A few challenges in large-rank matrix denoising and factorization

Antoine Maillard

Vittorio Erba, Florent Krzakala, Marc Mézard,

Emanuele Troiani, Lenka Zdeborová

Journal of Statistical Mechanics 2022

Mathematical and Scientific Machine Learning 2022

ETH zürich

NORDSTAT – June 22nd 2023

Setting

$\mathbf{F}^\star \sim P_F \quad \mathbf{X}^\star \sim P_X$

$$Y_{\mu l} \sim P_{\text{out}}\left(\cdot \,\Big|\, \frac{1}{\sqrt{n}} \sum_{i=1}^{n} F^\star_{\mu i} X^\star_{il}\right)$$

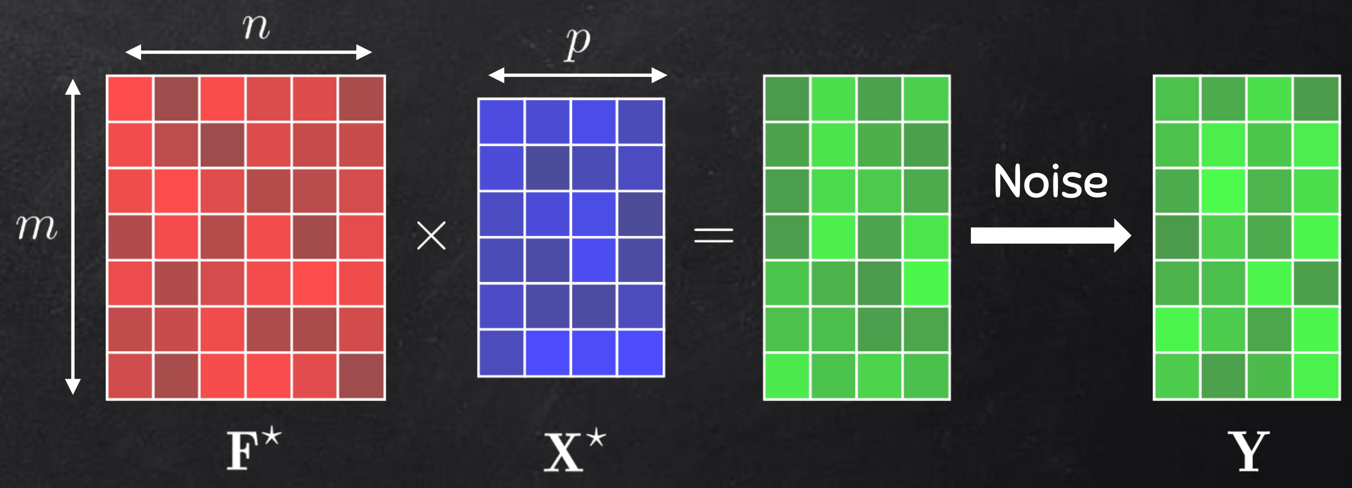

Goals

Recover $(\mathbf{F}^\star, \mathbf{X}^\star)$

Matrix factorization

Recover $\mathbf{S}^\star = \mathbf{F}^\star \mathbf{X}^\star$

Matrix denoising

Dictionary learning, sparse coding, sparse PCA, matrix completion...

Setting:

- High-dimensional: $m, p \to \infty$
- P_F, P_X, P_{out} are known ("Bayes-optimal").

Symmetric matrix factorization / denoising

$$Y_{\mu\nu} \sim P_{\text{out}}\left(\cdot \,\Big|\, \frac{1}{\sqrt{n}} \sum_{i=1}^{n} X^\star_{\mu i} X^\star_{\nu i}\right)$$

$\mathbf{X}^\star \sim P_X$

$m \to \infty$

The large - rank challenge

[M., Krzakala, Mézard & Zdeborová '22]

Symmetric matrix denoising $Y_{\mu\nu} \sim P_{\rm out}(\cdot|S^\star_{\mu\nu})$

$$\mathbf{S}^\star = \frac{1}{\sqrt{n}}\mathbf{X}^\star(\mathbf{X}^\star)^\top \quad \mathbf{X}^\star \in \mathbb{R}^{m\times n} \quad X_{\mu i} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$$

$n = \mathcal{O}(1)$

Statistical and algorithmic limits of the low-rank case are well understood.

