boldsymbol{u}$ from a noisy matrix $\boldsymbol{Y}$:

$$\boldsymbol{Y} = \beta \boldsymbol{u}\boldsymbol{u}^\top + \boldsymbol{W}, \tag{1}$$

  where
  - $\boldsymbol{u} \in \mathbb{R}^d$ is an unknown s-sparse unit vector,
  - $\boldsymbol{W} \sim \mathrm{GOE}(d)$ is distributed as the Gaussian orthogonal ensemble*,
  - $\boldsymbol{Y} \in \mathbb{R}^{d \times d}$,
  - $\beta > 0$ denotes the signal strength.

- Our goal is to estimate $\boldsymbol{u}$ from $\boldsymbol{Y}$ and characterize the regimes of (β, s, d) under which nontrivial recovery is information-theoretically and computationally possible.

* $\boldsymbol{W} = (\boldsymbol{A} + \boldsymbol{A}^\top)/\sqrt{2}$ with $\boldsymbol{A}$ having i.i.d. $\mathcal{N}(0, 1)$ entries.

Signal strength requirement

- Information-theoretic lower bound on signal strength is $\beta = \widetilde{\Omega}(\sqrt{s})$*.

- Existing polynomial-time algorithms fall into two main categories:
  - Spectral methods: vanilla and spectral projection (Brennan et al., 2018)
    - Vanilla: compute leading eigenpair of $\boldsymbol{Y}$
    - Spectral projection: projects leading eigenvector of $\boldsymbol{Y}$ onto sets of sparse vectors
    - Signal strength requirement: $\beta = \widetilde{\Omega}(\sqrt{d})$†
    - Complexity: $O(d^3)$
  - Thresholding methods: diagonal thresholding‡ and covariance thresholding§
    - Diagonal thresholding: estimate support by top s elements of diagonal of $\boldsymbol{Y}$
    - Covariance thresholding: soft/hard thresholding to $\boldsymbol{Y}$ and compute its leading eigenpair
    - Signal strength requirement: $\beta = \widetilde{\Omega}(s)$¶
    - Complexity: $O(d \log d + s^3)$ for diagonal thresholding and $O(d^3)$ for covariance thresholding

---

*Dia et al., 2016; Lelarge and Miolane, 2017; Banks et al., 2018; Perry et al., 2018, 2020

†Baik et al., 2005; Péché, 2006; Féral and Péché, 2007; Paul, 2007; Capitaine et al., 2009; Benaych-Georges and Nadakuditi, 2011; Brennan et al., 2018; Montanari et al., 2015

‡Johnstone and Lu, 2009

§Krauthgamer et al., 2015; Deshpande and Montanari, 2016

¶Hopkins et al., 2017; Brennan et al., 2018; Choo and d'Orsi, 2021

Recovery regimes

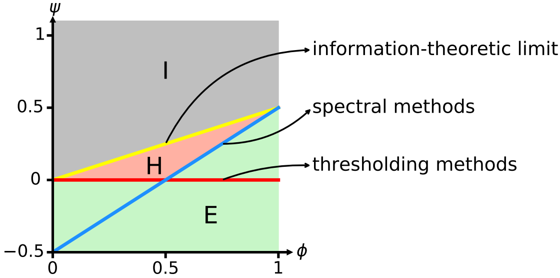

Figure 1: Recovery regimes in spiked Wigner model (Brennan et al., 2018), with $s = \widetilde{\Theta}(d^{\phi})$ and $s/\beta = \widetilde{\Theta}(d^{\psi})$.
Region I (Impossible): information-theoretically impossible regime
Region H (Hard): computationally hard regime
Region E (Easy): polynomial-time tractable regime
Yellow line: information-theoretic limit
Blue/Orange lines: computational boundaries achieved by spectral methods/thresholding methods

- Assuming planted-clique conjecture, for uniform amplitudes, statistical–computational gap in hard regime is believed to be fundamental and cannot be closed.

- We aim to identify a class of non-uniform spike which some polynomial-time algorithm can recover under signal strength $\beta = \widetilde{\Omega}(\sqrt{s})$.

