

PAPER • OPEN ACCESS

Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising

To cite this article: Antoine Maillard *et al* *J. Stat. Mech.* (2022) 083301

View the [article online](#) for updates and enhancements.

You may also like

- [Low rank matrix recovery with adversarial sparse noise](#)
Hang Xu, Song Li and Junhong Lin
- [Approximate survey propagation for statistical inference](#)
Fabrizio Antenucci, Florent Krzakala, Pierfrancesco Urbani et al.
- [A fault information-oriented weighted nuclear norm minimization method and its application to fault feature extraction in a rolling bearing](#)
Baoxiang Wang, Yuhe Liao, Yang Lv et al.



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

PAPER: Disordered systems, classical and quantum

Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising

Antoine Maillard^{1,2,*}, Florent Krzakala^{1,3},
Marc Mézard^{1,4} and Lenka Zdeborová⁵

¹ Laboratoire de Physique de l'École Normale Supérieure, PSL University, CNRS, Sorbonne Université, Université de Paris, France

² Institute for Mathematical Research (FIM), ETH Zurich, Switzerland

³ IdePHICS Laboratory, EPFL, Switzerland

⁴ Department of Computing Sciences, Bocconi University, Italy

⁵ SPOC Laboratory, EPFL, Switzerland

E-mail: antoine.maillard@math.ethz.ch

Received 4 February 2022

Accepted for publication 29 June 2022

Published 9 August 2022



CrossMark

Online at stacks.iop.org/JSTAT/2022/083301

<https://doi.org/10.1088/1742-5468/ac7e4c>

Abstract. Factorization of matrices where the rank of the two factors diverges linearly with their sizes has many applications in diverse areas such as unsupervised representation learning, dictionary learning or sparse coding. We consider a setting where the two factors are generated from known component-wise independent prior distributions, and the statistician observes a (possibly noisy) component-wise function of their matrix product. In the limit where the dimensions of the matrices tend to infinity, but their ratios remain fixed, we expect to be able to derive closed form expressions for the optimal mean squared error on the estimation of the two factors. However, this remains a very involved mathematical and algorithmic problem. A related, but simpler, problem is extensive-rank matrix denoising, where one aims to reconstruct a matrix with extensive but usually small rank from noisy measurements. In this paper, we approach both these problems using high-temperature expansions at fixed order parameters.

*Author to whom any correspondence should be addressed.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising

This allows to clarify how previous attempts at solving these problems failed at finding an asymptotically exact solution. We provide a systematic way to derive the corrections to these existing approximations, taking into account the structure of correlations particular to the problem. Finally, we illustrate our approach in detail on the case of extensive-rank matrix denoising. We compare our results with known optimal rotationally-invariant estimators, and show how exact asymptotic calculations of the minimal error can be performed using extensive-rank matrix integrals.

Keywords: learning theory, random matrix theory and extensions, statistical inference, spin glasses

Contents

1. Introduction	3
1.1. Motivation.....	3
1.2. Setting of the problem	4
1.3. A related problem: matrix denoising	6
1.4. Related works.....	7
1.5. Organization of the paper and main results.....	9
2. Derivation of the PGY expansion	12
2.1. PGY expansion and TAP equations in the symmetric setting.....	12
2.1.1. Sketch of the computation	12
2.1.2. The series at order 2 and the approximation of [KKM ⁺ 16].....	15
2.1.3. The series at third order: a non-trivial correction.....	16
2.2. Nature of the approximation in previous approaches.....	18
2.3. Non-symmetric extensive-rank matrix factorization	19
3. Probing our results in symmetric matrix denoising	21
3.1. The free entropy of factorization and denoising	21
3.1.1. The case $\alpha \leq 1$	22
3.1.2. The case $\alpha > 1$	26
3.2. The MMSE of denoising.....	27
3.3. Bayes-optimal estimator for denoising.....	28
3.3.1. From asymptotic results to estimation.....	28
3.3.2. Small- α expansion of the optimal denoiser	30
3.4. Denoising using the PGY expansion.....	31
3.4.1. PGY expansion for the denoising problem	31
3.4.2. Small- α expansion of the PGY denoiser.....	32
3.5. Numerical simulations	33
4. Conclusion and openings	38

Acknowledgments	39
Appendix A. Some definitions of probability theory.....	39
Appendix B. Technicalities of the PGY expansion.....	40
B.1. Setting of the PGY expansion.....	40
B.2. First orders of perturbation.....	40
B.3. PGY expansion in the non-symmetric model FX	44
B.3.1. The method.....	44
B.3.2. Order 3 of the expansion.....	45
B.3.3. Order 4 of the expansion.....	46
Appendix C. On the solutions to Matytsin’s equations	51
C.1. Derivation of result 3.1	51
C.2. Derivation of result 3.2	52
Appendix D. Small-α expansion of the optimal RIE	55
Appendix E. PGY expansion for denoising.....	56
Appendix F. Annealed calculations in the Gaussian setting.....	57
F.1. The symmetric case	57
F.2. The non-symmetric case.....	61
Appendix G. Other technicalities.....	65
G.1. The free entropy of [KKM ⁺ 16] in the Gaussian setting	65
G.2. Absence of contribution of diagonal terms to the MMSE.....	66
G.3. Proof of the I-MMSE theorem for denoising	67
References	67

1. Introduction

1.1. Motivation

Matrix factorization is a very generic problem that appears in a range of contexts. The goal is to reconstruct two matrices \mathbf{F}^* and \mathbf{X}^* , knowing only a noisy component-wise function of their product $\mathbf{F}^*\mathbf{X}^*$. In the Bayesian setting that we adopt here, one assumes some prior knowledge of the probability distributions from which \mathbf{F}^* and \mathbf{X}^* have been generated. Depending on the context, this prior will enforce some specific requirements like sparsity, or non-negativity. Applications that can be formulated as matrix factorization include dictionary learning or sparse

coding [OF96, OF97, KDMR⁺03, MBPS09], sparse principal component analysis [ZHT06], blind source separation [BAMCM97], low rank matrix completion [CR09, CT10], or robust principal component analysis [CLMW11].

Here, following exactly the problem setting of [KKM⁺16], we shall focus on the case of synthetic data, and inquire about the statistical limit of recovering matrices \mathbf{F}^* and \mathbf{X}^* from the observation of a function of their noisy product. While this problem has been solved in detail in the case in which \mathbf{F}^* , \mathbf{X}^* are low-rank (see e.g. [BBAP05, EAK18, LKZ17, DGJ18]), the case of extensive rank, where all dimensions of the matrices \mathbf{F}^* and \mathbf{X}^* go to infinity (with fixed ratios), is a much more challenging problem. In this limit, we are interested in two types of questions:

- (a) Finding the theoretical limits of the recovery of \mathbf{F}^* and \mathbf{X}^* , in particular understand when it is possible.
- (b) Designing algorithms for this inference, and understanding their own limits.

These challenges were addressed in [KKM⁺16], and related message-passing algorithms were developed in [PSC14a, PSC14b, ZZY21, LPPZ21]. We shall argue, however, that these previous works actually neglected some relevant correlations and do not provide a solution that is exact in the limit mentioned above. In this work we propose a systematic way to approach this problem with a high-temperature type of expansion at fixed order parameter. While the second order of this expansion gives back the previous result of [KKM⁺16], we show that higher order corrections are relevant and cannot be neglected.

We validate our approach on the special case corresponding to Gaussian priors on \mathbf{F}^* and \mathbf{X}^* and Gaussian additive noise on the observation of the product. This setting can be solved exactly using high-dimensional ‘HCIZ’ matrix integrals [HC57, IZ80, Mat94], or by leveraging known results on the denoising of rotationally invariant matrices [BABP16].

The high-temperature expansions presented in this paper open the way to a systematic construction of the mean-field ‘TAP-like’ equations for this problem, which could lead, with a proper iteration scheme, to an efficient message-passing algorithm. In future studies, it will also be interesting to make the contact between our high-temperature approach and the mean-field replica approach or the cavity method, which studies the statistical properties of the solutions to the TAP equations. All these approaches should eventually converge to a unified understanding of extensive-rank matrix factorization, and hopefully to efficient message-passing algorithms.

1.2. Setting of the problem

We consider the matrix factorization problem in the extensive-rank setting, which we define as the following inference problem:

Model FX (*Extensive-rank matrix factorization*)

Consider $n, m, p \geq 1$. Extensive-rank matrix factorization is the inference of the matrices $\mathbf{F}^* \in \mathbb{R}^{m \times n}$ and $\mathbf{X}^* \in \mathbb{R}^{n \times p}$ from the observation of $\mathbf{Y} \in \mathbb{R}^{m \times p}$, generated as:

$$Y_{\mu l} \sim P_{\text{out}} \left(\cdot \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n F_{\mu i}^* X_{il}^* \right. \right), \quad 1 \leq \mu \leq m, \quad 1 \leq l \leq p.$$

We assume that the matrix elements of \mathbf{F}^* and \mathbf{X}^* are both generated as i.i.d. random variables, with respective prior distributions P_F and P_X , both having zero mean and all moments finite. We consider this inference problem in the high-dimensional (or *thermodynamic*) limit, i.e. we assume $n, m, p \rightarrow \infty$ with finite limit ratios $m/n \rightarrow \alpha > 0$ and $p/n \rightarrow \psi > 0$.

In order to estimate \mathbf{F}^* and \mathbf{X}^* the statistician can use the *posterior distribution*, also referred to as *Gibbs measure* in the statistical physics literature⁶:

$$P_{\mathbf{Y},n}(\mathbf{dF}, \mathbf{dX}) \equiv \frac{1}{Z_{\mathbf{Y},n}} \prod_{\mu,i} P_F(\mathbf{dF}_{\mu i}) \prod_{i,l} P_X(\mathbf{dX}_{il}) \prod_{\mu,l} P_{\text{out}} \left(Y_{\mu l} \left| \frac{1}{\sqrt{n}} \sum_i F_{\mu i} X_{il} \right. \right). \quad (1)$$

We assume that she/he knows the distributions P_{out}, P_F, P_X ; this is known as the *Bayes-optimal setting*. In this case, it is well known that the mean under the posterior distribution of equation (1) is the information-theoretically optimal estimator of \mathbf{F}^* and \mathbf{X}^* .

