

## Phase Diagram of Extensive-Rank Symmetric Matrix Denoising beyond Rotational Invariance

Jean Barbier<sup>1,\*</sup>, Francesco Camilli<sup>1,†</sup>, Justin Ko<sup>2,‡</sup> and Koki Okajima<sup>3,§</sup><sup>1</sup>*The Abdus Salam International Centre for Theoretical Physics, Strada Costiera 11, 34151 Trieste, Italy*<sup>2</sup>*Department of Statistics and Actuarial Science, University of Waterloo,  
200 University Avenue West, Waterloo, Ontario N2L 3G1, Canada*<sup>3</sup>*Department of Physics, Graduate School of Science, The University of Tokyo,  
7-3-1 Hongo, Tokyo 113-0033, Japan*

(Received 12 November 2024; accepted 6 May 2025; published 6 June 2025)

Matrix denoising is central to signal processing and machine learning. Its statistical analysis when the matrix to infer has a factorized structure with a rank growing proportionally to its dimension remains a challenge, except when it is rotationally invariant. In this case, the information-theoretic limits and an efficient Bayes-optimal denoising algorithm, called the rotational invariant estimator, are known. Beyond this setting, few results can be found. The reason is that the model is not a usual spin system because of the growing rank dimension, nor a matrix model (as appearing in high-energy physics) due to the lack of rotation symmetry, but rather a hybrid between the two. In this paper, we make progress toward the understanding of Bayesian matrix denoising when the hidden signal is a factored matrix  $\mathbf{XX}^T$  that is not rotationally invariant. Monte Carlo simulations suggest the existence of a denoising-factorization transition separating a phase where denoising using the rotational-invariant estimator remains Bayes-optimal due to universality properties of the same nature as in random matrix theory, from one where universality breaks down and better denoising is possible, though algorithmically hard. We also argue that it is only beyond the transition that factorization, i.e., estimating  $\mathbf{X}$  itself, becomes possible up to irresolvable ambiguities. On the theoretical side, we combine mean-field techniques in an interpretable multiscale fashion in order to access the minimum mean-square error and mutual information. Interestingly, our alternative method yields equations reproducible by the replica approach of Sakata and Kabashima. Using numerical insights, we delimit the portion of phase diagram where we conjecture the mean-field theory to be exact and correct it using universality when it is not. Our complete *Ansatz* matches well the numerics in the whole phase diagram when considering finite-size effects.

DOI: [10.1103/PhysRevX.15.021085](https://doi.org/10.1103/PhysRevX.15.021085)

Subject Areas: Complex Systems, Statistical Physics

## I. INTRODUCTION

Statistical inference is the powerhouse of modern information processing systems as appearing, for instance, in signal processing, error correction for communication, and data cleaning in neuroscience, biology, or finance. Ultimately, the goal is always to separate the signal from the noise in the data at hand, where “signal” simply means what is considered the relevant part of the data for a specific application. Given the ubiquity of arrays to represent

modern data, a particularly general inference task is the one of matrix denoising. For example, in recommender systems, rows represent users and columns the products they have rated, or in medicine with patient versus measured physiological features. A basic reflex to make sense of a large data array is to write down its empirical covariance matrix, followed by a preprocessing denoising step.

However, denoising a large covariance matrix, modeled as a product of two factors, is a hard problem, both algorithmically and from the viewpoint of theoretical analysis, especially if its rank is proportional to its dimension as in the present contribution. The applications of matrix denoising range from dictionary learning [1], representation learning [2], or sparse coding [3–6] in machine learning and computational neuroscience to robust principal components analysis [7,8], submatrix localization [9], blind source separation [10], matrix completion, and recommender systems [11–13] or video processing [6,14] in signal processing, but also appear in the analysis of community detection in large graphs [15–17]. Matrix denoising

\*Contact author: [jbarbier@ictp.it](mailto:jbarbier@ictp.it)†Contact author: [fcamilli@ictp.it](mailto:fcamilli@ictp.it)‡Contact author: [justin.ko@uwaterloo.ca](mailto:justin.ko@uwaterloo.ca)§Contact author: [koki.okajima.25@gmail.com](mailto:koki.okajima.25@gmail.com)

problems can also be interpreted as “planted spin glasses,” and are therefore naturally linked to the statistical physics of disordered systems [18], a connection we will exploit extensively here.

In this paper, we focus on the denoising of symmetric positive definite signal matrices, which can therefore be interpreted as empirical covariance matrices, corrupted by a symmetric Gaussian matrix, which is a standard modeling assumption for noise. In order to study the fundamental limits of inference, i.e., the information theoretically optimal performance achievable independent of the computational cost to do so, the correct framework is the Bayesian one. Denoising is thus formulated as a Bayesian inference problem over an empirical covariance matrix  $\mathbf{X}\mathbf{X}^\top$  given a noisy data matrix  $\mathbf{Y} \in \mathbb{R}^{N \times N}$  generated as

$$\mathbf{Y} = \sqrt{\frac{\lambda}{N}} \mathbf{X}\mathbf{X}^\top + \mathbf{Z}. \quad (1)$$

The rectangular matrix  $\mathbf{X} \in \mathbb{R}^{N \times M}$  is the hidden factor [also called signal, a terminology we will also use for the square matrix  $\mathbf{X}\mathbf{X}^\top \in \mathbb{R}^{N \times N}$  of rank at most  $\min(N, M)$  if there is no ambiguity] and  $\mathbf{Z} \in \mathbb{R}^{N \times N}$  is a noise drawn from the Gaussian orthogonal ensemble (GOE), i.e., a symmetric matrix with independent Gaussian entries on its upper triangular part  $Z_{ij} = Z_{ji} \sim \mathcal{N}(0, 1 + \delta_{ij})$ . We consider the ideal setting where the signal  $\mathbf{X}$  has independent and identically distributed (i.i.d.) components  $(X_{i\mu})_{i \leq N, \mu \leq M}$  drawn from a sub-Gaussian prior distribution  $P_X$  with zero mean and unit variance.

When the prior is different than a centered Gaussian, the signal  $\mathbf{X}\mathbf{X}^\top$  (and thus the data) is *not* rotational invariant. Rotational invariance means in this context that the probability density  $P(\mathbf{X}\mathbf{X}^\top)$  would be equal to  $P(\mathbf{O}\mathbf{X}\mathbf{X}^\top\mathbf{O}^\top)$  for any  $N \times N$  orthogonal matrix  $\mathbf{O}$ . The setting where the prior  $P_X$  breaks this invariance, as well as the differences this induces with respect to the rotational-invariant Gaussian prior  $P_X = \mathcal{N}(0, 1)$ , will be of particular importance in this paper.

### A. The well known: Denoising a rank-one matrix

Model Eq. (1) may look rather special at first sight. However, a long line of work in the statistical physics of high-dimensional inference has proven that the analysis of prototypical models allow us to understand general mechanisms, similar to how dissecting the Ising model in physics has paved the way to understand phase transitions and critical phenomena way beyond it. The connection between the Ising model and matrix denoising is actually more than an analogy. Indeed, the rank-one version  $M = 1$  of matrix denoising (also called the spiked Wigner model) with signal entries taking binary values  $X_i = \pm 1$  is precisely the planted version of the celebrated Sherrington-Kirkpatrick mean-field spin glass [19] (see Refs. [18,20] for the mapping), whose

analysis is at the core of a fruitful 40-yr research field involving physicists and mathematicians [21,22].

The special features of this spin model inherent to the presence of a planted signal (or, equivalently, “on the Nishimori line”), first analyzed using spin-glass techniques by Nishimori [18], have triggered numerous works by various communities. Spiked matrix models (i.e., low-rank perturbations of much larger rank matrices, in the present case, the noise) are naturally a whole subject in random matrix theory. It was first understood in this context that the presence of the signal may yield a phase transition phenomenon now known as the Ben Arous-Baik-Péché (BBP) transition [23–25] at a critical value of the signal-to-noise ratio (SNR)  $\lambda$ . This is one of the most studied examples of phase transition in information processing. The top part of Fig. 1 illustrates the BBP transition: For SNR greater than a critical value  $\lambda_{\text{BBP}}$ , an outlier eigenvalue of the data  $\mathbf{Y}$  escapes the semicircular bulk of eigenvalues associated mostly with the noise component. A simple spectral estimator corresponding to the associated eigenvector then has a nontrivial overlap with  $\mathbf{X}$ . Before this point, however, no eigenvector of the data aligns with the signal, and spectral estimation is thus doomed. In the statistics literature, the spiked Wigner model was introduced as a statistical model for sparse principal components analysis (PCA) [26–29].

Concerning the Bayesian viewpoint on the model, as considered in the present paper, multiple approaches are now available for the computation of the main information-theoretic quantities, in particular, the mutual information between signal and data and the minimum achievable mean-square error. We mention the replica method [21,32,33] among the nonrigorous ones, and proofs based on the analysis of message-passing algorithms [34,35], spatial coupling [30,36], the interpolation method and its adaptive version [37–40], the cavity method [16,41], large deviation arguments [42], or partial differential equations approaches [43,44]. This influential series of works has led to a rich and general picture summarized by the bottom part of Fig. 1.

The phase diagram shown in the figure possesses three distinct regions. In the “impossible phase” at low  $\lambda$ , for a typical large realization of the probabilistic model Eq. (1) with  $M = 1$ , the data contain a vanishing information about the signal and thus inference, whether computationally efficient or not, is impossible. Consequently, even the Bayes-optimal estimator (corresponding to the Bayes posterior mean, black solid curve) yields trivially bad reconstruction error. Beyond the information-theoretic transition  $\lambda_{\text{IT}}$  the picture changes: here, the data typically contain information and inference becomes, in theory, possible. However, no practical algorithm running with a subexponential complexity in  $N$  is known. Those we have at hand, including the aforementioned spectral estimator and a Bayesian, prior-aware, approximate message-passing




FIG. 1. The top illustrates the Ben Arous-Baik-Péché transition occurring in the rank-one spiked matrix model. In orange is the histogram of eigenvalues of  $\mathbf{Y}$  when  $M = 1$ . Before the transition, it closely matches Wigner’s semicircular law (blue) and its top eigenvalue (red line) sticks to the end point of the bulk of eigenvalues. All eigenvectors of  $\mathbf{Y}$  have a vanishing  $o_N(1)$  overlap with the hidden signal  $\mathbf{X}$ . The BBP transition is marked by an outlier eigenvalue detaching from the bulk. The associated eigenvector aligns nontrivially with  $\mathbf{X}$  and can thus serve as spectral estimator. At the bottom is the phase diagram of the model considered from the Bayesian (information-theoretic) perspective, with sparse prior  $P_X = \rho\delta_0 + \tau\delta_{-a} + (1 - \tau - \rho)\delta_b$ , for a proper choice of parameters (see, e.g., Refs. [30,31]). Impossible, hard, and easy inference phases appear for typical realizations of Eq. (1) as the signal strength increases, delimited by information-theoretic and algorithmic transitions. The behavior of the Bayesian-optimal, spectral, and Bayesian approximate message-passing algorithm are shown. For this example of prior  $P_X$ , the transitions for the message-passing and spectral algorithms match, but the Bayesian algorithm outperforms the spectral one as it exploits the prior, while the latter does not.

(AMP) algorithm [33,45–47], are all failing. Finally, beyond the algorithmic threshold  $\lambda_{\text{Algo}}$  (which, for the choice of prior in Fig. 1, matches  $\lambda_{\text{BBP}}$ ), the signal is strong enough to be efficiently extracted from the data and, in particular, the Bayesian AMP algorithm suddenly becomes optimal in spite of its low computational cost, following a phase transition in its performance. In this region, inference is thus “easy”.

Such a phase diagram with a hard region (also called computational-to-statistical gap) and associated discontinuous transitions in algorithms performance is not specific to matrix denoising, but is actually typical in high-dimensional inference [20,48]. Their study is the central question in average-case complexity theory in modern computer science. Furthermore, web of reductions exist that prove that if the hard phase is fundamental (in the sense that there exists no polynomial-in- $N$  complexity algorithm for signal

extraction in this phase), then a plethora of other interconnected inference problems must also be hard [49,50]. The analyses of prototypical matrix denoising models therefore have a conceptual impact that goes way beyond their specific instance.

We mention that numerous generalizations of the spike Wigner model have recently been considered in Refs. [51–57].

## B. The challenge: Extensive-rank matrix denoising, beyond rotational invariance

In this paper, we aim at going beyond low-rank matrix inference [i.e.,  $M = O(1)$  while  $N \gg 1$ ] by considering the much more challenging extensive-rank setting; that is,  $N, M = M_N$  both large and verifying

$$\lim_{N \rightarrow +\infty} \frac{M}{N} = \alpha \in (0, +\infty).$$

In this scaling regime, much less is known; we refer to Table 1 in [58] for a compact account of the current status of the problem. This scaling is natural in random matrix theory. It is indeed the one where the signal is not anymore a “perturbation” (due to its low-rank nature) of the noise, but rather a “macroscopic deformation” of it. In this case, the BBP phenomenology is not valid anymore, because for any  $\lambda > 0$ , the spectrum of the data deviates macroscopically from the semicircular law. Many questions can then be answered using random matrix theory. For instance, the data asymptotic spectral density can be obtained using free probability [59].

In this paper, however, we are not interested in typical random matrix theory questions, but rather in information-theoretic ones, usually more akin to spin-glass approaches. A major difficulty is that, when interpreted in statistical physics terms, the model Eq. (1) is not a planted version of a well-known spin model as when  $M = 1$ , because we have to treat a whole matrix rather than a vector of variables as usual. This induces novel difficulties.

The standard strategy to access the information-theoretic quantities in statistical mechanics of inference is to cast the problem as the computation of a free entropy (i.e., log-partition function) associated with the Gibbs-Boltzmann distribution of the model, which is the posterior distribution given the data in the context of inference, see Eq. (3). In doing so, we will see in Sec. II that the partition function of extensive-rank matrix denoising resembles a “matrix model.” A prototypical matrix model is a partition function of the type

$$\mathcal{Z}_{\text{MM}}(\mathbf{A}) := \int_{\mathbb{H}_N} d\mathbf{F} \exp(-\text{Tr}\mathbf{F}^2 + \text{Tr}\mathbf{A}\mathbf{F}), \quad (2)$$

where  $\mathbb{H}_N$  is the space of Hermitian  $N \times N$  matrices,  $d\mathbf{F}$  denotes the Lebesgue measure over it, and the source

$\mathbf{A} \in \mathbb{H}_N$ . This one-matrix integral can be solved in the large- $N$  limit by the character expansion method [60,61] and by means of the Harish-Chandra–Itzykson–Zuber (HCIZ) integral [62,63] defined later in Eq. (8). Matrix models appear in high-energy physics, in particular, in Yang-Mills theories with internal symmetry groups of diverging dimension [64], in random matrix theory approaches to quantum gravity [65,66], quantum chromodynamics [67], statistical mechanics of planar random lattices [68–71], or when counting planar diagrams [72,73]. We refer the interested reader to the review in Ref. [74].

The known approaches for computing partition functions of matrix models require a Lebesgue measure over the matrix elements or a rotational-invariant one. This allows integration of the  $\Theta(N^2)$  rotational degrees of freedom using the HCIZ integral, leaving only  $N$  variables for the eigenvalues. If the exponent of the integrated matrix model stays  $\Theta(N^2)$ , then a final saddle point over eigenvalues (that are only  $N$ ) is justified to finish the computation.

In the present paper, however, we can no longer rely on rotational symmetry because it is, in general, broken by the prior measure  $P_X$  as we shall see. Therefore, the challenge that we tackle here is to provide a framework for analyzing information-theoretic limits of matrix denoising, cast as a statistical mechanics model looking like a matrix model (or vice versa). In doing so, we will not be able to rely on standard spin-glass approaches, as it is a matrix rather than a vector spin system, nor on those from matrix models due to the breaking of rotational invariance. Nevertheless, we will see that both types of techniques are key when properly adapted. This is the main contribution of the paper, with the ultimate goal of drawing the equivalent of the phase diagram in Fig. 1 but in the extensive-rank case  $M \approx \alpha N$ .

Finally, this scaling regime is interesting due not only to its challenging nature, but because it is also practically relevant. Many applications mentioned at the beginning in signal processing and learning naturally belong to this regime. Importantly, as it was realized that low-rank spiked matrix models play an important role in modern neural networks [75–77], so does their extensive-rank counterpart, see the recent works in Refs. [78,79], which exploit a deep connection to extensive-rank matrix denoising in order to analyze modern scaling regimes of neural networks.

## II. BAYESIAN SETTING, RELATED LITERATURE, AND MAIN CONTRIBUTIONS

Before properly defining the setting, let us introduce a few notations and information-theoretic basic definitions.

Symbol  $\sim$  expresses that a random variable is drawn from a certain law.  $X_{i\mu} \stackrel{\text{i.i.d.}}{\sim} P$  means that  $(X_{i\mu})_{i,\mu}$  are independent and identically distributed according to  $P$ . For a matrix  $\mathbf{x} = (x_{i\mu})_{i \leq N, \mu \leq M}$ , we denote  $dP_X(\mathbf{x}) = \prod_{i,\mu} dP_X(x_{i\mu})$ . Vectors are columns, and their transpose are rows.  $\mathbf{x}_i = (x_{i\mu})_{\mu \leq M} \in \mathbb{R}^M$  is the vector of the elements

of  $\mathbf{x}$  on row  $i$ . Consistently,  $\mathbf{x}_i$  will still be referred to as a column vector, even if it contains row elements. Analogously,  $\mathbf{x}_\mu = (x_{i\mu})_{i \leq N} \in \mathbb{R}^N$  is the vector of the elements of  $\mathbf{x}$  on column  $\mu$ . We denote the inner product between vectors equivalently as  $\mathbf{u} \cdot \mathbf{v} = \mathbf{u}^\top \mathbf{v} = \sum_i u_i v_i$ . The symbol  $\stackrel{D}{=}$  means equality in law. The norm  $\|\mathbf{A}\| = (\sum_{ij} A_{ij}^2)^{1/2}$  is the Frobenius norm. The limit  $\lim_{N \rightarrow \infty}$  means that both  $N$  and  $M$  diverge with  $M = \alpha N + o_N(1)$  and is called the “thermodynamic limit.”  $\mathcal{N}(m, \sigma^2)$  is the density of a Gaussian with mean  $m$  and variance  $\sigma^2$ .

The mutual information between two random variables  $X, Y$  with joint distribution  $P_{X,Y}$  and marginals  $P_X, P_Y$  is the Kullback-Leibler divergence  $I(X; Y) = I(Y; X) := D(P_{X,Y} \| P_X \otimes P_Y)$ . For our purposes,  $Y = \mathbf{Y}$  will be a continuous random variable, whereas  $X = \mathbf{X}$  can be discrete or continuous, but the conditional density  $P_{Y|X}(y|x)$  will always be accessible and well defined. In this regards, the mutual information (MI) can also be recast as

$$\begin{aligned} I(X; Y) &= \int dP_X(x) dy P_{Y|X}(y|x) \ln \frac{P_{Y|X}(y|x)}{P_Y(y)} \\ &= \mathcal{H}(Y) - \mathcal{H}(Y|X), \end{aligned}$$

where  $\mathcal{H}(Y)$  is the differential entropy and  $\mathcal{H}(Y|X)$  is the conditional differential entropy. The same symbol  $\mathcal{H}$  will also be used for discrete random variables, in which case it denotes the Shannon entropy. We refer the reader not familiar with these notions to Refs. [80,81] for more details.

### A. Bayesian matrix denoising: Setting

Consider the data generating model Eq. (1). In order to perform Bayesian inference, the first step is to write down the posterior distribution of the signal given the data using Bayes’s law. As the signal has a known factorized structure, the posterior is a probability distribution over the factor  $\mathbf{X}$  alone. Given observation Eq. (1) the posterior  $dP(\mathbf{x}|\mathbf{Y}) := \mathbb{P}(\mathbf{X} \in [\mathbf{x}, \mathbf{x} + d\mathbf{x}]|\mathbf{Y})$  reads

$$dP(\mathbf{x}|\mathbf{Y}) = \frac{\exp\left(-\frac{1}{4}\|\mathbf{Y} - \sqrt{\frac{\lambda}{N}}\mathbf{x}\mathbf{x}^\top\|^2\right)}{C(\mathbf{Y})} dP_X(\mathbf{x}), \quad (3)$$

with  $C(\mathbf{Y})$  a normalization. We denote the expectation with respect to the posterior using the bracket notation

$$\langle \cdot \rangle = \langle \cdot \rangle_{\mathbf{Y}} := \mathbb{E}[\cdot|\mathbf{Y}].$$

Expectation  $\mathbb{E}$  is with respect to the data  $\mathbf{Y}$  generated as Eq. (1) or, equivalently, with respect to the independent variables  $(\mathbf{X}, \mathbf{Z})$ . The posterior (3) gathers all information on the generating process of the signal and data: the factorized form of  $\mathbf{X}\mathbf{X}^\top$ , the size  $N \times M$  of the unknown factor  $\mathbf{X}$ , its prior law  $P_X$ , and the Gaussian nature of the symmetric noise, as well as the SNR  $\lambda$ . This setting where

the posterior is known exactly to the statistician is said to be ‘‘Bayes optimal.’’ This ideal case is the appropriate one for studying the fundamental performance limits in inference.

For the information-theoretical analysis, one of the main goals is the computation of the MI density, i.e., per signal element,

$$\frac{I(\mathbf{XX}^\top; \mathbf{Y})}{MN} = \frac{\mathcal{H}(\mathbf{Y})}{MN} - \frac{\mathcal{H}(\mathbf{Y}|\mathbf{XX}^\top)}{MN} = \frac{\mathcal{H}(\mathbf{Y})}{MN} - \frac{\mathcal{H}(\mathbf{Z})}{MN},$$

where given that the normalized distribution of the data is  $P(\mathbf{Y}) = C(\mathbf{Y}) / \prod_{i \leq j \leq N} \sqrt{2\pi(1 + \delta_{ij})}$ , the differential entropy of the data is

$$\begin{aligned} \mathcal{H}(\mathbf{Y}) &= \frac{1}{2} \sum_{i \leq j} \ln(2\pi(1 + \delta_{ij})) + \frac{1}{4} \mathbb{E} \|\mathbf{Y}\|^2 \\ &\quad - \mathbb{E} \ln \int_{\mathbb{R}^{N \times M}} dP_X(\mathbf{x}) e^{\frac{1}{2} \sqrt{\frac{\lambda}{N}} \text{Tr} \mathbf{Y} \mathbf{x} \mathbf{x}^\top - \frac{\lambda}{4N} \text{Tr}(\mathbf{x} \mathbf{x}^\top)^2} \end{aligned}$$

and the noise one is  $\mathcal{H}(\mathbf{Z}) = \frac{1}{2} \sum_{i \leq j} \ln(2\pi e(1 + \delta_{ij}))$ . Noticing that  $[1/(MN^2)] \mathbb{E} \|\mathbf{XX}^\top\|^2 = 1 + \alpha + O(1/N)$  for a centered unit variance prior, the MI density therefore reads

$$\begin{aligned} \frac{I(\mathbf{XX}^\top; \mathbf{Y})}{MN} &= \frac{\lambda}{4} (1 + \alpha) + O(1/N) \\ &\quad - \frac{1}{MN} \mathbb{E} \ln \int_{\mathbb{R}^{N \times M}} dP_X(\mathbf{x}) e^{\frac{1}{2} \sqrt{\frac{\lambda}{N}} \text{Tr} \mathbf{Y} \mathbf{x} \mathbf{x}^\top - \frac{\lambda}{4N} \text{Tr}(\mathbf{x} \mathbf{x}^\top)^2}. \end{aligned} \quad (4)$$

Observe that the MI between  $\mathbf{XX}^\top$  and  $\mathbf{Y}$  is the same as that between  $\mathbf{X}$  and  $\mathbf{Y}$  because  $\mathbf{Y}$  is a stochastic function of  $\mathbf{X}$  through only  $\mathbf{XX}^\top$ . In statistical mechanics of inference, the MI plays the same role as the free energy and is thus the thermodynamic potential of the problem. In the thermodynamic limit, its nonanalytic point(s) possibly encountered when tuning the control parameters  $(\alpha, \lambda)$  correspond (s) to the phase transition(s), whose characterization is a central goal in this paper.

The integral in the MI (4) is the logarithm of the partition function of a matrix model like Eq. (2) where  $\mathbf{F}$  is replaced by  $\mathbf{XX}^\top$ , which is symmetric positive semidefinite and where the data  $\mathbf{Y}$  play the role of source  $\mathbf{A}$ . A major difference is that the reference measure of the integral is not rotational invariant, in general.

The noise being additive Gaussian, the MI is linked to the minimum mean-square error (MMSE) on  $\mathbf{XX}^\top$  via a thermodynamic identity called the ‘‘I-MMSE relation’’ [82],

$$\frac{d}{d\lambda} \frac{I(\mathbf{XX}^\top; \mathbf{Y})}{MN} = \frac{1}{4MN^2} \mathbb{E} \|\mathbf{XX}^\top - \langle \mathbf{x} \mathbf{x}^\top \rangle\|^2. \quad (5)$$

The rhs is the MMSE =  $\text{MMSE}_N$  divided by 4, with

$$\begin{aligned} \text{MMSE} &:= \frac{1}{MN^2} \mathbb{E} \|\mathbf{XX}^\top - \langle \mathbf{XX}^\top \rangle\|^2 \\ &= \frac{1}{M} \mathbb{E} \left\| \frac{\mathbf{XX}^\top}{N} \right\|^2 - \frac{1}{M} \mathbb{E} \left\| \frac{\langle \mathbf{xx}^\top \rangle}{N} \right\|^2, \end{aligned} \quad (6)$$

using Proposition 2 in Appendix C for the second line.

## B. Related works

In the past 15 years, an interest in the statistical analysis of the challenging proportional scaling regime,  $M/N \rightarrow \alpha$  as  $N \rightarrow +\infty$ , has grown. Given the numerous conjectures revolving around extensive-rank matrix denoising, we now review relevant results.

### 1. Rotational-invariant setting

A major contribution came from random matrix theory. Ledoit and P ech e [83] and Bun *et al.* [84] introduced a rotational-invariant estimator (which could be more appropriately called rotational-equivariant estimator as noted in [85]) for optimal denoising of rotationally invariant signals. Given data as in Eq. (1), the rotational-invariant estimator (RIE) is the best estimator diagonalizing on the eigenbasis of the data, collected in an orthogonal matrix  $\mathbf{O}_Y$ . That given, it just remains to clean the eigenvalues by means of the ‘‘shrinkage formula,’’

$$\xi_i(\mathbf{Y}) := \frac{\gamma_{Y,i} - 2\pi \mathbf{H}[\rho_Y](\gamma_{Y,i})}{\sqrt{\lambda}}, \quad (7)$$

where  $(\gamma_{Y,i})$  are the eigenvalues of  $\mathbf{Y}/\sqrt{N}$ , the spectral density  $\rho_Y$  is the corresponding weak limit of  $(1/N) \sum_i \delta_{\gamma_{Y,i}}$  as  $N \rightarrow +\infty$ , and  $\mathbf{H}$  is the Hilbert transform  $\mathbf{H}[\rho_Y](x) := (1/\pi) \text{P.V.} \int d\rho_Y(y)/(x - y)$ , where P.V. denotes the principal value integral. The RIE is then composed as

$$\Xi(\mathbf{Y}) := \mathbf{O}_Y \text{diag} \left( (\xi_i(\mathbf{Y}))_{i \leq N} \right) \mathbf{O}_Y^\top.$$

The RIE exploits the factorized structure of the signal and its rank in (7) through its Marchenko-Pastur asymptotic spectral law  $\rho_{\mathbf{XX}^\top}$ , needed to compute  $\rho_Y = \rho_{\mathbf{XX}^\top} \boxplus \rho_Z$  by free convolution  $\boxplus$  with the noise semicircular spectral law  $\rho_Z$  (see Ref. [59]) and which depends on  $\alpha$  and the first and second moments of  $P_X$ ; higher-order moments are irrelevant to  $\rho_{\mathbf{XX}^\top}$  by spectral universality [86]. This is, however, *a priori* partial compared to the complete prior knowledge that the signal is of the form  $\mathbf{XX}^\top$  with  $X_{i\mu} \sim P_X$ , as exploited by a Bayesian estimator. For example, this washes out the dependence on local statistics of the eigenvalues or on the moments of  $P_X$  beyond the second. Yet, the RIE is optimal when the signal  $\mathbf{XX}^\top$  is rotational invariant [83,84].

The special case of Gaussian prior  $P_X = \mathcal{N}(0, 1)$  of the present model is the only one inducing rotational invariance of the signal and data. In this case, the log-partition function and MI of the model are exactly computable by means of the HCIZ integral defined as

$$\mathcal{Z}_{\text{HCIZ}}(\mathbf{A}, \mathbf{B}) := \int d\mu(\mathbf{O}) \exp\left(\frac{\beta N}{2} \text{Tr} \mathbf{O} \mathbf{A} \mathbf{O}^\dagger \mathbf{B}\right), \quad (8)$$

where  $\beta = 2$  if  $\mathbf{A}, \mathbf{B}$  are  $N \times N$  Hermitian matrices, and  $\beta = 1$  if they are real symmetric. Respectively, the integral is over the unitary group  $\mathbb{U}(N)$  when  $\beta = 2$  or the orthogonal one  $\mathbb{O}(N)$  when  $\beta = 1$ , with respect to the corresponding uniform Haar measure  $\mu$ . Remarkably, for  $\beta = 2$  this integral is solved in closed form [63] for any  $N$ , and its asymptotics was also computed first by Matytsin [87] and then rigorously proved in Refs. [88,89]. The asymptotic of this integral when  $\beta = 1$  is also known [90].

Starting from the HCIZ integral, the simplifications in Refs. [58,91] show that the MMSE with Gaussian prior reads

$$\lim_{N \rightarrow \infty} \text{MMSE} = \frac{1}{\lambda \alpha} \left( 1 - \frac{4\pi^2}{3} \int \rho_Y(y)^3 dy \right). \quad (9)$$

Integrating the I-MMSE relation Eq. (5) in the limit, one gets the asymptotic MI density  $\iota^G$  for matrix denoising when  $\mathbf{X}$  has i.i.d. Gaussian  $\mathcal{N}(0, 1)$  entries,

$$\begin{aligned} \iota^G(\alpha, \lambda) &:= \lim_{N \rightarrow +\infty} \frac{I(\mathbf{X}\mathbf{X}^\dagger; \mathbf{Y})}{MN} \\ &= \frac{1}{8\alpha} + \frac{1}{2\alpha} \int dx dy \rho_Y(x) \rho_Y(y) \ln |x - y|. \end{aligned} \quad (10)$$

For its numerical computation at SNR  $\lambda$ , a convenient approach is to use the I-MMSE relation, i.e., to numerically integrate 4 times the MMSE Eq. (9) with respect to  $\lambda' \in [0, \lambda]$  using  $\rho_Y$  given explicitly in Ref. [58] for a Wishart signal  $\mathbf{X}\mathbf{X}^\dagger / \sqrt{N}$ .

Optimal denoising and the RIE for the rectangular, nonsymmetric version of the problem was proposed recently in Refs. [92–95] based on the rectangular HCIZ integral [96]. The work in Ref. [97] provided an interesting approach to the factorization problem discussed in Sec. VI based on the RIE. The authors of Ref. [98] managed to exploit rotational invariance in order to analyze the dynamics of gradient flow for the denoising of a Wishart matrix.