[Rangan & Fletcher '12], [Deshpande & Montanari '14], [Lesieur, Krzakala & Zdeborová '15], [Perry, Wein , Bandeira & Moitra '16], [Lelarge & Miolane '19], [Aubin, Loureiro, M., Krzakala, Zdeborová '20], ...

Here: $m/n = \alpha \in (0,\infty)$; $m \to \infty$

- Posterior distribution $\mathbb{P}(\mathbf{S}|\mathbf{Y}) = \frac{1}{\mathcal{Z}(\mathbf{Y})} P^{(\alpha)}_{\rm Wish.}(\mathrm{d}\mathbf{S}) \prod_{1\leq\mu,\nu\leq m} P_{\rm out}(Y_{\mu\nu}|S_{\mu\nu})$
- Minimal Mean Squared Error estimator $\hat{\mathbf{S}}(\mathbf{Y}) = \mathbb{E}[\mathbf{S}|\mathbf{Y}]$ ⇒ $\mathrm{MMSE} = \mathbb{E}\|\mathbb{E}[\mathbf{S}|\mathbf{Y}] - \mathbf{S}^\star\|_F^2$

Large m limit ?

Reachable by efficient algorithms ?

Generalizations: Other priors P_X, rotationally-invariant models for $\mathbf{S}^\star$, matrix factorization.

Rotationally-invariant denoising (1)

[Bun, Allez, Bouchaud, Potters '16]

Assume noise is additive and rotationally-invariant

$$\mathbf{Y} = \mathbf{S}^\star + \sqrt{\Delta}\mathbf{Z}$$ → Gaussian (GOE) noise

Rotationally-invariant estimator (RIE) $\hat{\mathbf{S}}_{\rm RIE} = \sum_{\mu=1}^{m} \hat{\xi}_\mu \mathbf{u}_\mu \mathbf{u}_\mu^\top$ → Eigenvectors of $\mathbf{Y}$

$$\hat{\xi}_\mu = \arg\min_{\boldsymbol{\xi}\in\mathbb{R}^m} \|\mathbf{S}^\star - \hat{\mathbf{S}}(\boldsymbol{\xi})\|_F^2$$

$$\hat{\xi}_\mu = y_\mu - 2\Delta \times \mathrm{P.P.}\int \rho_{\mathbf{Y}}(\mathrm{d}t)\frac{1}{t - y_\mu}$$