Column thresholding

- Diagonal thresholding (Johnstone and Lu, 2009)
  - Statistical gap: $g_{\mathrm{diag}} := \min_{i \in \mathcal{T}} \mathbb{E}\left[\boldsymbol{Y}\right]_{ii} - \max_{i \in \mathcal{T}^c} \mathbb{E}\left[\boldsymbol{Y}\right]_{ii} = \beta \cdot \min_{i \in \mathcal{T}} |u_i|^2$.*
  - When g_{diag}/β is large, achieving any target gap sufficient for recovery requires less β.

- For a larger statistical gap, we propose the column thresholding algorithm†:
  1. $i_0 \leftarrow \arg\max_{i \in [d]} Y_{ii}$;
  2. $\widehat{\mathcal{T}} \leftarrow$ indices of s largest entries in $\left|\boldsymbol{Y}_{:,i_0}\right|$;
  3. $\boldsymbol{v} \leftarrow$ leading eigenvector of $\boldsymbol{Y}_{\widehat{\mathcal{T}}}$ with $\|\boldsymbol{v}\|_2 = 1$;
  4. Output $\hat{\boldsymbol{u}} \in \mathbb{R}^d$ with $\hat{\boldsymbol{u}}_{\widehat{\mathcal{T}}} \leftarrow \boldsymbol{v}$ and $\hat{\boldsymbol{u}}_{\widehat{\mathcal{T}}^c} \leftarrow \mathbf{0}$

  - Statistical gap: $g_{\mathrm{col}} := \min_{i \in \mathcal{T}} \left|\mathbb{E}\left[\boldsymbol{Y}\right]_{il}\right| - \max_{i \in \mathcal{T}^c} \left|\mathbb{E}\left[\boldsymbol{Y}\right]_{il}\right| = \beta\, |u_l| \min_{i \in \mathcal{T}} |u_i|$.
  - $g_{\mathrm{col}} \geq g_{\mathrm{diag}}$
  - Complexity: $O(d \log d + s^3)$ for column thresholding and $O(d \log d)$ for its variant

---

*$\mathcal{T}$: the support of $\boldsymbol{u}$

†In line with Jagatap and Hegde, 2019; Wu and Rebeschini, 2021; Cai et al., 2023, 2025

Theoretical results for column thresholding

- Noisy matrix: $\boldsymbol{Y} = \beta \boldsymbol{u}\boldsymbol{u}^\top + \boldsymbol{W}$ from (1).
- Estimation error: $\operatorname{dist}(\boldsymbol{u}, \hat{\boldsymbol{u}}) := \min\{\|\boldsymbol{u} - \hat{\boldsymbol{u}}\|_2, \|\boldsymbol{u} + \hat{\boldsymbol{u}}\|_2\}$.

Theorem (Estimation for column thresholding)

For any $\zeta \in (0,1]$, if $\beta \geq C_1 \zeta^{-1} \|\boldsymbol{u}\|_\infty^{-1} \sqrt{s \log d}$ for some universal constant $C_1 > 0$, then with probability at least $1 - 1.6d^{-1}$, the output $\hat{\boldsymbol{u}}$ of column thresholding satisfies $\operatorname{dist}(\boldsymbol{u}, \hat{\boldsymbol{u}}) \leq \zeta$.

Theorem (Support recovery)

Assume $\boldsymbol{u}$ is an s-sparse unit vector satisfying $|u_i| \geq \theta/\sqrt{s}$ for all $i \in \mathcal{T}$ and some constant $\theta > 0$. If $\beta \geq C_3 \theta^{-1} \|\boldsymbol{u}\|_\infty^{-1} \sqrt{s \log d}$ for some universal constant $C_3 > 0$, column thresholding recovers the support exactly (i.e., $\widehat{\mathcal{T}} = \mathcal{T}$) with probability at least $1 - 1.3d^{-1}$.

Truncated power method

- To improve estimation accuracy, we apply truncated power method (Yuan and Zhang, 2013).
  - Recovery problem: $\max_{\boldsymbol{w}} \boldsymbol{w}^\top \boldsymbol{Y} \boldsymbol{w}$, subject to $\|\boldsymbol{w}\|_2 = 1,\ \|\boldsymbol{w}\|_0 \le k$.
  - Initialization: our column thresholding
  - Refinement iteration: $\boldsymbol{u}^t = \mathcal{P}_{\mathbb{S}^{d-1}}(\mathcal{H}_k(\boldsymbol{Y}\boldsymbol{u}^{t-1}))$ *
- Column thresholding: a 'proper' initialization‡ with reduced signal strength requirement.
- Complexity: $O(ds + d\log d)$ per iteration