## 2. Statistical mechanics approaches and beyond

Moving away from the rotationally invariant setting, among the first works based on statistical physics we mention Ref. [99], which proposed a replica approach with an original ‘‘Gaussian Ansatz’’ (see Appendix F) for dictionary learning in an optimization setting. It was followed

by the concurrent works of Refs. [100–102] applying the same method in a Bayesian setting similar to ours. These works, however, do not take into account some correlations between variables that turn out to be relevant, in general, in the high-dimensional limit, as understood in Refs. [103,104] and further detailed by some of the previous authors in Ref. [91]. Although it is true that in Refs. [99–102] relevant correlations are neglected when considering the Gaussian prior, our numerical evidence suggests that this is not the case in the factorization phase we identify. The factorization phase cannot appear for the Gaussian prior and has thus eluded subsequent studies [104] that aimed at matching the formulas in Refs. [99–102] with the theoretical predictions for the Gaussian prior by means of the asymptotic expansion of Matytsin [87]. However, Ref. [104] was missing a justification of their final result, which was later provided in the concurrent works of Refs. [91,105] and rigorized in Ref. [58].

The authors of Ref. [91] provide a perturbative attempt to the computation of the MI based on Plefka-Georges-Yedidia (PGY) expansion [106,107]. The PGY formalism yields  $MN$  equations that should approximate to a certain perturbation order the Thouless-Anderson-Palmer (TAP) equations [108]. When comparing the PGY estimator, truncated at order 2 and 3 of the perturbative expansion, with the exact solution for the Gaussian prior given by the HCIZ integral, the authors still measure a significant discrepancy, a sign that higher perturbative orders may be non-negligible. The authors of Ref. [105] attempted an innovative spectral replica symmetric Ansatz, yet leading to an unsatisfactory solution as it is not possible to evaluate it and it is thus hard to extract predictions from it.

A step toward the phase diagram for generic prior was done in the series of works of Refs. [109–111], where the authors map the model into a sequence of neural networks for associative memory. They introduce a ‘‘decimation scheme,’’ greatly inspired by the replica computation for the Hopfield model [112] carried out by Amit *et al.* in Ref. [113] and by the unlearning mechanism [114–116]. Reference [110] provides a phase diagram for the decimation scheme and numerical evidence that, when the SNR is high,  $\alpha$  is low, and the prior is informative enough (such as the Rademacher case), decimation is able to outperform the RIE, despite requiring a computational time that is exponential in the size  $N$  of the system. The idea of decimation is that of retrieving the columns  $(\mathbf{X}_\mu)$  (i.e., patterns in an associative memory interpretation) one at a time. In order to do this, one needs to simplify the term  $-\lambda/(2N) \text{Tr}(\mathbf{x}\mathbf{x}^\dagger)^2$  in the Hamiltonian Eq. (17). The latter favors posterior samples whose columns  $(\mathbf{x}_\mu)$  remain quasiorthogonal to each other, a sort of ‘‘caging effect,’’ preventing pattern estimators to collapse onto the same ground truth patterns, thus boosting the joint recovery of all of them. This rigidity, which helps in the estimation, is completely given up on by decimation in favor of an analytically tractable model, paying the price of a brittle estimation: it works only for rather high SNR. On the

contrary, in extensive-rank matrix denoising, it is possible to give a nontrivial estimate of  $\mathbf{X}\mathbf{X}^\top$  as soon as  $\lambda > 0$  via the RIE. Therefore, a comprehensive theory should be able to take into account this caging effect.

Recently, Ref. [85] put forward a conjecture of “universality,” intended as independence of the MMSE and MI on the prior  $P_X$  (as long as it is centered and with unit variance), which can be formulated in a weak and a strong version. The weak one states that the RIE is optimal among finite-degree polynomial estimators in the data matrix elements, whereas the strong one would posit the universality to hold for all estimators. The strong version would imply that, under Gaussian additive noise, the MI is the one corresponding to the Gaussian prior case of Refs. [58,88,91,105]. However, this is in contrast to the evidence produced with decimation algorithms. From our analysis, a richer picture emerges. We disprove strong universality rigorously for any  $\alpha > 0$ . More precisely, we can prove that there exists an SNR beyond which the RIE cannot be optimal. Nevertheless, the conjecture is not to be completely discarded. In fact, we argue that universality of the MI and MMSE holds for low enough SNR, as supported by the next section’s numerics. A possible explanation for the breaking of validity of the analysis in Ref. [85] arguing strong universality is the following. The argument is based on the analysis of low-degree polynomials [117]. The author is able to show that no algorithm based on degree- $D$  polynomials of the data entries with  $D \leq 3$  can outperform the RIE when  $N \rightarrow +\infty$ . He also suggests that his argument extends to larger  $D$ . Then, assuming that the  $N \rightarrow +\infty$  and  $D \rightarrow +\infty$  can be exchanged, he concludes that no function of the data (approximated by a polynomial of arbitrary degree) can thus serve as better estimator. However, commutation of limits can break at a phase transition, as mentioned by the author. Therefore, it may be that the argument is valid up to a critical SNR, a point beyond which

more care should be taken with the large degree and size limits.

The universality-breaking threshold  $\lambda_c$  we identify in Sec. V depends on  $\alpha$ . In particular, recent literature supports the fact that when  $\alpha \rightarrow 0$  the universality phase should shrink below  $\lambda_c(\alpha = 0) =: \bar{\lambda}_c > 0$  (see Fig. 7). In Ref. [54], the authors prove that with a factorized prior, under some suitable hypothesis, the MI and MMSE reduce to those of the standard rank-one formulas when  $M = M_N$  grows sufficiently slowly, which is prior dependent and thus clearly not universal above the phase transition. Note that, for SNRs lower than the critical one, universality holds trivially at low rank, because the MMSE is that of the null estimator, i.e., 1. On top of it, here we show in Appendix E that the mean-square error (MSE) of the RIE when  $\alpha \rightarrow 0$  is that of naive PCA, which is worse than the Bayes-optimal performance for nonrotational-invariant priors beyond the transition. Furthermore, still for  $\alpha \rightarrow 0$ , decimation’s performance reduces to the one found via the low-rank formula for the MI [111] and is thus likely optimal in said limit. All this evidence supports that universality cannot hold for all  $\lambda$  when  $\alpha \rightarrow 0$ .

### C. Organization and main contributions

We now outline the organization of the remaining, while presenting our main results in order of appearance.

In Sec. III, we present the observations obtained from numerical experiments based on parallel tempering [118,119] and annealed importance sampling [120] used to compute the MI and MMSE.

- (i) Our numerical evidence in Sec. III A, see Fig. 2–4, suggests the existence of a phase transition occurring for structured enough priors, coined “denoising-factorization transition,” separating two regions of

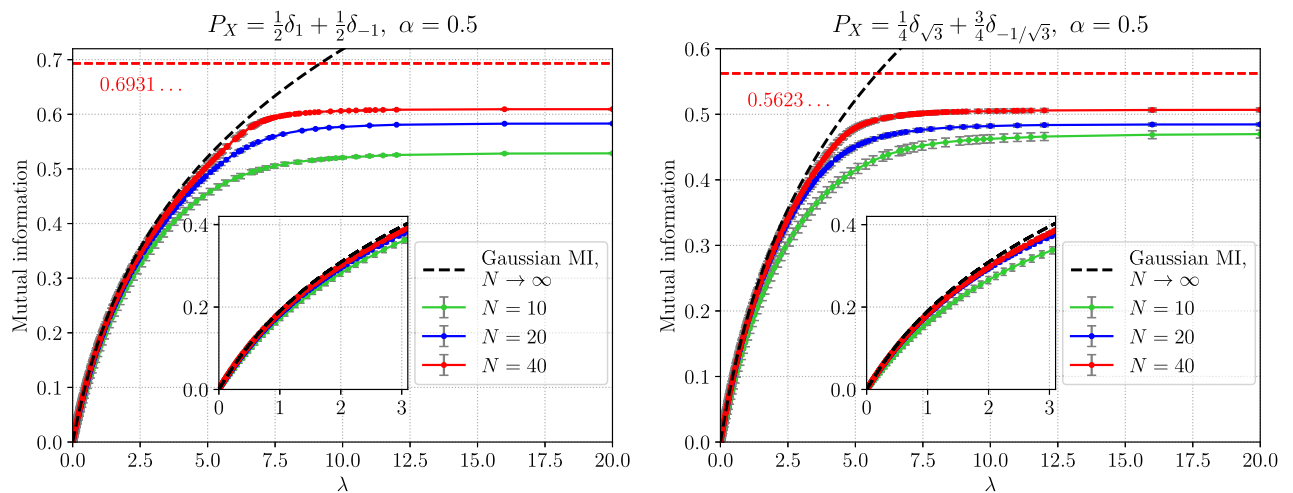


FIG. 2. Monte Carlo results for the mutual information with Rademacher prior (left) and discrete prior  $P_X = \frac{1}{4}\delta_{\sqrt{3}} + \frac{3}{4}\delta_{-1/\sqrt{3}}$  (right) for various sizes with  $\alpha = 0.5$ , i.e.,  $M = N/2$ . Error bars represent the standard error of the mean. The curves are compared to the exact infinite-size limit Eq. (10) in the case of standard Gaussian prior computed using the HCIZ integral. The red dashed curve corresponds to Shannon entropy of the prior considered in each panel, which bounds the MI for all sizes from above.

the diagram of different nature: a low-SNR region where prior universality holds and its complement where it does not.

- (ii) These numerics also support the fact that the rotational-invariant estimator discussed earlier is optimal for SNRs smaller than the critical value  $\lambda_c(\alpha)$  [later identified theoretically by Eq. (36)] and suboptimal for larger values. They also suggest in Fig. 5 that despite outperforming the RIE is information theoretically possible beyond  $\lambda_c(\alpha)$ , this is algorithmically hard for Monte Carlo methods without a sufficiently informative initialization.
- (iii) Importantly, the suboptimality of the RIE and the breaking of prior universality for discrete priors is proven in Sec. III B, see Eq. (11), thus discarding the strong version of the universality conjecture in Ref. [85] in a portion of the diagram. Yet, as previously mentioned, our numerics suggest that universality does hold before the transition, which we exploit for the analytical solution in that region, see result (iv).

All these findings are summarized by a qualitative phase diagram in  $(\alpha, \lambda)$  plane; see Fig. 7 which generalizes the low-rank picture of Fig. 1.

Section IV develops a “multiscale mean-field theory” able to match the numerical observations at high enough SNR. It relies on a combination of methods from spin glasses. Our approach yields equations characterizing the MMSE through Eq. (33) and the MI Eq. (35), also obtainable by an earlier replica calculation by Sakata and Kabashima Ref. [99], which was deemed incorrect in the whole phase diagram  $(\alpha, \lambda)$  for reasons discussed in Sec. II B. With our independent multiscale mean-field approach, we redeem their results, but only in a restricted region of the diagram.

- (iv) In Sec. VA, we identify this subregion of the phase diagram where these two *a posteriori* equivalent, mean-field approaches are conjecturally exact: the “factorization” or “nonuniversal phase” after the transition Eq. (36). Using our numerical findings (i, (iii) and the work in Ref. [85] while assuming the existence of a single transition point, we supersede the mean-field equations by the known exact ones for the Gaussian prior in the complement of the phase diagram based on the universality of the solution in  $P_X$ : This is the denoising or universal phase. The resulting complete *Ansatz* therefore takes into account these two scenarios (effective presence or lack of prior universality) and provides expressions for the MI between data and signal and for the MMSE, for all  $(\alpha, \lambda)$ : It corresponds to formula Eq. (37) paired with Eq. (35) for the MI after the identified transition Eq. (36), and Eq. (33) for the MMSE. Instead, for lower SNR than the transition, the HCIZ integral Eq. (10) is used for the

MI and Eq. (9) for the MMSE. This complete *Ansatz* together with (iii) are the main theoretical results, whose consequences are discussed in Sec. V B.

Our complete *Ansatz* Eq. (37) also brings forward the major conceptual result of this work: in extensive-rank matrix denoising, the model behaves in the two identified phases in fundamentally different ways, which require treatments based on very different techniques. The universal phase is akin to a matrix model, where universality properties hold and the HCIZ integral is the proper tool. In the factorization phase, strong concentration-of-measure effects take place and the model is instead akin to a planted spin glass treatable by mean-field methods. This mixed behavior in a single model is an interesting feature, in contrast to the low-rank regime, which remains always mean field in nature.

Section VI elaborates around the possibility to carry out factorization, namely, finding  $\mathbf{X}$  alone (rather than the covariance  $\mathbf{X}\mathbf{X}^\top$ ), up to unresolvable symmetries.

- (v) There we propose Algorithm 2, which is empirically shown to estimate  $\mathbf{X}$  in the factorization phase (thus its name) when given posterior samples obtained by Monte Carlo simulation. We also argue that, in contrast, its estimation is asymptotically impossible in the denoising phase where only nontrivial denoising of  $\mathbf{X}\mathbf{X}^\top$  is statistically possible.

Finally, the reader can find Appendixes after the Conclusions in the following order. In Appendix A, we discuss the irrelevance (from an information-theoretical perspective) of the diagonal of  $\mathbf{Y}$ . Appendix B gives the details about our numerical experiments carried out using Monte Carlo sampling. In Appendix C, we report the derivations of our multiscale mean-field theory of Sec. IV. In Appendix D, we discuss the main properties of the multiscale mean-field *Ansatz* for the MI. In Appendix E, we compute the  $\alpha \rightarrow 0$  limit of the estimation performance of the RIE. Appendix F contains the reduction of the approach of Ref. [99] to our setting. Appendix G explores potentially richer *Ansätze* for the MI, which end up yielding the same solution, thus supporting the validity of the proposed mean-field solution in the factorization phase.

The Monte Carlo simulations used in this paper are accessible at link in Ref. [121].

### III. EXPLORING THE PHASE DIAGRAM

We start by illustrating some important findings about the nature of extensive-rank matrix denoising that can be probed using numerics and simple rigorous arguments. To do so, we conducted replica-exchange Monte Carlo simulations, the details of which can be found in Appendix B. The left panel of Fig. 2 shows the results concerning the finite-size value of the MI with  $\alpha = 0.5$  for the Rademacher




FIG. 3. Monte Carlo results for the MMSE for  $N = 10$  (left),  $20$  (middle), and  $40$  (right) with Rademacher prior (solid colored lines) and  $\alpha = 0.5$ . They are compared to the MSE of the RIE for both Rademacher and Gaussian priors with corresponding  $N$ . Error bars for the RIE are omitted since they are too small to be visible. Error bars for the MMSE represent standard errors of the mean. Insets are in log scale.

prior  $P_X = \frac{1}{2}\delta_1 + \frac{1}{2}\delta_{-1}$ , which will be our main running example of discrete prior with same the mean and variance as the standard Gaussian prior throughout this paper. However, most of the conclusions we draw are not specific to the Rademacher prior; in fact, we also present the simulation results for the nonsymmetric distribution  $P_X = \frac{1}{4}\delta_{\sqrt{3}} + \frac{3}{4}\delta_{-1/\sqrt{3}}$  (again having the same mean and variance) in the right panel of Fig. 2 and demonstrate that quantitatively similar conclusions can be drawn. Because in the Gaussian case we know the exact asymptotic solution, we use as comparison the curve computed using the HCIZ integral corresponding to the  $N, M \rightarrow \infty$  limit of the MI, see Eqs. (9) and (10) and text below.

According to the I-MMSE relation Eq. (5), the derivative of the MI with respect to the SNR gives access to the MMSE. We plot it on Fig. 3 (where it was computed by Monte Carlo sampling, rather than by numerical differentiation of the MI) for the Rademacher prior. In addition, we plot the performance of the RIE for the same sizes. We recall that the RIE denoises the data  $\mathbf{Y}$  through eigenvalue shrinkage (keeping its eigenvectors untouched) in order to estimate  $\mathbf{X}\mathbf{X}^T$ , while exploiting its factorized structure and prior through only its asymptotic spectral law. This procedure is provably Bayes optimal when the signal is rotationally invariant, which corresponds to Gaussian prior  $P_X = \mathcal{N}(0, 1)$  in the present setting, as the MSE of the RIE matches the MMSE Eq. (9), see Refs. [58,122].

### A. Partial universality, a phase transition, and a statistical-to-computational gap

As mentioned before, the recent strong conjecture from Semerjian in Ref. [85] states that for any prior measure  $P_X$  for the i.i.d. entries of  $\mathbf{X}$ , which is centered with unit variance, the MI and MMSE are universal in the sense that, for a large system, their values match for all  $\lambda \geq 0$  and  $\alpha > 0$  the ones in the Gaussian case  $P_X = \mathcal{N}(0, 1)$  (up to vanishing, nonuniversal corrections). Moreover, this MMSE performance

should be reachable efficiently using the RIE. That this conjecture is wrong, in general, is clear from our finite-size numerical experiments. Indeed, a striking observation is that, beyond a certain SNR threshold, the MI curves for the Gaussian and discrete priors display an evident discrepancy in the behavior of their respective slopes, a point beyond which the MI for the discrete priors saturate quickly. This effect becoming sharper as the system size grows suggests a first-order discontinuous phase transition in the  $\lambda$  derivative of the MI, and thus the MMSE, which suddenly decreases beyond it in Fig. 3. The same effects suggesting a first-order transition are also observed for the nonsymmetric discrete prior, for which the MMSE is provided in Fig. 4, where the sharpening of the transition with increased system size becomes even more apparent. Based on the current numerics and later on our theory (see Sec. VB), we believe that in the thermodynamic limit this denoising-factorization transition  $\lambda_c(\alpha)$  is first order for  $\alpha > 0$ ; instead, for  $\alpha \rightarrow 0$ , which recovers the low-rank setting, the transition can be second order (continuous in the MMSE), like for the Rademacher prior.




FIG. 4. Monte Carlo results for the MMSE for  $N = 10, 20$ , and  $40$  with prior  $P_X = \frac{1}{4}\delta_{\sqrt{3}} + \frac{3}{4}\delta_{-1/\sqrt{3}}$  and  $\alpha = 0.5$ . Error bars represent standard errors of the mean.

Another important observation is the good agreement before the transition between the asymptotic MI for Gaussian prior and the finite-size MI for both discrete priors, which improves with increasing  $N$  and  $M$ . The numerics thus strongly suggest that, despite the universality not holding for all SNRs, it does before the transition in the thermodynamic limit  $N \rightarrow +\infty$ . Also, at the level of the MSE, we observe an almost perfect matching between the MMSE for the Rademacher prior computed by Monte Carlo sampling and the performance of the RIE already at  $N = 40$ , the latter being evaluated for both Rademacher and Gaussian priors (the same holds for the nonsymmetric prior of Fig. 4). As it should, the RIE performance is independent of the prior details as long as mean and variance match. However, while the RIE is Bayes optimal for all SNR in the Gaussian case, beyond  $\lambda_c$  it becomes suboptimal for the tested structured priors, as expected if universality breaks down, given that it cannot exploit the discrete nature of the priors.

A natural question is then whether in the region  $\lambda > \lambda_c$  there is a statistical-to-computational gap that prevents not only the RIE to reach the MMSE performance, but also a larger class of efficient algorithms including Bayesian ones like Monte Carlo methods. To address this, we also run standard Monte Carlo Metropolis-Hastings simulations (without replica exchange) with different initialization schemes and estimators summarized here:

Experimental setup in Fig. 5:

Red circle: A single Monte Carlo simulation with a random initialization, in which the sample average over the Markov chain is given as an estimator.




FIG. 5. MSE of the Metropolis-Hastings algorithm for estimating  $\mathbf{X}\mathbf{X}^T$  with  $(N, M) = (40, 28)$ , so  $\alpha = 0.7$ , with each estimator (red circle, blue triangle, and black diamond markers) designated in the main text under “Experimental setup.” Before the transition, all markers overlap. These are compared to the RIE performance in the large system limit (blue line). Experimental markers are averaged over 36 i.i.d. instances of the problem  $(\mathbf{X}, \mathbf{Z})$ .

Blue triangle: Results from 32 Monte Carlo simulations, each given a different random initialization. The sample average over all 32 Markov chains is the estimator.

Black diamond: A single Monte Carlo simulation initialized from the ground truth.

Further details are given in Appendix B. The numerics indicate that only with a strongly informative initialization about the signal (black diamond), outperforming the RIE becomes possible. Moreover, we observe that, beyond the transition where the “informative factorization state” appears ( $\lambda$  larger than  $\sim 7.5$ ), the performance of an estimator constructed from a single Monte Carlo chain with random initialization (red circle) is worse than the RIE (the name “factorization state” is linked to the possibility of solving the factorization task using samples associated with this thermodynamic state, see Sec. VI). The poor performance of Monte Carlo simulation is probably due to a dynamical glass phase transition [123] where exponentially many metastable states appear [124], preventing it to properly sample the equilibrium corresponding to the factorization state beyond  $\approx 7.5$ , recall Fig. 3. The metastable states are, however, correlated due to the data. Therefore, an additional average over random initializations (and thus over associated metastable states) should somehow “clean” the effect of glassiness. This hypothesis can be tested by running several Monte Carlo chains, each with an independent random initialization, and adopting the sample average over all of these chains as an estimator (blue triangle). The performance reached in this way matches precisely that of the RIE, so the aforementioned cleaning takes place, suggesting that the dynamical glass corresponds to a shattering of the RIE state. Yet, the equilibrium factorization state is not reached, which would require instead averaging over an exponentially (in  $N$ ) large number of Monte Carlo samples (or a smarter sampling procedure, like the replica-exchange Monte Carlo algorithm used for Fig. 3, which is not an inference algorithm). We conclude that the dynamical glass phase can be cleaned in polynomial time but the final outcome cannot outperform the RIE, which supports the existence of a hard phase. Algorithmic hardness was hypothesized in [85] as a possible consequence of the validity of the weak universality conjecture but not of the strong one.

To summarize, the picture that emerges is that of universality of the information-theoretic quantities effectively holding at low SNR, which suggests calling this region the universal phase. As proven in the next subsection, universality cannot hold for a discrete prior beyond a certain SNR. We claim that prior universality breaks down at larger SNR through a phase transition marking the starting point of the nonuniversal phase, which is also the hard phase for Monte Carlo algorithms.

The hard nonuniversal phase can also be alternatively called the factorization phase. The name comes from the fact that, in order to denoise  $\mathbf{X}\mathbf{X}^T$  optimally in this phase, exploiting all the prior information about the factorized structure of the signal and  $P_X$ , beyond simply  $\rho_{\mathbf{X}\mathbf{X}^T}$  as the RIE




FIG. 6. Monte Carlo results for the MMSE for  $N \in \{20, 40, 60\}$  with Rademacher prior, with  $\alpha = 0.5$  and SNR  $\lambda$  in the factorization phase. We also plot the rank-one MMSE and the MMSE obtained from our multiscale theory (see Sec. IV) for comparison. Error bars represent standard errors of the mean.

does, becomes necessary and, importantly, factorization is possible only there, see Sec. VI. The universal phase is instead algorithmically easy by optimality of the RIE for denoising there. Thus, we also call this region the denoising phase.

In the factorization phase, there is a close match between the MMSE and the one in the rank-one setting where factorization occurs, see Fig. 6. Quantitatively, this translates into an exponential decay of the MMSE as a function of the SNR (similar to what happens in dictionary learning, see Fig. 5 of [102]), whose exponent is the same as in rank-one matrix estimation, see Fig. 6. We also stress the even better agreement of our theory presented in Sec. VA with the  $N = 60$  points in this figure. In particular, despite the gap between the green points and the rank-one curve being very small, our theory can explain it. In contrast, in the denoising phase, the MMSE decay in  $\lambda$  is polynomial, see insets of Fig. 3. This change of polynomial to exponential decay of the MMSE will be captured by the theory, as well as the fact that the MMSE is close to the rank-one curve in the factorization phase, see Sec. VB.

### B. Breaking of universality beyond the transition

With all these observations in mind, we thus agree with the strong universality conjecture of Ref. [85] when the SNR is sufficiently low, but we can rigorously discard it beyond a threshold using a simple argument based on the definition of the MI. Consider a discrete prior  $P_X$  with same mean and variance as  $\mathcal{N}(0, 1)$ . Since  $\mathbf{X}$  is discrete, the MI can always be written as (see Theorem 3.4 in Ref. [81])

$$\frac{I(\mathbf{X}\mathbf{X}^\top; \mathbf{Y})}{MN} = \frac{\mathcal{H}(\mathbf{X}\mathbf{X}^\top)}{MN} - \frac{\mathcal{H}(\mathbf{X}\mathbf{X}^\top | \mathbf{Y})}{MN} \leq \frac{\mathcal{H}(\mathbf{X}\mathbf{X}^\top)}{MN},$$

where the upper bound follows from the fact that for discrete priors both entropies are non-negative. That is it: The MI cannot grow indefinitely as a function of the SNR for a discrete prior (the upper bound remains finite for any size), while it does so for the Gaussian prior (black dashed curve in Fig. 2). This discards the possibility of universality at least beyond the point where the MI density for the Gaussian case crosses the entropy density associated with the matrix  $\mathbf{X}\mathbf{X}^\top$ . We are going to show below that  $|\mathcal{H}(\mathbf{X}\mathbf{X}^\top) - \mathcal{H}(\mathbf{X})| = o(MN)$ . Therefore, we have a bound on the denoising-factorization transition. Let  $X \sim P_X$ , a centered discrete prior of unit variance. We have the following upper bound on the domain of possible prior universality:

$$\lambda_c(\alpha) \leq \inf\{\lambda \geq 0 : t^G(\alpha, \lambda) \geq \mathcal{H}(X)\}. \quad (11)$$

Because  $t^G(\alpha, \lambda)$  is known using the HCIZ integral via Eqs. (9) and (10), the bound is explicit [we have used that the prior is factorized over the signal matrix entries to get  $\mathcal{H}(\mathbf{X})/MN = \mathcal{H}(X)$ , but a similar bound can be written even if there is no such factorization]. In the case of Rademacher prior, its entropy equals  $\ln 2$ . We indeed see in the left panel of Fig. 2 that the bound holds (also for the other prior on the right panel), but it does not look sharp.

Interestingly, we can also explain the apparent gap between the value toward which the finite-size MI saturates and the entropy of the prior. In fact, the observations  $\mathbf{Y}$  are informed about the ground truth  $\mathbf{X}$  through only the combination  $\mathbf{X}\mathbf{X}^\top = \sum_{\mu \leq M} \mathbf{X}_\mu \mathbf{X}_\mu^\top$ . Hence, for large  $\lambda$ , we shall have  $I(\mathbf{X}\mathbf{X}^\top; \mathbf{Y}) \approx \mathcal{H}(\mathbf{X}\mathbf{X}^\top)$  at leading order, which is smaller than  $\mathcal{H}(\mathbf{X})$  by the data processing inequality. The inequality in this case is strict:  $\mathcal{H}(\mathbf{X}\mathbf{X}^\top) < \mathcal{H}(\mathbf{X})$  because  $\mathbf{X}\mathbf{X}^\top$  is a noninvertible function of  $\mathbf{X}$ . The degeneracy of the mapping  $\mathbf{X} \mapsto \mathbf{X}\mathbf{X}^\top$  explains accurately the gap. Specifically, for each  $\mathbf{X}\mathbf{X}^\top$  there can be  $M!$  permutations of the columns of  $\mathbf{X}$  that leave it invariant. Furthermore, if the prior is also symmetric, as in the Rademacher prior case in the left panel of Fig. 2, we must add possible sign flips of these columns to the invariance group of  $\mathbf{X}\mathbf{X}^\top$ . Of these  $2^M M!$ , some are actually redundant because columns of  $\mathbf{X}$  can match. The probability that any two columns match with a factorized Rademacher prior, however, is  $O(M^2 2^{-N})$ , which is negligible for large enough  $N$  with respect to the leading-order corrections that we compute. The probability of more columns matching is of even lower order. Hence, we can finally say that

$$\mathcal{H}(\mathbf{X}) = \mathcal{H}(\mathbf{X}\mathbf{X}^\top) + \ln(2^M M!) - o_N(1). \quad (12)$$

Putting these considerations together, we readily get

$$\frac{I(\mathbf{X}\mathbf{X}^\top; \mathbf{Y})}{MN} \xrightarrow{\lambda \rightarrow \infty} \frac{\mathcal{H}(\mathbf{X})}{MN} - \frac{\ln(2^M M!)}{MN} + o_N(1), \quad (13)$$

where the remainder can be estimated to be of order  $O(2^{-N}M \ln M)$ . It is now manifest that as  $M, N \rightarrow \infty$ , despite the degeneracy induced by signed permutations, the contributions yielded by the second and third terms on the rhs vanish. This explains why increasing  $N$  and  $M$  in Fig. 2 is sufficient to reduce the gap with the saturation value  $\ln 2$ . For example, for  $N = 20, M = 10$  the contribution of the signed permutations is approximately 0.1102 which summed to the saturation value of the blue curve at approximately 0.5826 yields approximately 0.6928, which is remarkably close to  $\ln 2 \approx 0.6931$ . The same holds for the other gaps, and the argument also explains the discrepancy for the nonsymmetric prior in the right panel of Fig. 2, where in this case one does not need to consider sign invariance. Note that we do not claim that beyond the transition the thermodynamic limit of the MI should saturate strictly to the prior entropy. Indeed, Fig. 3 shows that a small gap survives whose size decreases exponentially fast with  $\lambda$ .

The rhs of Eq. (11) is a rigorous bound above which universality of the MI and MMSE cannot hold for discrete priors in any possible scenario. We note that, given the extensive set of works on universality in random matrices in Refs. [59,86,125–131], it may come as a surprise that it breaks down. For example, at the level of the asymptotic spectrum of  $\mathbf{X}\mathbf{X}^\top$  and thus of  $\mathbf{Y}$ , or even at the level of more local statistics, the previous papers show universality. However, the mutual information depends on all the fine statistical properties of the data and  $\mathbf{X}$ , which become manifest at high enough SNR.

We also remark that the reasoning leading to the existence of a threshold beyond which universality with respect to the base measure of the matrix to infer cannot hold can be extended to a much broader set of matrix inference and learning models. In any such model, if the signal  $\mathbf{X}$  is discrete, the MI between data and signal cannot exceed the entropy of  $\mathbf{X}$  (as a function of the SNR, number of data, etc.). This establishes an upper bound on the region of possible validity of prior universality in models where replacing the signal by a Gaussian equivalent with matching moments yields an unbounded MI. Indeed, this MI saturation has already been shown to have deep implications in neural networks [79].

Whether this universality breaking holds for continuous (but non-Gaussian) priors is an unanswered question in this paper. We conjecture that, like the case with discrete priors, the nonrotational-invariant nature of the prior should be exploitable to construct a Bayesian estimator with performance well exceeding the RIE once the SNR is high enough. However, as the mutual information cannot be bounded from above in such settings, the previous arguments negating universality do not apply. Performing numerical experiments under continuous priors is also considerably more difficult due to the inability to exploit the binary nature of  $\mathbf{x}$  in the implementation for fast

computations. Moreover, the enlargement of the phase space from a discrete one to a continuous one may result in a longer relaxation time for the Monte Carlo chain to converge. Confirming whether this conjecture holds or not, either theoretically or numerically, is a challenging problem left open for future investigation.