In the thermodynamic limit, we can define the *single-instance free entropy* $\Phi_{\mathbf{Y},n}$ of the system as⁷:

$$\Phi_{\mathbf{Y},n} \equiv \frac{1}{n(m+p)} \ln \int \prod_{\mu,i} P_F(\mathbf{dF}_{\mu i}) \prod_{i,l} P_X(\mathbf{dX}_{il}) \prod_{\mu,l} P_{\text{out}} \left(Y_{\mu l} \left| \frac{1}{\sqrt{n}} \sum_i F_{\mu i} X_{il} \right. \right). \quad (2)$$

The averaged limit free entropy is denoted $\Phi \equiv \lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{Y}} \Phi_{\mathbf{Y},n}$. We finally define the asymptotic minimal mean squared error (or MMSE) of the matrices \mathbf{F} and \mathbf{X} as:

$$\begin{cases} \text{MMSE}_{\mathbf{F}} \equiv \lim_{n \rightarrow \infty} \frac{1}{nm} \sum_{\mu,i} \mathbb{E}_{\mathbf{Y}} \int P_F(\mathbf{dF}^*) P_X(\mathbf{dX}^*) (\langle F_{\mu i} \rangle - F_{\mu i}^*)^2, \\ \text{MMSE}_{\mathbf{X}} \equiv \lim_{n \rightarrow \infty} \frac{1}{np} \sum_{i,l} \mathbb{E}_{\mathbf{Y}} \int P_F(\mathbf{dF}^*) P_X(\mathbf{dX}^*) (\langle X_{il} \rangle - X_{il}^*)^2, \end{cases} \quad (3)$$

where the bracket denotes an average with respect to the measure $P_{\mathbf{Y},n}(\mathbf{dF}, \mathbf{dX})$ of equation (1).

Permutation symmetry—Note that since $\mathbf{F}^*, \mathbf{X}^*$ are generated i.i.d., there is a natural symmetry of the problem, in the sense that one can only recover them up to a common permutation $\{i \rightarrow \sigma(i)\}$ of the columns of \mathbf{F}^* and the rows of \mathbf{X}^* . Therefore a meaningful definition of the MMSEs of equation (3) requires to add an infinitely small side information to the system, breaking this symmetry. For lightness of the presentation, and since we will not evaluate numerically the factorization MMSEs, we do not introduce this technicality here. Note that this issue also arises in low-rank matrix factorization,

⁶ We will generically use Greek indices μ, ν for indices between 1 and m , while Latin i, j, k indices will run from 1 to n , and the l index between 1 and p .

⁷ We normalize $\Phi_{\mathbf{Y},n}$ by the total number $n(m+p)$ of variables to infer, while the normalization of [KKM⁺16] is n^2 .

Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising and is treated in the same way [LML⁺17]. Another way to deal with this symmetry is to compare quantities derived from \mathbf{F}^* , \mathbf{X}^* that are invariant under said symmetry.

We shall also consider a symmetric version of this problem, where the matrix has been generated as an empirical covariance matrix.

Model \mathbf{XX}^\top (*Symmetric extensive-rank matrix factorization*)

Consider $n, m \geq 1$. Symmetric extensive-rank matrix factorization is defined as the inference of the matrix $\mathbf{X}^* \in \mathbb{R}^{m \times n}$ from the observation of \mathbf{Y} generated as:

$$Y_{\mu\nu} \sim P_{\text{out}} \left(\cdot \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n X_{\mu i}^* X_{\nu i}^* \right. \right), \quad 1 \leq \mu < \nu \leq m.$$

We also assume that the elements of \mathbf{X}^* are generated i.i.d. from a *prior* distribution P_X . We consider this inference problem in the thermodynamic limit, where $n, m \rightarrow \infty$ with finite limit ratio $m/n \rightarrow \alpha > 0$.

As in the non-symmetric case, we can define both the posterior distribution of \mathbf{X} and the single-graph free entropy as:

$$\left\{ \begin{aligned} P_{\mathbf{Y},n}(\mathrm{d}\mathbf{X}) &\equiv \frac{1}{\mathcal{Z}_{\mathbf{Y},n}} \prod_{\mu,i} P_X(\mathrm{d}X_{\mu i}) \prod_{\mu < \nu} P_{\text{out}} \left(Y_{\mu\nu} \left| \frac{1}{\sqrt{n}} \sum_i X_{\mu i} X_{\nu i} \right. \right), \end{aligned} \right. \quad (4a)$$

$$\left\{ \begin{aligned} \Phi_{\mathbf{Y},n} &\equiv \frac{1}{nm} \ln \int \prod_{\mu,i} P_X(\mathrm{d}X_{\mu i}) \prod_{\mu < \nu} P_{\text{out}} \left(Y_{\mu\nu} \left| \frac{1}{\sqrt{n}} \sum_i X_{\mu i} X_{\nu i} \right. \right). \end{aligned} \right. \quad (4b)$$

The asymptotic MMSE on the matrix \mathbf{X} is defined as in the non-symmetric model (again assuming that we add an infinitely small side-information breaking the permutation symmetry between the columns of \mathbf{X}^*):

$$\text{MMSE}_{\mathbf{X}} \equiv \lim_{n \rightarrow \infty} \frac{1}{nm} \sum_{\mu i} \mathbb{E}_{\mathbf{Y}} \int P_X(\mathrm{d}\mathbf{X}^*) (\langle X_{\mu i} \rangle - X_{\mu i}^*)^2, \quad (5)$$

where the bracket denotes an average with respect to the measure $P_{\mathbf{Y},n}(\mathrm{d}\mathbf{X})$ of equation (4a).

The Gaussian setting—In the remaining of this work, we will denote *Gaussian setting* the specific models (both symmetric and non-symmetric) in which all the prior distributions and the channel distributions are zero-mean Gaussian. We consider the priors to have unit variance and the channel to have variance $\Delta > 0$.

1.3. A related problem: matrix denoising

Matrix denoising is a fundamental problem as well, with deep connection to the estimation of large covariance matrices in statistics and data analysis [BBP17]. In the problem of *denoising*, one aims at reconstructing a matrix \mathbf{Y}^* given by

$$Y_{\mu l}^* = \frac{1}{\sqrt{n}} \sum_{i=1}^n F_{\mu i}^* X_{il}^* \quad (6)$$

from noisy measurements of its elements, $Y_{\mu l} \sim P_{\text{out}}(Y_{\mu l} | Y_{\mu l}^*)$. The statistician knows that the matrix \mathbf{Y}^* was generated as a product of \mathbf{F}^* and \mathbf{X}^* , and while she/he knows the

prior distribution of \mathbf{F}^* and \mathbf{X}^* , she/he is only interested in reconstructing their product \mathbf{Y}^* . If we denote by \mathbf{G} the estimate of \mathbf{Y}^* , the posterior distribution $P_{\mathbf{Y},n}^{\text{den}}$ for denoising is

$$P_{\mathbf{Y},n}^{\text{den}}(d\mathbf{G}) \equiv \left\{ \int d\mathbf{F} d\mathbf{X} P_{\mathbf{Y},n}(d\mathbf{F}, d\mathbf{X}) \prod_{\mu,l} \delta \left(G_{\mu l} - \frac{1}{\sqrt{n}} \sum_i F_{\mu i} X_{il} \right) \right\} d\mathbf{G}. \quad (7)$$

This posterior is closely related to the posterior $P_{\mathbf{Y},n}(d\mathbf{F}, d\mathbf{X})$ of matrix factorization defined in equation (1). In particular, these two posterior distributions share the same partition function $\mathcal{Z}_{\mathbf{Y},n}$ and free entropy $\Phi_{\mathbf{Y},n}$. In a sense they are governed by the same measure, the difference being that the denoising statistician considers only observables that are built from the product $\mathbf{F}\mathbf{X}$.

Similarly, in the symmetric setting, the denoising problem aims at reconstructing the product $\mathbf{X}^*(\mathbf{X}^*)^\top$ so that the denoising posterior is:

$$P_{\mathbf{Y},n}^{\text{den}}(d\mathbf{G}) \equiv \left\{ \int d\mathbf{X} P_{\mathbf{Y},n}(d\mathbf{X}) \prod_{\mu < \nu} \delta \left(G_{\mu\nu} - \frac{1}{\sqrt{n}} \sum_i X_{\mu i} X_{\nu i} \right) \right\} d\mathbf{G}, \quad (8)$$

which again can be obtained from equation (4a). In the random matrix literature $\mathbf{X}^*(\mathbf{X}^*)^\top$ would often be called the Wishart matrix and we will hence denote this problem as denoising of a Wishart matrix. This case is particularly important as it corresponds to the problem of ‘cleaning’ empirical correlation matrices.

1.4. Related works

Previous approximations to extensive-rank matrix factorization—Beyond matrix factorization, many inference problems can be formulated in the statistical physics language. The large range of tools developed in the framework of statistical physics of disordered systems can then be used to tackle these models in the high-dimensional (or thermodynamic) limit. This has been attempted for the extensive rank matrix factorization problem with generic priors and output channels in a series of papers [SK13, KMZ13, KKM⁺16] applying the replica method to the problem. However, as we argue in detail in section 2.2, these works only provide an approximation, that is actually not exact in the thermodynamic limit due to some correlations between the variables that have been neglected.

A parallel line of work derived ‘approximate message-passing types’ of algorithms for the extensive rank matrix factorization problem [PSC14a, PSC14b]. By analogy with other graphical models where these algorithms are amenable to an exact analysis via an asymptotic description known as *state evolution* [BM11, BLM15, JM13, ZK16, GB21], the authors of [KKM⁺16] derived such a state evolution, but again making the same oversimplifying approximation. This implies that the state evolution stated in [KKM⁺16] is actually not following the behavior of the algorithm in the thermodynamic limit. Other recent works use related algorithms and their asymptotic analysis in a multi-layer setting [ZZY21, LPPZ21], relying essentially on the same approximation, which indicates that their analysis does not yield exact predictions in the thermodynamic limit.

Toward the exact solution of extensive rank matrix factorization and denoising—While the present authors realized the flaws in the solution presented in [KKM⁺16] back in 2017, works aiming at fixing this problem and finding the exact mean-field solution for the matrix factorization problems are scarce.

An important theoretical progress was made in the PhD thesis of Schmidt [Sch18], who realized that the approximation of [KKM⁺16] is not asymptotically exact when transposing it to the simpler case of matrix denoising with Gaussian additive noise and Gaussian prior on \mathbf{X}^* . The optimal error of matrix denoising for that case follows from a series of works on extensive-rank Harish–Chandra–Itzykson–Zuber (HCIZ) integrals [Mat94, GZ02] that are expressed via the solution of a hydrodynamical problem. In section 3, we show that in some cases this hydrodynamical problem can be solved exactly by the Dyson Brownian motion, and numerically evaluated, complementing the arguments of [Sch18]. The author of [Sch18] also attempted to correct the replica calculation of [KKM⁺16] including the missing correlations, but his proposition was also not leading to an exact solution.

An important algorithmic version of the optimal denoising has been worked out for rotationally-invariant estimators (RIE) in [BABP16]. In section 3 we confirm that indeed the algorithm of [BABP16] gives exactly the error predicted from the HCIZ matrix integrals.

Finally, another very recent work [BM21], came out during the completion of the present paper. An attempt of correcting the replica solution of [KKM⁺16] has already been provided in [Sch18]. Authors of [BM21] take this attempt further to develop a ‘spectral replica’ calculation for extensive-rank matrix factorization with non-Gaussian prior and (only) Gaussian channel. Its main result is an expression of the asymptotic free entropy, which is expressed as the function of a hydrodynamical HCIZ-like problem, making critical use of the results of [FG16, GH21] on the high-dimensional limit of ‘rectangular’ HCIZ integrals. While the expression obtained in [BM21] is closed and (conjectured to be) exact, a numerical solution of this hydrodynamical problem for cases of interest is yet to be provided.