Eigenvalues of $\mathbf{Y}$

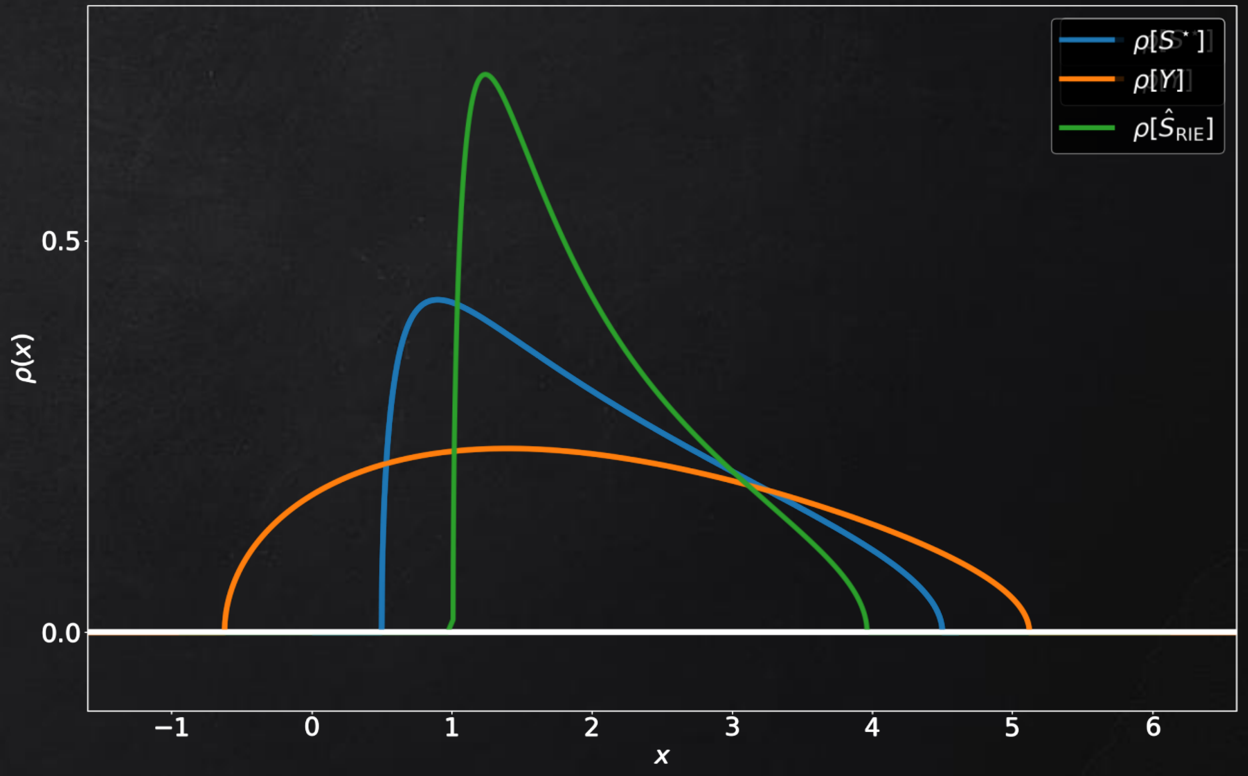

"Miraculous solution": only depends on the spectral properties of $\mathbf{Y}$

Easily computed with tools of free probability [Voiculescu, ...]

Rotationally-invariant denoising (2)

$$\mathbf{Y} = \mathbf{S}^\star + \sqrt{\Delta}\mathbf{Z}$$

$$\mathbb{P}(\mathbf{S}|\mathbf{Y}) = \frac{1}{\mathcal{Z}(\mathbf{Y})} P^{(\alpha)}_{\mathrm{Wish.}}(\mathbf{S}) \exp\left\{-\frac{1}{4\Delta}\|\mathbf{Y}-\mathbf{S}\|_F^2\right\} \propto P^{(\alpha)}_{\mathrm{Wish.}}(\mathbf{S}) \exp\left\{-\frac{\|\mathbf{S}\|_F^2}{4\Delta} + \frac{1}{2\Delta}\mathrm{Tr}[\mathbf{Y}\mathbf{S}]\right\}$$

$$f(\mathbf{O}\mathbf{D}\mathbf{O}^\top) = f(\mathbf{D}) \quad \Rightarrow \quad \mathbb{E}[f(\mathbf{S})|\mathbf{Y}] = \int_{\mathbb{R}^m} \mathrm{d}\mathbf{D}\, q(\mathbf{D})\, f(\mathbf{D}) \int_{\mathcal{O}(m)} \mathcal{D}\mathbf{O} \exp\left\{\frac{1}{2\Delta}\mathrm{Tr}\left[\mathbf{Y}\mathbf{O}\mathbf{D}\mathbf{O}^\top\right]\right\}$$

"HCIZ" integral [Harish-Chandra-Itzykson-Zuber]

$$\int_{\mathcal{O}(m)} \mathcal{D}\mathbf{O} \exp\left\{m\mathrm{Tr}\left[\mathbf{A}\mathbf{O}\mathbf{B}\mathbf{O}^\top\right]\right\}$$

Full-rank

known for $m \to \infty$

- When $\mathrm{rank}(\mathbf{A}) = o(m)$ [Guionnet '05]
- When $\mathrm{rank}(\mathbf{A}) = \Theta(m)$ [Matytsin '94, Guionnet&al '02]

[M., Krzakala, Mézard & Zdeborová '22]

- Analytical formula $\mathrm{MMSE} = \int \rho_{\mathbf{Y}}(\mathrm{d}t)\,(\cdots)$ cf also [Pourkamali, Barbier & Macris '23]
- Re-derivation of the optimal RIE estimator as $\hat{\mathbf{S}}_{\mathrm{opt.}} = \mathbb{E}[\mathbf{S}|\mathbf{Y}] \simeq \hat{\mathbf{S}}_{\mathrm{RIE}}$

Estimator and asymptotic characterization extend to the non-symmetric setting [Erba, Troiani, Krzakala, M. & Zdeborová '22]