Note that while  $N, M$  may seem small, the total number of variables itself is  $MN$ , which ranges from 50 to 1800 in our experiments. This is comparable with the extensive numerical studies done on the Edwards-Anderson Ising spin glass using parallel tempering [132], in which the system size ranges from 64 to 1000.

### 1. Summary and phase diagram

Our findings are condensed into a phase diagram in the  $(\alpha, \lambda)$  plane, see Fig. 7, which generalizes Fig. 1 to the extensive-rank regime. We have identified two phases for matrix denoising with discrete prior under the hypothesis that the denoising-factorization transition  $\lambda_c(\alpha)$  exists and is unique. In the denoising or universal phase  $\lambda \in [0, \lambda_c(\alpha))$ , treatable using matrix model techniques and the HCIZ integral, we conjecture that it is statistically not possible to




FIG. 7. Qualitative phase diagram for matrix denoising with discrete (Rademacher) prior. This picture generalizes that of the low-rank case in Fig. 1 to the extensive-rank regime. The thermodynamic phase transition  $\lambda_c(\alpha)$  (solid line) is believed to be first order for  $\alpha > 0$ . Increasing  $\lambda$ , it is preceded by a spinodal transition where the thermodynamic factorization state appears but is not yet the equilibrium (dashed line). The blue region is the denoising phase where prior universality holds, and optimal denoising is algorithmically easy using the RIE. The shaded region corresponds to the dynamical glass phase (we cannot assess whether it starts precisely at the spinodal transition as displayed). The red region is the factorization phase where prior universality does not hold. There, it is statistically possible but algorithmically hard to estimate  $\mathbf{X}$  and to denoise  $\mathbf{X}\mathbf{X}^\top$  better than the RIE. The line  $\alpha = 0$  recovers the results for the rank-one problem and is thus special:  $\bar{\lambda}_c = \lambda_c(0)$  is second order and there is no hard phase for Rademacher prior. For other priors (possibly discrete), a hard phase can appear in the low-rank case too, see Fig. 1.

outperform the RIE for denoising  $\mathbf{X}\mathbf{X}^\top$  despite it using only its asymptotic spectral distribution as prior. It is therefore algorithmically easy to reach the MMSE performance (i.e., in polynomial complexity in  $N$ ). Universality in the prior of the MI and MMSE holds and may be linked to the quasiuniformity of the eigenvectors of  $\mathbf{X}\mathbf{X}^\top$  [133–136] and the universality of its spectrum [86,130] when  $\mathbf{X}$  is made of i.i.d. entries. We argue in Sec. VI that, in this region, factorization, namely, estimating  $\mathbf{X}$  up to a unresolvable ambiguities inherent to the mapping  $\mathbf{X} \mapsto \mathbf{X}\mathbf{X}^\top$ , is statistically impossible.

Beyond the transition point, we find the factorization or nonuniversal phase  $\lambda \in (\lambda_c(\alpha), +\infty)$  where techniques from mean-field spin glasses become correct. In this phase, exploiting the complete prior knowledge about  $\mathbf{X}\mathbf{X}^\top$  beyond its spectrum becomes essential: Information theoretically, a Bayesian algorithm, such as Monte Carlo sampling, can outperform the RIE for denoising. Moreover, factorization is also possible on top of denoising. Both tasks are, however, algorithmically hard due to the presence of a dynamical glass phase appearing through a spinodal transition.

Finally, we can also interpret the transition as marking the starting point of a feature learning regime, where it becomes information theoretically possible to detect and exploit fine features in the data for achieving better denoising and factorization. In the denoising phase, on the contrary, even knowing in advance the features (factorized and discrete structure) does not yield any improvement compared to assuming no particular structure (i.e., a Gaussian prior).

Notice two major differences compared to the phase diagram in Fig. 1 for  $M = 1$ . First, in the extensive-rank case, the diagram is defined on a plane  $(\alpha, \lambda)$  rather than parametrized by the SNR  $\lambda$  only. The transition is thus now a line  $\lambda_c(\alpha)$  and its nature (order) can change along it. Indeed, between  $\alpha = 0$  or strictly positive, it does change, for instance, in the Rademacher prior case. Second, at fixed  $\alpha > 0$ , no easy regime appears once the transition is crossed, even at very large SNR. Extensive-rank matrix denoising is thus algorithmically much more demanding than in the low-rank case for large  $\lambda$ .

#### IV. MULTISCALE MEAN-FIELD THEORY

We now outline a theory matching the numerical observations discussed earlier. As already mentioned, we cannot employ methods from matrix models such as the HCIZ integral. Instead, we rely on spin-glass techniques concatenated in an original way. The proposed approach is called multiscale mean-field theory because it uses multiple mean-field techniques in a multiscale fashion. The details of the derivations of this section are deferred to Appendix C.

To begin with, as proved in Appendix A, we observe that the diagonal of the data  $(Y_{ii})_{i \leq N}$  does not contribute to the MI in the thermodynamic limit. In fact,  $Y_{ii} = \sqrt{\lambda/N} \|\mathbf{X}_i\|^2 + Z_{ii} \approx \sqrt{\lambda/N} \|\mathbf{X}_i\|^2$ . However,  $\|\mathbf{X}_i\|^2/N$

can be already estimated precisely without knowing  $Y_{ii}$ : Its value is known to a Bayes-optimal statistician, as it concentrates on the second moment of the prior (i.e., 1) by the law of large numbers. We thus now work with

$$Y_{ij} = Y_{ji} = \sqrt{\frac{\lambda}{N}} \mathbf{X}_i^\top \mathbf{X}_j + Z_{ij}, \quad 1 \leq i < j \leq N. \quad (14)$$

The MI for this channel then reads

$$\frac{I(\mathbf{X}\mathbf{X}^\top; \mathbf{Y})}{MN} = \frac{\lambda}{4} - \phi_N + O(N^{-1}) \quad (15)$$

with “free entropy”

$$\phi_N := \frac{1}{MN} \mathbb{E}_{\mathbf{Y}} \ln \mathcal{Z}_{N,M}(\mathbf{Y}). \quad (16)$$

Using the language of statistical mechanics, we introduced the partition function of the model

$$\mathcal{Z}_{N,M}(\mathbf{Y}) := \int_{\mathbb{R}^{N \times M}} dP_{\mathbf{X}}(\mathbf{x}) e^{-H_N(\mathbf{x}; \mathbf{Y})}$$

with (minus) Hamiltonian (which is the log-likelihood)

$$-H_N(\mathbf{x}; \mathbf{Y}) = \sqrt{\frac{\lambda}{N}} \sum_{i < j, 1}^N Y_{ij} \mathbf{x}_i^\top \mathbf{x}_j - \frac{\lambda}{2N} \sum_{i < j, 1}^N (\mathbf{x}_i^\top \mathbf{x}_j)^2. \quad (17)$$

The associated Boltzmann-Gibbs average, i.e., the posterior average, will also be denoted by  $\langle \cdot \rangle$  in this section.

The MMSE of the inference problem Eq. (14) does not include any error on the reconstruction of the diagonal part of the signal  $\mathbf{X}\mathbf{X}^\top$  and is thus defined (with a little abuse of notation) as

$$\begin{aligned} \text{MMSE} &= \frac{1}{MN^2} \mathbb{E} \sum_{i \neq j, 1}^N [(\mathbf{X}_i^\top \mathbf{X}_j)^2 - \langle \mathbf{x}_i^\top \mathbf{x}_j \rangle^2] \\ &= 1 - \frac{1}{NM} \mathbb{E} \sum_{i=1}^{N-1} \langle \mathbf{x}_i^\top \mathbf{x}_N \rangle^2 + O(N^{-1}). \end{aligned} \quad (18)$$

One can show that the above is equal to Eq. (6) up to  $o_N(1)$ , which is why it has been denoted in the same way. This  $o_N(1)$  discrepancy is precisely due to the diagonal contributions that would appear in Eq. (6) (see Appendix A). The second line is instead obtained using the exchangeability of the rows, selecting  $j = N$  for convenience.

The free entropy  $\phi_N$  can be viewed as that of a nonstandard spin system made of  $N$  spins [the rows of  $\mathbf{x}$ ,  $(\mathbf{x}_i)_{i \leq N}$ ] which are  $M$  dimensional (i.e., have extensive dimensionality instead of being scalar or low-dimensional variables) or, instead,  $M$  spins [the columns of  $\mathbf{x}$ ,  $(\mathbf{x}_\mu)_{\mu \leq M}$ ] which are  $N$  dimensional. The two viewpoints are equivalent, but the first turns out to be more amenable for a cavity

computation, which would imply the extraction of a row of  $\mathbf{x}$ . The reason is that the part of the Hamiltonian depending on row  $\mathbf{x}_i \in \mathbb{R}^M$  carries a noise variable  $\mathbf{Z}_i \in \mathbb{R}^N$  independent of the rest. More precisely, the posterior can be decomposed in terms of rows  $\mathbf{x}_i$  and  $\mathbf{Y}_i = (Y_{ij})_{j \leq i-1}$ ,  $i = 1, \dots, N$ , as

$$\begin{aligned} dP(\mathbf{x}|\mathbf{Y}) &\propto dP_X(\mathbf{x}_N)P(\mathbf{Y}_N|\mathbf{x}_N, (\mathbf{x}_j)_{j \leq N-1}) \\ &\quad \times \prod_{i \leq N-1} dP_X(\mathbf{x}_i)P(\mathbf{Y}_i|\mathbf{x}_i, (\mathbf{x}_j)_{j \leq i-1}) \\ &\propto dP(\mathbf{x}_N|\mathbf{Y}_N, (\mathbf{x}_j)_{j \leq N-1}) \\ &\quad \times dP((\mathbf{x}_j)_{j \leq N-1} | (\mathbf{Y}_j)_{j \leq N-1}), \end{aligned} \quad (19)$$

where we have dropped normalizations. This decomposition following from Eq. (14) emphasizes that, conditional on the ‘‘bulk’’  $(\mathbf{x}_j)_{j \leq N-1}$ , the ‘‘cavity’’  $\mathbf{x}_N$  depends on a single data row  $\mathbf{Y}_N$ , and the bulk marginal posterior measure is independent of the cavity. Note that the same argument holds when extracting any other cavity row after a proper permutation of row indices. This decomposition is crucial as, first, it makes manifest that the posterior of the cavity conditional on the bulk corresponds to a lower-dimensional Bayes-optimal inference problem. Second, the bulk measure can be simplified without having to modify the cavity measure. Such a decomposition cannot be written if we dig a column cavity.

### A. Computing the limiting MMSE

The perspective introduced above is particularly useful when combined with the MMSE in Eq. (18). The latter is, in fact, the so-called YMMSE of a random linear estimation (RLE) problem with noisy design (see Refs. [137,138]), where samples  $\bar{\mathbf{x}} := (\mathbf{x}_i)_{i \leq N-1}$  are used for the estimation of the uncertain design matrix  $\bar{\mathbf{X}} \in \mathbb{R}^{(N-1) \times M}$ , while samples  $\mathbf{x}_N$  from the cavity posterior measure  $dP(\mathbf{x}_N|\mathbf{Y}_N, (\mathbf{x}_j)_{j \leq N-1})$  are used to estimate  $\mathbf{X}_N$ .

#### 1. First scale reduction

Mapping onto the RLE task has the advantage of converting a matrix model into a vector inference problem over the cavity (still with an annealed bulk matrix), potentially tractable using known approaches. This has the effect of ‘‘reducing the scale of the problem’’: we go from a partition function associated with a Hamiltonian of magnitude  $O(MN)$  to one with an  $O(M)$  Hamiltonian. This  $O(MN) \rightarrow O(M)$  reduction performed by means of the cavity method corresponds to the first scale reduction in our scheme. Contrary to standard spin or inference models (like the Hopfield model or the perceptron treated with the cavity method in Refs. [139,140]), the RLE problem that will appear remains high dimensional because  $M \rightarrow +\infty$  with  $N$ .

Let us now identify the correct RLE problem. We denote the cavity variable  $\mathbf{x}_N$  by  $\boldsymbol{\eta} \in \mathbb{R}^M$ ,  $\mathbf{X}_N$  by  $\mathbf{H} \in \mathbb{R}^M$ , and the bulk (i.e., all but the cavity) variables as  $\bar{\mathbf{x}}, \bar{\mathbf{X}} \in \mathbb{R}^{(N-1) \times M}$ . The RLE under study then corresponds to the two following observation channels:

$$\tilde{\mathbf{Y}}(\lambda) = \sqrt{\frac{\lambda}{N}} \bar{\mathbf{X}} \mathbf{H} + \tilde{\mathbf{Z}}, \quad (20)$$

$$Y'_{ij}(\zeta) = \sqrt{\frac{\zeta}{N}} \mathbf{X}_i^\dagger \mathbf{X}_j + Z_{ij}, \quad 1 \leq i < j \leq N-1. \quad (21)$$

The first one is the linear observation of the cavity ground truth variable  $\mathbf{H}$  through the uncertain design matrix  $\bar{\mathbf{X}}$ , whereas the second is additional information about  $\bar{\mathbf{X}}$  only.  $\zeta$  is an auxiliary SNR that should be taken equal to  $\lambda$ , but we consider a more general model where it may differ, in which case we denote the data  $\mathbf{Y}'(\zeta) = (Y'_{ij}(\zeta))_{i < j \leq N-1}$ . Denoting by  $\langle \cdot \rangle_\circ$  the posterior average acting on  $(\bar{\mathbf{x}}, \boldsymbol{\eta})$  of this Bayes-optimal inference problem over  $(\bar{\mathbf{X}}, \mathbf{H})$ , we define the YMMSE given Eqs. (20) and (21) as

$$\begin{aligned} \text{YMMSE}(\lambda, \zeta) &:= \frac{1}{NM} \mathbb{E} \|\bar{\mathbf{X}} \mathbf{H} - \langle \bar{\mathbf{x}} \boldsymbol{\eta} \rangle_\circ\|^2 \\ &= 1 - \frac{1}{NM} \mathbb{E} \|\langle \bar{\mathbf{x}} \boldsymbol{\eta} \rangle_\circ\|^2, \end{aligned}$$

where the second equality follows from the Nishimori identity Proposition 2 in Appendix C. This is the correct RLE problem because for  $\zeta = \lambda$  the above is equal to Eq. (18) up to vanishing errors (and also  $\langle \cdot \rangle = \langle \cdot \rangle_\circ$ ). The YMMSE is obtained by an I-MMSE relation similar to Eq. (5) through a  $\lambda$  derivative of the following conditional MI:

$$\begin{aligned} I_N^{\text{RLE}}(\lambda, \zeta) &:= \frac{1}{M} I((\bar{\mathbf{X}}, \mathbf{H}); \tilde{\mathbf{Y}}(\lambda) | \mathbf{Y}'(\zeta)) \\ &= \frac{1}{M} \mathcal{H}(\tilde{\mathbf{Y}}(\lambda) | \mathbf{Y}'(\zeta)) - \frac{1}{M} \mathcal{H}(\tilde{\mathbf{Y}}(\lambda) | \mathbf{Y}'(\zeta), \bar{\mathbf{X}}, \mathbf{H}). \end{aligned} \quad (22)$$

Specifically, we have

$$\frac{d}{d\lambda} I_N^{\text{RLE}}(\lambda, \zeta) = \frac{1}{2} \text{YMMSE}(\lambda, \zeta). \quad (23)$$

The reason we introduced  $\zeta$  in the first place is to be able to write down this relation. Taking  $\zeta = \lambda$  from the beginning would prevent accessing  $\text{MMSE} \approx \text{YMMSE}(\lambda, \zeta = \lambda)$  through an I-MMSE relation with SNR  $\lambda$  influencing a single (cavity) row. This is thus needed to relate the calculation of MMSE Eq. (18) to the RLE problem. The second entropy in Eq. (22) is just the one of the noise  $\tilde{\mathbf{Z}}$  and is thus simple. The first is instead related to the free entropy

$$\begin{aligned} \phi_N^{\text{RLE}}(\lambda, \zeta) &= \phi_N^{\text{RLE}} \\ &:= \frac{1}{M} \mathbb{E}_{\bar{\mathbf{Y}}(\lambda), \mathbf{Y}'(\zeta)} \ln \int_{\mathbb{R}^M} dP_X(\boldsymbol{\eta}) \left\langle e^{-H_{N,M}^{\text{row}}(\boldsymbol{\eta}; \bar{\mathbf{X}}, \lambda)} \right\rangle' \end{aligned} \quad (24)$$

by a constant shift and a sign. Here,  $H_{N,M}^{\text{row}}$  given by Eq. (C3) in Appendix C 1 is the ‘‘cavity Hamiltonian,’’ namely, the term capturing all the dependence on the extracted row  $\boldsymbol{\eta}$  (and its interaction with the bulk) in the total Hamiltonian Eq. (17). What remains of Eq. (17) is the ‘‘bulk Hamiltonian,’’ see Eq. (C2). The latter defines the posterior average  $\langle \cdot \rangle'$  given  $\mathbf{Y}'(\zeta)$  only, thus acting on the bulk  $\bar{\mathbf{X}}$  but not the cavity  $\boldsymbol{\eta}$ .  $\phi_N^{\text{RLE}}$  clarifies the connection between our mapping to the RLE setting and the cavity method. It actually mirrors a term appearing in the Aizenman-Sims-Starr representation of the free energy in standard spin glasses and inference [16,41].

## 2. Second scale reduction

The previous reduction by the cavity method is exact. However, the complexity due the inherent matrix nature of the problem has not disappeared: It is now segregated into the uncertainty of the design  $\bar{\mathbf{x}} = (\mathbf{x}_i)_{i \leq N-1}$  measured by  $\langle \cdot \rangle'$ . In order to integrate this bulk, we will resort to a mean-field *Ansatz*, yielding a model that is tractable by means of rigorous replica formulas [48], in which one can identify scalar inference problems with  $O(1)$  Hamiltonians. This  $O(M) \rightarrow O(1)$  magnitude reduction of the Hamiltonian corresponds to the second scale reduction in our approach.

Before doing this, let us make a remark. We could have tried to simplify the problem using an *Ansatz* directly at the level of the original Hamiltonian in order to go from  $O(MN)$  degrees of freedom to  $O(1)$  effective ones in one step, e.g., using the replica method as in Refs. [91,99,102,105]. However, having already reduced the scale down to  $O(M)$  by the cavity method in an exact manner, our assumptions will only be at the level of the bulk (and thus potentially more interpretable): the ‘‘direct’’ interaction between cavity and bulk captured by channel Eq. (20), which defines the first posterior in Eq. (19), is kept untouched, with approximations only made through the simplification of  $\langle \cdot \rangle'$ , i.e., the second posterior law in Eq. (19). In essence, our mean-field *Ansatz* assumes that the interactions among bulk variables have a lower-order effect on the free entropy  $\phi_N^{\text{RLE}}$  obtained by the cavity method compared to the direct cavity-bulk interaction.

In addition to being physically meaningful, these assumptions allow us to integrate the bulk variables. We achieve this by replacing channel Eq. (21) by an effective, more manageable one,

$$\mathbf{Y}'_{\text{eff}}(\sigma) = \sqrt{\sigma} \bar{\mathbf{X}} + \mathbf{Z}'_{\text{eff}}, \quad (25)$$

where  $\mathbf{Z}'_{\text{eff}} = (Z'_{i\mu})_{i \leq N-1, \mu \leq M}$  is a matrix of i.i.d. standard Gaussian variables,  $\bar{\mathbf{X}}$  is the ground truth bulk as before,

and  $\sigma$  is a tunable SNR for the effective observations  $\mathbf{Y}'_{\text{eff}}(\sigma) = (\mathbf{Y}'_i \in \mathbb{R}^M)_{i \leq N-1} = (Y'_{i\mu})_{i \leq N-1, \mu \leq M}$ . This is the anticipated mean-field *Ansatz*, i.e., a parametrization of the bulk measure in terms of marginals only. Its potential validity in this extensive-rank setting will be thoroughly discussed later. For now, let us just notice that this form of marginal is asymptotically exact for standard low-rank models [16,30,48], provided  $\sigma$  is suitably tuned.

Within this effective factorized mean-field *Ansatz*, the simplified bulk measure reads

$$\langle \cdot \rangle'_{\text{eff}} = \prod_{i,\mu=1}^{N-1,M} \frac{1}{\mathcal{Z}'(Y'_{i\mu})} \int dP_X(x_{i\mu}) e^{\sqrt{\sigma} Y'_{i\mu} x_{i\mu} - \frac{1}{2} \sigma x_{i\mu}^2}(\cdot)$$

with proper normalization  $\mathcal{Z}'$ . We can now leverage on known results [48] in order to derive a variational expression for the limit of  $\phi_N^{\text{RLE}}$  after trading Eq. (21) for Eq. (25) (see Appendix C),

$$\begin{aligned} \phi_{\text{eff}}^{\text{RLE}}(\lambda, \sigma) &= \text{extr}_+ \left\{ \mathbb{E}_{\xi, H} \ln \int dP_X(\boldsymbol{\eta}) e^{(\xi \sqrt{r} + H r) \boldsymbol{\eta} - \frac{1}{2} r \boldsymbol{\eta}^2} \right. \\ &\quad \left. - \frac{r q}{2} + \frac{\lambda}{2} - \frac{1}{2\alpha} \ln(1 + \lambda \alpha (1 - J(\sigma) q)) \right\}, \end{aligned} \quad (26)$$

where  $\xi \sim \mathcal{N}(0, 1)$ ,  $H \sim P_X$ , and  $\text{extr}_+$  denotes the extremization operation with respect to  $(r, q)$ . The  $+$  subscript indicates that, in case of multiple extremal points, the one that realizes the largest value of the potential should be considered. Finally, we called this free entropy  $\phi_{\text{eff}}^{\text{RLE}}$  and not  $\phi^{\text{RLE}}$  to distinguish our *Ansatz* based on Eq. (25) from the true asymptotic free entropy. Through our *Ansatz*,  $\zeta$  has disappeared in favor of a properly chosen parameter  $\sigma$  (we will later see how to choose it) that appears in a function  $J(\sigma)$  defined as

$$J(\sigma) := \mathbb{E}_{X, Z'} \left[ X \frac{\int dP_X(x) x e^{(\sqrt{\sigma} Z' + X \sigma) x - \frac{1}{2} \sigma x^2}}{\int dP_X(x) e^{(\sqrt{\sigma} Z' + X \sigma) x - \frac{1}{2} \sigma x^2}} \right], \quad (27)$$

where  $X \sim P_X$  and  $Z' \sim \mathcal{N}(0, 1)$ . This is the overlap of a Gaussian channel of SNR  $\sigma$ . Based on definition Eq. (22), we can relate the free entropy  $\phi_{\text{eff}}^{\text{RLE}}$  for RLE with uncertain design with effective observations Eq. (25) to the corresponding approximation  $t_{\text{eff}}^{\text{RLE}}(\lambda, \sigma)$  of  $\lim_{N \rightarrow \infty} t_N^{\text{RLE}}(\lambda, \zeta)$ ,

$$t_{\text{eff}}^{\text{RLE}}(\lambda, \sigma) = \frac{\lambda}{2} - \phi_{\text{eff}}^{\text{RLE}}(\lambda, \sigma). \quad (28)$$

Note that, under the effective bulk observations Eq. (25), this equality is essentially rigorous. This is because, in this case, the problem can be mapped onto a standard RLE problem with perfect knowledge of the design matrix but with higher measurement noise (as we exploit in Appendix C 2), which is covered by the theorems in Ref. [48].

Before moving to the saddle point equations, we observe that in the Bayes-optimal setting the extremizer  $q_* = q_*(\lambda, \sigma)$  in Eq. (26) corresponds to the asymptotic overlap between the ground truth row  $\mathbf{H}$  and a sample  $\boldsymbol{\eta}$  from the posterior associated with the RLE problem with noisy covariates Eqs. (20) and (25) or between two conditionally independent posterior samples [by the Nishimori identity, see Eq. (C1)], which concentrate onto their mean given by

$$\begin{aligned} q_* &= \lim_{N \rightarrow \infty} \mathbb{E} \frac{1}{M} \mathbf{H} \cdot \mathbb{E}[\boldsymbol{\eta} | \tilde{\mathbf{Y}}(\lambda), \mathbf{Y}'_{\text{eff}}(\sigma)] \\ &= \lim_{N \rightarrow \infty} \mathbb{E} \frac{1}{M} \|\mathbb{E}[\boldsymbol{\eta} | \tilde{\mathbf{Y}}(\lambda), \mathbf{Y}'_{\text{eff}}(\sigma)]\|^2, \end{aligned}$$

where  $\mathbb{E}[\boldsymbol{\eta} | \tilde{\mathbf{Y}}, \mathbf{Y}'_{\text{eff}}]$  equals the posterior mean of the cavity given direct observations  $\tilde{\mathbf{Y}}$  about it and the effective observations for the bulk. Instead, the extremizer  $r_* = r_*(\lambda, \sigma)$  is the effective SNR that controls the cavity fields underlying the marginals of the cavity coordinates under model Eqs. (20) and (25): asymptotically, the un-normalized random marginal probability measure of any of the  $\eta_i$  reads

$$dP(\eta_i \in [\eta, \eta + d\eta]) \propto dP_X(\eta) e^{(\xi_i \sqrt{r_*} + H_i r_*) \eta - \frac{1}{2} r_* \eta^2}, \quad (29)$$

with randomness  $\xi_i \sim \mathcal{N}(0, 1)$  and  $H_i \sim P_X$ . In other words, in the thermodynamic limit, a cavity coordinate has marginal corresponding to a random posterior distribution  $\mathbb{P}(H_i \in |\sqrt{r_*} H_i + \xi_i)$  associated with the effective scalar observation  $\sqrt{r_*} H_i + \xi_i$  with prior  $P_X$ .

## B. Saddle point equations and mutual information

The saddle point equations are obtained by equating the gradient of  $\{\cdot\}$  in Eq. (26) to zero,

$$q = \mathbb{E}_{\xi, H} [H \langle \eta \rangle_r], \quad r = \frac{\lambda J(\sigma)}{1 + \lambda \alpha (1 - J(\sigma) q)}, \quad (30)$$

where

$$\langle \cdot \rangle_r = \langle \cdot \rangle_r(\xi, H) := \frac{1}{\mathcal{Z}_r} \int dP_X(\eta) e^{(\xi \sqrt{r} + H r) \eta - \frac{1}{2} r \eta^2}(\cdot),$$

and  $\mathcal{Z}_r(\xi, H)$  is the normalization. The solutions  $(r_*, q_*)$  depend on  $\alpha, \lambda$ , and  $J(\sigma)$ .

The main question now is how to set the effective bulk SNR  $\sigma$  in order to capture the effect of the original bulk observations Eqs. (21) when  $\zeta = \lambda$ . We argue that it can be determined with a consistency argument.  $\sigma$  parametrizes the marginals of the bulk variables  $\bar{\mathbf{x}} = (x_{i\mu})_{i \leq N-1, \mu \leq M}$  through Eq. (25). In the original problem or, equivalently, the model Eqs. (20) and (21) with  $\zeta = \lambda$ , rows of  $\mathbf{x}$  are completely exchangeable under the law  $\mathbb{E}[\langle \cdot \rangle]$ . Therefore, the random cavity marginals Eq. (29) corresponding to SNR  $\lambda$  should have the same law as the bulk marginals. This is

achieved by matching the SNRs in the effective scalar observation channels for the cavity coordinates,  $\sqrt{r_*(\lambda, \sigma)} H_i + \xi_i$ , and bulk coordinates,  $\sqrt{\sigma} X_{i\mu} + Z'_{i\mu}$ . This yields the third missing consistency equation needed for  $\sigma_*(\lambda)$  in addition to the two we already have for  $(q, r)$ ,

$$r_* = \frac{\lambda J(\sigma_*)}{1 + \lambda \alpha (1 - J(\sigma_*) q_*)} = \sigma_*. \quad (31)$$

By Eq. (27),  $J$  is the overlap that any bulk row in  $\bar{\mathbf{X}}$  has with the corresponding ground truth row in  $\bar{\mathbf{X}}$ . As intuition may suggest, matching the laws of the cavity and bulk marginals also equates the two overlaps  $J$  and  $q_*$ . To see this, it is sufficient to realize that the function  $F(r) = \mathbb{E} H \langle \eta \rangle_r$  is monotonic in  $r \geq 0$  and thus invertible. Then, by a simple application of it to both sides of Eq. (31), we readily get  $q_* = J(\sigma_*)$  as predicted. The two conditions  $\sigma_* = r_*$  and  $J(\sigma_*) = q_*$  are thus equivalent.

Using this knowledge, we recast Eq. (30) as

$$q = \mathbb{E}_{\xi, H} [H \langle \eta \rangle_r], \quad r = \frac{\lambda q}{1 + \lambda \alpha (1 - q^2)}. \quad (32)$$

These are solved via fixed-point iterations. Their solutions shall be denoted by  $q_*(\alpha, \lambda)$  and  $r_*(\alpha, \lambda)$ . Using the relation Eq. (28) and that the solution is extremizing  $\phi_{\text{eff}}^{\text{RLE}}$ , we finally get an *Ansatz* for the MMSE by using the I-MMSE relation Eq. (23) but for the model with the mean-field bulk *Ansatz* [here approximate equalities are up to vanishing corrections  $O(1/N)$ ],

$$\begin{aligned} \text{MMSE} \approx \text{YMMSE}(\lambda, \lambda) &\approx 2 \frac{\partial}{\partial \lambda} \left( \frac{\lambda}{2} - \phi_{\text{eff}}^{\text{RLE}}(\alpha, \lambda, \sigma_*) \right) \\ &= \frac{1 - q_*(\alpha, \lambda)^2}{1 + \lambda \alpha (1 - q_*(\alpha, \lambda)^2)}. \end{aligned} \quad (33)$$

Note that  $q_* = r_* = \sigma_* = 0$  is always a solution for any  $(\alpha, \lambda)$ . We call it the “paramagnetic solution” as it corresponds to no alignment between posterior sample and ground truth. However, the MMSE for this solution,

$$\text{MMSE}_{\text{para}} = \frac{1}{1 + \alpha \lambda},$$

is a nontrivial decreasing function of the SNR. This is because denoising is still possible, for any  $\lambda > 0$ , when the matrix  $\mathbf{X}$  has no macroscopic overlap with a posterior sample [ $q_* = 0$  which is both the row or column overlap under *Ansatz* Eq. (25)]. We have checked that, apart from very low SNR, this solution is unstable in the sense that any small positive initialization for the iterative solution of the saddle point equations will yield another solution.