TAP equations and Plefka–Georges–Yedidia (PGY) expansion—An important strategy of theoretical statistical physics, pioneered in the context of spin glasses by Thouless–Anderson–Palmer [TAP77], is to write mean-field equations that are exact for weakly coupled systems with infinite-range interactions. This was achieved in the Sherrington–Kirpatrick (SK) model [SK75] by finding a free energy at fixed order parameters: one constrains the average magnetization of each spin i to take a given value m_i , and writes a TAP free-energy as function of all the m_i ’s. This can then be optimized in order to precisely describe the physics of the system. The cavity method [MPV86] can be used both to derive the TAP free energy, and to perform a statistical analysis of its solutions (by going back to cavity fields). In more general models where the variables are not binary, the generalization of the TAP approach [Méz89] aims at computing the free entropy of the system (e.g. equations (2) and (4a)), while constraining the first and second moments of the fields (e.g. of $\{X_{\mu i}\}$ in equation (4a)). These moments thus become

parameters of the resulting ‘TAP’ free entropy, and extremizing over these variables leads to the celebrated *TAP equations*. These equations can be turned into algorithms, in which one seeks a solution of the TAP equations by iterating them, using appropriate time-iteration schemes [Bol14].

The TAP line of approach has led to prolific developments in increasingly complex models. For instance, when considering correlated data structures, such improvements include adaptive TAP (adaTAP) [OW01a, OW01b], expectation-consistency (EC) [Min01, OW05], as well as recent algorithmic approaches based on approximate message-passing algorithms [Ran11, DJM13, RSF17, SRF16, OCW16].

In a recent work [MFC⁺19], all these approaches have been shown to lead to equivalent approximation schemes for rotationally-invariant data structures in finite-rank models. Moreover, it was shown that they could be understood in terms of another technique in the theoretical physics toolbox, namely *high-temperature expansions at fixed order parameters*. These expansions were introduced by Plefka [Ple82] for the SK model, and then amply generalized by Georges and Yedidia [GY91]. In the following of this paper we shall denote them as PGY expansions. The core of this method is to systematically compute high-temperature expansions of the free entropy of the system, while constraining the first and second moments of the fields. This provides a systematic and controllable way to derive the TAP free energy and the TAP equations. An important aim of the present work is to apply these approaches to the extensive-rank matrix factorization problem. Importantly, we will see that in model XX^\dagger , this PGY expansion turns out to be an expansion in powers of the parameter $\alpha = m/n$.

1.5. Organization of the paper and main results

Let us describe the structure of the paper.

- Section 2 is dedicated to the detailed derivation of the PGY ‘high-temperature’ expansion at fixed order parameters for extensive-rank matrix factorization. We show that the second order of the expansion gives back the results of the previous approach of [KKM⁺16], but relevant new terms appear already at order 3 for the symmetric problem, and order 4 for the non-symmetric. We also comment in detail on the type of correlations that were neglected in [KKM⁺16]. In the majority of this section we focus on the symmetric model, and we end the section by detailing the PGY expansion for the non-symmetric matrix factorization problem.
- In section 3, we consider the simpler problem of extensive-rank matrix denoising with Gaussian additive noise. First, we show that exact calculations of the minimal error can be performed using extensive-rank spherical integrals [Har57, IZ80], and that they match the error achieved by the RIE of [BABP16]. As far as we know the explicit link between these approaches (one algorithmic, one asymptotic) has not been made previously. Secondly, we compare these results to the predictions of the PGY expansion developed in section 2, adapted for the denoising problem. We show that the order-3 corrections found in the PGY approach and neglected in [KKM⁺16] indeed match the optimal estimator, providing evidence for the correctness of our PGY-based approach.

Finally, note that while we focus in most of the following on model \mathbf{XX}^\top , in which many of our results can be stated in a lighter manner, all our techniques can be transposed to model \mathbf{FX} , as we show for instance in section 2.3.

The main contributions of this paper can be summarized as follows.

Extensive-rank matrix factorization for generic prior and noise channel—

We develop the PGY expansion for model \mathbf{XX}^\top and model \mathbf{FX} , for any priors P_X, P_F and observation channel P_{out} . For the symmetric model \mathbf{XX}^\top , the second order provides the approximation of [KKM⁺16]. However, we show that the third order of the PGY expansion provides additional non-negligible terms. Going further, we conjecture that the PGY approach at all orders would provide exact solution to the extensive rank matrix factorization and denoising. The advantage of this approach with respect to the replica method stems from its relation to the TAP equations and message-passing algorithms: because of this connection, it is directly related to an algorithm that can be used to solve practical instances of the problem. On the other hand, the recent replica result of [BM21] might help understand the large- n behavior of the solution to the TAP equations, what is called *state evolution* in the literature [BM11, BLM15, JM13, ZK16, GB21]. Another important aspect is that the TAP equations are conjectured for any type of output channel P_{out} , while recent replica results [BM21] are restricted to Gaussian additive noise.

Extensive-rank matrix denoising—

We probe our theoretical calculations on the simpler problem of matrix denoising described in section 1.3. Our contribution are as follows:

- We derive an analytical expression for the asymptotic free entropy of the problem by using HCIZ integrals [Mat94]. We provide an explicit solution to the optimization problem resulting from these integrals, and show both analytically and numerically that it agrees with the performance of the rotationally-invariant denoising algorithm of [BABP16].
- We evaluate the MMSE of matrix denoising, both from our asymptotic formula based on HCIZ integrals (red lines in figure 1), and from the algorithmic performance of the RIE of [BABP16] (green points in figure 1). We compare this MMSE to the solution of [KKM⁺16], that agrees with the second order of the PGY expansion developed in the present paper (blue points in figure 1), and show that it falls short of the actual MMSE.
- We truncate PGY expansion at third order (orange points in figure 1) and show that, while we do not reach the MMSE at this order, at small α we obtain a significant improvement with respect to the second-order truncation.

In figure 1 we illustrate some of the main results of this paper on the denoising problem. The presented plots are for denoising of Wishart matrices as defined in equation (8), with Gaussian output channel of zero mean and variance Δ . The factors \mathbf{X}^* are generated with i.i.d. $\mathcal{N}(0, 1)$ components. The three panels correspond to different aspects ratios $\alpha = m/n$ of \mathbf{X}^* . $\alpha \gg 1$ corresponds to denoising a small-rank matrix, a regime which we call *undercomplete*. On the other hand, the regime $\alpha \ll 1$ is called *overcomplete*, and corresponds to denoising a matrix close to a Wigner matrix, i.e. with random independent components. The four lines plotted in figure 1 are the following:

J. Stat. Mech. (2022) 083301

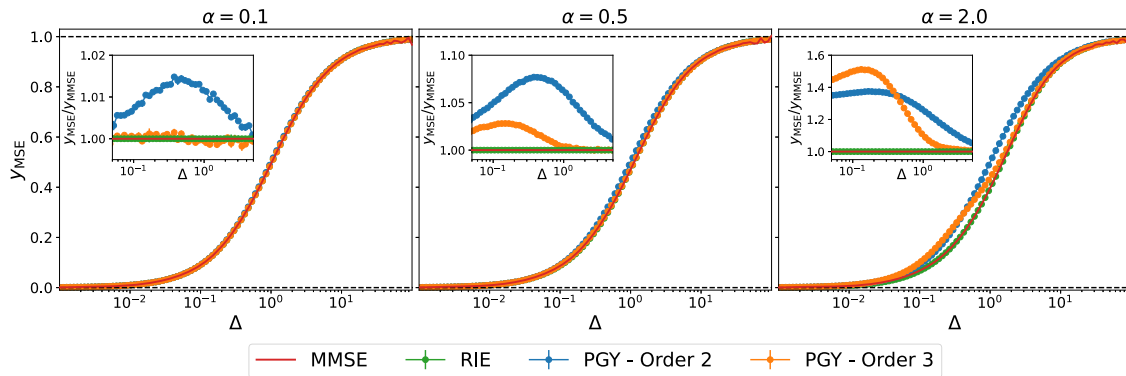


Figure 1. Denoising of a Wishart matrix as defined in equation (8) observed with Gaussian noise of zero mean and variance Δ . The factors \mathbf{X}^* are generated with i.i.d. $\mathcal{N}(0, 1)$ components and $\alpha = m/n$ is their aspect ratio. We compare the MMSE obtained from the analytic prediction using HCIZ integrals, the RIE estimator, and the estimator obtained from the PGY expansion truncated at orders 2 and 3. Note that the PGY expansion is an expansion at small α . In fact one sees that the third-order truncation works very well at small α . The truncation at order 2 corresponds to taking the limit $\alpha \rightarrow 0$, and therefore to assuming that the underlying matrix is a Wigner matrix.

- *MMSE*: the exact MMSE of the matrix denoising obtained analytically (as described in section 3.2, using the matrix integrals presented in section 3.1). For small α (in the overcomplete regime) the denoised matrix has almost independent elements, and thus the MMSE is very close to the MMSE of Wigner denoising.
- *RIE*: the MSE obtained by the RIE of [BABP16], presented in section 3.3. The algorithm was run on single instances of the problem with $m = 3000$ and the data were averaged over three samples (error bars are invisible). The agreement with the MMSE is perfect.
- *PGY order 2*: as explained in section 2.2 the second order of the PGY expansion for matrix denoising corresponds to ‘Wigner denoising’, that is it implicitly assumes that the components of the ground-truth matrix are independent. The approximation of [KKM⁺16] agrees exactly with the second order of the PGY expansion as we explain in section 2.1.2, it thus corresponds to Wigner (scalar) denoising. As in the overcomplete regime $\alpha \ll 1$ a Wishart matrix behaves like a Wigner matrix, this approximation is good in this domain, and deteriorates as α grows, as shown in figure 1
- *PGY order 3*: this corresponds to truncating the PGY expansion at order 3, as explained in section 3.4.1. It can be shown that increasing the order of the PGY expansion is equivalent to increasing the order considered in a small- α expansion of the denoiser, which can be analytically computed, see equation (81).

2. Derivation of the PGY expansion

2.1. PGY expansion and TAP equations in the symmetric setting

For a given instance of the problem, a generic way to compute the single-graph free entropy of equation (2) is to follow the formalism of [GY91] to perform the high-temperature expansion at fixed order parameters pioneered in [Ple82]. Here we show how this PGY method can be applied the symmetric model \mathbf{XX}^\top .