Theorem (Estimation for truncated power method)

For any $\zeta \in (0,1)$, if $\beta \ge C_4 \max\left\{\|\boldsymbol{u}\|_\infty^{-1}, \zeta^{-1}\right\}\sqrt{s\log d}$ with some universal constant $C_4 > 0$, then with probability at least $1 - 1.5d^{-1}$, $\{\boldsymbol{u}^t\}_{t\ge 1}$ produced by truncated power method, initialized with $\boldsymbol{u}^0$ from column thresholding and using $k = C_5 s$ with some universal constant $C_5 > 0$, satisfies $\mathrm{dist}(\boldsymbol{u}, \boldsymbol{u}^t) \le \eta^t\, \mathrm{dist}(\boldsymbol{u}, \boldsymbol{u}^0) + h\zeta$, where $\eta, h \in (0,1)$ are constants about C_5.

---

* $\mathcal{P}_{\mathbb{S}^{d-1}} : \mathbb{R}^d \setminus \{\mathbf{0}\} \to \mathbb{S}^{d-1}$ defined by $\mathcal{P}_{\mathbb{S}^{d-1}}(\boldsymbol{z}) = \boldsymbol{z}/\|\boldsymbol{z}\|_2$; k: sparsity level used in recovery problem; $\mathcal{H}_k : \mathbb{R}^d \to \mathbb{R}^d$ is hard thresholding operator that retains top k entries (in absolute value) and zeros out the rest.

‡ 'Proper' initialization: satisfy the convergence condition of truncated power method in Yuan and Zhang, 2013.

Empirical validation of signal thresholds

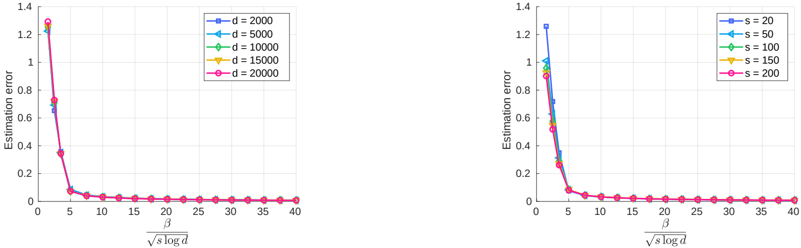

Figure 2: Estimation error vs. scaled signal strength for our column thresholding.

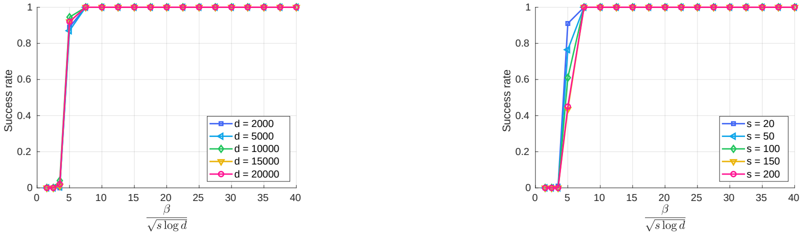

Figure 3: Support-recovery success rate vs. scaled signal strength for our column thresholding.

- Column thresholding empirically achieves the signal strength requirement $\beta = \Omega(\sqrt{s \log d})$.

Statistical and computational comparison

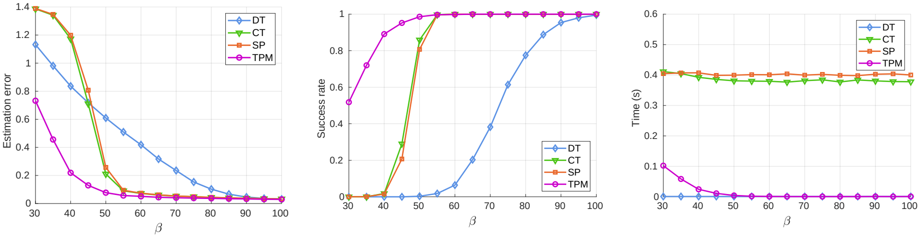

Figure 4: Estimation error, support-recovery success rate and runtime as functions of the signal strength β.