Given the MMSE, one could be tempted to use the I-MMSE formula Eq. (5) to integrate it with respect to the

SNR  $\lambda$  in order to obtain the MI of matrix denoising. However, there could be multiple solutions  $q_*(\alpha, \lambda)$  associated with given values of  $\lambda$  and  $\alpha$ , which impeaches this naive reasoning. A more careful approach is required. Given Eqs. (32), we reconstruct in Appendix D the “mutual information potential” that generates them, and we shall always choose solutions that minimize it. Indeed, there is a unique potential function  $\iota(r, q; \alpha, \lambda)$  that yields the system Eq. (32) when extremized over  $(r, q)$  and that at the same time satisfies the I-MMSE relation Eq. (5) with the previously found solution Eq. (33) as MMSE, namely, such that

$$\frac{d}{d\lambda} \iota(r_*, q_*; \alpha, \lambda) = \frac{1}{4} \frac{1 - q_*(\alpha, \lambda)^2}{1 + \lambda\alpha(1 - q_*(\alpha, \lambda)^2)}.$$

This potential is

$$\begin{aligned} \iota(r, q; \alpha, \lambda) := & \frac{rq}{2} + \frac{1}{4\alpha} \ln(1 + \lambda\alpha(1 - q^2)) \\ & - \mathbb{E}_{\xi \sim \mathcal{N}(0,1), H \sim P_X} \ln \int dP_X(\eta) e^{(\xi\sqrt{r} + Hr)\eta - \frac{1}{2}r\eta^2}. \end{aligned} \quad (34)$$

The mean-field prediction for the MI of the original problem is then

$$\iota(\alpha, \lambda) := \text{extr}_- \{ \iota(r, q; \alpha, \lambda) \}, \quad (35)$$

where extremization is with respect to  $(r, q)$ , and  $\text{extr}_-$  selects the solution of the saddle point Eqs. (32) minimizing the potential. Let us stress again that this formula will not be our final prediction for the MI in the whole phase diagram  $(\alpha, \lambda)$ , but just for a subregion of it as explained in the next section.

This formula recovers the known replica symmetric potential for the rank-one case [16,30] (which is always correctly described by mean-field theory) when  $\alpha \rightarrow 0$ ,

$$\begin{aligned} \iota(r, q; 0, \lambda) = & \frac{rq}{2} + \frac{\lambda}{4} (1 - q^2) \\ & - \mathbb{E}_{\xi, H} \ln \int dP_X(\eta) e^{(\xi\sqrt{r} + Hr)\eta - \frac{1}{2}r\eta^2}. \end{aligned}$$

We observe that, in the converse limit  $\alpha \rightarrow \infty$ , the formula trivializes and is optimized at  $r = q = 0$ . This is due to our choice of normalizations in channel Eq. (14). We refer to Refs. [85,91] for a different scaling where this limit is meaningful and thoroughly investigated. In addition, Appendix D 2 shows that  $\iota(\alpha, \lambda \rightarrow \infty)$  approaches the entropy of the prior  $\mathcal{H}(X)$  as it should; it is also apparent from Fig. 9 appearing later on.

From the more general derivation in Appendix C, it is clear that our multiscale mean-field theory can handle generic observation channels as long as the observations

remain independent conditionally on the signal, namely,  $Y_{ij} \sim P_{\text{out}}(\cdot | \mathbf{X}_i \cdot \mathbf{X}_j / \sqrt{N})$  conditionally independent. The fixed-point equation  $q = \mathbb{E}H\langle \eta \rangle_r$  is unchanged, whereas  $r = g_{\text{out}}(q)$  with a channel-dependent  $g_{\text{out}}$ . The consistency  $r_* = \sigma_*$  still holds, as well as  $J(\sigma_*) = q_*$ .

### 1. Relation to the replica method with Gaussian Ansatz of Sakata and Kabashima

The same formula for the MI Eq. (35) can be produced using the earlier approach of Ref. [99], also used in Refs. [100–102] in a Bayesian context. We have adapted their method to our model in Appendix F for completeness, which is close to what is done in Ref. [102]. This alternative derivation does not go through a cavity computation but instead uses a nonstandard replica method at the scale  $O(MN)$ . From our understanding, the first Gaussian Ansatz they use is to approximate  $(z_{ij}^a = \mathbf{x}_i^a \cdot \mathbf{x}_j^a / \sqrt{N})_{i,j=1,\dots,N, a=0,\dots,n}$ , where  $a$  is the replica index, as a jointly Gaussian family. The second Ansatz is on their covariance. Somewhat counterintuitively, they consider  $(z_{ij}^a)$  independent (up to symmetry,  $z_{ij}^a = z_{ji}^a$ ). Overall, this is equivalent to consider  $(z_{ij}^a)_{i,j}$  as a GOE matrix, with a semicircular spectral density. This is hard to justify *a priori*, given that this matrix is of Wishart type, whose spectrum follows the Marchenko-Pastur law. This assumption may, however, be more justifiable as  $\alpha \approx M/N$  gets large because, in this case, a Wishart matrix resembles more a GOE one with a mean. As discussed in Ref. [91], this approximation is supposedly more accurate in this case (as confirmed by our numerics) and even possibly exact in the  $\alpha \rightarrow \infty$  limit. Interestingly, Eq. (34) also yields the rank-one formula when  $\alpha \rightarrow 0$ , which signals that the approach is actually taking into account both  $\alpha$  limits.

As mentioned in Sec. II B, Eqs. (32) were deemed incorrect for a long time [91,103]. In fact, in the past they were tested against the exact MI for Gaussian prior only [91]. However, this is precisely the case in which they are not supposed to hold for all  $(\alpha, \lambda)$ , because the denoising phase never ends for the Gaussian prior and so the HCIZ integral is the right tool. With our contribution, we redeem the replica equations found more than a decade ago in Refs. [99–102] using an independent, original, mean-field approach, by showing that they actually yield excellent results but in the factorization phase only (and good ones in the denoising phase for large  $\alpha$ ), a phase which exists only for non-Gaussian priors. Despite leading to the same equations, the multiscale theory we have presented is based on different, interpretable hypotheses. The fact that the earlier replica method and our cavity-based approach yield the same equations is not evident *a priori* given the nonstandard Gaussian Ansatz used in the former. This is an indication of their validity when the concentration-of-measure effects needed to reduce the high-dimensional integrals entering the free entropy to saddle point equations

over order parameters take place. This, we claim, happens only in the factorization phase, which was not understood until now.

## V. COMPLETE ANSATZ TAKING INTO ACCOUNT THE PRIOR UNIVERSALITY IN THE DENOISING PHASE

### A. Complete Ansatz for the whole phase diagram

We have confirmed numerically that the strong prior universality conjecture of Ref. [85] holds but only before the denoising-factorization transition. This should be incorporated to correct the inaccuracies of the mean-field approximations in this region. Indeed, the MMSE predicted by the mean-field theory does not match the known exact one for Gaussian prior, nor for discrete priors in the denoising phase. This is somehow expected. The reason is that our *Ansatz* for the bulk measure, corresponding to taking the posterior associated with the channel  $\mathbf{Y}'_{\text{eff}} = \sqrt{\sigma}\mathbf{X} + \mathbf{Z}'_{\text{eff}}$ , breaks an “effective rotational symmetry” that seems to hold in the denoising phase as we argue in Sec. VI. The rotational symmetry breaking comes from the fact that this channel corresponds to a direct observation of  $\mathbf{X}$  instead of  $\mathbf{X}\mathbf{X}^\top$  which “mixes” the entries. Therefore, there is no reason for the theory to be correct in the denoising phase. Yet, in the factorization phase, the mean-field prediction for the MMSE is in agreement with the exponential decay in Fig. 6, due to the strong alignment between posterior “patterns”  $(\mathbf{X}_\mu)_{\mu \leq M}$  and the planted ones  $(\mathbf{X}_\mu)_{\mu \leq M}$  (see Sec. VI). This implies that in this phase the behavior of a row  $\boldsymbol{\eta}$  is indeed dominated by its direct interaction with the associated planted row  $\mathbf{H}$  (mediated by  $H_{N,M}^{\text{row}}$ ), while the effect of the bulk simplifies greatly and is captured by the mean-field *Ansatz* Eq. (25). We conclude that the denoising phase is a “matrix model regime” which needs to be analyzed by its own set of techniques (i.e., proofs of universality and the HCIZ integral). In contrast, the factorization phase is a “mean-field regime” amenable to statistical physics methods. We are thus confident that the mean-field solution is accurate in this region, as backed up by the numerics and its robustness to attempts on improving it, see Appendix G, as well as the matching with the replica approach of Sakata and Kabashima [99].

From Fig. 5 and the related discussion, we deduce that there are only three possible thermodynamic states: One corresponds to the universal RIE branch, another to the nonuniversal factorization branch reached from informative initialization, and finally the dynamical glass state. Given our numerical evidence as well as the work of Semerjian [85], we select the universal RIE branch as the equilibrium state for low SNRs, which corresponds to the Gaussian MI Eq. (10). After the transition yet to be identified, we instead select the factorization state described by the mean-field formula Eq. (35).

The question now is how to locate the denoising-factorization phase transition, assuming there is only one transition. The single transition scenario is backed up by our numerical experiments. We know three key properties of the asymptotic MI that can be used to locate the transition between the universal and mean-field branches: (i) The MI is continuous in  $\lambda$ ; (ii) it is also concave in  $\lambda$  (its second derivative is the first derivative of the MMSE, which must decrease when  $\lambda$  increases); and (iii) for a discrete prior, we know from Eq. (13) that the MI saturates to  $\mathcal{H}(X)$  when  $\lambda \rightarrow +\infty$ . Therefore, the only possibility to link the Gaussian solution at low SNR to the mean-field one at higher values [which does approach  $\mathcal{H}(X)$  exponentially fast in  $\lambda$  by Appendix D 2] is to locate the transition as the point where they cross, namely,

$$\lambda_c(\alpha) = \sup\{\lambda \geq 0: \iota^G(\alpha, \lambda) \leq \iota(\alpha, \lambda)\}. \quad (36)$$

As we show in the next section, this is in good agreement with our finite-size experiments once we properly take into account finite-size effects. Consistently, our final *Ansatz*  $\iota^*$  for the MI is then

$$\iota^*(\alpha, \lambda) = \min\{\iota^G(\alpha, \lambda), \iota(\alpha, \lambda)\}. \quad (37)$$

The choice of selecting the Gaussian MI in the denoising phase for an informative prior like Rademacher is also well motivated by the following argument. An informative (structured) prior carries more information than the Gaussian one, and thus the associated MMSE cannot be larger than in the Gaussian case given the same data. Therefore, for such a prior, the mean-field theory cannot be correct when it predicts an MMSE (33) larger than the exact one for the Gaussian prior Eq. (9). The Gaussian MMSE is thus at least a better approximation to the true one in such a region. Figure 8 shows that there is indeed a low-SNR region where this happens and therefore the mean-field solution is incorrect. We claim that in this region the Gaussian curve is not just a better approximation to the truth, but is actually exact, in agreement with [85] and our numerics. On the other hand, as illustrated by Fig. 6 and further discussed in the following section, the choice of the mean-field solution beyond a point is validated by the excellent agreement with our numerical experiments. Then, from the single transition hypothesis and the aforementioned properties (i)–(iii) of the MI, the transition Eq. (36) emerges.

The form of the final *Ansatz*  $\iota^*$  shares features of classical limit formulas in spin glasses such as the Parisi formula [141,142]. In the Parisi formula, the limit free entropy  $F^*(\beta)$  at inverse temperature  $\beta$  is also given by a minimization over different candidate “branches,”

$$F^*(\beta) = \min\{F_{\text{RS}}(\beta), F_{\text{1RSB}}(\beta), F_{\text{2RSB}}(\beta), \dots\},$$

where  $F_{\text{RS}}(\beta)$ ,  $F_{\text{IRSB}}(\beta)$ ,  $\dots$ , etc., are the respective replica symmetric, 1-replica symmetry-breaking step,  $\dots$ , approximations of the free entropy. However, there are important differences too. First, our model is replica symmetric as it lives on its Nishimori line. However, more importantly, in contrast to our setting, a key difference in spin glasses and usual inference problems is that, in all their phases (i.e., regimes of symmetry breaking), these models remain mean-field in nature, and thus the different branches of symmetry breaking can be described by one formalism yielding an overarching formula. This is not the case for extensive-rank matrix denoising: the current techniques to compute the MI in the denoising and factorization phases are very different, which reflects their fundamentally distinct nature. We believe it is unlikely there exists a mean-field theory describing both concurrently, unlike what happens in the Parisi formula. This is the main reason why we had to employ matrix model techniques in the denoising phase.

### B. Theoretical predictions and consistency with the numerical experiments

We plot in Fig. 8 the mean-field solution for the MMSE Eq. (33) obtained by iterating Eq. (32) (dark blue curves) and the associated asymptotic MI Eq. (35) in Fig. 9. When solving Eq. (32) by fixed-point iteration, we call the initialization  $q^{t=0} = 1$  “informative”. The “uninformative” initialization is instead  $q^{t=0} = \varepsilon \ll 1$  but nonzero. From these two obtained solutions, the equilibrium  $(q_*, r_*)$  is the one that minimizes Eq. (34). The other solution (we always observed at most two) is called “metastable.” It matches the equilibrium solution if unique, as for  $\alpha = 0.5$  in Fig. 9, while they can differ for  $\alpha = 0.7$ .

Recall, however, that before the transition Eq. (36) we claim that the correct branch is instead the universal one. For all plots displayed, the complete *Ansatz* taking into account universality at low SNR  $t^*(\alpha, \lambda)$  Eq. (37) (cyan curves for both MI and MMSE) correctly describes the key features of the problem evidenced by our numerical experiments. In particular, we stress the very good agreement between 4 times the  $\lambda$  derivative of  $t^*$  and the finite-size MMSEs in the factorization phase, see Figs. 6 and 8, but also in the denoising phase, see Figs. 3 and 8. Hereby, we describe our theoretical findings focusing on the complete *Ansatz*  $t^*$  if not specified otherwise, comparing them with those of Sec. III.

- (i) Agreement with the mutual information and transition upper bounds.  $t^*$  remains upper bounded by the entropy  $\mathcal{H}(X)$  of the prior as it should, see Appendix D 2 and Fig. 9. This implies that the theoretical transition Eq. (36) occurs before the rigorous upper bound Eq. (11), which becomes tighter as  $\alpha$  increases, see Fig. 10. This nontrivial consistency check entails that, for high SNR, the theoretical MMSE decays exponentially with it, as

supported by the numerical experiments, see Figs. 6 and 8. This exponential decay in the factorization phase, which is asymptotically in  $\lambda$  the same as for the rank-one case with Rademacher prior, hints at the fact that in this phase the posterior patterns  $(\mathbf{x}_\mu)$  are weakly coupled, and the problem resembles (but is not equivalent to)  $M$ -independent rank-one problems with properly tuned SNRs.

- (ii) Universality at low SNR. Despite the mean-field solution  $t$  being only approximate in the denoising phase, at low SNR it captures universality: the mean-field curves  $t(\alpha, \lambda)$  for Gaussian and Rademacher priors do match, see Fig. 9; it is also evident at the level of the MMSE in Fig. 8. This is in agreement with experiments shown in Figs. 2 and 3. When  $\alpha$  increases, this matching lasts until close to the crossing of  $t$  and the Rademacher entropy: This signature of universality becomes more pronounced. It shows that  $t$  becomes more accurate and consistent at larger  $\alpha$  (but it also is exact at  $\alpha = 0$ ). Obviously,  $t^*$  is universal before the transition by construction.
- (iii) Denoising-factorization transition.  $t^*$  encodes the denoising-factorization transition observed in our numerical experiments. Indeed, for all  $\alpha$ 's tested we see a change of behavior in the  $\lambda$  derivative of  $t^*(\alpha, \lambda)$  separating the regime where the MI is prior independent from the factorization phase where the MI for discrete prior approaches the entropy of the prior exponentially fast in  $\lambda$ . The MI for the Gaussian prior instead continues to increase smoothly. The algebraic versus exponential decay of the MMSE with the SNR observed numerically is quantitatively caught by the theory, see Fig. 8.

For what concerns the transition location, it seems from Fig. 8 that there is a non-negligible gap with what  $t^*(\alpha, \lambda)$  predicts. We argue here that this is due to the finite-size of the simulated system and can be accounted for. In fact, our arguments in Sec. III make it clear that for finite sizes the MI for Rademacher prior must not saturate to  $\ln 2$  due to the degeneracy induced by column permutations and signs, which explains precisely the gap to  $\ln 2$  in the left panel of Fig. 2. Since the empirical MI curves closely follow  $t^G(\alpha, \lambda)$  at low SNR, they must start saturating earlier, namely, at a smaller  $\lambda$  than the asymptotic one  $\lambda_c$ , and toward a lower value than  $\ln 2$ . For instance, for  $\alpha = 0.5$  and  $N = 40$  saturation occurs at  $\text{MI} \approx 0.6228$ , which is also reported as the lowest dotted horizontal line in the left panel of Fig. 9. The curve  $t^G$  crosses that value at  $\lambda \approx 7.2$ , which gives us a rough upper bound on the transition location for that size, in good agreement with what is observed in Figs. 2 and 3 (even if formally we cannot talk about phase transition at finite size). However, we see that the change in the behavior of the derivative of the finite-size MI is not sudden and starts even before 7.2.

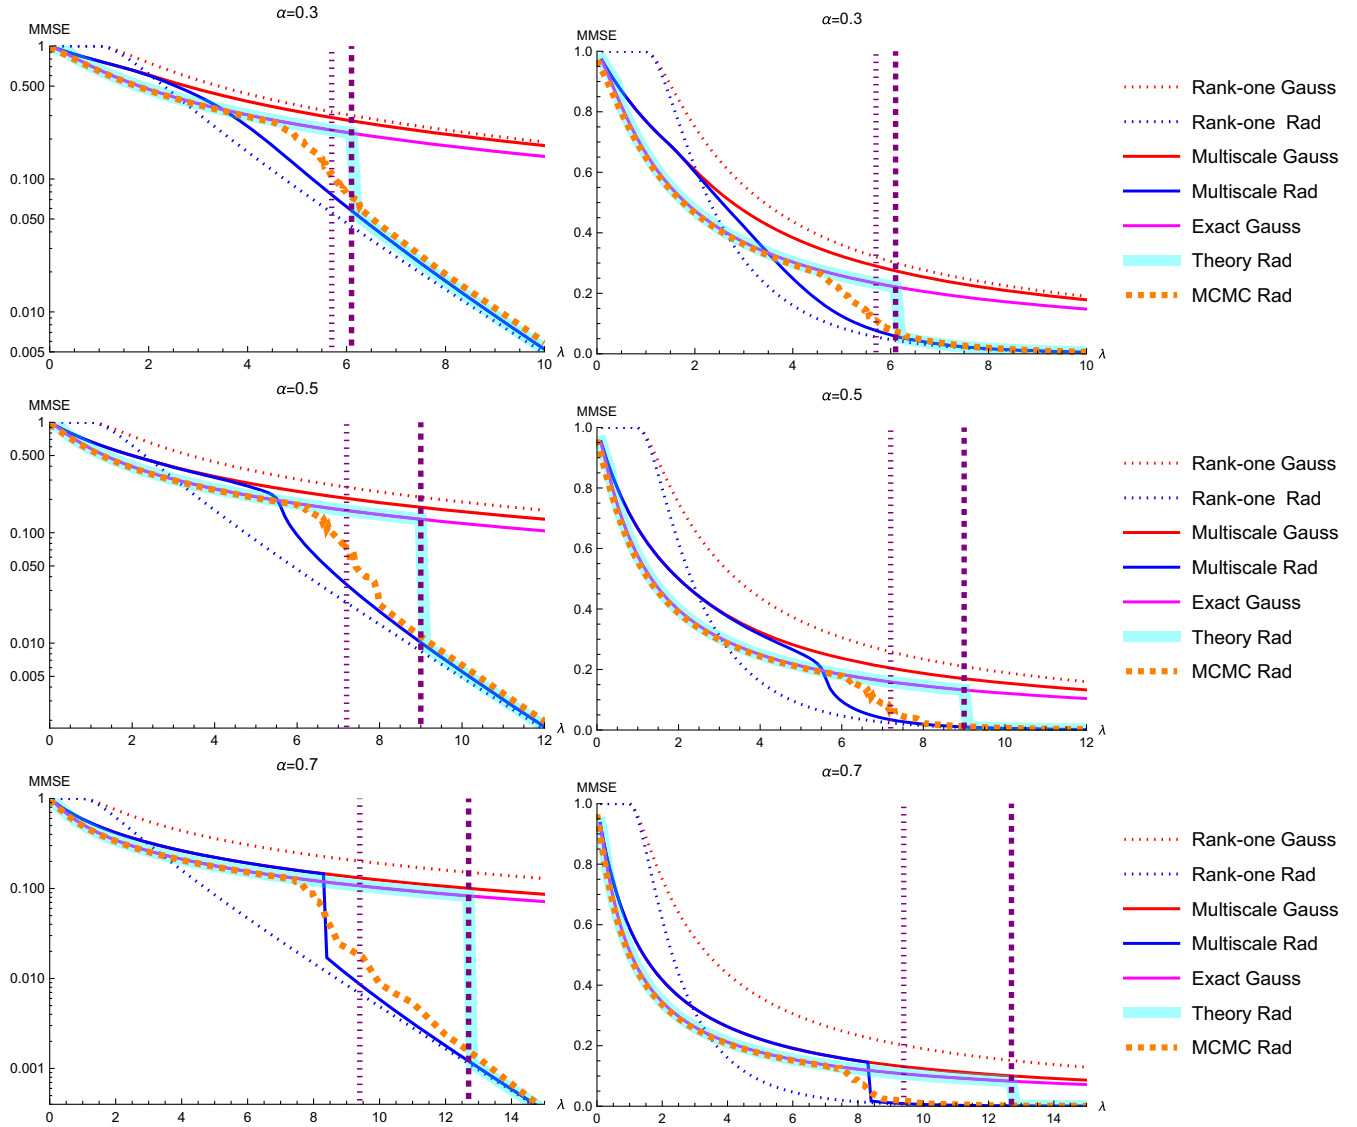


FIG. 8. Thermodynamic limit of the MMSE as function of the SNR for  $\alpha \in \{0.3, 0.5, 0.7\}$  with Gaussian and Rademacher priors. Logarithmic scale on the left, linear scale on the right. We plot (i) the rank-one MMSEs, (ii) the “multiscale” curves correspond to the mean-field Eq. (33) with the equilibrium solution selected, (iii) the “exact Gauss” curve is the true Gaussian MMSE obtained via Eq. (9), which also corresponds to the performance of the RIE for any prior with zero mean and unit variance, (iv) the cyan “theory Rad” is our theoretical prediction for the MMSE with Rademacher prior corresponding to 4 times the derivative of  $\iota^*(\alpha, \lambda)$  Eq. (37) with respect to  $\lambda$ , and (v) “MCMC Rad” are the Monte Carlo results with Rademacher prior (see Tables I and II for details). Finally, (vi) the rightmost vertical line pinpoints the location of the infinite  $N$  transition  $\lambda_c$  Eq. (36), while the one on the left is its approximate location once taking into account finite-size effects; see the caption of Fig. 9 and the text for an explanation.

This is consistent with Fig. 8 for  $\alpha = 0.5$ , where the Markov chain Monte Carlo (MCMC) curve detaches from the Gaussian MI around  $\lambda \approx 6.5$ . Nevertheless, observe that the finite-size-corrected approximation to the transition at 7.2 falls nicely in the middle of the “crossover region” between the two phases for the MMSE. The same correcting approach works for the other values of  $\alpha$ .

Let us explain in more detail the way we computed corrections to the transition location. The procedure starts by computing the finite-size entropy

$\mathcal{H}(\mathbf{X}\mathbf{X}^\top)/(NM)$  using Eq. (12). Then we find the SNR where  $\iota^G$  attains this value: This gives an upper bound on the finite-size “transition,” which we expect to be close to it given the numerical and theoretical evidence. This implies that the larger  $\alpha$ , the smaller the slope of  $\iota^G(\alpha, \lambda)$  in  $\lambda$ , and thus the larger the gap between the asymptotic transition point  $\lambda_c$  and its approximate location for finite size, see Fig. 9.

On top of this gap increasing with  $\alpha$ , the crossover region for the transition should also get wider. To see it, we notice that the finite-size corrections to the MI,




FIG. 9. Thermodynamic limit of the mutual information as function of the SNR for  $\alpha \in \{0.5, 0.7\}$  with Gaussian and Rademacher priors. We plot (i) the rank-one MI, (ii) the multiscale curves referring to  $\iota(\alpha, \lambda)$  Eq. (35) with the equilibrium solution selected, (iii) the “multiscale met.” curve corresponds to the metastable solution (which exists only for Rademacher prior), (iv) the exact Gauss curve is the exact MI for Gaussian prior  $\iota^G(\alpha, \lambda)$  Eq. (10), (v) the theory Rad is the complete Ansatz  $\iota^*(\alpha, \lambda)$  Eq. (37) for Rademacher prior once taking into account universality in the denoising phase, (vi) the entropy density  $\mathcal{H}(\mathbf{X}\mathbf{X}^T)/(NM)$  for  $N \rightarrow +\infty$ , equal to  $\ln 2$ , and its finite-size approximation Eq. (12) given by  $\ln 2 - \ln(2^M M!)/(MN)$  with  $N = 40, M = \alpha N$ . Finally, (vii) the rightmost vertical line pinpoints the crossing of the exact Gaussian MI with the multiscale mean-field prediction, i.e., the infinite-size transition  $\lambda_c$  Eq. (36), while the leftmost vertical line shows the crossing of the Gaussian MI with  $\ln 2 - \ln(2^M M!)/(MN)$  which approximates the location of the denoising-factorization transition taking into account finite-size effects.

estimated through Eq. (13) as  $\ln(2^M M!)/(MN)$  ( $\approx 5 \times 10^{-2}$  in all our experiments), are large given the size of the system ( $MN \approx 1000$  in all our experiments). As comparison, in the rank-one spiked Wigner model, finite-size corrections to the MI instead vanish as  $1/N$ , which is faster relative to the system size. For size 1000, this yields a correction roughly 2 orders of magnitude smaller than here. Notice also that finite-size corrections increase in  $\alpha$ , since they are directly linked to the number of column permutations. In summary, thanks to the aforementioned correction procedure, the predicted finite-size transition does fall in the middle of the crossover region for the MMSE, which gets wider with larger  $\alpha$ . Out of these crossover regions, the theory proves to be extremely accurate. Despite the exponential decays and the modest sizes that can be simulated, our predictions for the MMSEs deep in the factorization phase remain precise even at a  $y$  scale of  $10^{-3}$ .

- (iv) Spinodal transition and statistical-to-computational gap. According to our theory  $\iota^*$ , the thermodynamic denoising-factorization phase transition at  $\lambda_c(\alpha)$  is of first order, i.e., with a discontinuity in its  $\lambda$  derivative, which is the MMSE over 4 by Eq. (5), for any  $\alpha > 0$ . We have already discussed based on Fig. 5 that, at low SNR, the only thermodynamic state present is the universal RIE state. This implies that, before the first order transition, there must be a spinodal transition linked to the appearance of the informative factorization state in addition to the RIE state, a point around which the RIE state seems to shatter into a dynamical glassy state observed in Fig. 5. A spinodal transition followed by the coexistence of informative and glassy states is a typical mechanism at the root of

algorithmic hardness in high-dimensional inference [123,143]. This connection should, however, be considered with some caution, as it has been noticed in various settings that the onset of glassiness does not always prevent inference algorithms, including sampling based, to perform well [123,144,145].  $\iota^*$  is a pure equilibrium theory and is therefore unable to detect the location of the spinodal transition. One would need to modify it in order to take into account out-of-equilibrium glassy effects linked to replica symmetry breaking [21,123], which is out of the scope of the present paper. Given that we are in the Bayes-optimal setting, the thermodynamic equilibrium should anyway be described by a replica symmetric theory [18,146,147].




FIG. 10. Thermodynamic denoising-factorization phase transition  $\lambda_c(\alpha)$  given by Eq. (36) and its upper bound given by the rhs of Eq. (11), for Rademacher prior.

It is worth noting that, even though the mean-field MI  $t(\alpha, \lambda)$  is inexact in the denoising phase, the potential Eq. (34) captures the presence of a spinodal transition for  $\alpha \geq 0.7$ . Indeed, when  $\alpha = 0.7$  the mean-field theory exhibits a metastable branch around  $\lambda \approx 7$ ; see the discontinuity in the green curve in Fig. 9. It detaches from the equilibrium mean-field curve through a spinodal transition where the informative fixed point appears. We thus interpret the mean-field theory as an approximation in the denoising phase (which becomes exact after the transition, while  $t^*$  is deemed to be exact everywhere), in the sense of the critical phenomenology it implies that cannot be probed from the equilibrium theory  $t^*$  alone.

## VI. PROBING THE PHASES THROUGH THE OVERLAP MATRIX AND FACTORIZATION

So far, we have focused on the denoising task of recovering the square signal  $\mathbf{X}\mathbf{X}^\top$ . We have divided the phase diagram in two thermodynamic phases: the universal denoising phase where the RIE is optimal and the factorization phase where it is not, because exploiting the prior and factorized structure of the signal beyond its asymptotic spectrum  $\rho_{\mathbf{X}\mathbf{X}^\top}$  is needed there to better denoise. It is thus natural to wonder if optimal denoising is related to the factorization task, i.e., estimating the single factor  $\mathbf{X}$  nontrivially. Nontrivial estimation means finding an estimator  $\hat{\mathbf{X}} = \hat{\mathbf{X}}(\mathbf{Y}) \in \mathbb{R}^{N \times M}$  verifying

$$\lim_{N \rightarrow +\infty} \frac{1}{NM} \max_{\mathbf{\Pi} \in \bar{\pi}_M} \text{Tr}(\mathbf{X}^\top \hat{\mathbf{X}} \mathbf{\Pi}) > 0, \quad (38)$$

where  $\bar{\pi}_M$  is the set of  $M \times M$  signed permutation matrices. It is not clear that it is possible given the numerous symmetries and invariances of the model when  $M \rightarrow +\infty$  with  $N$ . When  $M$  is independent of  $N$ , if denoising is possible, then factorization is, because the number of symmetries to resolve in order to infer  $\mathbf{X}$  given an estimator  $\widehat{\mathbf{X}\mathbf{X}^\top}$  of  $\mathbf{X}\mathbf{X}^\top$  remains generally finite; e.g., for  $M = 1$ , the eigenvector of  $\widehat{\mathbf{X}\mathbf{X}^\top}$  with largest eigenvalue estimates  $\mathbf{X}$  up to a global sign.

We will explore here this question by getting a more detailed description of the phases through the analysis of a richer order parameter than the MMSE, namely, the overlap matrix,

$$\mathbf{Q} := \frac{\mathbf{X}^\top \mathbf{x}}{N} = \left( \frac{\mathbf{X}_\mu^\top \mathbf{x}_\nu}{N} \right)_{\mu, \nu \leq M} \in \mathbb{R}^{M \times M},$$

where  $\mathbf{x} \sim P(\cdot | \mathbf{Y})$  is a posterior sample, and  $(\mathbf{x}_\mu)$  are the columns, or patterns, of  $\mathbf{x}$ . The MMSE can be written directly in terms of the overlap,

$$\text{MMSE} = \frac{1}{M} \mathbb{E} \left\| \frac{\mathbf{X}\mathbf{X}^\top}{N} \right\|^2 - \frac{1}{M} \mathbb{E} \langle \|\mathbf{Q}\|^2 \rangle, \quad (39)$$

which follows from the Nishimori identity. From Eq. (38), a posterior sample  $\mathbf{X}$  is a nontrivial estimator for the factorization task if  $(1/M)\text{Tr}(\mathbf{Q}\mathbf{\Pi}^*) > o_N(1)$  for an  $\mathbf{x}$ -dependent permutation  $\mathbf{\Pi}^*$ . The overlap can also be defined in terms of two conditionally i.i.d. samples from the posterior rather than using the ground truth:  $\mathbf{x}^{(1)\top} \mathbf{x}^{(2)}/N$ . By Bayes optimality and the Nishimori identities, the conclusions of this section based on  $\mathbf{Q}$  would still hold.