2.1.1. Sketch of the computation. In order to formulate our problem in a formalism adapted to high-temperature expansions, we first introduce an auxiliary field $\hat{\mathbf{H}} \equiv \mathbf{XX}^\top/\sqrt{n}$ in equation (4a), which defines the free entropy:

$$e^{nm\Phi_{\mathbf{Y},n}} = \int \prod_{\mu,i} P(dX_{\mu i}) \prod_{\mu<\nu} \left[\int d\hat{H}_{\mu\nu} P_{\text{out}}(Y_{\mu\nu}|\hat{H}_{\mu\nu}) \delta\left(\hat{H}_{\mu\nu} - \frac{1}{\sqrt{n}} \sum_i X_{\mu i} X_{\nu i}\right) \right]. \quad (9)$$

Using the Fourier transform of the delta function $\delta(x) = (2\pi)^{-1} \int dh e^{ihx}$, we reach an effective free entropy in terms of two fields \mathbf{X}, \mathbf{H} , with $\mathbf{H} \equiv \{H_{\mu\nu}\}_{\mu<\nu}$ a symmetric $m \times m$ matrix (whose diagonal can be taken equal to 0):

$$e^{nm\Phi_{\mathbf{Y},n}} = \int \prod_{\mu<\nu} P_{H,Y}^{\mu\nu}(dH_{\mu\nu}) \prod_{\mu,i} P_X(dX_{\mu i}) e^{-H_{\text{eff}}[\mathbf{X},\mathbf{H}]}, \quad (10)$$

in which we introduced an *effective Hamiltonian* $H_{\text{eff}}[\mathbf{F}, \mathbf{X}, \mathbf{H}]$ and un-normalized probability distributions $\{P_{H,Y}^{\mu\nu}\}$ defined as:

$$H_{\text{eff}}[\mathbf{X}, \mathbf{H}] \equiv \frac{1}{\sqrt{n}} \sum_{\mu<\nu} \sum_i (iH)_{\mu\nu} X_{\mu i} X_{\nu i}, \quad (11)$$

$$P_{H,Y}^{\mu\nu}[dH] \equiv \int \frac{d\hat{H}}{2\pi} e^{iH\hat{H}} P_{\text{out}}(Y_{\mu\nu}|\hat{H}). \quad (12)$$

The PGY formalism allows to compute the free entropy of the system, constraining the means and variances of each variable $\{X_{\mu i}, H_{\mu\nu}\}$ by ‘tilting’ the Gibbs measure of equation (10). We thus impose:

$$\begin{cases} \langle X_{\mu i} \rangle = m_{\mu i}, & \langle (iH)_{\mu\nu} \rangle = -g_{\mu\nu}, \\ \langle X_{\mu i}^2 \rangle = v_{\mu i} + (m_{\mu i})^2, & \langle (iH)_{\mu\nu}^2 \rangle = -r_{\mu\nu} + g_{\mu\nu}^2. \end{cases} \quad (13)$$

Here, $\langle \cdot \rangle$ denotes an average over the (now tilted) Gibbs measure. Note that we will sometimes symmetrize the quantities involved, e.g. we write $g_{\mu\nu} \equiv g_{\nu\mu}$ for $\mu > \nu$, and moreover we adopt the convention $g_{\mu\mu} = r_{\mu\mu} = 0$. The resulting free entropy is a function of the means and variances $\{\mathbf{m}, \mathbf{v}, \mathbf{g}, \mathbf{r}\}$, on which we will then have to extremize. The conditions of equation (13) will be imposed via Lagrange multipliers, which we denote respectively by $\{\boldsymbol{\lambda}, \boldsymbol{\gamma}, \boldsymbol{\omega}, \mathbf{b}\}$. The free entropy can now be expressed as a function of

all the parameters $\{\mathbf{m}, \mathbf{v}, \mathbf{g}, \mathbf{r}, \boldsymbol{\lambda}, \boldsymbol{\gamma}, \boldsymbol{\omega}, \mathbf{b}\}$ (we still denote it $\Phi_{\mathbf{Y},n}$, with a slight abuse of notations):

$$nm\Phi_{\mathbf{Y},n} = \sum_{\mu,i} \left[\lambda_{\mu i} m_{\mu i} + \frac{\gamma_{\mu i}}{2} (v_{\mu i} + (m_{\mu i})^2) \right] + \sum_{\mu < \nu} \left[-\omega_{\mu\nu} g_{\mu\nu} - \frac{b_{\mu\nu}}{2} (-r_{\mu\nu} + g_{\mu\nu}^2) \right] + \ln \int P_{H,Y}(\mathbf{dH}) P_X(\mathbf{dX}) e^{-S_{\text{eff}}[\mathbf{X},\mathbf{H}]}, \quad (14)$$

in which we introduced an *effective action* S_{eff} :

$$S_{\text{eff}}[\mathbf{X}, \mathbf{H}] \equiv \sum_{\mu,i} \left[\lambda_{\mu i} X_{\mu i} + \frac{\gamma_{\mu i}}{2} X_{\mu i}^2 \right] + \sum_{\mu < \nu} \left[\omega_{\mu\nu} (iH)_{\mu\nu} - \frac{b_{\mu\nu}}{2} (iH_{\mu\nu})^2 \right] + H_{\text{eff}}[\mathbf{X}, \mathbf{H}]. \quad (15)$$

We note that this effective action is a sum of decoupled terms, apart from a single ‘cubic interaction’ term which is H_{eff} defined in equation (11).

We now apply the standard PGY expansion by introducing a parameter $\eta > 0$ in front of the interaction Hamiltonian, i.e. replacing $H_{\text{eff}} \rightarrow \eta H_{\text{eff}}$ in equation (15). We shall then expand in powers of η , and take $\eta = 1$ in the end. We denote the corresponding free entropy $\Phi_{\mathbf{Y},n}(\eta)$. Importantly, for $\eta = 0$ all the fields $\{X_{\mu i}, H_{\mu\nu}\}$ are independent, which allows for an efficient perturbative computation.

In order to control the PGY expansion, we introduce two natural assumptions on the structure of the correlations of the variables:

- H.1 At their physical value, the variables $\{m_{\mu i}\}$ should be uncorrelated, in coherence with the fact that the elements of \mathbf{X}^* are drawn i.i.d. Importantly, this is not true for \mathbf{H} : although the elements $H_{\mu\nu}$ are independent by equation (12), their distribution depends on (μ, ν) and therefore their statistics (e.g. the variables $\{g_{\mu\nu}\}$) might be correlated through the correlation of the variables $\{Y_{\mu\nu}\}$.
- H.2 Recall that $g_{\mu\nu}$ is the average of $(iH)_{\mu\nu}$, the conjugate variable to $\hat{H}_{\mu\nu} \equiv \sum_i X_{\mu i} X_{\nu i} / \sqrt{n}$. In particular, as one can see from a simple calculation (see equations (20) and (21) that follow), $\{g_{\mu\nu}\}$ takes the form of an independent-component estimator of the error achieved by $\boldsymbol{\omega}$, which is the estimator of the Wishart matrix $\mathbf{X}\mathbf{X}^\top / \sqrt{n}$. In order to estimate this error, we use the observed value of $Y_{\mu\nu}$, built as an independent component-wise non-linearity applied to the Wishart matrix $\mathbf{X}^*(\mathbf{X}^*)^\top / \sqrt{n}$. Similarly, we assume that *the correlations of the $\mathbf{g} = \{g_{\mu\nu}\}$ scale as the ones of a matrix built as the $f(\mathbf{M}\mathbf{M}^\top / \sqrt{n})$, for \mathbf{M} a matrix with i.i.d. zero-mean components, and f a—possibly random—component-wise function.*

Assuming H.1 and H.2, we derive in appendix B the first three terms of the small η expansion. Note that H.2 is actually not necessary for result 2.1, as we stop our computation of the series at order 3. However we believe that understanding the structure of $\{g_{\mu\nu}\}$ might prove critical when tackling higher-order terms in the expansion, similarly as [MFC⁺19] for finite-rank problems. We leave for future work such a derivation of a general expression at all orders, and possibly a resummation.

Result 2.1 (*First three orders of the PGY expansion for model \mathbf{XX}^\top*)

At leading order as $n, m \rightarrow \infty$:

$$nm[\Phi_{\mathbf{Y},n}(\eta) - \Phi_{\mathbf{Y},n}(0)] = \frac{\eta}{\sqrt{n}} \sum_{\mu < \nu} g_{\mu\nu} m_{\mu i} m_{\nu i} - \frac{\eta^2}{2n} \sum_{\mu < \nu} r_{\mu\nu} [v_{\mu i} v_{\nu i} + v_{\mu i} m_{\nu i}^2 + m_{\mu i}^2 v_{\nu i}] + \frac{\eta^2}{4n} \sum_{\mu, \nu, i} g_{\mu\nu}^2 v_{\mu i} v_{\nu i} + \frac{\eta^3}{6n^{3/2}} \sum_i \sum_{\substack{\mu_1, \mu_2, \mu_3 \\ \text{pairwise distinct}}} g_{\mu_1 \mu_2} g_{\mu_2 \mu_3} g_{\mu_3 \mu_1} v_{\mu_1 i} v_{\mu_2 i} v_{\mu_3 i} + \mathcal{O}(\eta^4). \quad (16)$$

The zero-order term $\Phi_{\mathbf{Y},n}(0)$ is given by (with implicit extremization over the Lagrange multipliers)

$$nm\Phi_{\mathbf{Y},n}(0) = \sum_{\mu, i} \left[\lambda_{\mu i} m_{\mu i} + \frac{\gamma_{\mu i}}{2} (v_{\mu i} + (m_{\mu i})^2) + \ln \int P_X(dx) e^{-\frac{\gamma_{\mu i}}{2} x^2 - \lambda_{\mu i} x} \right] + \sum_{\mu < \nu} \left[-\omega_{\mu\nu} g_{\mu\nu} - \frac{b_{\mu\nu}}{2} (-r_{\mu\nu} + g_{\mu\nu}^2) + \ln \int dz \frac{e^{-\frac{1}{2b_{\mu\nu}}(z-\omega_{\mu\nu})^2}}{\sqrt{2\pi b_{\mu\nu}}} P_{\text{out}}(Y_{\mu\nu}|z) \right]. \quad (17)$$

We now consider the fixed point equations that result from extremizing the free entropy of result 2.1. Note that the Lagrange multipliers only appear in the $\eta = 0$ part of the free entropy (this is a general consequence of the formalism of [GY91]), so we can easily write the maximization over these parameters:

$$\begin{cases} m_{\mu i} = \mathbb{E}_{P(\lambda_{\mu i}, \gamma_{\mu i})}[X], & v_{\mu i} = \mathbb{E}_{P(\lambda_{\mu i}, \gamma_{\mu i})}[(X - m_{\mu i})^2], \\ g_{\mu\nu} = g_{\text{out}}(Y_{\mu\nu}, \omega_{\mu\nu}, b_{\mu\nu}), & r_{\mu\nu} = -\partial_{\omega} g_{\text{out}}(Y_{\mu\nu}, \omega_{\mu\nu}, b_{\mu\nu}). \end{cases} \quad (18)$$

In equation (18) we defined $P_{\lambda, \gamma}$ and g_{out} as:

$$\begin{cases} P(\lambda, \gamma)(x) \equiv \frac{P_X(x) e^{-\frac{1}{2}\gamma x^2 - \lambda x}}{\int dx' P_X(x') e^{-\frac{1}{2}\gamma x'^2 - \lambda x'}}, \\ g_{\text{out}}(y, \omega, b) \equiv \frac{1}{b} \frac{\int dz P_{\text{out}}(y|z) (z - \omega) e^{-\frac{(z-\omega)^2}{2b}}}{\int dz P_{\text{out}}(y|z) e^{-\frac{(z-\omega)^2}{2b}}}. \end{cases} \quad (19)$$

We shall then focus on the maximization over the parameters $\{\mathbf{m}, \mathbf{v}, \mathbf{g}, \mathbf{r}\}$. It is important to first understand the physical interpretation of $g_{\mu\nu}$ and $r_{\mu\nu}$.