Beyond rotation invariant denoising

[M., Krzakala, Mézard & Zdeborova '22]

$$Y_{\mu\nu} \sim P_{\text{out}}(\cdot|S^\star_{\mu\nu})$$

$$\mathbf{S}^\star = \frac{1}{\sqrt{n}}\mathbf{X}^\star(\mathbf{X}^\star)^\top$$

- Exact characterization of the MMSE ✗
- Efficient optimal estimator ✗

Unknown

Proposed Approximate Message Passing (AMP) algorithms

→ Convergence problems + hard-to-control assumptions

[Kabashima & al '16, Parker &al '14, Zou & al '21, Lucibello & al '21]

This talk: sketch a perturbative approach to clarify these difficulties, and lay a path for improvement.

$$S_{\mu\nu} = \frac{1}{\sqrt{n}}\sum_{i=1}^{n} x_{\mu i}x_{\nu i}$$

$H_{\mu\nu}$: conjugate field

$$\mathbb{P}(\mathbf{X}|\mathbf{Y}) = \frac{1}{\mathcal{Z}(\mathbf{Y})}\prod_{\mu,i} p_X(x_{\mu i})\prod_{\mu,\nu} P_{\text{out}}\left(Y_{\mu\nu}\middle|\frac{1}{\sqrt{n}}\sum_{i=1}^{n} x_{\mu i}x_{\nu i}\right)$$

$$P(\mathbf{X},\mathbf{H}) \propto \left[\prod_{\mu,i} p_X(x_{\mu i})\right]\left[\prod_{\mu,\nu} Q(Y_{\mu\nu},H_{\mu\nu})\right]\exp\left\{\frac{1}{\sqrt{n}}\sum_{i,\mu,\nu} H_{\mu\nu}x_{\mu i}x_{\nu i}\right\}$$

"Effective" distribution

Plefka–Georges–Yedidia expansion

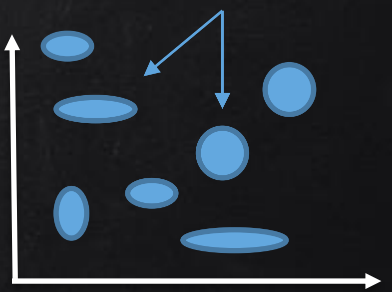

Idea: $P(\mathbf{X},\mathbf{H}) \propto \left[\prod_{\mu,i} p_X(x_{\mu i})\right]\left[\prod_{\mu,\nu} Q(Y_{\mu\nu},H_{\mu\nu})\right] \exp\left\{\frac{\not{1}\,\eta}{\sqrt{n}} \sum_{i,\mu,\nu} H_{\mu\nu} x_{\mu i} x_{\nu i}\right\}$

Thouless–Anderson–Palmer approximation [TAP77]

There is a function $\Phi_{\rm TAP}(\mathbf{m},\boldsymbol{\sigma})$ whose maxima give the "pure states" in which $P(\mathbf{X},\mathbf{H})$ concentrates its mass.

TAP free energy

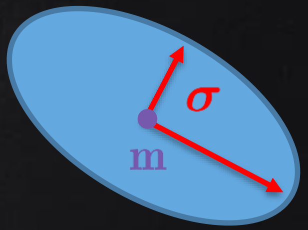

Worked out in spin glass models and simpler statistical inference models [Parisi&Potters '95], [M.&al '19]

$$\Phi_{\rm TAP}(\mathbf{m},\boldsymbol{\sigma}) = \sum_{k=0}^{\infty} \frac{\partial_\eta^k \Phi_{\rm TAP}(\mathbf{m},\boldsymbol{\sigma})[\eta=0]}{k!}$$

- $\partial_k^\eta \Phi_{\rm TAP}(\mathbf{m},\boldsymbol{\sigma})[\eta=0]$ can be recursively computed by the "PGY" method [Plefka '82, Georges&Yedidia '91]
- It turns out that (at least for the first orders) $\eta \Leftrightarrow \sqrt{m/n} = \sqrt{\alpha}$: "overcomplete" limit $\mathbf{S}^\star \simeq \mathrm{I}_m + \varepsilon(\alpha)$.