### A. Behavior of the overlap in the two phases: Can $\mathbf{X}$ be inferred in addition to $\mathbf{X}\mathbf{X}^\top$ ?

It is insightful to distinguish between diagonal and off-diagonal contributions of the overlap to the MMSE. Note that, from the independence of the data on the permutation of the columns of  $\mathbf{X}$  (and their signs for symmetric priors), the induced invariance of the posterior implies that the columns of posterior samples  $\mathbf{x}$  obtained from the Monte Carlo sequence must be sorted appropriately to match the ones of the ground truth in order to access a meaningful overlap. In other words, we need an efficient way to approximate the solution of the optimization problem Eq. (38). Numerically, this is achieved via iteration with respect to the columns of  $\mathbf{X}$ . In the first iteration, one finds the column of  $\mathbf{X}$  showing the largest absolute inner product with  $\mathbf{X}_1$ , labeling its index as  $\pi(1)$ . For iterations  $\mu = 2, \dots, M$ , choose  $\pi(\mu)$  as  $\pi(\mu) = \text{argmax}_{\nu \in S_\mu} |\mathbf{X}_\mu^\top \mathbf{x}_\nu|$ , where  $S_\mu$  is the set of indices from 1 to  $M$  excluding  $\pi(1), \dots, \pi(\mu - 1)$ , see Algorithm 1 below.

By executing this procedure for *each* posterior sample, we can distinguish between diagonal and off-diagonal contributions to the MMSE, with the diagonal one representing the error made from misalignment between matched patterns and the off-diagonals representing the contributions from alignment between unmatched patterns. For all the rest of this section, every time a posterior sample  $\mathbf{x}$  appears, it is assumed that the signs and ordering of its columns has already been resolved according to the above greedy procedure as  $\mathbf{x} \leftarrow (\pi \circ \sigma)(\mathbf{x})$ , and the Monte Carlo approximations of expectations  $\langle \cdot \rangle$  are taking into account this per-sample operation.

Figure 11 shows the diagonal and off-diagonal overlap contributions to the MMSE Eq. (39) for finite sizes

Algorithm 1. Greedy column and sign alignment.

---

**Input:** Reference matrix  $\mathbf{X}$ , ordered matrix  $\mathbf{x}$   
**Output:** Permutation  $\pi$ , signs  $\sigma$   
**Initialize:**  $S \leftarrow \{1, 2, \dots, M\}$  (Set of available indices)  
**Initialize:**  $\pi$  and  $\sigma$  as empty arrays of length  $M$

- 1: **for**  $\mu = 1$  to  $M$  **do**
- 2:      $\pi(\mu) \leftarrow \text{argmax}_{\nu \in S} |\mathbf{X}_\mu^\top \mathbf{x}_\nu|$
- 3:      $\sigma(\mu) \leftarrow \text{sign}(\mathbf{X}_\mu^\top \mathbf{x}_{\pi(\mu)})$
- 4:      $S \leftarrow S \setminus \{\pi(\mu)\}$
- 5: **end for**

---

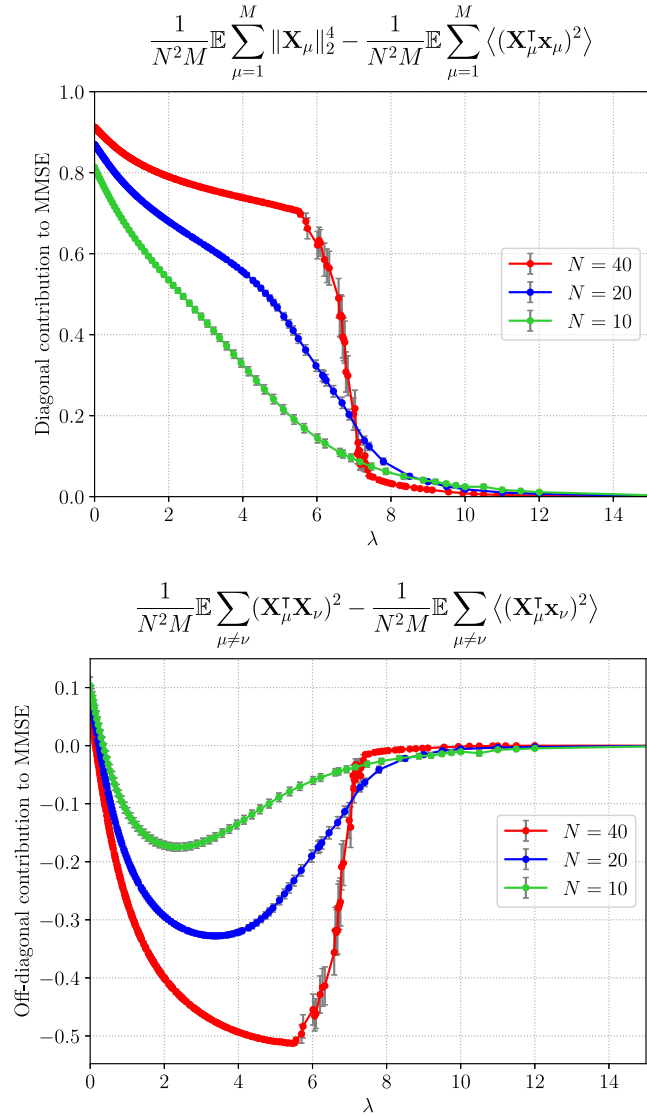


FIG. 11. Monte Carlo results for the diagonal (upper) and off-diagonal (lower) contributions to the MMSE Eq. (39) after reordering and sign selection of the columns of each posterior sample according to Algorithm VI A, with Rademacher prior, for various  $N$  with  $\alpha = 0.5$ . Error bars are standard errors of the mean.

computed by Monte Carlo sampling. First, looking at the diagonal contribution for SNRs in the denoising phase ( $\lambda \lesssim 6$  for this size), the curve seems to approach  $1 = \mathbb{E} \sum_{\mu=1}^M \|\mathbf{X}_\mu\|_2^4 / (N^2 M)$  as  $N$  increases, indicating that no strong “pattern matching” is happening. Looking at the bottom part of Fig. 11, we observe that in this phase it is instead the off diagonal of the overlap that contributes the most in decreasing the MMSE, e.g., in Fig. 3. The strong negative value of  $(\mathbb{E} \sum_{\mu \neq \nu} (\mathbf{X}_\mu^\top \mathbf{X}_\nu)^2 - \mathbb{E} \sum_{\mu \neq \nu} \langle (\mathbf{X}_\mu^\top \mathbf{x}_\nu)^2 \rangle) / (N^2 M)$  indicates that  $(Q_{\mu\nu})_{\mu \neq \nu}$  fluctuate more around 0 than in the case of i.i.d. patterns  $(\mathbf{X}_\mu^\top \mathbf{X}_\nu / N)_{\mu \neq \nu}$ . This is even more evident when plotting

the histogram of the overlap entries in the denoising phase, see Fig. 14. The broader distributions compared to the gray one indicate larger fluctuations than in the random case (where  $\mathbf{X}$  and  $\mathbf{x}$  are i.i.d.). Each  $\mathbf{X}_\mu$  has thus a small projection on most  $(\mathbf{X}_\nu)$  of order  $O(1/\sqrt{N})$ ; yet, as  $\lambda$  increases, a large fraction of these projections become typically larger than between independent patterns (but of the same order).

Beyond the transition, however, the diagonal contribution to the MMSE drops to zero, meaning that posterior patterns pair with ground truth patterns which are retrieved in a synchronized way (lower part of Fig. 12). Consistently, the off-diagonal contribution approaches 0, as posterior patterns “orthogonalize” with respect to the planted patterns, except for matched pairs that contribute to the reduction in the diagonal MMSE contribution. This shows that, in the factorization phase, nontrivial estimation of  $\mathbf{X}$  is possible, up to the unavoidable signs and permutation ambiguities of the columns.

Let us discuss a concrete algorithm to do so. Once the posterior samples’ column permutation and sign symmetry are broken in the factorization phase, their average is expected to possess a nontrivial alignment with  $\mathbf{X}$  up to the aforementioned ambiguity. The posterior samples can be synchronized efficiently by taking a single posterior sample as reference matrix (instead of  $\mathbf{X}$ ) and then aligning each remaining sample against this using again Algorithm 1, see Algorithm 2 below.

This procedure outputs an estimator  $\langle \mathbf{x} \rangle_{\text{alg}}$  of  $\mathbf{X}$  verifying criterion Eq. (38). If in addition one is given an oracle




FIG. 12. Average squared overlap  $N^{-2} \langle (\mathbf{X}^\top \mathbf{x})^2 \rangle$  (the square applies entrywise) for various SNRs obtained by Monte Carlo sampling after reordering and sign selection of the columns of each posterior sample according to Algorithm VI A, for a single instance  $(\mathbf{X}, \mathbf{Z})$  with Rademacher prior and  $(N, M) = (60, 30)$ .

Algorithm 2. Synchronized posterior average.

---

**Input:** Posterior samples  $\{\mathbf{x}_t\}_{t=1}^{N_{\text{MC}}}$ .  
**Output:**  $\langle \mathbf{x} \rangle_{\text{alg}}$ .  
**Initialize:**  $\langle \mathbf{x} \rangle_{\text{alg}} = \mathbf{x}_1 / N_{\text{MC}}$   
 1: **for**  $t = 2$  to  $N_{\text{MC}}$  **do**  
      $\pi, \sigma \leftarrow$  Algorithm 1 with Inputs  $\mathbf{x}_1$  and  $\mathbf{x}_t$   
      $\langle \mathbf{x} \rangle_{\text{alg}} \leftarrow \langle \mathbf{x} \rangle_{\text{alg}} + (\pi \circ \sigma)(\mathbf{x}_t) / N_{\text{MC}}$   
 4: **end for**

---

permutation  $\hat{\pi}(\mathbf{X})$  and signs  $\hat{\sigma}(\mathbf{X})$ , the remaining ambiguities can be broken by setting  $\langle \mathbf{x} \rangle_{\text{alg}} \leftarrow (\hat{\pi} \circ \hat{\sigma})(\langle \mathbf{x} \rangle_{\text{alg}})$ . Figure 13 depicts the recovery performance of  $\langle \mathbf{x} \rangle_{\text{alg}}$  from Algorithm 2 (in which the oracle permutation and signs are exploited for the sake of illustration). As can be seen clearly from both the Frobenius norm and trace overlap between  $\langle \mathbf{x} \rangle_{\text{alg}}$  and  $\mathbf{X}$ , information on  $\mathbf{X}$  itself can be retrieved in the factorization phase. Notice that Algorithm 2 requires as input posterior samples, which take a time exponential in  $N$  to sample by Monte Carlo simulation beforehand. Thus, we do not claim that factorization can be done efficiently, but rather that it is statistically possible.

Whether a nontrivial estimation of  $\mathbf{X}$  is possible in the denoising phase is less clear. The averaged entrywise squared overlap is displayed in Fig. 12 for various SNRs and, in particular, close to the crossover region (three plots at the bottom). The homogeneous background and weaker diagonal before the transition (upper part) signals that the information about the ground truth patterns is somehow diluted in all the posterior patterns, consistent with the bottom panel of Fig. 11. The diagonal then gets suddenly stronger and the background much smaller around the transition, indicating again a strong pattern-to-pattern matching between posterior samples and the signal in the factorization phase. Overall, these figures illustrate a clear change of behavior between the two phases, where the




FIG. 13. MSE (blue line, left axis) and overlap with the ground truth  $\mathbf{X}$  (red, right) for estimator  $\langle \mathbf{x} \rangle_{\text{alg}}$  from Algorithm 2 with  $5.10^5$  Monte Carlo samples for  $(N, M) = (60, 30)$  and a single data instance.

quasibasis of posterior patterns ( $\mathbf{x}_\mu$ ) (i.e., a set of  $N$ -dimensional vectors whose pairwise inner products are approximately  $\sqrt{N}$ ) “aligns” much more with the quasibasis of planted patterns ( $\mathbf{X}_\mu$ ) beyond the transition.

We have made two empirical observations that may suggest at first sight that factorization is possible in the denoising phase, too: (i) the diagonal contribution to the MMSE (upper part of Fig. 11) seems bounded away from one, and (ii) the diagonals displayed in Fig. 12 seem  $O(1)$  in the denoising phase. The question we thus aim at elucidating is whether these are artifacts of finite-size effects combined with the column reordering procedure applied to the posterior samples (which necessarily amplifies the overlap diagonal) or not. In other words, is the overlap  $|\mathbf{x}_\mu \cdot \mathbf{X}_\mu|/N$  between matched pairs vanishing or not in the thermodynamic limit in the denoising phase (still assuming the permutation ambiguity is resolved)?

In order to answer this question, we can heuristically estimate finite-size corrections, helping ourselves with the histograms at the top panel in Fig. 14. Similar histograms were obtained from overlap matrices to which we removed the diagonal for a fair comparison (even if the histogram almost does not change if keeping the diagonal, which is not true anymore beyond the transition as clear from the bottom panel in Fig. 14). The diagonal is removed since it has different statistical properties than the off-diagonal entries, and we need to understand the statistics of the complement of the diagonal in order to assess if such random variables can induce the observed diagonal just through the reordering procedure. From Fig. 14, we deduce that the marginal law of off-diagonal elements rescaled by  $\sqrt{N}$  behaves as centered Gaussian variables with variance  $\sigma^2(\lambda)$  tuned by the SNR, evaluated numerically from the histograms. We thus make the simplifying assumption that the correlations among the  $(Q_{\mu\nu})_{\mu, \nu \leq M}$  are sufficiently weak to consider them independent. This allows us to estimate the average of the diagonal of the entrywise squared overlap matrix under the null hypothesis that it is just due to the reordering.

According to the aforementioned independence and Gaussianity hypotheses on the overlap entries, the law of the  $\mu$ th diagonal squared overlap element in Fig. 12 can be modeled as that of

$$\chi_\mu := \max_{\mu \leq \nu \leq M} Q_{\mu\nu}^2.$$

The cumulative distribution function of  $\chi_\mu$  is explicit,

$$\mathbb{P}(\chi \leq t) = \mathbb{P}(|Q_{11}| \leq \sqrt{t})^{M_\mu} = \text{erf}\left(\sqrt{\frac{Nt}{2\sigma^2(\lambda)}}\right)^{M_\mu}$$

for any  $t \geq 0$ , where  $M_\mu := M - \mu + 1$ , and  $\text{erf}(z) := (2/\sqrt{\pi}) \int_0^z e^{-x^2} dx$ . Furthermore,  $\chi_\mu$  being a positive

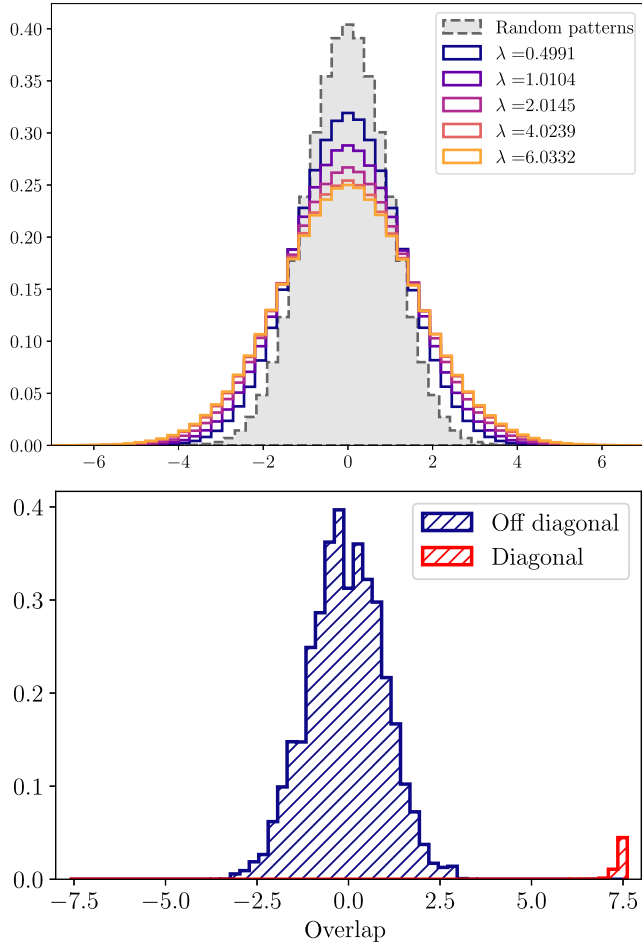


FIG. 14. Upper: Histogram of the entries of  $\sqrt{N}\mathbf{Q}$  in the denoising phase for a single instance of  $(\mathbf{X}, \mathbf{Z})$  with size  $(N, M) = (60, 30)$  and Rademacher prior, obtained from  $5 \times 10^3$  Monte Carlo samples after reordering of each sample column. All entries are accumulated into each histogram, resulting in a sample size of  $5 \times 10^3 M^2$  each. The gray histogram indicates the distribution of the overlap between two random Rademacher patterns of size  $N$ , which approaches a standard normal distribution as  $N \rightarrow +\infty$ . Lower: Histogram of  $\sqrt{N}\mathbf{Q}$  for  $\lambda = 7.96$  which lies at the start of the factorization phase given that size, see Fig. 8. In that case, outliers emerge and so we have separated the diagonal and off-diagonal contributions. The diagonal one is peaked around  $\sqrt{N} \approx 7.7$ . The off-diagonal histogram resembles again a normal distribution, and its variance is very close to 1.

random variable, the expectation of any function  $f$  of it can be computed via tail integration as

$$\begin{aligned} \mathbb{E}f(\chi_\mu) &= f(0) + \int_0^\infty f'(t)\mathbb{P}(\chi_\mu \geq t)dt \\ &= f(0) + \int_0^\infty f'(t) \left(1 - \operatorname{erf}\left(\sqrt{\frac{Nt}{2\sigma^2(\lambda)}}\right)^{M_\mu}\right) dt. \end{aligned}$$

We are interested in evaluating the mean and order of the fluctuations of the average of the diagonal elements

$$\bar{\chi} := \frac{1}{M} \sum_{\mu=1}^M \chi_\mu.$$

Following our assumptions, the  $\chi_\mu$ 's are independent random variables, so the variance  $\mathbb{V}(\bar{\chi}) = \sum_{\mu=1}^M \mathbb{V}(\chi_\mu)/M^2$ . Therefore, if the empirical dominant diagonal in Fig. 12 was just an artifact of the reordering, its average value should be compatible with

$$\bar{\chi} \approx \mathbb{E}\bar{\chi} \pm \frac{2}{M} \sqrt{\sum_{\mu=1}^M \mathbb{V}(\chi_\mu)},$$

considering two standard deviations as confidence interval. Recall that both values on the rhs depend on the SNR trough  $\sigma^2(\lambda)$ .

Figure 15 compares the value of  $\bar{\chi}$  with its uncertainty (orange) to the arithmetic mean of the diagonal elements of the averaged squared overlap  $N^{-2}\langle(\mathbf{X}^\top \mathbf{X})^{\otimes 2}\rangle$  (blue). There is an evident agreement in the denoising phase: the curve obtained via MCMC never detaches from  $\bar{\chi}$  by more than 2 standard deviations. This suggests that our assumptions on the overlap elements  $Q_{\mu\nu}$ 's are effective, and above all that the diagonals in Fig. 12 for  $\lambda \leq 6.723$ , i.e., in the denoising phase for that size, are compatible with the diagonal-

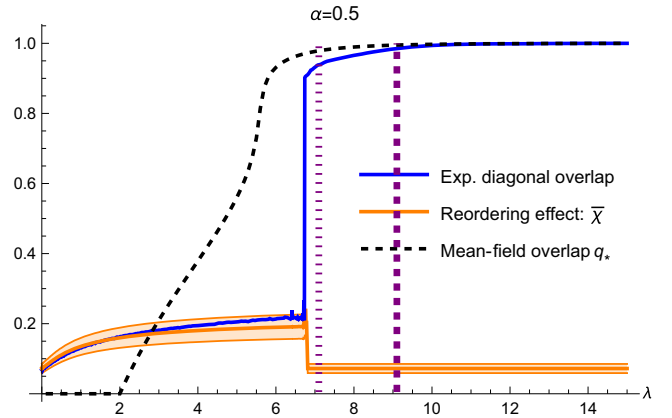


FIG. 15. Experimental average diagonal of the squared overlap as a function of  $\lambda$  obtained via Monte Carlo sampling after reordering of each sample column for a single instance  $(\mathbf{X}, \mathbf{Z})$  with Rademacher prior and  $(N, M) = (60, 30)$  (blue), as in Fig. 12. It is compared to the effect of the reordering of the columns of the posterior samples estimated by  $\mathbb{E}\bar{\chi}$  (orange curve), with 2 standard deviations confidence interval (band). We also display the equilibrium diagonal overlap  $q_*$  given by the mean-field theory Eq. (32). The complete Ansatz Eq. (37) predicts a transition at  $\lambda_c(\alpha = 0.5) \approx 9.1$  (vertical purple dashed curve), which can be corrected for finite-size effects yielding  $\approx 7.1$ , see caption of Fig. 9.

amplifying effect of the column reordering procedure. This entails that the posterior patterns ( $\mathbf{x}_\mu$ ) are delocalized in the quasibasis of the ground truth patterns ( $\mathbf{X}_\mu$ ). On the contrary, after the transition, the values of  $\bar{\chi}$  are incompatible (further than 5 standard deviations) with the empirical curve, signaling that the values in the diagonals of the bottom center and right panels in Fig. 12 are due not only to reordering, but to a real effect happening in the factorization phase. Note that the sudden decrease in  $\mathbb{E}\bar{\chi}$  in Fig. 15 is related to the collapse of the blue histogram at the bottom of Fig. 14 onto a variance one normal law right after the transition, while it was more spread just before it, i.e., it is a manifestation of the transition through  $\sigma^2(\lambda)$ . As a last comment, we can show that the effect  $\mathbb{E}\bar{\chi}$  of the reordering vanishes as  $O(2\sigma^2(\lambda)\ln(2M)/N)$  (as it is a mean of maxima of i.i.d. Gaussians, see Ref. [148]).

In conclusion, from the evidence we gathered, we expect that the seemingly bright diagonals in Fig. 12 for  $\lambda \leq 6.723$  disappear in the thermodynamic limit. Consequently, in the denoising phase, it is more plausible that factorization is not possible, in contrast with the factorization phase, and only denoising of  $\mathbf{X}\mathbf{X}^\top$  is. This aligns with the fact that one cannot extract an estimator of  $\mathbf{x}$  from the RIE, which we think is Bayes optimal in that phase (or at least we do not see how to do so).

Note that Figs. 11–13 and 15 once more suggest that the phase transition is discontinuous for discrete priors.

We have argued that nontrivial estimation of  $\mathbf{X}$  is impossible in the denoising phase. In order to clarify what is inferred about it that allows denoising of  $\mathbf{X}\mathbf{X}^\top$ , consider the following rotation for a given posterior sample:

$$\mathbf{O}_*(\mathbf{X}) := \underset{\mathbf{O} \in \mathbb{R}^{M \times M}: \mathbf{O}^\top \mathbf{O} = I_M}{\operatorname{argmin}} \|\mathbf{X} - \mathbf{x}\mathbf{O}\|^2.$$

In Fig. 16, we plot the overlap matrix between “optimally rotated” posterior samples and the ground truth, averaged over Monte Carlo samples, i.e.,  $\mathbf{X}^\top \langle \mathbf{x}\mathbf{O}_*(\mathbf{x}) \rangle / N$ . Clearly, in the denoising phase the profile of this matrix is diagonal, but with the diagonal now remaining order one as  $N$  increases and whose intensity grows with  $\lambda$ . Note that there is a finite diagonal overlap even when  $\lambda = 0$ , as two random  $M$ -dimensional subspaces of  $\mathbb{R}^N$  always have a common subspace of dimension  $O(N)$ . This demonstrates two points: On one hand, the denoising phase is the regime where the ground truth factor is estimated but only up to an unknown rotation of the basis in which the patterns ( $\mathbf{X}_\mu$ ) are expressed. On the other hand, this also reveals that it is not sufficient to break the permutation symmetry of the columns in order to see a nontrivial overlap appear (as was done in Fig. 12). One instead needs to break the much larger symmetry group of rotations, which further requires an oracle already knowing the planted matrix. This “effective rotational invariance” is what we believe to be the source of




FIG. 16. Average overlap  $N^{-1}\mathbf{X}^\top \langle \mathbf{x}\mathbf{O}_*(\mathbf{x}) \rangle$  computed by Monte Carlo sampling after reordering of each sample columns for a single instance of  $(\mathbf{X}, \mathbf{Z})$  with Rademacher prior and size  $(N, M) = (60, 30)$ .

universality in the denoising phase and what prevents factorization. We refer to Ref. [149] for a related discussion.

As a last comment, solving the saddle point Eqs. (32), we get that the mean-field overlap  $q_*$  detaches from 0 around  $\lambda = 2$  in the denoising phase, see Fig. 15. We have verified that this has an effect on the multiscale curves in Figs. 8 and 9, but it is just too small to be visible by eye. This may seem in contradiction with our negative conclusions concerning factorization in this phase. However, we claim that this is an artifact of the mean-field theory being applied in the “matrix model and denoising regime” where the necessary concentration-of-measure effects for it to hold are not present. As discussed in Sec. V, mean-field equations are reliable only after the transition  $\lambda_c(\alpha = 0.5) \approx 9.1$ , as manifest from the figure. We therefore do not interpret this detachment from 0 as a true transition, nor as a signal of possible factorization in this phase.

## VII. CONCLUSION AND PERSPECTIVES

In this paper, we have first conducted numerical experiments to grasp features of the phase diagram of matrix denoising in the challenging regime of extensive-rank beyond the rotational-invariant setting. We have clarified the rich phenomenology of the model. In particular, we uncovered a first-order phase transition separating a regime akin to a matrix model where strong universality properties hold and the RIE is a Bayes-optimal denoising algorithm, from a fundamentally different phase where it is not anymore due to universality breaking and which is more amenable to statistical physics techniques. The free entropy Eq. (16) seen as a function of the SNR  $\lambda$  thus interpolates between a

matrix model solvable by the HCIZ integral (the denoising phase) and a planted mean-field spin model (the factorization phase).

We have then proposed a mean-field theory based on a combination of spin-glass approaches to compute the mutual information and minimum mean-square error. Interestingly, it yields the same equations as the replica approach of Refs. [99–102]. The latter was believed to be incorrect for a long time. Through our independent analysis, we have instead clarified in which regime it is reliable—the factorization phase—or not—the denoising phase—and have shown how to correct it when not accurate thanks to universality, as conjectured by Semerjian [85], but for the whole phase diagram. Our complete *Ansatz* for the two phases is in good quantitative agreement with the experiments, up to unusually large finite-size corrections which we were able to quantify.

The numerical evidence suggests that optimal denoising is algorithmically hard only beyond the transition point, due to the presence of the universal RIE state at all signal-to-noise ratios. We have experimentally confirmed that this state prevents Monte Carlo algorithms from outperforming the RIE.

There are many directions to pursue from here. Concerning our present setting, we have proven that beyond a certain SNR universality cannot hold. It would be interesting to rigorously confirm that the universality indeed holds at low SNR and breaks precisely at the predicted transition  $\lambda_c(\alpha)$ , and that the mean-field theory does become asymptotically exact beyond that point.

Concerning generalizations, the most natural one is the nonsymmetric version of the problem, with more general noises, as considered in Refs. [99–102].

In addition, our setting shares several similarities with the Hopfield model [112] of associative memory. The phenomenology of the latter is well understood (at least at the replica symmetric level) thanks to the seminal work in Ref. [113]. The key idea is that samples from the Boltzmann-Gibbs measure cannot have extensive projections on all memorized patterns, but only on a finite subset and typically  $O(1/\sqrt{N})$  projections onto the others. These infinitesimal contributions, being extensively many, give rise to a contribution that can be treated as effective noise called pattern interference. Its intensity depends on the load of the memory  $M/N \approx \alpha$ . There exists a retrieval phase for low  $\alpha$  and temperatures, where an algorithm sampling from the Boltzmann-Gibbs distribution with informative initialization “recalls” a pattern. Out of this region, the samples have instead  $O(1/\sqrt{N})$  overlap with all the patterns.

We observe that for matrix denoising too, in the denoising phase, the overlap matrix presents  $O(1/\sqrt{N})$  entries. However, instead of treating these elements as noise as done in Hopfield, here they actually give a fundamental contribution to decreasing the MMSE as illustrated in Fig. 11. This is one of the main differences with the

Hopfield model: The sum of spurious alignments of the posterior sample columns  $\mathbf{x}_\mu$  onto other patterns of the ground truth is exploitable as information when denoising  $\mathbf{X}\mathbf{X}^\top$ , and it is not a disturbance. The term  $\text{Tr}(\mathbf{x}\mathbf{x}^\top)^2$  in the Hamiltonian Eq. (17) contributes to this phenomenology, generating a caging effect that prevents posterior patterns from overlapping and thus helps. As a consequence, we find that the “load”  $\alpha$  that the memory can withstand is higher. Our theory supports the fact that it is possible to retrieve *all* the patterns even for  $\alpha$ ’s sensibly larger than the Hopfield critical load  $\alpha_c \approx 0.138$  for binary patterns, provided an informative initialization is available and the signal-to-noise ratio is sufficiently high.

Furthermore, our theory and its agreement with experiments in the factorization phase indicates that this caging effect becomes strong enough to almost completely cleanse the system from pattern interference, yielding an MMSE akin to that of several quasidecoupled rank-one problems.

Another major difference is that in matrix denoising one attempts to recall all patterns (columns of  $\mathbf{X}$ ) at once (see the recent related work in Ref. [150] where multiple, but finitely many, patterns are jointly recovered), different from the approach of Ref. [110] for instance.

To conclude this analogy, the study of the memorization properties of this model should be carried out in a non-Bayes-optimal setting. One issue is that the universal branch of the matrix denoising problem is known exactly only thanks to the Nishimori identities which allow one to simplify the HCIZ integral [58,91]. Bayes optimality also prevents replica symmetry breaking [21]. These simplifying symmetries break down in the absence of Bayes optimality, and solving the model without relying on them could be much more challenging.

A further connection that is drawing attention is that with information-theoretical limits of neural networks. It was recently pinpointed by some of the present authors that a theory similar to ours can be used to describe learning transitions in shallow neural networks. More specifically, in [78] the authors exploited a mapping between the learning of an extensive-width neural network with quadratic activation function in a teacher-student setting and a matrix denoising problem, where the hidden matrix follows a Wishart distribution (and thus is exactly rotationally invariant). Based on the ideas of the present paper, Ref. [79] goes beyond by tackling any activation function but also possibly nonrotationally invariant weight matrices. Given the similar nature of the problem treated in Ref. [79], the remarkable agreement between theoretical predictions and numerical experiments in the said paper also supports our claims of exactness of the theory proposed here. Hence, our theory (together with the one of Refs. [99–102]) and, in particular, the rich phenomenological picture we have identified has the potential to become a valuable tool for describing a wide variety of matrix models going much beyond matrix denoising.