General channels—Recall that we introduced the field $\hat{\mathbf{H}}$ in equation (9), and then its conjugate field \mathbf{H} in equation (10). Recall as well that we defined $\{\mathbf{g}, \mathbf{r}\}$ as the first and second moments that we constraint on the Gibbs measure, cf equation (13). Because the variables $\{H_{\mu\nu}\}$ have been introduced as auxiliary parameters, the physical interpretation of $g_{\mu\nu}$ and $r_{\mu\nu}$ is not totally obvious, but it can be found through simple integration by parts, starting from equation (13). We get, with $\hat{\mathbf{H}} \equiv \mathbf{XX}^\top / \sqrt{n}$:

$$g_{\mu\nu} = \left\langle \frac{\partial_{\hat{H}} P_{\text{out}}(Y_{\mu\nu} | \hat{H}_{\mu\nu})}{P_{\text{out}}(Y_{\mu\nu} | \hat{H}_{\mu\nu})} \right\rangle_{\mathbf{Y}}; \quad r_{\mu\nu} = \left\langle \frac{\partial_{\hat{H}}^2 P_{\text{out}}(Y_{\mu\nu} | \hat{H}_{\mu\nu})}{P_{\text{out}}(Y_{\mu\nu} | \hat{H}_{\mu\nu})} \right\rangle_{\mathbf{Y}} - g_{\mu\nu}^2. \quad (20)$$

Simplification for Gaussian channels—For a Gaussian channel (as is the case in the denoising problem), we have $P_{\text{out}}(Y|\cdot) = \mathcal{N}(Y, \Delta)$, and thus one gets:

$$\begin{cases} g_{\mu\nu} = \frac{1}{\Delta} \left[Y_{\mu\nu} - \frac{1}{\sqrt{n}} \left\langle \sum_{i=1}^n X_{\mu i} X_{\nu i} \right\rangle \right], \\ r_{\mu\nu} = \frac{1}{\Delta} - \frac{1}{\Delta^2} \left[\left\langle \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n X_{\mu i} X_{\nu i} \right)^2 \right\rangle - \left\langle \frac{1}{\sqrt{n}} \sum_{i=1}^n X_{\mu i} X_{\nu i} \right\rangle^2 \right]. \end{cases} \quad (21)$$

Notice that, from these equations, the PGY estimator for the denoising of \mathbf{Y} takes a simple form in the case of a Gaussian channel:

$$\hat{Y}_{\mu\nu} \equiv Y_{\mu\nu} - \Delta g_{\mu\nu}. \quad (22)$$

2.1.2. *The series at order 2 and the approximation of [KKM⁺16].* We first examine the second order of the perturbation series of result 2.1. The free entropy reads:

$$\begin{aligned} nm[\Phi_{\mathbf{Y},n}(\eta) - \Phi_{\mathbf{Y},n}(0)] &= \frac{\eta}{\sqrt{n}} \sum_{\substack{\mu < \nu \\ i}} g_{\mu\nu} m_{\mu i} m_{\nu i} - \frac{\eta^2}{2n} \sum_{\substack{\mu < \nu \\ i}} r_{\mu\nu} [v_{\mu i} v_{\nu i} + v_{\mu i} m_{\nu i}^2 \\ &+ m_{\mu i}^2 v_{\nu i}] + \frac{\eta^2}{4n} \sum_{\mu, \nu, i} g_{\mu\nu}^2 v_{\mu i} v_{\nu i} + \mathcal{O}(\eta^3). \end{aligned} \quad (23)$$

Since they are taken at $\eta = 0$, the fixed point equations of equation (18) are unchanged. The maximization over the physical order parameters (the means and variances) can be done and yields (we indicate on the right of the equation the corresponding parameter over which we maximized):

$$\begin{cases} b_{\mu\nu} = \frac{\eta^2}{n} \sum_i [v_{\mu i} v_{\nu i} + v_{\mu i} (m_{\nu i})^2 + (m_{\mu i})^2 v_{\nu i}], & (r_{\mu\nu}) \\ \omega_{\mu\nu} = \frac{\eta}{\sqrt{n}} \sum_i m_{\mu i} m_{\nu i} - g_{\mu\nu} \left[b_{\mu\nu} - \frac{\eta^2}{n} \sum_i v_{\mu i} v_{\nu i} \right], & (g_{\mu\nu}) \\ \gamma_{\mu i} = \frac{\eta^2}{n} \sum_{\nu} [r_{\mu\nu} v_{\nu i} + r_{\mu\nu} (m_{\nu i})^2 - g_{\mu\nu}^2 v_{\nu i}], & (v_{\mu i}) \\ \lambda_{\mu i} = -\frac{\eta}{\sqrt{n}} \sum_{\nu} g_{\mu\nu} m_{\nu i} + m_{\mu i} \left[-\gamma_{\mu i} + \frac{\eta^2}{n} \sum_{\nu} r_{\mu\nu} v_{\nu i} \right] & (m_{\mu i}). \end{cases} \quad (24)$$

Taking $\eta = 1$, the reader can check easily that the combined equations (18) and (24) are actually completely equivalent to the GAMP equations derived in [KKM⁺16], replacing

Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising the notations of the variables as follows⁸:

$$\begin{aligned}
 b_{\mu\nu} &\rightarrow V_{\mu\nu} & \omega_{\mu\nu} &\rightarrow \omega_{\mu\nu} & g_{\mu\nu} &\rightarrow g_{\text{out}}(Y_{\mu\nu}, \omega_{\mu\nu}, V_{\mu\nu}) & r_{\mu\nu} &\rightarrow -\partial_{\omega} g_{\text{out}}(Y_{\mu\nu}, \omega_{\mu\nu}, V_{\mu\nu}) \\
 \gamma_{\mu i} &\rightarrow Z_{\mu i}^{-1} & \lambda_{\mu i} &\rightarrow -\frac{W_{\mu i}}{Z_{\mu i}} & m_{\mu i} &\rightarrow \hat{f}_{\mu i} & v_{\mu i} &\rightarrow s_{\mu i}.
 \end{aligned}$$

In conclusion, the PGY expansion truncated to order η^2 gives back exactly the stationary limit of the GAMP algorithm of [KKM⁺16]. However, as we will see below, the higher-order corrections of order η^3 (and most probably beyond) cannot be neglected: this shows explicitly how the GAMP equations of [KKM⁺16] (and also the BiGAMP equations of e.g. [PSC14a, PSC14b, ZZY21], which are based on the same approximation) are missing relevant terms.

2.1.3. The series at third order: a non-trivial correction. While the order η^2 truncation of result 2.1 yields back the approximation of [KKM⁺16], we have computed via the PGY expansion the order η^3 , which we recall:

$$\frac{1}{3!nm} \frac{\partial^3 \Phi_{\mathbf{Y},n}}{\partial \eta^3}(\eta = 0) = \frac{1}{6mn^{5/2}} \sum_i \sum_{\substack{\mu_1, \mu_2, \mu_3 \\ \text{pairwise distinct}}} g_{\mu_1 \mu_2} g_{\mu_2 \mu_3} g_{\mu_3 \mu_1} v_{\mu_1 i} v_{\mu_2 i} v_{\mu_3 i} + \mathcal{O}_n(1). \quad (25)$$

A crucial observation is that in general the rhs of equation (25) is not negligible as $n \rightarrow \infty$. This is a consequence of the diagrammatic analysis of [MFC⁺19]. As an example, consider the case of a Gaussian additive noise for the observation channel. Given that \mathbf{Y} is the sum of a Wishart and a Wigner matrix, it is natural to assume that \mathbf{g} behaves like a rotationally-invariant matrix, see H.2. Assuming moreover for simplicity that the variances are all equal $v_{\mu i} = v$, theorem 1 of [MFC⁺19] (under the rotation-invariance hypothesis we described) yields:

$$\frac{1}{3!nm} \frac{\partial^3 \Phi_{\mathbf{Y},n}}{\partial \eta^3}(\eta = 0) = \frac{v^3 \alpha^{3/2}}{6} c_3(\mathbf{g}/\sqrt{m}) + \mathcal{O}_n(1),$$

in which the coefficients $\{c_p(\mathbf{g}/\sqrt{m})\}$ are the *free cumulants* of random matrix theory, and are functions of the spectrum of \mathbf{g}/\sqrt{m} , cf appendix A. Therefore, assuming that the bulk of eigenvalues of \mathbf{g}/\sqrt{m} stays of order 1 as $n \rightarrow \infty$ (which is a natural scaling given equation (22)), this order-3 term gives a non-negligible contribution to the free entropy. This indicates that the approximation of [KKM⁺16] breaks down, even in the simple case of Gaussian noise.

⁸[KKM⁺16] considers primarily the non-symmetric model FX, but one can very easily transfer their equations to model XX[†].