The PGY expansion

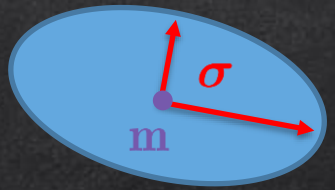

$\mathbf{m} = (m_{\mu\nu})$

$\boldsymbol{\sigma} = \mathrm{Diag}(\{\sigma_{\mu\nu}\})$

PGY

$$\Phi_{\mathrm{TAP}}(\mathbf{m}, \boldsymbol{\sigma}) = \sum_{\mu,\nu} \Big[\operatorname*{extr}_{\boldsymbol{\omega},\mathrm{b}} \Big\{ -\omega_{\mu\nu} m_{\mu\nu} - \frac{b_{\mu\nu}}{2}\Big(-\sigma_{\mu\nu}^2 + m_{\mu\nu}^2 \Big) + \ln \int \mathrm{d}z \, \frac{e^{-\frac{1}{2b_{\mu\nu}}(z-\omega_{\mu\nu})^2}}{\sqrt{2\pi b_{\mu\nu}}} P_{\mathrm{out}}(Y_{\mu\nu}|z) \Big\} \Big]$$

$$+ \frac{\eta^2}{2} \sum_{\mu,\nu} [m_{\mu\nu}^2 - \sigma_{\mu\nu}^2] + \frac{\eta^3}{6n^{1/2}} \sum_{\substack{\mu_1,\mu_2,\mu_3 \\ \text{pairwise distinct}}} m_{\mu_1\mu_2} m_{\mu_2\mu_3} m_{\mu_3\mu_1} + \mathcal{O}(\eta^4)$$

- Iterative equations to find the maxima of Φ_{TAP} can be turned into an algorithm [M., Foini, &al '19]
- Truncating at order η^2 ⟺ "AMP" algorithms of [Kabashima & al'16, ...]

  → We explicit their approximation
- However, order η^3 and above are not negligible

  → [Kabashima & al'16, ...] effectively neglect some 3rd order correlations

Numerics for Gaussian denoising

$$\mathbf{Y} = \frac{\mathbf{X}^\star (\mathbf{X}^\star)^\intercal}{\sqrt{n}} + \sqrt{\Delta}\mathbf{Z} \quad ; \quad m = \alpha n$$