Another stretched, yet potentially fruitful comparison is that with Anderson localization [151]. The phenomenon is different as being quantum by nature and is not a thermodynamic phase transition. Yet, the similarity is striking due to both transitions being linked to localization, i.e., macroscopic pairwise alignment of two bases of vectors: the eigenvectors of the energy operator of the quantum system at hand localize in lattice basis when tuning the amount of disorder, while in the present problem the quasibasis of posterior patterns ( $\mathbf{x}_\mu$ ) localizes at the transition in the quasibasis of planted ones ( $\mathbf{X}_\mu$ ), which manifests itself by the abrupt increase of the overlap diagonal in Fig. 12. It would be interesting to explore how far this analogy can be pushed.

A final direction to pursue would be to compare known Bayesian algorithms for matrix inference to our results, such as BiGAMP [152], the AMP algorithm of Ref. [102], or Unitary AMP [153]. Given our results on the hardness of denoising using Monte Carlo sampling, we think that these algorithms adapted to the present setting may, too, face a hard phase and be outperformed by the RIE without additional information on  $\mathbf{X}$ . We plan to look into this in a future work.

### ACKNOWLEDGMENTS

J. B. and F. C. were funded by the European Union (ERC, CHORAL, Project No. 101039794). J. K. was

funded by the Natural Sciences and Engineering Research Council of Canada and the Canada Research Chairs programme (RGPIN-2020-04597, DGEGR-2020-00199). K. O. was funded by JSPS KAKENHI Grant No. 22KJ1074, JST CREST Grant No. JPMJCR1912, and Grant-in-Aid for Transformative Research Areas (A), ‘‘Foundation of Machine Learning Physics’’ (22H05117). Computational resources of the AI Bridging Cloud Infrastructure (ABCI) provided by the National Institute of Advanced Industrial Science and Technology (AIST) were used. Part of the work of K. O. was done while visiting the International Center for Theoretical Physics in Trieste. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

### APPENDIX A: IRRELEVANCE OF THE DIAGONAL OF THE DATA

*Proposition 1: Information irrelevance of the data diagonal components.* The diagonal part of the data  $(Y_{ii})_{i \leq N}$  does not contribute to the MI in the high-dimensional limit. Specifically, the inference problem

$$\tilde{Y}_{ij} = \sqrt{\frac{\lambda}{N}} \sum_{\mu=1}^M X_{i\mu} X_{j\mu} + \tilde{Z}_{ij}, \quad \tilde{Z}_{ij} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \quad 1 \leq i < j \leq N \quad (\text{A1})$$

has the same asymptotic mutual information density between data and signal as Eq. (1).

*Proof.* The proof is based on standard interpolation. The interpolating inference problem is

$$Y_{ij}^t = \begin{cases} \tilde{Y}_{ij} & i < j \\ \sqrt{\frac{t\lambda}{N}} \sum_{\mu=1}^M X_{i\mu}^2 + Z_{ii} & i = j \end{cases}$$

where  $Z_{ii} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 2)$ . The MI density for this problem is

$$\begin{aligned} \frac{1}{MN} I(\mathbf{Y}^t; \mathbf{X}) &= \frac{\lambda}{4MN^2} \sum_{i \neq j} \sum_{\mu, \nu=1}^M \mathbb{E} X_{i\mu} X_{j\mu} X_{i\nu} X_{j\nu} + \frac{t\lambda}{4MN^2} \sum_{i=1}^N \sum_{\mu, \nu=1}^M \mathbb{E} X_{i\mu}^2 X_{i\nu}^2 \\ &\quad - \frac{1}{MN} \mathbb{E} \ln \int dP_{\mathbf{X}}(\mathbf{x}) \exp \sum_{i < j} \left[ \sqrt{\frac{\lambda}{N}} \tilde{Y}_{ij} \mathbf{x}_i \cdot \mathbf{x}_j - \frac{\lambda}{2N} (\mathbf{x}_i \cdot \mathbf{x}_j)^2 \right] \\ &\quad \times \exp \frac{1}{2} \sum_{i=1}^N \left[ \sqrt{\frac{t\lambda}{N}} Y_{ii}^t \|\mathbf{x}_i\|^2 - \frac{t\lambda}{2N} \|\mathbf{x}_i\|^4 \right] \\ &= \frac{\lambda(1 + \alpha t)}{4} - \frac{1}{MN} \mathbb{E} \ln \int dP_{\mathbf{X}}(\mathbf{x}) \exp \sum_{i < j} \left[ \sqrt{\frac{\lambda}{N}} \tilde{Y}_{ij} \mathbf{x}_i \cdot \mathbf{x}_j - \frac{\lambda}{2N} (\mathbf{x}_i \cdot \mathbf{x}_j)^2 \right] \\ &\quad \times \exp \frac{1}{2} \sum_{i=1}^N \left[ \sqrt{\frac{t\lambda}{N}} Y_{ii}^t \|\mathbf{x}_i\|^2 - \frac{t\lambda}{2N} \|\mathbf{x}_i\|^4 \right] + O(N^{-1}). \end{aligned}$$

The last error term is uniform in  $t \in [0, 1]$ . The  $t$  derivative of this interpolating MI thus reads

$$\frac{d}{dt} \frac{1}{MN} I(\mathbf{X}; \mathbf{Y}^t) = \frac{\lambda\alpha}{4} - \frac{1}{MN} \sum_{i=1}^N \mathbb{E} \left\langle \frac{1}{4} \sqrt{\frac{\lambda}{tN}} Z_{ii} \|\mathbf{x}_i\|^2 + \lambda \frac{\|\mathbf{X}_i\|^2 \|\mathbf{x}_i\|^2}{2N} - \frac{\lambda}{4N} \|\mathbf{x}_i\|^4 \right\rangle_t + O(N^{-1}).$$

Now we can integrate the  $Z_{ii}$ 's by parts, taking into account that their variance is 2,

$$\begin{aligned} \frac{d}{dt} \frac{1}{MN} I(\mathbf{X}; \mathbf{Y}^t) &= \frac{\lambda\alpha}{4} - \frac{1}{MN} \sum_{i=1}^N \mathbb{E} \left\langle \frac{\lambda}{2N} \|\mathbf{x}_i^{(1)}\|^4 - \frac{\lambda}{2N} \|\mathbf{x}_i^{(1)}\|^2 \|\mathbf{x}_i^{(2)}\|^2 + \lambda \frac{\|\mathbf{X}_i\|^2 \|\mathbf{x}_i^{(1)}\|^2}{2N} - \frac{\lambda}{4N} \|\mathbf{x}_i^{(1)}\|^4 \right\rangle_t + O(N^{-1}) \\ &= \frac{\lambda\alpha}{4} - \frac{\lambda}{4MN^2} \sum_{i=1}^N \mathbb{E} \|\mathbf{X}\|^4 + O(N^{-1}) = O(N^{-1}), \end{aligned}$$

where in the last step we also used the Nishimori identity

$$\mathbb{E} \langle \|\mathbf{x}_i^{(1)}\|^2 \|\mathbf{x}_i^{(2)}\|^2 \rangle_t = \mathbb{E} \langle \|\mathbf{X}_i\|^2 \|\mathbf{x}_i^{(1)}\|^2 \rangle_t,$$

and superscripts  $\mathbf{x}^{(k)}$  denote replica indices, namely, conditionally (on  $\mathbf{Y}$ ) independent samples from the posterior measure. The error term is again uniform in  $t$  so the proof is complete by noticing that  $\mathbf{Y}^{t=1}$  are the observations in Eq. (1) and  $\mathbf{Y}^{t=0}$  is Eq. (A1). ■

The above proposition holds also if  $\alpha$  is vanishing in the high-dimensional limit, i.e., for low-rank settings.

## APPENDIX B: DETAILS OF THE MONTE CARLO PROCEDURE

In this appendix, we elaborate on the numerical procedures necessary to evaluate the MI and MMSE for finite-size systems. The target distribution is the posterior measure given by

$$\begin{aligned} P_\lambda(\mathbf{x}; \mathbf{X}, \mathbf{Z}) &= \frac{P_X(\mathbf{x}) e^{-H_\lambda(\mathbf{x}; \mathbf{X}, \mathbf{Z})}}{\mathcal{Z}_\lambda(\mathbf{X}, \mathbf{Z})}, \\ -H_\lambda(\mathbf{x}; \mathbf{X}, \mathbf{Z}) &= \sqrt{\frac{\lambda}{N}} \sum_{i < j, 1}^N (\mathbf{X}_i \cdot \mathbf{X}_j + Z_{ij}) \mathbf{x}_i \cdot \mathbf{x}_j \\ &\quad - \frac{\lambda}{2N} \sum_{i < j, 1}^N (\mathbf{x}_i \cdot \mathbf{x}_j)^2, \end{aligned}$$

where  $\mathcal{Z}_\lambda$  is the partition function for SNR  $\lambda$ . Here, all expressions are explicitly rewritten using the quenched random variables  $\mathbf{X}, \mathbf{Z}$ , and the subscripts for  $N, M$  are dropped for brevity.

Evaluating the MI numerically requires one to calculate  $\ln \mathcal{Z}_\lambda(\mathbf{X}, \mathbf{Z})$  with high precision, which is difficult to perform with standard MCMC samplers. Therefore, we use bridge sampling [154,155], more popularly known as ‘‘annealed importance sampling’’ [120]. Given a strictly increasing sequence of SNRs  $0 = \lambda_0 < \lambda_1 < \dots < \lambda_R$ , the logarithm of the partition function at  $\lambda_R$  can alternatively be expressed as a telescopic sum,

$$\begin{aligned} \ln \mathcal{Z}_{\lambda_R}(\mathbf{X}, \mathbf{Z}) &= \ln \prod_{r=0}^{R-1} \frac{\mathcal{Z}_{\lambda_{r+1}}(\mathbf{X}, \mathbf{Z})}{\mathcal{Z}_{\lambda_r}(\mathbf{X}, \mathbf{Z})} \\ &= \sum_{r=0}^{R-1} \ln \langle \exp[H_{\lambda_r}(\mathbf{x}; \mathbf{X}, \mathbf{Z}) - H_{\lambda_{r+1}}(\mathbf{x}; \mathbf{X}, \mathbf{Z})] \rangle_{\lambda_r}, \end{aligned} \tag{B1}$$

where  $\langle \cdot \rangle_\lambda$  is the expectation with respect to  $P_\lambda(\mathbf{x}; \mathbf{X}, \mathbf{Z})$ . By taking  $\lambda_{r+1}$  and  $\lambda_r$  sufficiently close, and consequently taking  $H_{\lambda_r}(\mathbf{x}; \mathbf{X}, \mathbf{Z}) - H_{\lambda_{r+1}}(\mathbf{x}; \mathbf{X}, \mathbf{Z})$  close to 0, the value in the brackets would yield low variance, allowing one to estimate its mean with high precision using samples generated from  $P_\lambda(\mathbf{x}; \mathbf{X}, \mathbf{Z})$ . To sample from a sequence of posterior measures with SNRs  $\{\lambda_r\}_{r=0}^R$  in an efficient manner, we employ parallel tempering, or the replica-exchange Monte Carlo method [118,119]. This approach involves simultaneous execution of ‘‘local’’ MCMC samplers specifically simulating each of the  $R + 1$  systems. In addition, configurations between adjacent SNRs are occasionally exchanged while satisfying the detailed balance condition for fast mixing time. More explicitly, given an adjacent pair of SNRs  $\lambda$  and  $\lambda'$ , with its current configuration of their corresponding sampler being given by  $\mathbf{x}$  and  $\mathbf{x}'$ , respectively, the exchange probability of the two configurations can be given using Metropolis's method,

$$\min\{1, \exp[H_{\lambda'}(\mathbf{x}'; \mathbf{X}, \mathbf{Z}) - H_{\lambda'}(\mathbf{x}; \mathbf{X}, \mathbf{Z}) + H_\lambda(\mathbf{x}; \mathbf{X}, \mathbf{Z}) - H_\lambda(\mathbf{x}'; \mathbf{X}, \mathbf{Z})]\}.$$

In our experiments, we use a single spin-flip Metropolis-Hastings algorithm as the local MCMC sampler for each SNR. A single MCMC step is defined as  $30NM$  local spin-flip moves on all  $R + 1$  systems, followed by a replica-exchange move between all adjacent pairs of SNR. Before each replica-exchange move, the configuration of  $\mathbf{x}\mathbf{x}^\top$  of each MCMC chain is recorded. After the Monte Carlo procedure has finished, the thermal average appearing in the expression for the log-partition function Eq. (B1) needed for the MI and in the MMSE Eq. (6) is obtained by the empirical average over these samples. This whole

TABLE I. Specific parameters used to run each Monte Carlo experiment. Experiments (1)–(3) are conducted for both the Rademacher prior case and the nonsymmetric prior case ( $P_X = \frac{1}{4}\delta_{\sqrt{3}} + \frac{3}{4}\delta_{-1/\sqrt{3}}$ ) with the same parameters. The rest of the experiments are conducted only for the Rademacher prior case.

Experiment no.	$(N, M)$	$\alpha$	Burn-in MCMC steps	Sampling MCMC steps	$R + 1$	No. of random instances
(1)	(10, 5)	0.5	2000	8000	72	128
(2)	(20, 10)	0.5	10 000	40 000	135	128
(3)	(40, 20)	0.5	10 000	40 000	261	32
(4)*	(20, 10)	0.5	10 000	40 000	29	128
(5)*	(40, 20)	0.5	10 000	40 000	29	128
(6)*	(60, 30)	0.5	40 000	160 000	29	128
(7)*	(40, 28)	0.7	10 000	40 000	29	128
(8)	(60, 30)	0.5	500 000	500 000	313	1
(9)	(60, 18)	0.3	40 000	160 000	256	32
(10)	(40, 28)	0.7	40 000	160 000	261	32

Monte Carlo procedure is executed for several instances of  $(\mathbf{X}, \mathbf{Z})$  to then evaluate the average over the quenched randomness.

The specific number of MCMC steps and the number of replicas  $R$ , as well as the number of random instances of  $(\mathbf{X}, \mathbf{Z})$  used for the simulations are given in Table I. Table II specifies which data were used to generate each figure in this paper. The precise values of  $\{\lambda_r\}_{r=0}^R$  are chosen such that the exchange rate between adjacent SNRs is always finite (typically larger than 0.3) by using the adaptive algorithm given by Ref. [156]. This is with the exception of experiments designated with the star in Table I, whose SNR values are given by equispaced values in the range 8–15.

As indicated in Table II, the MCMC curves in Fig. 8 are obtained by augmenting data from two different types of experiments: one from the usual replica-exchange MCMC [Eqs. (3) and (11)] and the other from MCMC specifically run under high SNR, Eqs. (7) and (8). Note that with only the results from Eqs. (3) and (11), the mean of the MMSE will become hidden in the statistical error from a finite number of random instances. The results from high SNR, Eqs. (7) and (8), are given to clarify the exponential decay of the MMSE, which have significantly lower statistical error due to the larger number of random instances we can afford to average over.

TABLE II. Table of corresponding experiments used to generate each figure in this paper.

Figure ref.	Experiment no.
2, 3, 4	(1,2,3)
6	(4,5,6)
8 ( $\alpha = 0.3$ )	(9)
8 ( $\alpha = 0.5$ )	(3) (for $\lambda < 8$ ) (6) (for $\lambda > 8$ )
8 ( $\alpha = 0.7$ )	(10) (for $\lambda < 12$ ) (7) (for $\lambda > 12$ )
11	(1,2,3)
12–16	(8)

### 1. Experiments for Fig. 5

The experimental procedure to obtain Fig. 5 is very different from that to obtain the other ones. Unlike the other experiments, here we consider only MCMC without replica-exchange moves, i.e., independent MCMC chains are simply run for each value of SNR. The informative points (black markers) are obtained by initializing a MCMC chain on the ground truth and performing  $50\,000 \times 30NM$  local spin-flip moves for each SNR. This is repeated on 36 random realizations of  $(\mathbf{X}, \mathbf{Z})$  to evaluate the disordered average. The uninformative points (red markers) are obtained by initializing each MCMC chain to a random configuration and repeating the same process as that to obtain the informative points. Last, the averaged uninformative points, indicated by the blue markers, are obtained by running 32 individual MCMC chains with random initialization for each SNR and  $(\mathbf{X}, \mathbf{Z})$ . The estimator for  $\mathbf{X}\mathbf{X}^T$  is obtained by averaging over all 32 runs with different initializations. More concretely, for a given realization of  $(\mathbf{X}, \mathbf{Z})$ , if we let  $\mathbf{x}_i(t) \in \{-1, +1\}^{N \times M}$  be the configuration of the MCMC chain after  $30NMt$  local spin-flip moves were performed from a random initialization indexed by  $i$ , the estimate of  $\mathbf{X}\mathbf{X}^T$  is given by

$$\langle \mathbf{x}\mathbf{x}^T \rangle_{\text{uninfo averaged}} = \frac{1}{32 \times 4000} \sum_{i=1}^{32} \sum_{t=1001}^{5000} \mathbf{x}_i(t)\mathbf{x}_i(t)^T.$$

This procedure is then repeated over 36 random realizations of  $(\mathbf{X}, \mathbf{Z})$  to evaluate the disordered average. Note that, for each initialization, only  $5000 \times 30NM$  local spin-flip moves were performed in total, since we found after performing the experiments for uninformative points (red markers) that the MCMC chains get stuck to a random spin configuration typically after  $\sim 100 \times 30NM$  local spin-flip moves for SNR larger than 7. Note that this is different from the procedure to obtain uninformative points (red markers), where for a given realization of  $(\mathbf{X}, \mathbf{Z})$ , the estimate is given only by a single initialization,

$$\langle \mathbf{x}\mathbf{x}^\top \rangle_{\text{unifinfo}} = \frac{1}{40\,000} \sum_{t=10\,001}^{50\,000} \mathbf{x}_1(t)\mathbf{x}_1(t)^\top.$$

## 2. Initialization of the MCMC chains

In principle, given sufficient amount of Monte Carlo steps, the MCMC simulations are able to sample from the true equilibrium measure, irrespective of the configuration in which each chain was initialized. Here, the ground truth  $\mathbf{X}$  is chosen as the initial configuration, with the obvious exception of experiment Eq. (3). This is done for the following reason. As seen in Fig. 5, for low SNR the system converges to the RIE irrespective of initialization, indicating a benign energy landscape explorable by local spin flips. On the other hand, under high SNR, an MCMC chain with random initialization encounters a possibly glassy phase, where mixing time increases exponentially with respect to  $N$ . This resembles a “golf-course”-type of energy landscape [157], where it consists of a large number of excited glassy states and a small number of ground states, which correspond to the ground truth  $\mathbf{X}$ . An important observation is that, while it is difficult to reach the ground truth from the excited states by using replica-exchange and local spin-flip moves, one can easily reach excited states from a ground state by exchanging configurations with replicas in low SNR. Therefore, given that the replicas exchange with finite probability in our simulations, replicas with a paramagnetic equilibrium will not remain trapped in the initial ground state for exponentially long periods; these replicas can access equilibrium states through exchanges with replicas in low SNR or via local spin-flip moves if the landscape is benign (which is the case in the denoising phase). Simultaneously, for the replicas within the factorization phase, initializing on  $\mathbf{X}$  eliminates the effort to search for rare ground states across the entire configuration

space. This enables efficient investigation of the problem without introducing significant bias toward informative states in the numerical results.

## APPENDIX C: DETAILS OF THE MULTISCALE MEAN-FIELD COMPUTATION

The subsequent analysis will rely on special identities inherent to the Bayes optimality of the setting. As a simple consequence of the tower rule for conditional expectations, one can prove the Nishimori identities that will induce numerous simplifications along the way. With a little abuse of notation, we denote by  $\langle \cdot \rangle$  also the infinite product measure  $\otimes_{n=1}^{\infty} \mathbb{E}[\cdot | \mathbf{Y}]$ . Hence, replicas, namely, conditionally (on  $\mathbf{Y}$ ) independent samples from the posterior (3), are averaged jointly by the measure  $\langle \cdot \rangle$ . Now we can state the following.

*Proposition 2: Nishimori identity.* Let  $f$  be a bounded function of the observations  $\mathbf{Y}$ , the ground truth factor  $\mathbf{X}$ , and of replicas  $(\mathbf{x}^{(k)})_{k \leq n}$  drawn independently from the posterior (3). Then,

$$\mathbb{E}_{\mathbf{X}, \mathbf{Y}} \langle f(\mathbf{Y}; \mathbf{X}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n)}) \rangle = \mathbb{E}_{\mathbf{Y}} \langle f(\mathbf{Y}; \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(n)}) \rangle. \quad (\text{C1})$$

An elementary proof of this fact can be found in Ref. [16].

### 1. First scale reduction: $\mathcal{O}(MN) \rightarrow \mathcal{O}(M)$ reduction by the cavity method

Here we detail more precisely how to devise properly the RLE problem mentioned in Sec. IV via the cavity method. Recall the Hamiltonian Eq. (17) with  $N$  rows and  $M$  columns and where  $Y_{ij}$  has been expanded in signal and noise contributions,

$$-H_{N,M}(\mathbf{x}; \mathbf{Z}, \mathbf{X}) = \sum_{i < j}^N \left[ \sqrt{\frac{\lambda}{N}} Z_{ij} \mathbf{X}_i \cdot \mathbf{X}_j + \frac{\lambda}{N} (\mathbf{X}_i \cdot \mathbf{X}_j) \mathbf{x}_i \cdot \mathbf{x}_j - \frac{\lambda}{2N} (\mathbf{x}_i \cdot \mathbf{x}_j)^2 \right]$$

with  $\mathbf{x}_i \cdot \mathbf{x}_j = \sum_{\mu \leq M} x_{i\mu} x_{j\mu}$ . Recall also the notation for the row cavity  $\mathbf{x}_N \equiv \boldsymbol{\eta} \in \mathbb{R}^M$ ,  $\mathbf{X}_N \equiv \mathbf{H}$ , and the bulk (i.e., all but the cavity) variables  $\bar{\mathbf{x}}, \bar{\mathbf{X}} \in \mathbb{R}^{N-1 \times M}$ . Now we segregate the contribution of  $\boldsymbol{\eta}$  from the Hamiltonian,

$$-H_{N,M}(\bar{\mathbf{x}}, \boldsymbol{\eta}; \lambda) = -H'_{N,M}(\bar{\mathbf{x}}; \lambda) - H_{N,M}^{\text{row}}(\boldsymbol{\eta}; \bar{\mathbf{x}}, \lambda),$$

with bulk Hamiltonian

$$-H'_{N,M}(\bar{\mathbf{x}}; \lambda) = \sum_{i < j}^{N-1} \left[ \sqrt{\frac{\lambda}{N}} Z_{ij} \mathbf{x}_i \cdot \mathbf{x}_j + \frac{\lambda}{N} (\mathbf{X}_i \cdot \mathbf{X}_j) \mathbf{x}_i \cdot \mathbf{x}_j - \frac{\lambda}{2N} (\mathbf{x}_i \cdot \mathbf{x}_j)^2 \right] \quad (\text{C2})$$

and cavity Hamiltonian

$$-H_{N,M}^{\text{row}}(\boldsymbol{\eta}; \bar{\mathbf{x}}, \lambda) = \sum_{i=1}^{N-1} \left[ \sqrt{\frac{\lambda}{N}} Z_{iN} \mathbf{x}_i \cdot \boldsymbol{\eta} + \frac{\lambda}{N} (\mathbf{X}_i \cdot \mathbf{H}) \mathbf{x}_i \cdot \boldsymbol{\eta} - \frac{\lambda}{2N} (\mathbf{x}_i \cdot \boldsymbol{\eta})^2 \right]. \quad (\text{C3})$$

In order to lighten the notation, we have dropped dependencies on the ground truth and the noise. From  $-H_{N,M}^{\text{low}}$ , we recognize (up to a constant) the log-likelihood of an RLE with uncertain design inference problem over both  $(\tilde{\mathbf{X}}, \mathbf{H})$  with data  $\tilde{\mathbf{Y}}(\lambda) = \mathbf{Y}_N = (\tilde{Y}_i = Y_{Ni})_{i \leq N-1}$  generated as in Eq. (20) or, more generally, as

$$\tilde{Y}_i \sim P_{\text{out}}\left(\cdot \left| \frac{\mathbf{X}_i \cdot \mathbf{H}}{\sqrt{N}} \right.\right), \quad i \leq N-1, \quad (\text{C4})$$

with output channel

$$P_{\text{out}}(y|x) := \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(y-\sqrt{\lambda}x)^2}. \quad (\text{C5})$$

In addition, we also have side information about the bulk variables coming from the observations Eq. (21).

We are going to work in this general setting and specialize to the Gaussian linear channel Eq. (C5) at the end. Note that, with these notations, the mutual information  $I(\tilde{\mathbf{Y}}(\lambda), \mathbf{Y}'(\zeta); \tilde{\mathbf{X}}, \mathbf{H})$  is equal to the original one Eq. (15) when  $\zeta = \lambda$ .

Let us also introduce

$$\langle \cdot \rangle' = \frac{1}{\mathcal{Z}'(\mathbf{Y}'(\zeta))} \int_{\mathbb{R}^{N-1 \times M}} dP_X(\tilde{\mathbf{x}}) e^{-H_{N,M}^{\text{low}}(\tilde{\mathbf{x}}; \zeta)}(\cdot)$$

with proper normalization  $\mathcal{Z}'$ .  $\langle \cdot \rangle'$  is the posterior measure associated with the channel Eq. (21) with SNR  $\zeta$ , and it is the one averaging over the bulk degrees of freedom.

## 2. Second scale: $\mathcal{O}(M) \rightarrow \mathcal{O}(1)$ reduction by bulk *Ansatz* and replica formulas

We recall that within the mean-field *Ansatz* Eq. (25) the simplified  $\sigma$ -dependent bulk measure reads

$$\langle \cdot \rangle'_{\text{eff}} = \prod_{i,\mu=1}^{N-1,M} \frac{1}{\mathcal{Z}'(Y'_{i\mu})} \int dP_X(x_{i\mu}) e^{\sqrt{\sigma} Y'_{i\mu} x_{i\mu} - \frac{1}{2} \sigma x_{i\mu}^2}(\cdot),$$

where  $\mathcal{Z}'(Y'_{i\mu})$  is the proper normalization per bulk entry. Under this simplified bulk measure, we define the free entropy  $\phi_{\text{eff},N}^{\text{RLE}}$  similar to Eq. (24),

$$\phi_{\text{eff},N}^{\text{RLE}}(\lambda, \sigma) = \phi_{\text{eff},N}^{\text{RLE}} := -\frac{1}{\alpha} \mathbb{E} \ln P_{\text{out}}(\tilde{Y}_1|0) + \frac{1}{M} \mathbb{E}_{\tilde{\mathbf{Y}}(\lambda), \mathbf{Y}'_{\text{eff}}(\sigma)} \ln \int_{\mathbb{R}^M} dP_X(\boldsymbol{\eta}) \left\langle \prod_{i=1}^{N-1} P_{\text{out}}\left(\tilde{Y}_i \left| \frac{\mathbf{x}_i \cdot \boldsymbol{\eta}}{\sqrt{N}} \right.\right) \right\rangle'_{\text{eff}}.$$

For the first term, we have used the fact that  $P_{\text{out}}(\tilde{\mathbf{Y}}|0)$  factorizes over the elements of  $\tilde{\mathbf{Y}}$ , giving a sum of contributions that are all equal in average. Again, the different symbol for the above quantity stresses the fact that we have used our mean-field *Ansatz*. Notice also that we are using the more generic definition Eq. (C4). There is an immediate simplification due the factorization of the bulk measure,

$$\phi_{\text{eff},N}^{\text{RLE}} = -\frac{1}{\alpha} \mathbb{E} \ln P_{\text{out}}(\tilde{Y}_1|0) + \frac{1}{M} \mathbb{E}_{\tilde{\mathbf{Y}}(\lambda), \mathbf{Y}'_{\text{eff}}(\sigma)} \ln \int dP_X(\boldsymbol{\eta}) \prod_{i=1}^{N-1} \left\langle P_{\text{out}}\left(\tilde{Y}_i \left| \frac{\mathbf{x}_i \cdot \boldsymbol{\eta}}{\sqrt{N}} \right.\right) \right\rangle'_{\text{eff}},$$

where we were able to factorize the  $i$  product.

Under the factorized measure  $\langle \cdot \rangle'_{\text{eff}}$ , the rescaled sums  $\mathbf{x}_i \cdot \boldsymbol{\eta} / \sqrt{N}$  behave as Gaussians by the central limit theorem, for any fixed  $\boldsymbol{\eta}$ . Therefore, it is now equivalent to think of the  $x_{i\mu}$ 's as independent Gaussians with mean and variance

$$m(Y'_{i\mu}) := \langle x_{i\mu} \rangle'_{\text{eff}}, \quad V(Y'_{i\mu}) := \langle x_{i\mu}^2 \rangle'_{\text{eff}} - (\langle x_{i\mu} \rangle'_{\text{eff}})^2$$

or, equivalently,

$$x_{i\mu} \stackrel{\text{D}}{=} m(Y'_{i\mu}) + \xi_{i\mu} \sqrt{V(Y'_{i\mu})}, \quad (\text{C6})$$

with  $\xi_{i\mu} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$  the new annealed bulk degrees of freedom. It is now the moment to define  $J(\sigma)$  in Eq. (27) as the overlap for the bulk variables

$$J = J(\sigma) := \mathbb{E} m(Y'_{i\mu}) X_{i\mu} = \mathbb{E}_{X, Z'} \left[ X \frac{\int dP_X(x) x e^{(\sqrt{\sigma} Z' + X\sigma)x - \frac{1}{2} \sigma x^2}}{\int dP_X(x) e^{(\sqrt{\sigma} Z' + X\sigma)x - \frac{1}{2} \sigma x^2}} \right]$$

with  $Z' \sim \mathcal{N}(0, 1)$ ,  $X \sim P_X$ . By the Nishimori identity, it is also equal to  $J = \mathbb{E} m(Y'_{i\mu})^2$ . Hence, using the properties of the prior, we have that  $(m(Y'_{i\mu}))_{i,\mu}$  are i.i.d. sub-Gaussian random variables with

$$\mathbb{E}m(Y'_{i\mu}) = 0, \quad \mathbb{E}m(Y'_{i\mu})^2 = J. \quad (\text{C7})$$

Rewriting  $\phi_{\text{eff},N}^{\text{RLE}}$  by means of Eq. (C6) we get

$$\phi_{\text{eff},N}^{\text{RLE}} = \frac{1}{M} \mathbb{E}_{\tilde{\mathbf{Y}}, \mathbf{Y}'_{\text{eff}}} \ln \int dP_X(\boldsymbol{\eta}) \prod_{i=1}^{N-1} \mathbb{E}_{\xi_i} P_{\text{out}} \left( \tilde{Y}_i \left| \frac{m(\mathbf{Y}'_i) \cdot \boldsymbol{\eta}}{\sqrt{N}} + \frac{(\sqrt{V(\mathbf{Y}'_i)} \circ \boldsymbol{\xi}_i) \cdot \boldsymbol{\eta}}{\sqrt{N}} \right. \right) - \frac{1}{\alpha} \mathbb{E} \ln P_{\text{out}}(\tilde{Y}_1|0) + o_N(1)$$

with  $\boldsymbol{\xi}_i = (\xi_{i\mu})_{\mu \leq M}$ ,  $m(\mathbf{Y}'_i) = (m(Y'_{i\mu}))_{\mu}$  and similar for  $V(\mathbf{Y}'_i)$ , and  $\circ$  denotes the componentwise product. The remainder is due to the fact that  $i$  runs to  $N-1$  and not to  $N$ . It will be dropped in the following.