The third-order correction changes the order-2 TAP equation (24) as follows:

$$\begin{cases} b_{\mu\nu} = \frac{\eta^2}{n} \sum_i [v_{\mu i} v_{\nu i} + v_{\mu i} (m_{\nu i})^2 + (m_{\mu i})^2 v_{\nu i}], & (r_{\mu\nu}) \\ \omega_{\mu\nu} = \frac{\eta}{\sqrt{n}} \sum_i m_{\mu i} m_{\nu i} - g_{\mu\nu} \left[b_{\mu\nu} - \frac{\eta^2}{n} \sum_i v_{\mu i} v_{\nu i} \right] + \frac{\eta^3}{2n^{3/2}} \sum_{\rho} g_{\nu\rho} g_{\rho\mu} \sum_i v_{\mu i} v_{\nu i} v_{\rho i}, & (g_{\mu\nu}) \\ \gamma_{\mu i} = \frac{\eta^2}{n} \sum_{\nu} [r_{\mu\nu} v_{\nu i} + r_{\mu\nu} (m_{\nu i})^2 - g_{\mu\nu}^2 v_{\nu i}] - \frac{\eta^3}{n^{3/2}} \sum_{\nu, \rho} g_{\mu\nu} g_{\nu\rho} g_{\rho\mu} v_{\nu i} v_{\rho i}, & (v_{\mu i}) \\ \lambda_{\mu i} = -\frac{\eta}{\sqrt{n}} \sum_{\nu} g_{\mu\nu} m_{\nu i} + m_{\mu i} \left[-\gamma_{\mu i} + \frac{\eta^2}{n} \sum_{\nu} r_{\mu\nu} v_{\nu i} \right] & (m_{\mu i}). \end{cases} \quad (26)$$

It is interesting to see the typical scaling of the correction terms of order η^3 in the thermodynamic limit:

- As we saw in the argument above, we expect the correction to the free entropy in equation (25) to scale as $\eta^3 \alpha^{3/2}$ in the thermodynamic limit.
- The correction to $\omega_{\mu\nu}$ scales as $(\eta^3/n^{3/2}) \sum_{\rho} g_{\nu\rho} g_{\rho\mu} \sum_i v_{\mu i} v_{\nu i} v_{\rho i}$, while the correction to $\gamma_{\mu i}$ scales as $(\eta^3/n^{3/2}) \sum_{\nu, \rho} g_{\mu\nu} g_{\nu\rho} g_{\rho\mu} v_{\nu i} v_{\rho i}$. Using similar arguments as above in the case of Gaussian additive noise, we expect that these corrections scale respectively as $\eta^3 \alpha^{1/2}$ and $\eta^3 \alpha^{3/2}$ in the thermodynamic limit.

Therefore the corrections at order η^3 are relevant and cannot be neglected, except in the small α limit where they become small. We end this discussion by a few remarks on the results presented here:

- **Similarity with finite-rank problems**—One may notice that the first orders of result 2.1 are very similar the TAP free entropy of simple symmetric models with rotationally-invariant couplings, that is derived in [MFC⁺19], see equation (25) of the mentioned paper. A crucial difference is that here the role of the coupling matrix is played by \mathbf{g} , which is itself a parameter of the TAP free entropy. While this limits to the first order presented in result 2.1, this similarity is already striking. Importantly, the resummation of all orders of the PGY expansion in [MFC⁺19] then suggests that all orders $p \geq 3$ of the free entropy might perhaps be expressible solely in terms of the *eigenvalue distribution* of \mathbf{g}/\sqrt{m} .
- **Nature of the order parameter**—The previous remark suggests that the order parameter in extensive-rank matrix factorization would be a probability distribution of eigenvalues⁹. This is an important difference from the finite-rank case: for rank- k matrix factorization (or in general for rank- k recovery problems), the state of the system is governed by a $k \times k$ overlap matrix.
- **Small α series**—As we have seen, the corrections at order η^3 to the free entropy are of order $\alpha^{3/2}$. Here we remind that α is the ratio between the size of the observed matrix m and the size of the factor n . Large α corresponds to the low-rank matrix estimation (or undercomplete regime) and small α leads to the observed matrix

⁹This hint is strengthened by recent results in [BM21] using the spectral replica method: in the special case of Gaussian channels P_{out} , it analytically computes the asymptotic mean free energy with the replica method, and its results strongly suggest that the proper order parameter is indeed a probability measure of eigenvalues.

having elements close to i.i.d. (or overcomplete regime). It seems to us that the expansion in powers of η , which is the core of the PGY approach, actually turns into an expansion of the TAP equations in powers of α . While this is easily seen up to order 3 from result 2.1, we have no formal proof of this statement for all orders, and leave it as a conjecture.

- **Iterating the equations**—As thoroughly discussed in [MFC⁺19], the TAP equations are in general not sufficient to obtain an algorithm with good convergence properties. A classical example is given by the generalized vector approximate passing algorithm (G-VAMP) [SRF16], for which the corresponding TAP equations are derived in [MFC⁺19] using PGY expansions. As highlighted there, the TAP equations correspond to the stationary limit of G-VAMP, however, *there is no obvious iterative resolution scheme of the TAP equations that automatically gives the GVAMP algorithm*. This indicates that even if one obtains all the orders of perturbations in result 2.1, turning them into an efficient algorithm may require further work¹⁰.

2.2. Nature of the approximation in previous approaches

As we have found that the conjecture of [KKM⁺16] about exactness of their solution is incorrect, it is useful to go back to it and understand where that approach failed. In [KKM⁺16] the free entropy of the matrix factorization problem is estimated in the Bayes-optimal setting, in two different ways, namely via the replica method and via belief propagation (BP) equations. In both cases there are somewhat hidden hypotheses which we think are not valid. These hypotheses are also present in the derivation of the BiGAMP (and BiG-VAMP) algorithm (cf e.g. [PSC14a, PSC14b, ZZY21]). We therefore believe that these algorithms are also not able to give an exact asymptotic computation of the marginal probabilities in this problem.

Let us now describe both approaches taken in [KKM⁺16], and explain how the assumptions behind them fail, focusing primarily on the replica analysis performed in section V.B of [KKM⁺16]. The main idea behind the replica method is to compute the *quenched* free entropy from the evaluation of the moments of the partition function, using the relation [MPV87]:

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \mathbb{E}_{\mathbf{Y}} \ln \mathcal{Z}_{\mathbf{Y},n} = \frac{\partial}{\partial r} \left[\lim_{n \rightarrow \infty} \frac{1}{n^2} \ln \mathbb{E}_{\mathbf{Y}} \mathcal{Z}_{\mathbf{Y},n}^r \right]_{r=0}.$$

The computation of the quenched free entropy then reduces to the evaluation of the integer moments of the partition function, by analytically expanding the expression of the r th moment to any $r > 0$. When writing out $\mathbb{E}_{\mathbf{Y}} \mathcal{Z}_{\mathbf{Y},n}^r$ there naturally appears $(r + 1)$ *replicas* of the system, that interact via the channel distribution term, as represented in the following equation (here we consider model FX, to be in the same setting as [KKM⁺16]):

¹⁰ We unsuccessfully tried naive iteration schemes to solve equations (18) and (26), as we found the algorithm to always either diverge or converge to a trivial solution $\mathbf{m} = 0$. We leave a more precise investigation of the numerical properties of these iterations to a future work.

$$\mathbb{E}_{\mathbf{Y}} Z_{\mathbf{Y},n}^r = \int d\mathbf{Y} \left[\prod_{a=0}^r P_F(d\mathbf{F}^a) P_X(d\mathbf{X}^a) \right] \prod_{a=0}^r \left[\prod_{\mu,l} P_{\text{out}} \left(Y_{\mu l} \left| \frac{1}{\sqrt{n}} \sum_i F_{\mu i}^a X_{i l}^a \right. \right) \right].$$

A key step in the calculation of [KKM+16] is the assumption that

$$Z_{\mu l}^a \equiv \frac{1}{\sqrt{n}} \sum_i F_{\mu i}^a X_{i l}^a \tag{27}$$

are multivariate Gaussian random variables. However, although $\{\mathbf{F}^a\}$ and $\{\mathbf{X}^a\}$ follow statistically independent distributions, with zero mean and finite variances, this is not enough to guarantee the Gaussianity of the $Z_{\mu l}^a$ variables in the high-dimensional limit. Indeed, there is a number $\mathcal{O}(n^2)$ of variables $Z_{\mu l}^a$, and therefore classical central limit results cannot conclude on the asymptotic Gaussianity of the joint distribution of such variables. In general, this Gaussianity is actually false, as can be seen, e.g. by considering the following quantity, for a single replica:

$$L_4 \equiv \lim_{n \rightarrow \infty} \mathbb{E} \left[\frac{1}{n^3} \sum_{\mu_1 \neq \mu_2} \sum_{l_1 \neq l_2} Z_{\mu_1 l_1} Z_{\mu_1 l_2} Z_{\mu_2 l_2} Z_{\mu_2 l_1} \right]. \tag{28}$$

Let us assume that both P_F and P_X are standard Gaussian distributions for the simplicity of the argument. The computation of L_4 then yields $L_4 = \alpha^2 \psi^2$. However, should the *joint* distribution of the $\{Z_{\mu l}\}$ converge to a multivariate (zero-mean) Gaussian distribution, such a distribution would satisfy by definition

$$\mathbb{E}[Z_{\mu_1 l_1} Z_{\mu_2 l_2}] = \frac{1}{n} \sum_{i,j} \mathbb{E}[F_{\mu_1 i} F_{\mu_2 j}] \mathbb{E}[X_{l_1 i} X_{l_2 j}] = \delta_{\mu_1 \mu_2} \delta_{l_1 l_2}, \tag{29}$$

and Wick’s theorem would give, wrongly, $L_4 = 0$. Therefore the joint distribution of $\{Z_{\mu l}^a\}$ is not Gaussian.

The message-passing approach—Another approach to the problem are the BP equations [MM09], also referred to as the *cavity method* in the physics literature [MPV86]. The goal of BP is to compute the *marginal* distributions of each variable in the system, by solving iterative equations involving probability distributions over each single variable. These probability distributions are called *messages* in the BP language, and the fixed point of the iterative equations yields an estimate of the marginal distributions. While a detailed treatment of the BP derivation of [KKM+16] is beyond the scope of this paper, we have seen in the previous section 2.1 that the message-passing approach of [KKM+16] is equivalent to neglecting (wrongly) some higher order terms in the PGY expansion.

2.3. Non-symmetric extensive-rank matrix factorization

Performing the PGY expansion for model **FX** is extremely similar to the calculation we have done for model **XXT**: one can introduce a field $\tilde{\mathbf{H}} \equiv \mathbf{FX}/\sqrt{n}$, and then perform the same calculations via the Fourier transform of the delta function. In the following of section 2.3, we give the results of our derivation, while more details are given in

appendix B.3. We adopt similar notations to the ones of section 2.1, adding additional X, F subscripts to differentiate the two fields. More precisely, we impose the first and second moment constraints:

$$\begin{cases} \langle F_{\mu i} \rangle = m_{\mu i}^F, & \langle X_{i l} \rangle = m_{i l}^X, \\ \langle F_{\mu i}^2 \rangle = v_{\mu i}^F + (m_{\mu i}^F)^2, & \langle X_{i l}^2 \rangle = v_{i l}^X + (m_{i l}^X)^2, \\ \langle (iH)_{\mu l} \rangle = -g_{\mu l}, & \langle (iH)_{\mu l}^2 \rangle = -r_{\mu l} + g_{\mu l}^2. \end{cases} \quad (30)$$

The PGY expansion at order η^4 —Similarly to result 2.1, we obtain the first orders of the free entropy as:

Result 2.2 (First orders of the PGY expansion for model FX)