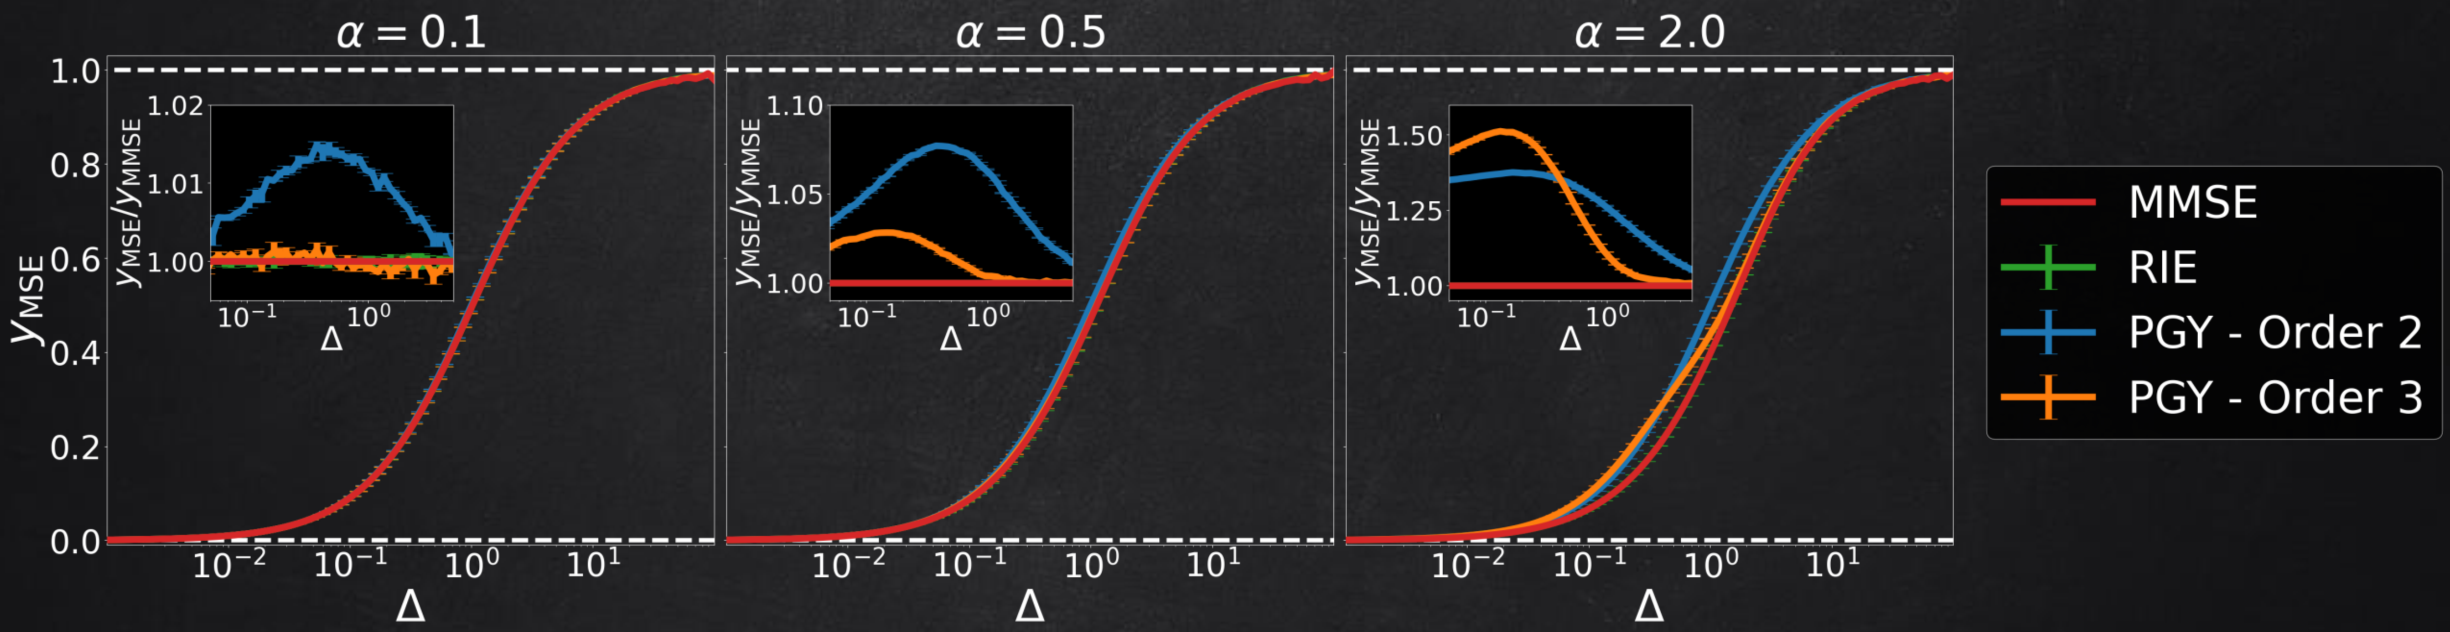

- "PGY order 3" significantly improves over order 2, in the <u>overcomplete</u> regime $\alpha \ll 1$.
- Analytical check that $\hat{\mathbf{S}}_{\mathrm{PGY}} \simeq \mathbb{E}[\mathbf{S}|\mathbf{Y}]$ up to order $(\sqrt{\alpha})^3$ ✓

<u>Limitation of the PGY method</u> ⚠

Orders 1, 2, 3, ... of the expansion ➡ Educated conjecture about arbitrary orders

For orders ≥ 4, PGY expansion becomes very tedious, need more investigation !

Conclusion

Some (of the many) open directions

- PGY expansion at orders ≥ 4 ? Arbitrary orders ? Possible resummation of the series ?
- Efficient denoising/factorization algorithms when $n = \Theta(m)$ and for non-RI noise ? $Y_{\mu\nu} \sim P_{\rm out}(\cdot|\sqrt{m}S^\star_{\mu\nu})$
- Transition between low-rank and extensive-rank regimes when rotationally-invariant :

$$I(\mathbf{A}) = \frac{1}{m^2} \log \int_{\mathcal{O}(m)} \mathcal{D}\mathbf{O} \exp\left\{m \mathrm{Tr}\left[\mathbf{AOBO}^\top\right]\right\}$$

$\mathrm{rank}(A) = o(m)$ ⟷ ? $\mathrm{rank}(A) = \Theta(m)$

[Alice Guionnet, "Rare events in Random Matrix theory", ICM 2022]

Other recent works: [Camilli & Mézard '23, Barbier & Macris '23, Pourkamali & Macris '23, Landau, Mel & Ganguli '23]

THANK YOU !