Being in a Bayes-optimal setting, strong concentration effects allow us to simplify  $((\sqrt{V(\mathbf{Y}'_i)}/N \circ \boldsymbol{\xi}_i) \cdot \boldsymbol{\eta})_i$ . Indeed, conditionally on  $\boldsymbol{\eta}$  these are independent centered Gaussian random variables with variance

$$\frac{1}{N} \sum_{\mu=1}^M V(Y'_{i\mu}) \eta_{\mu}^2 = \alpha \mathbb{E}V(\sqrt{\sigma}X + Z') + o_N(1),$$

where the equality holds with overwhelming probability when  $N, M \rightarrow \infty$ . We stress that we have used  $\|\boldsymbol{\eta}\|^2/M = 1 + o_N(1)$ .  $\boldsymbol{\eta}$  is a sample from a posterior, so this replacement is *a priori* not guaranteed, but thanks to Bayes-optimality  $\|\boldsymbol{\eta}\|^2/M$  concentrates onto its mean, which is the same as for the ground truth by the Nishimori identity. Still using the Nishimori identity, one gets directly that  $\mathbb{E}V(Y'_{i\mu}) = 1 - J$ . Therefore, by defining the noisier channel

$$\tilde{P}_{\text{out}}^{(\sigma, \alpha)}(y|x) := \mathbb{E}_{\xi \sim \mathcal{N}(0,1)} P_{\text{out}}(y|x + \xi \sqrt{\alpha(1-J(\sigma))})$$

and collecting all the above observations, up to negligible corrections the free entropy is rewritten as

$$\phi_{\text{eff},N}^{\text{RLE}} = -\frac{1}{\alpha} \mathbb{E} \ln P_{\text{out}}(\tilde{Y}_1|0) + \frac{1}{M} \mathbb{E}_{\tilde{\mathbf{Y}}, \mathbf{Y}'_{\text{eff}}} \ln \int dP_X(\boldsymbol{\eta}) \prod_{i=1}^{N-1} \tilde{P}_{\text{out}}^{(\sigma, \alpha)} \left( \tilde{Y}_i \left| \frac{m(\mathbf{Y}'_i) \cdot \boldsymbol{\eta}}{\sqrt{N}} \right. \right).$$

The first contribution is simple to compute since  $\tilde{Y}_1 \sim P_{\text{out}}(\cdot|\mathbf{X}_1 \cdot \mathbf{H}/\sqrt{N})$  is asymptotically equal in law to  $\tilde{Y}_1 \sim P_{\text{out}}(\cdot|\sqrt{\alpha}Z')$  with  $Z' \sim \mathcal{N}(0, 1)$ . A final simplification is to replace the i.i.d. centered sub-Gaussian variables  $(m(Y'_{i\mu}))$  playing the role of quenched covariates by Gaussians  $(\sqrt{J}\tau_{i\mu})$  with  $\tau_{i\mu} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ , using Eq. (C7). This leaves the free entropy unchanged at leading order because the replica computation for the generalized linear model and its rigorous proof simply leverage on the fact that  $(m(\mathbf{Y}'_i) \cdot \boldsymbol{\eta}/\sqrt{N})_{i \leq N}$  behave as Gaussian random variables under the quenched expectation with respect to  $\mathbf{Y}'$ , and these have the same law in the large system limit as  $(\boldsymbol{\tau}_i \cdot \boldsymbol{\eta} \sqrt{J/N})_i$ . Therefore, we have

$$\phi_{\text{eff},N}^{\text{RLE}} = -\frac{1}{\alpha} \mathbb{E} \ln P_{\text{out}}(\tilde{Y}_1|0) + \frac{1}{M} \mathbb{E}_{\tilde{\mathbf{Y}}, \boldsymbol{\tau}} \ln \int dP_X(\boldsymbol{\eta}) \prod_{i=1}^{N-1} \tilde{P}_{\text{out}}^{(\sigma, \alpha)} \left( \tilde{Y}_i \left| \sqrt{J} \frac{\boldsymbol{\tau}_i \cdot \boldsymbol{\eta}}{\sqrt{N}} \right. \right)$$

at leading order. It is finally in the familiar form of free entropy of a generalized linear model with quenched Gaussian covariates matrix, whose rigorous limiting value is found in Ref. [48],

$$\begin{aligned} \phi_{\text{eff}}^{\text{RLE}} = \lim_{N \rightarrow \infty} \phi_{\text{eff},N}^{\text{RLE}} = \text{extr}_+ \left\{ -\frac{rq}{2} + \mathbb{E}_{\xi, H} \ln \int dP_X(\boldsymbol{\eta}) e^{(\xi\sqrt{r} + Hr)\boldsymbol{\eta} - \frac{1}{2}r\boldsymbol{\eta}^2} - \frac{1}{\alpha} \mathbb{E}_{\xi} \int dy P_{\text{out}}(y|\sqrt{\alpha}\xi) \ln P_{\text{out}}(y|0) \right. \\ \left. + \frac{1}{\alpha} \mathbb{E}_{\xi, u_0} \int dy \tilde{P}_{\text{out}}^{(\sigma, \alpha)}(y|\sqrt{\alpha J(\sigma)(1-q)}u_0 + \sqrt{\alpha J(\sigma)q\xi}) \ln \mathbb{E}_u \tilde{P}_{\text{out}}^{(\sigma, \alpha)}(y|\sqrt{\alpha J(\sigma)(1-q)}u + \sqrt{\alpha J(\sigma)q\xi}) \right\}, \end{aligned}$$

where extremization is intended with respect to  $(r, q)$ , and  $(\xi, u_0, u)$  are i.i.d. standard Gaussian variables,  $H \sim P_X$ .  $\text{extr}_+$  selects the solution of the saddle point equations, obtained by equating to zero the gradient of the replica potential  $\{\dots\}$ , which maximizes it. If we specify the above to the Gaussian output channel Eq. (C5), we readily obtain Eq. (26).

## APPENDIX D: PROPERTIES OF THE MEAN-FIELD MUTUAL INFORMATION

### 1. Reconstruction of Eq. (34) from the saddle point equations

We are looking for an MI potential, denoted in the following as  $\iota(r, q; \alpha, \lambda)$ , generating the saddle point Eqs. (32). We start by imposing the I-MMSE relation Eq. (5),

$$\frac{d}{d\lambda} \iota(r_*, q_*; \alpha, \lambda) = \frac{1}{4} \frac{1 - q_*(\alpha, \lambda)^2}{1 + \lambda\alpha(1 - q_*(\alpha, \lambda)^2)}. \quad (\text{D1})$$

We observe that Eq. (32) defines a vector field

$$v_r(r, q) := q - \mathbb{E}H\langle \eta \rangle_r, \quad v_q(r, q) := r - \frac{\lambda q}{1 + \lambda\alpha(1 - q^2)}.$$

It is conservative, because it is irrotational on a simply connected domain,

$$\frac{\partial}{\partial q} v_r(r, q) = 1 = \frac{\partial}{\partial r} v_q(r, q).$$

Therefore, any line integral of this vector field starting from (0,0) and ending at an arbitrary point  $(r, q)$  will yield the potential, up to a constant shift and an arbitrary common rescaling of the vector field components. We choose the integration path  $(0,0) \rightarrow (r,0) \rightarrow (r,q)$ ,

$$\begin{aligned} \iota(r, q; \alpha, \lambda)K + C &= \int_0^r v_r(s, 0)ds + \int_0^q v_q(r, p)dp \\ &= -2\mathbb{E}_{\xi, H} \ln \int dP_X(\eta) e^{(\xi\sqrt{r}+Hr)\eta - \frac{1}{2}r\eta^2} \\ &\quad + rq + \frac{1}{2\alpha} \ln(1 + \lambda\alpha(1 - q^2)), \end{aligned}$$

where we have introduced two constants  $K$  and  $C$  to be fixed later. The term  $\mathbb{E}H\langle \eta \rangle_s = 1 - \text{MMSE}(H|\sqrt{s}H + \xi)$  has been integrated with respect to  $s$  using the I-MMSE relation for a scalar Gaussian channel [82]. It is easy to check that this potential generates the correct saddle point equations. Now we evaluate it at their solutions. This will yield an MI whose basic properties will help us to fix  $K$  and  $C$ . In particular, at  $\lambda = 0$  one has  $q_* = r_* = 0$  and  $\iota(r_*, q_*; \alpha, 0) = 0$ , which directly entails  $C = 0$ . In order to fix  $K$ , we compute the derivative

$$\frac{d}{d\lambda} \iota(r_*, q_*; \alpha, \lambda) = \frac{1}{2K} \frac{1 - q_*(\alpha, \lambda)^2}{1 + \lambda\alpha(1 - q_*(\alpha, \lambda)^2)}.$$

The I-MMSE relation Eq. (D1) then requires  $K = 2$ . This determines Eq. (34).

### 2. Large SNR limit of Eq. (35) for discrete priors

With a discrete prior, one can see that  $1 - q^2 = o_\lambda(\lambda^{-1}) = 1 - q$ . This must be verified in order for  $r$  in Eq. (32) to diverge as  $\sim \lambda$ , so that  $q$  indeed saturates to 1 in the infinite SNR limit. Furthermore, the  $\mathbb{E}_{\xi, H} \dots$  term of Eq. (34) is related to a Bayes-optimal scalar Gaussian channel with SNR  $r$ . Hence, when  $r \sim \lambda \rightarrow \infty$ , the integral over the prior can be done by saddle point and must select  $\eta = H$ ,

$$\begin{aligned} \mathbb{E}_{\xi, H} \ln \int dP_X(\eta) e^{(\xi\sqrt{r}+Hr)\eta - \frac{1}{2}r\eta^2} \\ = \frac{r}{2} + \mathbb{E}_H \ln \int dP_X(\eta) \mathbb{1}(\eta = H) + o_\lambda(1). \end{aligned}$$

Here,  $\mathbb{1}(\eta = H)$  is the indicator function of the event  $\eta = H$  (recall that  $H$  is discrete in this derivation). The integral on the second line is thus precisely  $-\mathcal{H}(X)$ . Plugging this into Eq. (34), we get

$$\iota(\alpha, \lambda) = \frac{r(q-1)}{2} + o_\lambda(1) + \frac{1}{4\alpha} \ln(1 + o_\lambda(1)) + \mathcal{H}(X).$$

When  $\lambda \rightarrow \infty$ ,  $r \sim \lambda$  at leading order and both  $\lambda(1 - q^2)$ ,  $\lambda(1 - q)$  are  $o_\lambda(1)$ . Hence,

$$\lim_{\lambda \rightarrow \infty} \iota(\alpha, \lambda) = \mathcal{H}(X).$$

## APPENDIX E: MEAN-SQUARE ERROR OF THE ROTATIONAL-INVARIANT ESTIMATOR WHEN $\alpha \rightarrow 0$

Consider the usual problem,

$$\frac{\mathbf{Y}}{\sqrt{N}} = \frac{\sqrt{\lambda}}{N} \mathbf{X} \mathbf{X}^\top + \frac{\mathbf{Z}}{\sqrt{N}},$$

where we rescaled the observations so that they have  $O(1)$  eigenvalues and a well-defined, finitely supported, asymptotic spectral density  $\rho_Y$ . Standard random matrix theory [59] shows that  $\rho_Y = \rho_{\text{s.c.}} \boxplus \rho_{\sqrt{\lambda}\text{MP}}(\cdot; \alpha)$  is the free convolution between the semicircular density and the spectral density of a Wishart matrix (Marchenko-Pastur distribution of parameter  $1/\alpha$ ) multiplied by  $\sqrt{\lambda}$ . We recall once more the shrinking procedure of Ref. [84] needed to get the eigenvalues of the RIE,

$$\xi_i = \xi_i(\mathbf{Y}) := \frac{\gamma_{\mathbf{Y},i} - 2\pi\mathcal{H}[\rho_Y](\gamma_{\mathbf{Y},i})}{\sqrt{\lambda}}, \quad (\text{E1})$$

where  $(\gamma_{\mathbf{Y},i})$  are the eigenvalues of  $\mathbf{Y}/\sqrt{N}$ , and

$$\pi\mathcal{H}[\rho_Y](x) = \text{P.V.} \int dy \frac{\rho_Y(y)}{x-y} = \lim_{\epsilon \rightarrow 0^+} \text{Reg}_Y(x + i\epsilon).$$

$g_Y$  is the Stieltjes transform of the spectral density  $\rho_Y$  [59],

$$g_Y(z) := \int dy \frac{\rho_Y(y)}{z-y}, \quad z \in \mathbb{C} \setminus \text{Supp}(\rho_Y).$$

We consider the low-rank limit  $\alpha \rightarrow 0$ . For simplicity, we set  $M = 1$ , but the argument below can be extended to any  $M$  finite, and we claim it can be extended to sublinear rank regimes  $M = \alpha N^\gamma$  with  $\gamma < 1$ . Under this hypothesis,  $g_Y = g_Z$ , since low-rank perturbations are not enough to modify the spectral density of the noise.  $\mathbf{Y}$  then turns into a low-rank perturbation of a full-rank Wigner matrix. Hence, some eigenvalues (one for  $M = 1$ ) will pop out of the bulk spectrum of the noise.

The eigenvalue(s) that pops out, for  $\lambda > 1$ , is located at  $z^* = \sqrt{\lambda} + (1/\sqrt{\lambda})$  according to our notation [23]. The Stieltjes transform of the observations evaluated at  $z^*$  is thus

$$\begin{aligned} g_Y(z^*) \xrightarrow{\alpha \rightarrow 0} g_Z(z^*) &= \frac{1}{2} \left( z^* - z^* \sqrt{1 - 4/z^{*2}} \right) \\ &= \frac{1}{2} \left( \sqrt{\lambda} + \frac{1}{\sqrt{\lambda}} - \sqrt{\lambda - 2 + \frac{1}{\lambda}} \right) = \frac{1}{\sqrt{\lambda}}, \end{aligned}$$

and hence

$$\pi \mathbf{H}[\rho_Y](z^*) \xrightarrow{\alpha \rightarrow 0} \frac{1}{\sqrt{\lambda}}.$$

From the above, we can infer the shrinkage of the leading eigenvalue,

$$\xi_1 = \frac{1}{\sqrt{\lambda}} \left( \sqrt{\lambda} + \frac{1}{\sqrt{\lambda}} - \frac{2}{\sqrt{\lambda}} \right) = 1 - \frac{1}{\lambda}.$$

All other eigenvalues are flushed to 0. In fact, we can safely assume that, under low-rank perturbations of the noise bulk, all eigenvalues except the leading one lie inside the bulk. In that case, thanks to the identity

$$\text{P.V.} \int dy \frac{\rho_Z(y)}{x-y} = \frac{x}{2}, \quad x \in \mathbb{R},$$

it is straightforward to see that Eq. (E1) yields 0 for all such eigenvalues.

Therefore, the estimator given by the RIE in the low-rank limit (recall  $M = 1$ ) would read

$$\left\langle \frac{\mathbf{X}\mathbf{X}^\top}{N} \right\rangle_{\mathbf{Y}} = \mathbf{u}_{\mathbf{Y}} \mathbf{u}_{\mathbf{Y}}^\top \left( 1 - \frac{1}{\lambda} \right),$$

where  $\mathbf{u}_{\mathbf{Y}}$  is the unit eigenvector corresponding to the leading eigenvalue  $z^*$  of  $\mathbf{Y}$ . This estimator is also Bayes optimal when the prior on  $\mathbf{X}$  is rotational invariant, and the MSE attained is

$$\text{MSE} = 1 - \left( 1 - \frac{1}{\lambda} \right)^2.$$

This MSE is attained by the RIE regardless of the prior and is equal to that of PCA [16], which is suboptimal for general prior other than Gaussian. This argument thus supports the suboptimality of the RIE in the limit  $\alpha \rightarrow 0$ , unless there is a discontinuous behavior in  $\alpha$ .

## APPENDIX F: THE SAKATA AND KABASHIMA REPLICAS APPROACH TO EQ. (35)

Here we adapt to our model the method of Ref. [99] developed originally for dictionary learning in an optimization setting, then concurrently extended in Refs. [100–102] to the Bayesian case. In order to match notations, we consider the generic output channel

$$Y_{ij} \sim P_{\text{out}} \left( \cdot \mid \frac{\mathbf{X}_i \cdot \mathbf{X}_j}{\sqrt{N}} \right), \quad 1 \leq i < j \leq N.$$

The free entropy we are interested in computing, to be later connected to the MI, is

$$\tilde{\phi}_N := \frac{1}{MN} \mathbb{E} \ln \tilde{\mathcal{Z}}(\mathbf{Y}) - \frac{1}{MN} \mathbb{E} \ln \prod_{i < j} P_{\text{out}}(Y_{ij} | 0),$$

where

$$\tilde{\mathcal{Z}}(\mathbf{Y}) := \int dP_{\mathbf{X}}(\mathbf{x}) \prod_{i < j, 1}^N P_{\text{out}} \left( Y_{ij} \mid \frac{\mathbf{x}_i \cdot \mathbf{x}_j}{\sqrt{N}} \right).$$

We proceed to the computation of the log-partition function with the replica method,

$$\mathbb{E} \tilde{\mathcal{Z}}^n(\mathbf{Y}) = \int \prod_{i < j}^N dY_{ij} \int \prod_{a=0}^n dP_{\mathbf{X}}(\mathbf{x}^a) \prod_{i < j, 1}^N P_{\text{out}} \left( Y_{ij} \mid \frac{\mathbf{x}_i^a \cdot \mathbf{x}_j^a}{\sqrt{N}} \right),$$

where we set  $\mathbf{X}^0 = \mathbf{X}$ . We now introduce a constraint on the overlaps

$$\begin{aligned} \mathbb{E} \tilde{\mathcal{Z}}^n(\mathbf{Y}) &\propto \int \prod_{i < j}^N dY_{ij} \int \prod_{a=0}^n dP_{\mathbf{X}}(\mathbf{x}^a) \\ &\times \int \prod_{a \leq b=0}^n dQ_{ab} \delta(NMQ_{ab} - \text{Tr}[\mathbf{x}^a \mathbf{x}^{b\top}]) \\ &\times \prod_{a=0}^n \prod_{i < j, 1}^N P_{\text{out}} \left( Y_{ij} \mid \frac{\mathbf{x}_i^a \cdot \mathbf{x}_j^a}{\sqrt{N}} \right). \end{aligned}$$

Now we aim at identifying the joint law of the variables

$$z_{ij}^a = \frac{\mathbf{x}_i^a \cdot \mathbf{x}_j^a}{\sqrt{N}}, \quad 1 \leq i < j \leq N.$$

More precisely, let us introduce a  $\delta$  function to isolate these variables

$$\int d\mathbf{Y} \int \prod_{a \leq b=0}^n dQ_{ab} \int \prod_{a=0}^n dP_X(\mathbf{x}^a) \prod_{a \leq b=0}^n \delta(NMQ_{ab} - \text{Tr}[\mathbf{x}^a \mathbf{x}^{b\top}]) \prod_{i < j, 1}^N \prod_{a=0}^n \delta\left(z_{ij}^a - \frac{\mathbf{x}_i^a \cdot \mathbf{x}_j^a}{\sqrt{N}}\right) P_{\text{out}}(Y_{ij} | z_{ij}^a)$$

and let us isolate their non-normalized probability measure for a given value of the overlaps  $Q_{ab}$ ,

$$\int \prod_{a=0}^n dP_X(\mathbf{x}^a) \prod_{a \leq b=0}^n \delta(NMQ_{ab} - \text{Tr}[\mathbf{x}^a \mathbf{x}^{b\top}]) \prod_{i < j, 1}^N \prod_{a=0}^n \delta\left(z_{ij}^a - \frac{\mathbf{x}_i^a \cdot \mathbf{x}_j^a}{\sqrt{N}}\right).$$

The key assumption is then to consider the  $z$ 's as a Gaussian family with zero mean and a covariance equal to

$$\mathbb{E}_z z_{ij}^a z_{kl}^b = \alpha Q_{ab}^2 \delta_{ik} \delta_{jl}.$$

This Gaussian *Ansatz* allows us to move forward with the replica approach. With the above covariance structure, coupled only in replica space, we can write

$$P_z((z_{ij}^a)_{a \leq n} | (Q_{ab})_{a,b=0}^n) = \frac{1}{\sqrt{(2\pi)^{n+1} \det \mathcal{T}}} \exp\left[-\frac{1}{2} \sum_{a,b=0}^n z_{ij}^a z_{ij}^b (\mathcal{T}^{-1})_{ab}\right]$$

with  $\mathcal{T} := (\alpha Q_{ab}^2)_{a,b=0}^n$ . Hence, we can replace

$$\int \prod_{a=0}^n dP_X(\mathbf{x}^a) \prod_{a \leq b=0}^n \delta(NMQ_{ab} - \text{Tr}[\mathbf{x}^a \mathbf{x}^{b\top}]) \prod_{i < j, 1}^N \prod_{a=0}^n \delta\left(z_{ij}^a - \frac{\mathbf{x}_i^a \cdot \mathbf{x}_j^a}{\sqrt{N}}\right) \approx P_z((z_{ij}^a)_{a \leq n} | (Q_{ab})_{a,b=0}^n) e^{\ln V(Q)},$$

where

$$V(Q) := \int \prod_{a=0}^n dP_X(\mathbf{x}^a) \prod_{a \leq b=0}^n \delta(NMQ_{ab} - \text{Tr}[\mathbf{x}^a \mathbf{x}^{b\top}])$$

is a normalization. Therefore, the replicated partition function reads

$$\mathbb{E} \tilde{Z}^n(\mathbf{Y}) \propto \int \prod_{a \leq b=0}^n dQ_{ab} \exp(\ln V(Q)) \left[ \mathbb{E}_{\mathbf{Z} \sim \mathcal{N}(0; \mathcal{T})} \int dY \prod_{a=0}^n P_{\text{out}}(Y | z^a) \right]^{N^2/2}.$$

Letting  $N$  and  $M$  diverge, we get the following expression by saddle point with respect to the integration variables  $(Q_{ab})_{0 \leq a, b \leq n}$ :

$$\frac{1}{MN} \ln \mathbb{E} \tilde{Z}^n(\mathbf{Y}) = \text{extr}_+ \left\{ \frac{1}{MN} \ln V(Q) + \frac{1}{2\alpha} \ln \mathbb{E}_{\mathbf{Z} \sim \mathcal{N}(0; \mathcal{T})} \int dY \prod_{a=0}^n P_{\text{out}}(Y | z^a) \right\} + o_N(1).$$

Now, we specify our saddle point to the replica symmetric *Ansatz*:  $\mathcal{T}_{RS} = \alpha(1 - q^2)\delta_{ab} + \alpha q^2$ , and we then send  $n \rightarrow 0$ . Here, we have further imposed that  $Q_{a0} = q = Q_{ab}$  for any  $a \neq b$ , and  $v = Q_{aa}$  for all  $a = 0, \dots, n$ , which amounts to imposing the Nishimori identities on the solution of the saddle point. This is justified by the Bayes-optimal setting. The computation of  $V(Q)$  also requires the introduction of Fourier  $\delta$  representations. These steps are standard, and we defer the reader to Ref. [100] for this computation, since for where  $V(Q)$  is concerned, it is identical. In the end, one gets

$$\frac{1}{MN} \ln V(Q) = -n \text{extr}_+ \left\{ \frac{rq}{2} + \mathbb{E} \ln \int dP_X(x) \exp\left[(\sqrt{r}Z + rX)x - \frac{r}{2}x^2\right] \right\} + O(n^2),$$

where extremization is with respect to  $r$ . The other term is easily dealt with,

$$\frac{1}{\alpha} \ln \mathbb{E}_{\mathbf{Z} \sim \mathcal{N}(0; \mathcal{T})} \int dY \prod_{a=0}^n P_{\text{out}}(Y|z^a) = \frac{n}{\alpha} \mathbb{E}_{u_0, \zeta} \int dY P_{\text{out}}\left(Y \left| \sqrt{\alpha(v^2 - q^2)}u_0 + \sqrt{\alpha q^2 \zeta} \right.\right) \ln \mathbb{E}_u P_{\text{out}}\left(Y \left| \sqrt{\alpha(v^2 - q^2)}u + \sqrt{\alpha q^2 \zeta} \right.\right) + O(n^2).$$

We reach

$$\frac{1}{MNn} \ln \mathbb{E} \tilde{Z}^n(\mathbf{Y}) = \text{extr}_+ \left\{ -\frac{rq}{2} + \mathbb{E} \ln \int dP_X(x) \exp \left[ (\sqrt{r}Z + rX)x - \frac{r}{2}x^2 \right] + \frac{1}{2\alpha} \mathbb{E}_{u_0, \zeta} \int dY P_{\text{out}}\left(Y \left| \sqrt{\alpha(1 - q^2)}u_0 + \sqrt{\alpha q^2 \zeta} \right.\right) \ln \mathbb{E}_u P_{\text{out}}\left(Y \left| \sqrt{\alpha(1 - q^2)}u + \sqrt{\alpha q^2 \zeta} \right.\right) \right\} + o_N(1) + O(n),$$

where extremization is always intended with respect to  $r, q$ . Specifying to Gaussian channel

$$P_{\text{out}}(y|x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(y - \sqrt{\lambda}x)^2\right),$$

one gets

$$\text{extr}_+ \left\{ -\frac{rq}{2} + \mathbb{E} \ln \int dP_X(x) \exp \left[ (\sqrt{r}Z + rX)x - \frac{r}{2}x^2 \right] - \frac{1}{4\alpha} \ln(1 + \lambda\alpha(1 - q^2)) - \frac{1}{4\alpha} \log 2\pi e \right\} + O(n)$$

and we can now subtract the term

$$\frac{1}{MN} \mathbb{E} \ln \prod_{i < j}^N P_{\text{out}}(Y_{ij}|0) = \frac{1}{2\alpha} \mathbb{E} \ln P_{\text{out}}(Y|0) + o_N(1) = -\frac{1}{4\alpha} \log 2\pi e - \frac{\lambda}{4}.$$

Putting it all together, we get the free entropy

$$\tilde{\phi} = \lim_{N \rightarrow \infty} \tilde{\phi}_N = \text{extr}_+ \left\{ \frac{\lambda}{4} - \frac{rq}{2} + \mathbb{E} \ln \int dP_X(x) \exp \left[ (\sqrt{r}Z + rX)x - \frac{r}{2}x^2 \right] - \frac{1}{4\alpha} \ln(1 + \lambda\alpha(1 - q^2)) \right\}.$$

For Gaussian channel, in order to obtain the MI it is sufficient to subtract  $\lambda/4$  from the above and flip the sign,

$$\iota(\alpha, \lambda) = \text{extr}_- \left\{ \frac{rq}{2} + \frac{1}{4\alpha} \ln(1 + \lambda\alpha(1 - q^2)) - \mathbb{E} \ln \int dP_X(x) \exp \left[ (\sqrt{r}Z + rX)x - \frac{r}{2}x^2 \right] \right\},$$

where  $r$  and  $q$  have to be extremized. This formula matches Eq. (35).

### APPENDIX G: CAN THE MEAN-FIELD THEORY $\iota(\alpha, \lambda)$ BE IMPROVED?

We have shown that the mean-field Eqs. (35) based on the bulk Ansatz Eq. (25) produce excellent results in the factorization phase, where an MMSE lower than the Gaussian one is possible; they are, however, only approximate in the denoising phase where  $i^*$  correct them. One may wonder if it is possible to do better by generalizing Ansatz Eq. (25). We have performed several attempts, keeping the bulk measure factorized over the indices  $(i, \mu)$  as needed to justify the computation of the RLE free entropy.

As first trial, one may think that each entry  $(i, \mu)$  comes with a different SNR  $\sigma_{i\mu}$ , whose values are i.i.d. drawn from a certain distribution that becomes the new order parameter. In that case, Eqs. (30) simply contain  $\mathbb{E}_\sigma J(\sigma)$  in place of  $J(\sigma)$ . In particular, the equation for  $r$  forces the distribution of  $\sigma$  to be a  $\delta_{\sigma_*}$  where again  $\sigma_* = r_*$ , since the rhs is deterministic. Hence, this approach recovers the proposed solution.

Another approach is to introduce an SNR profile  $(\sigma_\mu)$  such that  $Y'_{i\mu} = \sqrt{\sigma_\mu} X_{i\mu} + Z'_{i\mu}$ , labeled only by  $\mu$  to detect possible inhomogeneities at equilibrium in the ‘‘column index space.’’ The  $(\sigma_\mu)$  are then assumed to be divided in an arbitrary number  $C$  of blocks to be sent to  $+\infty$  after the thermodynamic limit, in a similar fashion to spatial coupling, see Refs. [158,159] for similar computations.

Specifically,  $\sigma_\mu = \sigma_c$  if  $\mu \in \Lambda_c \subset \{1, \dots, M\}$ . The computation in Sec. IV (and in Appendix C) still works, except that this time we need to introduce more order parameters  $(r_c, q_c)_{c=1, \dots, C}$ , where  $q_c$  corresponds to the overlap  $(C/M) \sum_{\mu \in \Lambda_c} \eta_\mu H_\mu$ . These order parameters then satisfy the equations

$$q_c = \mathbb{E}H\langle \eta \rangle_{r_c}, \quad r_c = \frac{\lambda J(\sigma_c)}{1 + \lambda \alpha (1 - \frac{1}{C} \sum_{c=1}^C J(\sigma_c) q_c)},$$

together with the cavity-bulk consistency equation  $\sigma_c = r_c$ . This system always admits the paramagnetic solution  $\sigma_c = r_c = q_c = 0$  for any given set of  $c$ 's and the homogeneous solution  $(\sigma_c, r_c, q_c) = (\sigma_*, r_*, q_*)$  for all  $c$ 's, which also coincides with the solution given by the Ansatz Eq. (25). No other solution seems to exist. Interestingly, the latter seems to be always the most convenient in terms of MMSE, since we see numerically that in that case all the SNR values  $\sigma_c$  get attracted to a higher SNR  $\sigma_*$  with respect to the case when some  $\sigma_c$  are set to 0. Furthermore, if one interprets the above system of equations as a fixed-point equation for  $\sigma_c$ ,

$$\sigma_c = \frac{\lambda J(\sigma_c)}{1 + \lambda \alpha (1 - \frac{1}{C} \sum_{c=1}^C J(\sigma_c)^2)},$$

the solutions where some  $\sigma_c = 0$  are unstable under such iteration which discards them. One can also run a population dynamics [124] using the above equation interpreted as a distributional fixed-point equation, which thus looks for a solution where the law of the  $(\sigma_\mu)$  is the same for all  $\mu$ , rather than to look for solutions for the ordered profile  $(\sigma_c)$  shared by all rows. Again, the final histogram of  $(\sigma_\mu)$  ends up peaked around the value of the homogeneous solution  $\sigma_*$ .