We have, at leading order as $n \rightarrow \infty$:

$$\begin{aligned} n(m+p)[\Phi_{\mathbf{Y},n}(\eta) - \Phi_{\mathbf{Y},n}(0)] &= \frac{\eta}{\sqrt{n}} \sum_{\mu,i,l} g_{\mu l} m_{\mu i}^F m_{i l}^X - \frac{\eta^2}{2n} \sum_{\mu,i,l} r_{\mu l} [v_{\mu i}^F (m_{i l}^X)^2 + (m_{\mu i}^F)^2 v_{i l}^X] \\ &+ \frac{\eta^2}{2n} \sum_{\mu,i,l} (g_{\mu l}^2 - r_{\mu l}) v_{\mu i}^F v_{i l}^X + \frac{\eta^4}{4n^2} \sum_i \sum_{\mu_1 \neq \mu_2} \sum_{l_1 \neq l_2} g_{\mu_1 l_1} g_{\mu_2 l_1} g_{\mu_1 l_2} g_{\mu_2 l_2} v_{\mu_1 i}^F v_{\mu_2 i}^F v_{i l_1}^X v_{i l_2}^X + \mathcal{O}(\eta^5). \end{aligned}$$

Recall that the term $\Phi_{\mathbf{Y},n}(0)$ contains the dependency on the channel and priors contributions, as well as the Lagrange multipliers introduced to enforce the conditions of equation (30). Its precise form is:

$$\begin{aligned} n(m+p)\Phi_{\mathbf{Y},n}(0) &= \sum_{\mu,i} \left[\lambda_{\mu i}^F m_{\mu i}^F + \frac{\gamma_{\mu i}^F}{2} (v_{\mu i}^F + (m_{\mu i}^F)^2) + \ln \int P_F(df) e^{-\frac{\gamma_{\mu i}^F}{2} f^2 - \lambda_{\mu i}^F f} \right] \\ &+ \sum_{i,l} \left[\lambda_{i l}^X m_{i l}^X + \frac{\gamma_{i l}^X}{2} (v_{i l}^X + (m_{i l}^X)^2) + \ln \int P_X(dx) e^{-\frac{\gamma_{i l}^X}{2} x^2 - \lambda_{i l}^X x} \right] \\ &+ \sum_{\mu,l} \left[-\omega_{\mu l} g_{\mu l} - \frac{b_{\mu l}}{2} (-r_{\mu l} + g_{\mu l}^2) + \ln \int dz \frac{e^{-\frac{1}{2b_{\mu l}}(z-\omega_{\mu l})^2}}{\sqrt{2\pi b_{\mu l}}} P_{\text{out}}(Y_{\mu l}|z) \right]. \end{aligned} \quad (31)$$

Higher-order terms, and breakdown of previous approximations—As in the symmetric case, the approximation of [KKM⁺16, PSC14a, PSC14b, ZZY21] amounts to truncating the perturbation series of result 2.2 at order η^2 . However, exactly as for model XX[†], the higher-order terms in result 2.2 are in general non-negligible. Under a similar hypothesis as H.2, they are related to the asymptotic singular value distribution of \mathbf{g}/\sqrt{n} , similarly to what happened in the symmetric case. The limit of such terms has been worked out in [MFC⁺19]. Assuming, e.g. $v_{\mu i}^F = v^F$ and $v_{i l}^X = v^X$, we have at $\eta = 0$:

$$\frac{1}{4!n(m+p)} \frac{\partial^4 \Phi_{\mathbf{Y},n}}{\partial \eta^4}(\eta=0) = \frac{\psi}{\alpha + \psi} \frac{(v^F)^2 (v^X)^2}{4} \Gamma_2 \left(\frac{\alpha}{\psi}, \frac{\mathbf{g}^\top \mathbf{g}}{n} \right) + \mathcal{O}_n(1). \quad (32)$$

The functions Γ_p are related to the rectangular free cumulants introduced in [BG11], and their precise definition is given in [MFC⁺19]. Again, in general these coefficients are non-negligible in the limit $n \rightarrow \infty$.

For completeness, let us write the TAP equations for this problem, with the corrections arising from the order η^4 . As in the symmetric case the extremization with respect to the Lagrange multipliers only depends on the zeroth order term, and yields:

$$\begin{cases} m_{\mu i}^F = \mathbb{E}_{P_F(\lambda_{\mu i}^F, \gamma_{\mu i}^F)}[f], & v_{\mu i}^F = \mathbb{E}_{P_F(\lambda_{\mu i}^F, \gamma_{\mu i}^F)}[(f - m_{\mu i}^F)^2], \\ m_{i l}^X = \mathbb{E}_{P_X(\lambda_{i l}^X, \gamma_{i l}^X)}[x], & v_{i l}^X = \mathbb{E}_{P_X(\lambda_{i l}^X, \gamma_{i l}^X)}[(x - m_{i l}^X)^2], \\ g_{\mu l} = g_{\text{out}}(Y_{\mu l}, \omega_{\mu l}, b_{\mu l}), & r_{\mu l} = -\partial_{\omega} g_{\text{out}}(Y_{\mu l}, \omega_{\mu l}, b_{\mu l}). \end{cases} \quad (33)$$

The fourth-order corrections affect the remaining TAP equations, which read:

$$\begin{cases} b_{\mu l} = \frac{\eta^2}{n} \sum_i [v_{\mu i}^F v_{i l}^X + v_{\mu i}^F (m_{i l}^X)^2 + (m_{\mu i}^F)^2 v_{i l}^X], \\ \omega_{\mu l} = \frac{\eta}{\sqrt{n}} \sum_i m_{\mu i}^F m_{i l}^X - g_{\mu l} \left[b_{\mu l} - \frac{\eta^2}{n} \sum_i v_{\mu i}^F v_{i l}^X \right] + \frac{\eta^4}{n^2} \sum_{i, \mu', l'} g_{\mu' l} g_{\mu' l'} g_{\mu l} v_{\mu' i}^F v_{\mu' i}^F v_{i l}^X v_{i l}^X, \\ \gamma_{\mu i}^F = \frac{\eta^2}{n} \sum_l [r_{\mu l} v_{i l}^X + r_{\mu l} (m_{i l}^X)^2 - g_{\mu l}^2 v_{i l}^X] - \frac{\eta^4}{n^2} \sum_{\mu'} \sum_{l_1 \neq l_2} g_{\mu' l_1} g_{\mu' l_2} g_{\mu l} v_{\mu' i}^F v_{\mu' i}^F v_{i l_1}^X v_{i l_2}^X, \\ \lambda_{\mu i}^F = -\frac{\eta}{\sqrt{n}} \sum_l g_{\mu l} m_{i l}^X + m_{\mu i}^F \left[-\gamma_{\mu i}^F + \frac{\eta^2}{n} \sum_l r_{\mu l} v_{i l}^X \right], \\ \gamma_{i l}^X = \frac{\eta^2}{n} \sum_{\mu} [r_{\mu l} v_{\mu i}^F + r_{\mu l} (m_{\mu i}^F)^2 - g_{\mu l}^2 v_{\mu i}^F] - \frac{\eta^4}{n^2} \sum_{l'} \sum_{\mu_1 \neq \mu_2} g_{\mu_1 l} g_{\mu_2 l} g_{\mu_2 l'} g_{\mu_1 l'} v_{\mu_1 i}^F v_{\mu_2 i}^F v_{i l}^X, \\ \lambda_{i l}^X = -\frac{\eta}{\sqrt{n}} \sum_{\mu} g_{\mu l} m_{\mu i}^F + m_{i l}^X \left[-\gamma_{i l}^X + \frac{\eta^2}{n} \sum_{\mu} r_{\mu l} v_{\mu i}^F \right]. \end{cases} \quad (34)$$

3. Probing our results in symmetric matrix denoising

3.1. The free entropy of factorization and denoising

We consider the symmetric problem, specializing to the Gaussian setup in which the prior $P_X = \mathcal{N}(0, 1)$, and the measurement channel $P_{\text{out}} = \mathcal{N}(0, \Delta)$. Then the partition function of both matrix factorization and denoising, defined in equations (4a) and (8), is given by¹¹:

¹¹ Note that we added the diagonal terms $\mu = \nu$ with respect to equation (4a). However, this number $\mathcal{O}(n)$ of terms does not affect the thermodynamic free entropy.

$$\mathcal{Z}_{\mathbf{Y},n} = \int_{\mathbb{R}^{m \times n}} \frac{d\mathbf{X}}{(2\pi)^{\frac{mn}{2}} (2\pi\Delta)^{\frac{m(m-1)}{4}}} e^{-\frac{1}{2} \sum_{\mu,i} X_{\mu i}^2 - \frac{1}{4\Delta} \sum_{\mu,\nu} \left(Y_{\mu\nu} - \frac{1}{\sqrt{n}} \sum_i X_{\mu i} X_{\nu i} \right)^2}. \quad (35)$$

The corresponding free entropy $\Phi_{\mathbf{Y},n} \equiv (nm)^{-1} \ln \mathcal{Z}_{\mathbf{Y},n}$ can be computed for a given signal \mathbf{Y} , using random matrix techniques. We shall compute here its thermodynamic limit $\Phi_{\mathbf{Y}} \equiv \lim_{n \rightarrow \infty} \Phi_{\mathbf{Y},n}$, in the generic case where the asymptotic eigenvalue distribution of \mathbf{Y}/\sqrt{m} is well defined, and we denote it $\rho_{\mathbf{Y}}$. This computation relies on a sophisticated result on the celebrated HCIZ integrals, that emerged in the theoretical physics literature [Mat94], before being rigorously established [GZ02]. This approach was first discussed in the context of extensive-rank matrix factorization in [Sch18]. Here we elaborate on this work, filling in gaps in its arguments, and most importantly we work out how to evaluate analytically and numerically the corresponding equations, in order to provide explicit comparison to previous works and the PGY expansion proposed here. The recent and independent work [BM21] also derives asymptotic results in matrix factorization using Matytsin’s results [Mat94], but does not evaluate them or compare them to other approaches (beyond scalar denoising).

Beyond Wishart matrices—While we focus on the denoising of Wishart matrices (since we are initially interested in the matrix factorization problem), we emphasize that one can easily generalize all the conclusions of sections 3.1, 3.2 and 3.3 to the more general denoising of a rotationally-invariant matrix. This allows to compare our results with the general ones on symmetric matrix denoising of [BABP16, BBP17] in the case of Gaussian noise¹².