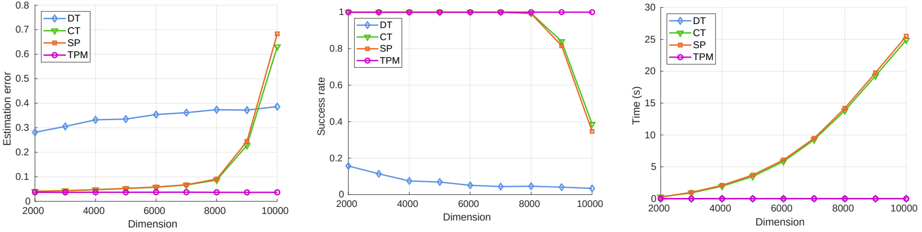

Figure 5: Estimation error, support-recovery success rate, and runtime as a function of the dimension d.

- Our TPM outperforms existing methods in estimation, support recovery and runtime.

*DT:diagonal thresholding;CT:covariance thresholding (Krauthgamer et al., 2015);SP:spectral projection (Brennan et al., 2018)

References I

Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.

Jess Banks, Cristopher Moore, Roman Vershynin, Nicolas Verzelen, and Jiaming Xu. Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *IEEE Transactions on Information Theory*, 64(7):4872–4894, 2018.

Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.

Matthew Brennan, Guy Bresler, and Wasim Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. In *Proceedings of the 31st Conference On Learning Theory*, volume 75, pages 48–166, 2018.

Jian-Feng Cai, Jingyang Li, and Juntao You. Provable sample-efficient sparse phase retrieval initialized by truncated power method. *Inverse Problems*, 39(7):075008, 2023.

Jian-Feng Cai, Zhuozhi Xian, and Jiaxi Ying. Fast and provable algorithms for sparse PCA with improved sample complexity. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267, pages 6319–6340, 2025.

Mireille Capitaine, Catherine Donati-Martin, and Delphine Féral. The largest eigenvalues of finite rank deformation of large Wigner matrices: Convergence and nonuniversality of the fluctuations. *The Annals of Probability*, 37(1):1–47, 2009.

Davin Choo and Tommaso d'Orsi. The complexity of sparse tensor PCA. *Advances in Neural Information Processing Systems*, 34: 7993–8005, 2021.

Yash Deshpande and Andrea Montanari. Information-theoretically optimal sparse PCA. In *2014 IEEE International Symposium on Information Theory*, pages 2197–2201, 2014.

Yash Deshpande and Andrea Montanari. Sparse PCA via covariance thresholding. *Journal of Machine Learning Research*, 17(141): 1–41, 2016.

Mohamad Dia, Nicolas Macris, Florent Krzakala, Thibault Lesieur, Lenka Zdeborová, et al. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. *Advances in Neural Information Processing Systems*, 29, 2016.

Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large Wigner matrices. *Communications in Mathematical Physics*, 272(1):185–228, 2007.

References II

Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 720–731, 2017.

Gauri Jagatap and Chinmay Hegde. Sample-efficient algorithms for recovering structured signals from magnitude-only measurements. *IEEE Transactions on Information Theory*, 65(7):4434–4456, 2019.

Iain M Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009.

Robert Krauthgamer, Boaz Nadler, and Dan Vilenchik. Do semidefinite relaxations solve sparse PCA up to the information limit? *The Annals of Statistics*, 43(3):1300–1322, 2015.

Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65, pages 1297–1301, 2017.

Andrea Montanari, Daniel Reichman, and Ofer Zeitouni. On the limitation of spectral methods: From the Gaussian hidden clique problem to rank-one perturbations of Gaussian tensors. In *Advances in Neural Information Processing Systems*, volume 28, 2015.

Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4): 1617–1642, 2007.

S. Péché. The largest eigenvalue of small rank perturbations of Hermitian random matrices. *Probability Theory and Related Fields*, 134(1):127–173, 2006.

Amelia Perry, Alexander S. Wein, Afonso S. Bandeira, and Ankur Moitra. Optimality and sub-optimality of PCA I: Spiked random matrix models. *The Annals of Statistics*, 46(5):2416–2451, 2018.

Amelia Perry, Alexander S. Wein, and Afonso S. Bandeira. Statistical limits of spiked tensor models. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 56(1):230–264, 2020.

Fan Wu and Patrick Rebeschini. Hadamard wirtinger flow for sparse phase retrieval. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130, pages 982–990, 2021.

Xiao-Tong Yuan and Tong Zhang. Truncated power method for sparse eigenvalue problems. *Journal of Machine Learning Research*, 14(4):899–925, 2013.