A final attempt we made consists of admitting a block structure also in the  $i$  index, i.e., the row index:  $\sigma_{i\mu} = \sigma_{\ell c}$  for  $i \in \Lambda_{\ell'}, \mu \in \Lambda_c$ , with  $\ell' = 1, \dots, L$  and  $c$  as before. With this block structure, one can recast the fixed-point equations in terms of the same order parameters as above with some slight differences,

$$q_{\ell'c} = \mathbb{E}H\langle \eta \rangle_{r_{\ell'c}},$$

$$r_{\ell'c} = \frac{1}{L} \sum_{\ell'=1}^L \frac{\lambda J(\sigma_{\ell'c})}{1 + \lambda \alpha (1 - \frac{1}{C} \sum_{c=1}^C J(\sigma_{\ell'c}) q_c)},$$

where  $\ell'$  corresponds to the index of the row block from which the cavity was extracted. They are completed with the cavity-bulk consistency  $\sigma_{\ell'c} = r_{\ell'c}$ . All this must be true for all  $\ell' \in \{1, \dots, L\}$  as the cavity can be extracted from any block. Because the rhs of the  $r_{\ell'c} = \dots$  equation is  $\ell'$  independent, solutions must be homogeneous in the row-block indices and we are back to the previous case.

In conclusion, none of these possible generalizations of Eq. (25) yield an improvement with respect to our current Ansatz. This strengthens our belief that the mean-field theory is exact in the factorization phase, as it is robust to richer Ansätze as long as the bulk measure is assumed to be factorized.

- [1] I. Tošić and P. Frossard, *Dictionary learning*, *IEEE Signal Process. Mag.* **28**, 27 (2011).
- [2] Y. Bengio, A. Courville, and P. Vincent, *Representation learning: A review and new perspectives*, *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1798 (2013).
- [3] B. A. Olshausen and D. J. Field, *Emergence of simple-cell receptive field properties by learning a sparse code for natural images*, *Nature (London)* **381**, 607 (1996).
- [4] B. A. Olshausen and D. J. Field, *Sparse coding with an overcomplete basis set: A strategy employed by v1?*, *Vision Res.* **37**, 3311 (1997).
- [5] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski, *Dictionary learning algorithms for sparse representation*, *Neural Comput.* **15**, 349 (2003).
- [6] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09 (Association for Computing Machinery, New York, NY, 2009), pp. 689–696.
- [7] E. J. Candès, X. Li, Y. Ma, and J. Wright, *Robust principal component analysis?*, *J. ACM* **58**, 1 (2011).
- [8] A. Perry, A. S. Wein, A. S. Bandeira, and A. Moitra, *Optimality and sub-optimality of PCA I: Spiked random matrix models*, *Ann. Stat.* **46**, 2416 (2018).
- [9] B. Hajek, Y. Wu, and J. Xu, *Information limits for recovering a hidden community*, *IEEE Trans. Inf. Theory* **63**, 4729 (2017).
- [10] A. Belouchrani, K. Abed-Meraim, J.-F. Cardoso, and E. Moulines, *A blind source separation technique using second-order statistics*, *IEEE Trans. Signal Process.* **45**, 434 (1997).
- [11] E. J. Candès and B. Recht, *Exact matrix completion via convex optimization*, *Found. Comput. Math.* **9**, 717 (2009).
- [12] E. J. Candès and T. Tao, *The power of convex relaxation: Near-optimal matrix completion*, *IEEE Trans. Inf. Theory* **56**, 2053 (2010).
- [13] A. Mnih and R. R. Salakhutdinov, *Advances in Neural Information Processing Systems*, edited by J. Platt, D. Koller, Y. Singer, and S. Roweis (Curran Associates, Inc., 2007), Vol. 20, [https://papers.nips.cc/paper\\_files/paper/2007/file/d7322ed717dedf1eb4e6e52a37ea7bcd-Paper.pdf](https://papers.nips.cc/paper_files/paper/2007/file/d7322ed717dedf1eb4e6e52a37ea7bcd-Paper.pdf).
- [14] J. Mairal, M. Elad, and G. Sapiro, *Sparse representation for color image restoration*, *IEEE Trans. Image Process.* **17**, 53 (2008).
- [15] E. Abbe, *Community detection and stochastic block models: Recent developments*, *J. Mach. Learn. Res.* **18**, 6446 (2017), <https://jmlr.org/papers/v18/16-480.html>.

- [16] M. Lelarge and L. Miolane, *Fundamental limits of symmetric low-rank matrix estimation*, *Probab. Theory Relat. Fields* **173**, 859 (2019).
- [17] F. Caltagirone, M. Lelarge, and L. Miolane, *Recovering asymmetric communities in the stochastic block model*, *IEEE Trans. Netw. Sci. Eng.* **5**, 237 (2017).
- [18] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Clarendon Press, Oxford, 2001), <https://academic.oup.com/book/5185>.
- [19] D. Sherrington and S. Kirkpatrick, *Solvable model of a spin-glass*, *Phys. Rev. Lett.* **35**, 1792 (1975).
- [20] L. Zdeborová and F. Krzakala, *Statistical physics of inference: Thresholds and algorithms*, *Adv. Phys.* **65**, 453 (2016).
- [21] M. Mézard, G. Parisi, and M.-A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific Publishing Co., Inc., Pergamon Press, New York, 1990).
- [22] D. Panchenko, *The Sherrington-Kirkpatrick Model*, Springer Monographs in Mathematics (Springer, New York, NY, 2013).
- [23] J. Baik, G. B. Arous, and S. Pécché, *Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices*, *Ann. Probab.* **33**, 1643 (2005).
- [24] J. Baik and J. W. Silverstein, *Eigenvalues of large sample covariance matrices of spiked population models*, *J. Multivariate Anal.* **97**, 1382 (2006).
- [25] F. Benaych-Georges and R. R. Nadakuditi, *The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices*, *Adv. Math.* **227**, 494 (2011).
- [26] I. Johnstone, *On the distribution of the largest eigenvalue in principal components analysis*, *Ann. Stat.* **29**, 295 (2001).
- [27] I. Johnstone and A. Lu, *Sparse principal components analysis* (unpublished).
- [28] H. Zou, T. Hastie, and R. Tibshirani, *Sparse principal component analysis*, *J. Comput. Graph. Stat.* **15**, 265 (2006).
- [29] I. Johnstone and A. Lu, *On consistency and sparsity for principal components analysis in high dimensions*, *J. Am. Stat. Assoc.* **104**, 682 (2012).
- [30] J. Barbier, M. Dia, N. Macris, F. Krzakala, T. Lesieur, and L. Zdeborová, *Proceedings of the 30th International Conference on Neural Information Processing Systems*, NIPS'16 (Curran Associates, Inc., Red Hook, New York, 2016), pp. 424–432.
- [31] L. Miolane, *Fundamental Limits of Inference: A Statistical Physics Approach*, Ph.D. thesis, Ecole normale supérieure-ENS PARIS, Inria, Paris (2019).
- [32] T. Lesieur, F. Krzakala, and L. Zdeborová, *MMSE of probabilistic low-rank matrix estimation: Universality with respect to the output channel*, *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)* (IEEE, New York, NY, 2015), pp. 680–687.
- [33] T. Lesieur, F. Krzakala, and L. Zdeborová, *Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications*, *J. Stat. Mech.* (2017) 073403.
- [34] Y. Deshpande and A. Montanari, *Information-theoretically optimal sparse PCA*, in *2014 IEEE International Symposium on Information Theory, Honolulu, HI* (IEEE, New York, 2014), pp. 2197–2201, [10.1109/ISIT.2014.6875223](https://doi.org/10.1109/ISIT.2014.6875223).
- [35] Y. Deshpande, E. Abbe, and A. Montanari, *Asymptotic mutual information for the balanced binary stochastic block model*, *Inf. Infer.* **6**, 125 (2017).
- [36] J. Barbier, M. Dia, N. Macris, F. Krzakala, and L. Zdeborová, *Rank-one matrix estimation: Analysis of algorithmic and information theoretic limits by the spatial coupling method*, [arXiv:1812.02537](https://arxiv.org/abs/1812.02537) (2018).
- [37] S. Korada and N. Macris, *Exact solution of the gauge symmetric p-spin glass model on a complete graph*, *J. Stat. Phys.* **136**, 205 (2009).
- [38] J. Barbier and N. Macris, *The adaptive interpolation method: A simple scheme to prove replica formulas in Bayesian inference*, *Probab. Theory Relat. Fields* **174**, 1133 (2019).
- [39] J. Barbier and N. Macris, *The adaptive interpolation method for proving replica formulas. Applications to the Curie-Weiss and Wigner spike models*, *J. Phys. A* **52**, 294002 (2019).
- [40] G. Reeves, *Information-theoretic limits for the matrix tensor product*, *IEEE J. Sel. Areas Inf. Theory* **1**, 777 (2020).
- [41] M. Aizenman, R. Sims, and S. L. Starr, *Extended variational principle for the Sherrington-Kirkpatrick spin-glass model*, *Phys. Rev. B* **68**, 214403 (2003).
- [42] A. El Alaoui and F. Krzakala, *Estimation in the spiked Wigner model: A short proof of the replica formula*, in *2018 IEEE International Symposium on Information Theory (ISIT)* (IEEE, New York, 2018), pp. 1874–1878, [10.1109/ISIT.2018.8437810](https://doi.org/10.1109/ISIT.2018.8437810).
- [43] J.-C. Mourrat, *Hamilton-Jacobi equations for finite-rank matrix inference*, *Ann. Appl. Probab.* **30**, 2234 (2020).
- [44] H. Chen, J.-C. Mourrat, and J. Xia, *Statistical inference of finite-rank tensors*, *Ann. Henri Lebesgue* **5**, 1161 (2022).
- [45] Y. Kabashima, *A cdma multiuser detection algorithm on the basis of belief propagation*, *J. Phys. A* **36**, 11111 (2003).
- [46] D. L. Donoho, A. Maleki, and A. Montanari, *Message-passing algorithms for compressed sensing*, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 18914 (2009).
- [47] S. Rangan and A. Fletcher, *Iterative estimation of constrained rank-one matrices in noise*, in *IEEE International Symposium on Information Theory* (IEEE, New York, NY, 2012), pp. 1246–1250.
- [48] J. Barbier, F. Krzakala, N. Macris, L. Miolane, and L. Zdeborová, *Optimal errors and phase transitions in high-dimensional generalized linear models*, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 5451 (2019).
- [49] M. Brennan, G. Bresler, and W. Huleihel, *Reducibility and computational lower bounds for problems with planted sparse structure*, in *Conference on Learning Theory* (PMLR, Stockholm, Sweden, 2018), pp. 48–166.
- [50] M. Brennan and G. Bresler, *Reducibility and statistical-computational gaps from secret leakage*, in *Conference on Learning Theory* (PMLR, Graz, Austria, 2020), pp. 648–847.

- [51] J. Barbier, F. Camilli, M. Mondelli, and M. Sáenz, *Fundamental limits in structured principal component analysis and how to reach them*, *Proc. Natl. Acad. Sci. U.S.A.* **120**, e2302028120 (2023).
- [52] J. Barbier, F. Camilli, Y. Xu, and M. Mondelli, *Information limits and Thouless-Anderson-Palmer equations for spiked matrix models with structured noise*, *Phys. Rev. Res.* **7**, 013081 (2025).
- [53] J. Barbier, T. Hou, M. Mondelli, and M. Sáenz, *The price of ignorance: How much does it cost to forget noise structure in low-rank matrix estimation?*, *Adv. Neural Inf. Process. Syst.* **35**, 36733 (2022), [https://papers.nips.cc/paper\\_files/paper/2022/hash/ee74a6ade401e200985e242-1b20bbae4-Abstract-Conference.html](https://papers.nips.cc/paper_files/paper/2022/hash/ee74a6ade401e200985e242-1b20bbae4-Abstract-Conference.html).
- [54] J. Barbier, J. Ko, and A. A. Rahman, *Information-theoretic limits for sublinear-rank symmetric matrix factorization*, in International Zurich Seminar on Information and Communication (IZS 2024) (ETH Zurich, Zurich, Switzerland, 2024) p. 16.
- [55] D. Alberici, F. Camilli, P. Contucci, and E. Mingione, *The multi-species mean-field spin-glass on the Nishimori line*, *J. Stat. Phys.* **182**, 2 (2021).
- [56] D. Alberici, F. Camilli, P. Contucci, and E. Mingione, *The solution of the deep Boltzmann machine on the Nishimori line*, *Commun. Math. Phys.* **387**, 1 (2021).
- [57] Y. Zhang and M. Mondelli, *Matrix denoising with doubly heteroscedastic noise: Fundamental limits and optimal spectral methods*, *Adv. Neural Inf. Process. Syst.* **37**, 93060 (2024), [https://proceedings.neurips.cc/paper\\_files/paper/2024/hash/a96368eb38bce0956a1132154d70d72d-Abstract-Conference.html](https://proceedings.neurips.cc/paper_files/paper/2024/hash/a96368eb38bce0956a1132154d70d72d-Abstract-Conference.html).
- [58] F. Pourkamali, J. Barbier, and N. Macris, *Matrix inference in growing rank regimes*, *IEEE Trans. Inf. Theory* **70**, 8133 (2024).
- [59] M. Potters and J.-P. Bouchaud, *A First Course in Random Matrix Theory: For Physicists, Engineers and Data Scientists* (Cambridge University Press, Cambridge, England, 2020).
- [60] A. A. Migdal, *Recursion equations in gauge theories*, *Sov. Phys. JETP* **42**, 413 (1975).
- [61] V. A. Kazakov, M. Staudacher, and T. Wynter, *Exact solution of discrete two-dimensional  $r_2$  gravity*, *Nucl. Phys.* **B471**, 309 (1996).
- [62] Harish-Chandra, *Differential operators on a semisimple lie algebra*, *Am. J. Math.* **79**, 87 (1957).
- [63] C. Itzykson and J.-B. Zuber, *The planar approximation. II*, *J. Math. Phys. (N.Y.)* **21**, 411 (1980).
- [64] V. A. Kazakov and A. A. Migdal, *Induced gauge theory at large  $n$* , *Nucl. Phys.* **B397**, 214 (1993).
- [65] J. Ambjørn, B. Durhuus, and T. Jonsson, *Quantum Geometry: A Statistical Field Theory Approach*, Cambridge Monographs on Mathematical Physics (Cambridge Univ. Press, Cambridge, England, 2005).
- [66] P. Francesco, P. Ginsparg, and J. Zinn-Justin, *2d gravity and random matrices*, *Phys. Rep.* **254**, 1 (1995).
- [67] J. Verbaarschot and T. Wettig, *Random matrix theory and chiral symmetry in QCD*, *Annu. Rev. Nucl. Part. Sci.* **50**, 343 (2000).
- [68] V. A. Kazakov, *Ising model on a dynamical planar random lattice: Exact solution*, *Phys. Lett. A* **119**, 140 (1986).
- [69] V. A. Kazakov, M. Staudacher, and T. Wynter, *Almost flat planar diagrams*, *Commun. Math. Phys.* **179**, 235 (1996).
- [70] D. V. Boulatov and V. A. Kazakov, *The Ising model on a random planar lattice: The structure of the phase transition and the exact critical exponents*, *Phys. Lett. B* **186**, 379 (1987).
- [71] V. A. Kazakov, *Exactly solvable Potts models, bond- and tree-like percolation on dynamical (random) planar lattice*, *Nucl. Phys. B, Proc. Suppl.* **4**, 93 (1988).
- [72] E. Brezin, C. Itzykson, G. Parisi, and J. B. Zuber, *Planar diagrams*, *Commun. Math. Phys.* **59**, 35 (1978).
- [73] A. Zvonkin, *Matrix integrals and map enumeration: An accessible introduction*, *Math. Comp. Model.* **26**, 281 (1997).
- [74] V. A. Kazakov, *Solvable matrix models*, [arXiv:hep-th/0003064](https://arxiv.org/abs/hep-th/0003064) (2000).
- [75] J. Ba, M. A. Erdogdu, T. Suzuki, Z. Wang, D. Wu, and G. Yang, *High-dimensional asymptotics of feature learning: How one gradient step improves the representation*, *Adv. Neural Inf. Process. Syst.* **35**, 37932 (2022), [https://papers.nips.cc/paper\\_files/paper/2022/hash/f7e7fabd73b3df96c-54a320862afcb78-Abstract-Conference.html](https://papers.nips.cc/paper_files/paper/2022/hash/f7e7fabd73b3df96c-54a320862afcb78-Abstract-Conference.html).
- [76] Y. Dandi, F. Krzakala, B. Loureiro, L. Pesce, and L. Stephan, *How two-layer neural networks learn, one (giant) step at a time*, *J. Mach. Learn. Res.* **25**, 1 (2024), <https://www.jmlr.org/papers/volume25/23-1543/23-1543.pdf>.
- [77] Z. Wang, D. Wu, and Z. Fan, *Nonlinear spiked covariance matrices and signal propagation in deep neural networks*, in *The Thirty Seventh Annual Conference on Learning Theory* (PMLR, Edmonton, Canada, 2024), pp. 4891–4957.
- [78] A. Maillard, E. Troiani, S. Martin, F. Krzakala, and L. Zdeborová, *Bayes-optimal learning of an extensive-width neural network from quadratically many samples*, [arXiv:2408.03733](https://arxiv.org/abs/2408.03733) (2024).
- [79] J. Barbier, F. Camilli, M.-T. Nguyen, M. Pastore, and R. Skerk, *Optimal generalisation and learning transition in extensive-width shallow neural networks near interpolation*, [arXiv:2501.18530](https://arxiv.org/abs/2501.18530) (2025).
- [80] T. M. Cover, *Elements of Information Theory* (John Wiley & Sons, New York, 1999).
- [81] Y. Polyanskiy and Y. Wu, *Information Theory: From Coding to Learning* (Cambridge University Press, Cambridge, England, 2025).
- [82] D. Guo, S. Shamai, and S. Verdú, *Mutual information and minimum mean-square error in gaussian channels*, *IEEE Trans. Inf. Theory* **51**, 1261 (2005).
- [83] O. Ledoit and S. Péché, *Eigenvectors of some large sample covariance matrix ensembles*, *Probab. Theory Relat. Fields* **151**, 233 (2011).
- [84] J. Bun, R. Allez, J.-P. Bouchaud, and M. Potters, *Rotational invariant estimator for general noisy matrices*, *IEEE Trans. Inf. Theory* **62**, 7475 (2016).
- [85] G. Semerjian, *Matrix denoising: Bayes-optimal estimators via low-degree polynomials*, *J. Stat. Phys.* **191**, 139 (2024).
- [86] N. S. Pillai and J. Yin, *Universality of covariance matrices*, *Ann. Appl. Probab.* **24**, 935 (2014).
- [87] A. Matytsin, *On the large- $n$  limit of the Itzykson-Zuber integral*, *Nucl. Phys.* **B411**, 805 (1994).

- [88] A. Guionnet and O. Zeitouni, *Large deviations asymptotics for spherical integrals*, *J. Funct. Anal.* **188**, 461 (2002).
- [89] A. Guionnet, *First order asymptotics of matrix integrals: A rigorous approach towards the understanding of matrix models*, *Commun. Math. Phys.* **244**, 527 (2004).
- [90] J.-B. Zuber, *The large- $n$  limit of matrix integrals over the orthogonal group*, *J. Phys. A* **41**, 382001 (2008).
- [91] A. Maillard, F. Krzakala, M. Mézard, and L. Zdeborová, *Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising*, *J. Stat. Mech.* (2022) 083301.
- [92] E. Troiani, V. Erba, F. Krzakala, A. Maillard, and L. Zdeborová, *Optimal denoising of rotationally invariant rectangular matrices*, in *Mathematical and Scientific Machine Learning* (PMLR, Beijing, China, 2022), pp. 97–112.
- [93] F. Pourkamali and N. Macris, *Rectangular rotational invariant estimator for general additive noise matrices*, in *2023 IEEE International Symposium on Information Theory (ISIT)* (IEEE, New York, 2023), pp. 2081–2086, 10.1109/ISIT54713.2023.10206784.
- [94] F. Pourkamali and N. Macris, *Rectangular rotational invariant estimator for high-rank matrix estimation*, arXiv:2403.04615 (2024).
- [95] I. D. Landau, G. C. Mel, and S. Ganguli, *Singular vectors of sums of rectangular random matrices and optimal estimation of high-rank signals: The extensive spike model*, *Phys. Rev. E* **108**, 054129 (2023).
- [96] A. Guionnet and J. Huang, *Asymptotics of rectangular spherical integrals*, *J. Funct. Anal.* **285**, 110144 (2023).
- [97] F. Pourkamali and N. Macris, *Bayesian extensive-rank matrix factorization with rotational invariant priors*, *Adv. Neural Inf. Process. Syst.* **36** (2024), [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/4b8afc-47273c746662a96dfdf562f87f-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/4b8afc-47273c746662a96dfdf562f87f-Paper-Conference.pdf).
- [98] A. Bodin and N. Macris, *Gradient flow on extensive-rank positive semi-definite matrix denoising*, in *2023 IEEE Information Theory Workshop (ITW)* (IEEE, New York, 2023), pp. 365–370.
- [99] A. Sakata and Y. Kabashima, *Statistical mechanics of dictionary learning*, *Europhys. Lett.* **103**, 28008 (2013).
- [100] A. Sakata and Y. Kabashima, *Sample complexity of Bayesian optimal dictionary learning*, *2013 IEEE International Symposium on Information Theory*, (IEEE, New York, 2013) p. 669, 10.1109/ISIT.2013.6620310.
- [101] F. Krzakala, M. Mézard, and L. Zdeborová, *Phase diagram and approximate message passing for blind calibration and dictionary learning*, in *2013 IEEE International Symposium on Information Theory* (IEEE, New York, 2013), pp. 659–663, 10.1109/ISIT.2013.6620308.
- [102] Y. Kabashima, F. Krzakala, M. Mézard, A. Sakata, and L. Zdeborová, *Phase transitions and sample complexity in Bayes-optimal matrix factorization*, *IEEE Trans. Inf. Theory* **62**, 4228 (2016).
- [103] F. Krzakala, M. Mézard, A. Manuel, and L. Zdeborová (unpublished).
- [104] H. C. Schmidt, *Statistical Physics of Sparse and Dense Models in Optimization and Inference*, Ph.D. thesis, Université Paris Saclay (COMUE) (2018).
- [105] J. Barbier and N. Macris, *Statistical limits of dictionary learning: Random matrix theory and the spectral replica method*, *Phys. Rev. E* **106**, 024136 (2022).
- [106] T. Plefka, *Convergence condition of the tap equation for the infinite-ranged Ising spin glass model*, *J. Phys. A* **15**, 1971 (1982).
- [107] A. Georges and J. S. Yedidia, *How to expand around mean-field theory using high-temperature expansions*, *J. Phys. A* **24**, 2173 (1991).
- [108] D. J. Thouless, P. W. Anderson, and R. G. Palmer, *Solution of solvable model of a spin glass*, *Philos. Mag.* **35**, 593 (1977).
- [109] F. Camilli, *New Perspectives in Statistical Mechanics and High-Dimensional Inference*, Ph.D. thesis, University of Bologna and École Normale Supérieure—PSL (2023).
- [110] F. Camilli and M. Mézard, *Matrix factorization with neural networks*, *Phys. Rev. E* **107**, 064308 (2023).
- [111] F. Camilli and M. Mézard, *The decimation scheme for symmetric matrix factorization*, *J. Phys. A* **57**, 085002 (2024).
- [112] J. J. Hopfield, *Neural networks and physical systems with emergent collective computational abilities*, *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554 (1982).
- [113] D. J. Amit, H. Gutfreund, and H. Sompolinsky, *Storing infinite numbers of patterns in a spin-glass model of neural networks*, *Phys. Rev. Lett.* **55**, 1530 (1985).
- [114] J. J. Hopfield, D. I. Feinstein, and R. G. Palmer, “Unlearning” has a stabilizing effect in collective memories, *Nature (London)* **304**, 158 (1983).
- [115] J. Van Hemmen, L. Ioffe, R. Kühn, and M. Vaas, *Increasing the efficiency of a neural network through unlearning*, *Physica (Amsterdam)* **163A**, 386 (1990).
- [116] M. Benedetti, E. Ventura, E. Marinari, G. Ruocco, and F. Zamponi, *Supervised perceptron learning vs unsupervised Hebbian unlearning: Approaching optimal memory retrieval in Hopfield-like networks*, *J. Chem. Phys.* **156**, 104107 (2022).
- [117] D. Kunisky, A. S. Wein, and A. S. Bandeira, *Notes on computational hardness of hypothesis testing: Predictions using the low-degree likelihood ratio*, in *ISAAC Congress (International Society for Analysis, its Applications and Computation)* (Springer, New York, 2019), pp. 1–50.
- [118] E. Marinari and G. Parisi, *Simulated tempering: A new Monte Carlo scheme*, *Europhys. Lett.* **19**, 451 (1992).
- [119] K. Hukushima and K. Nemoto, *Exchange Monte Carlo method and application to spin glass simulations*, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- [120] R. M. Neal, *Annealed importance sampling*, *Stat. Comput.* **11**, 125 (2001). [https://github.com/darjeeling-0/MCMC\\_for\\_RademacherMF](https://github.com/darjeeling-0/MCMC_for_RademacherMF) (2024).
- [122] J. Bun, J.-P. Bouchaud, and M. Potters, *Cleaning large correlation matrices: Tools from random matrix theory*, *Phys. Rep.* **666**, 1 (2017).
- [123] F. Antenucci, S. Franz, P. Urbani, and L. Zdeborová, *Glassy nature of the hard phase in inference problems*, *Phys. Rev. X* **9**, 011020 (2019).

- [124] M. Mezard and A. Montanari, *Information, Physics and Computation* (Oxford University Press, Oxford, England, 2009).
- [125] G.W. Anderson, A. Guionnet, and O. Zeitouni, *An Introduction to Random matrices* (Cambridge University Press, Cambridge, England, 2010) p. 118.
- [126] Z. Bai and J. W. Silverstein, *Spectral Analysis of Large Dimensional Random Matrices* (Springer, New York, 2010), Vol. 20.
- [127] T. Tao and V. Vu, *Random matrices: Universality of local eigenvalue statistics up to the edge*, *Commun. Math. Phys.* **298**, 549 (2010).
- [128] L. Erdős, S. Péché, J. A. Ramírez, B. Schlein, and H.-T. Yau, *Bulk universality for Wigner matrices*, *Commun. Pure Appl. Math.* **63**, 895 (2010).
- [129] L. Erdős, B. Schlein, and H.-T. Yau, *Universality of random matrices and local relaxation flow*, *Inventiones Mathematicae* **185**, 75 (2011).
- [130] T. Tao and V. Vu, *Random covariance matrices: Universality of local statistics of eigenvalues*, *Ann. Probab.* **40**, 1285 (2012).
- [131] O. H. Ajanki, L. Erdős, and T. Krüger, *Universality for general Wigner-type matrices*, *Probab. Theory Relat. Fields* **169**, 667 (2017).
- [132] W. Wang, J. Machta, and H. G. Katzgraber, *Comparing Monte Carlo methods for finding ground states of Ising spin glasses: Population annealing, simulated annealing, and parallel tempering*, *Phys. Rev. E* **92**, 013303 (2015).
- [133] Z. D. Bai, B. Q. Miao, and G. M. Pan, *On asymptotics of eigenvectors of large sample covariance matrix*, *Ann. Probab.* **35**, 1532 (2007).
- [134] L. Erdős, B. Schlein, and H.-T. Yau, *Local semicircle law and complete delocalization for Wigner random matrices*, *Commun. Math. Phys.* **287**, 641 (2009).
- [135] F. Benaych-Georges, *Eigenvectors of Wigner matrices: Universality of global fluctuations*, arXiv:1104.1219 (2011).
- [136] S. O'Rourke, V. Vu, and K. Wang, *Eigenvectors of random matrices: A survey*, *J. Comb. Theory Ser. A* **144**, 361 (2016).
- [137] J. Barbier, N. Macris, M. Dia, and F. Krzakala, *Mutual information and optimality of approximate message-passing in random linear estimation*, *IEEE Trans. Inf. Theory* **66**, 4270 (2020).
- [138] F. Krzakala, M. Mézard, and L. Zdeborová, *Compressed sensing under matrix uncertainty: Optimum thresholds and robust approximate message passing*, in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (IEEE, New York, 2013), pp. 5519–5523, 10.1109/ICASSP.2013.6638719.
- [139] M. Talagrand, *Mean Field Models for Spin Glasses: Volume I: Basic Examples* (Springer, New York, 2010).
- [140] M. Talagrand, *Mean Field Models for Spin Glasses: Volume II: Advanced Replica-Symmetry and Low Temperature* (Springer, New York 2011), Vol. 55.
- [141] G. Parisi, *Infinite number of order parameters for spin-glasses*, *Phys. Rev. Lett.* **43**, 1754 (1979).
- [142] G. Parisi, *A sequence of approximated solutions to the SK model for spin glasses*, *J. Phys. A* **13**, L115 (1980).
- [143] H. Sompolinsky, N. Tishby, and H. S. Seung, *Learning from examples in large neural networks*, *Phys. Rev. Lett.* **65**, 1683 (1990).
- [144] M. C. Angelini and F. Ricci-Tersenghi, *Limits and performances of algorithms based on simulated annealing in solving sparse hard inference problems*, *Phys. Rev. X* **13**, 021011 (2023).
- [145] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, *Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications*, *Phys. Rev. E* **84**, 066106 (2011).
- [146] J. Barbier, *Overlap matrix concentration in optimal Bayesian inference*, *Inf. Infer.* **10**, 597 (2019).
- [147] J. Barbier and D. Panchenko, *Strong replica symmetry in high-dimensional optimal Bayesian inference*, *Commun. Math. Phys.* **393**, 1199 (2022).
- [148] P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling extremal events: For insurance and finance* (Springer Science & Business Media, New York, 2013), Vol. 33.
- [149] G. Semerjian, *Some observations on the ambivalent role of symmetries in Bayesian inference problems*, *C.R. Phys.* **26**, 199 (2025).
- [150] E. Agliari, A. Alessandrelli, A. Barra, M. S. Centonze, and F. Ricci-Tersenghi, *Generalized hetero-associative neural networks*, *J. Stat. Mech.*(2025) 013302.
- [151] P. W. Anderson, *Absence of diffusion in certain random lattices*, *Phys. Rev.* **109**, 1492 (1958).
- [152] J. T. Parker, P. Schniter, and V. Cevher, *Bilinear generalized approximate message passing—part i: Derivation*, *IEEE Trans. Signal Process.* **62**, 5839 (2014).
- [153] Z. Yuan, Q. Guo, Y. C. Eldar, and Y. Li, *Unitary approximate message passing for matrix factorization*, in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, New York, 2024), pp. 9576–9580, 10.1109/ICASSP48485.2024.10447745.
- [154] Y. Ogata, *A Monte Carlo method for an objective Bayesian procedure*, *Ann. Inst. Stat. Math.* **42**, 403 (1990).
- [155] M. Jerrum and A. Sinclair, *Polynomial-time approximation algorithms for the Ising model*, *SIAM J. Comput.* **22**, 1087 (1993).
- [156] S. Syed, A. Bouchard-Côté, G. Deligiannidis, and A. Doucet, *Non-reversible parallel tempering: A scalable highly parallel MCMC scheme*, *J. R. Stat. Soc. Ser. B* **84**, 321 (2022).
- [157] J. Machta, *Strengths and weaknesses of parallel tempering*, *Phys. Rev. E* **80**, 056706 (2009).
- [158] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová, *Probabilistic reconstruction in compressed sensing: Algorithms, phase diagrams, and threshold achieving matrices*, *J. Stat. Mech.* (2012) P08009.
- [159] J. Barbier and F. Krzakala, *Approximate message-passing decoder and capacity achieving sparse superposition codes*, *IEEE Trans. Inf. Theory* **63**, 4894 (2017).