3.1.1. *The case $\alpha \leq 1$.* We shall first study the case $\alpha \leq 1$ (i.e. $m \leq n$), so that the distribution of $\mathbf{S} \equiv \mathbf{X}\mathbf{X}^\top/\sqrt{nm}$ is non-singular. We can rewrite equation (35) as:

$$\exp\{nm\Phi_{\mathbf{Y},n}\} = \int P_S(d\mathbf{S}) \frac{e^{-\frac{1}{4\Delta} \sum_{\mu,\nu} (Y_{\mu\nu} - \sqrt{m}S_{\mu\nu})^2}}{(2\pi\Delta)^{\frac{m(m-1)}{4}}}, \quad (36)$$

in which $P_S(d\mathbf{S})$ is the Wishart distribution given by (with Γ_p the multivariate Γ function):

$$P_S(d\mathbf{S}) \equiv \frac{m^{\frac{mn}{2}}}{(2\sqrt{\alpha})^{\frac{mn}{2}} \Gamma_m(n/2)} (\det \mathbf{S})^{\frac{n-m-1}{2}} e^{-\frac{\sqrt{nm}}{2} \text{Tr} \mathbf{S}} d\mathbf{S}. \quad (37)$$

We can decompose the integration in equation (36) on the eigenvalues and eigenvectors of $\mathbf{S} = \mathbf{O}\mathbf{L}\mathbf{O}^\top$, where the diagonal matrix \mathbf{L} contains the m eigenvalues $\{l_\mu\}$. This change of variables yields a combinatorial factor, that can be found, e.g. using equation (C.4) of [Nic14] or propositions 4.1.1 and 4.1.14 of [AGZ10]. All in all, equation (36) turns into:

¹²Note that [BABP16, BBP17] also consider generic rotationally-invariant additive noise, beyond the Gaussian assumption. While this is beyond our scope here, adapting the spherical integral techniques presented hereafter to this more generic setting could be possible, and we leave it as an open direction of research.

$$\exp\{nm\Phi_{\mathbf{Y},n}\} = \frac{m^{\frac{mn}{2}} e^{-\frac{1}{4\Delta}\sum_{\mu,\nu} Y_{\mu\nu}^2}}{(2\sqrt{\alpha})^{\frac{mn}{2}} \Gamma_m(n/2)(2\pi\Delta)^{\frac{m(m-1)}{4}}} \times \frac{\pi^{\frac{m(m+1)}{4}}}{2^{m/2}\Gamma(m+1)\prod_{\mu=1}^m \Gamma(\mu/2)} \quad (38)$$

$$\times \int_{\mathbb{R}_+^m} d\mathbf{L} \prod_{\mu < \nu} |l_\mu - l_\nu| e^{-\frac{m}{2}\sum_{\mu=1}^m \left(\frac{l_\mu}{\sqrt{\alpha}} + \frac{l_\mu^2}{2\Delta}\right) + \frac{n-m-1}{2}\sum_{\mu=1}^m \ln l_\mu} \int_{\mathcal{O}(m)} \mathcal{D}\mathbf{O} e^{\frac{\sqrt{m}}{2\Delta} \text{Tr}[\mathbf{Y}\mathbf{O}\mathbf{L}\mathbf{O}^\top]}.$$

The last integral is given by the large- n limit of the extensive-rank HCIZ integral [Mat94, GZ02]:

$$\int_{\mathcal{O}(m)} \mathcal{D}\mathbf{O} e^{\frac{\sqrt{m}}{2\Delta} \text{Tr}[\mathbf{Y}\mathbf{O}\mathbf{L}\mathbf{O}^\top]} \simeq \exp\left\{\frac{m^2}{2} I_\Delta[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}]\right\}.$$

The function I_Δ depends only on the asymptotic eigenvalue distribution $\rho_{\mathbf{Y}}$ of \mathbf{Y}/\sqrt{m} , and on the asymptotic distribution of the eigenvalues of \mathbf{S} . In principle, using the Laplace method, the result should be expressed as a supremum over the asymptotic distribution of these eigenvalues, and one would need to assume the existence of this asymptotic distribution (this is justified as there is a number $\mathcal{O}(n)$ of eigenvalues, and we consider asymptotics in the scale $\exp\{\Theta(n^2)\}$). It turns out that in our case we do not need to perform this complicated optimization on the spectrum of \mathbf{S} . This is a consequence of *Bayes optimality*. Indeed, $\rho_{\mathbf{S}}$ is the limit of the eigenvalue distribution of \mathbf{S} under the posterior measure. But by the Nishimori identity [Nis01], for any bounded function ϕ we have:

$$\mathbb{E}\left\langle \int \phi(\lambda) \rho_{\mathbf{S}}(d\lambda) \right\rangle = \mathbb{E} \int \phi(\lambda) \rho_{\mathbf{S}}^*(d\lambda), \quad (39)$$

in which $\rho_{\mathbf{S}}^*$ is the limit eigenvalue distribution of the ground-truth signal \mathbf{S}^* . In other terms, the Nishimori identity implies that the eigenvalue distribution of \mathbf{S} under the distribution of equation (36) concentrates on the Marchenko–Pastur law.

Therefore the function I_Δ depends on two known spectra, $\rho_{\mathbf{S}}$ which is the Marcenko–Pastur spectrum of the matrix $\mathbf{X}^*(\mathbf{X}^*)^\top/\sqrt{nm}$, and $\rho_{\mathbf{Y}}$ which is the spectrum of the signal, equal to a Wishart matrix plus a Gaussian noise. We moreover have from [Mat94, GZ02]:

$$I_\Delta[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}] = \frac{\ln \Delta}{2} - \frac{3}{4} + \frac{1}{2\Delta} \left[\int \rho_{\mathbf{Y}}(dx) x^2 + \int \rho_{\mathbf{S}}(dx) x^2 \right] \quad (40)$$

$$- \frac{1}{2} \int \rho_{\mathbf{Y}}(dx) \rho_{\mathbf{Y}}(dy) \ln |x - y| - \frac{1}{2} \int \rho_{\mathbf{S}}(dx) \rho_{\mathbf{S}}(dy) \ln |x - y|$$

$$- \frac{1}{2} \left\{ \int_0^\Delta dt \int dx \rho(x, t) \left[\frac{\pi^2}{3} \rho(x, t)^2 + v(x, t)^2 \right] \right\},$$

where the function $f(x, t) \equiv v(x, t) + i\pi\rho(x, t)$ satisfies a complex Burgers' equation with prescribed boundary conditions:

$$\begin{cases} \partial_t f + f\partial_x f = 0, \\ \rho(x, t = 0) = \rho_S(x), \\ \rho(x, t = \Delta) = \rho_Y(x). \end{cases} \quad (41)$$

We now evaluate the thermodynamic limit of the free entropy of equation (38), when $m, n \rightarrow \infty$ with a fixed limit of $\alpha = m/n$:

$$\Psi[\rho_Y, \rho_S] \equiv \lim_{n \rightarrow \infty} \Phi_{Y,n} = \Psi_1 + \Psi_2 + \Psi_3, \quad (42)$$

where

$$\begin{cases} \Psi_1 \equiv \lim_{m \rightarrow \infty} \left[\frac{1}{2} \ln m - \frac{1}{nm} \ln \Gamma_m(n/2) - \frac{1}{nm} \sum_{\mu=1}^m \Gamma(\mu/2) \right], \\ \Psi_2 \equiv -\frac{\ln 2}{2} - \frac{\ln \alpha}{4} - \frac{3\alpha}{8} - \frac{\alpha}{4} \ln 2, \\ \Psi_3 \equiv \frac{\alpha}{2} \int dx dy \rho_S(x)\rho_S(y) \ln|x-y| - \frac{\alpha}{2} \int dx \rho_S(x) \left[\frac{x}{\sqrt{\alpha}} + \frac{x^2}{2\Delta} \right] \\ \quad + \frac{1-\alpha}{2} \int dx \rho_S(x) \ln x + \frac{\alpha}{2} I_\Delta[\rho_Y, \rho_S]. \end{cases} \quad (43)$$

Ψ_1 can be computed using the asymptotic expansion

$$\frac{1}{nm} \ln \Gamma_m(n/2) = \frac{\alpha}{4} \ln \pi + \frac{1}{nm} \sum_{\mu=1}^m \ln \Gamma\left(\frac{n+1-\mu}{2}\right) + \mathcal{O}_n(1), \quad (44)$$

and Stirling's formula $\ln \Gamma(z) = z \ln z - z + \mathcal{O}(\ln z)$ as $z \rightarrow \infty$. We get:

$$\begin{aligned} & \frac{1}{2} \ln m - \frac{1}{nm} \ln \Gamma_m(n/2) - \frac{1}{nm} \sum_{\mu=1}^m \ln \Gamma(\mu/2) \\ &= \frac{\alpha + (\alpha - 1)^2 \ln(1 - \alpha) - (\alpha - 2)\alpha \ln(\alpha)}{4\alpha} - \frac{\alpha}{4} \ln \pi + \frac{1 + \ln 2}{2} + \mathcal{O}_n(1). \end{aligned} \quad (45)$$

Collecting all pieces, we get the final expression for the free entropy density of extensive-rank symmetric matrix factorization in the Gaussian setting:

$$\begin{aligned} \Psi[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}] &= \frac{1}{8} \left\{ \frac{2(\alpha - 1)^2 \ln(1 - \alpha)}{\alpha} - \alpha(3 + \ln 4) + 2 \ln \alpha - 2\alpha \ln \pi\alpha + 6 \right\} \\ &+ \frac{\alpha}{4} \int \rho_{\mathbf{S}}(dx) \rho_{\mathbf{S}}(dy) \ln |x - y| - \frac{\sqrt{\alpha}}{2} \int \rho_{\mathbf{S}}(dx) x \\ &- \frac{\alpha - 1}{2} \int \rho_{\mathbf{S}}(dx) \ln x - \frac{\alpha}{4} \int \rho_{\mathbf{Y}}(dx) \rho_{\mathbf{Y}}(dy) \ln |x - y| \\ &- \frac{\alpha}{4} \int_0^\Delta dt \int dx \rho(x, t) \left[\frac{\pi^2}{3} \rho(x, t)^2 + v(x, t)^2 \right]. \end{aligned} \quad (46)$$

As stated above, the function $f(x, t) = v(x, t) + i\pi\rho(x, t)$ satisfies the complex Burgers' equation (41).

It turns out that we can greatly simplify equation (46). This simplification is a consequence of two results, which we state hereafter. While the first of these two results is not directly used in the simplification, we believe it also bears independent interest that allows to understand the reasons allowing for such a simplification.

The first one is the analytical solution to equation (41), using the Dyson Brownian motion. It was given in [Sch18], but incompletely justified.

Result 3.1 (*Matytsin's solution and Dyson Brownian motion*)

Recall that $\mathbf{Y}/\sqrt{m} = \mathbf{S} + \sqrt{\Delta}\mathbf{Z}/\sqrt{m}$, with \mathbf{Z} a Wigner matrix. For any \mathbf{S} , the solution to equation (41) is given by the Stieltjes transform¹³ of the Dyson Brownian motion $\mathbf{Y}(t)/\sqrt{m} \stackrel{d}{=} \mathbf{S} + \sqrt{t}\mathbf{Z}/\sqrt{m}$, with $t \in [0, \Delta]$. More precisely, if $g_{\mathbf{Y}(t)}(z) \equiv \int \rho_{\mathbf{Y}(t)}(dt)/(t - z)$, the solution is given by:

$$\begin{cases} \rho(x, t) &= \frac{1}{\pi} \lim_{\epsilon \downarrow 0} \{ \text{Im}[g_{\mathbf{Y}(t)}(x + i\epsilon)] \}, \\ v(x, t) &= \lim_{\epsilon \downarrow 0} \{ -\text{Re}[g_{\mathbf{Y}(t)}(x + i\epsilon)] \}. \end{cases} \quad (47)$$

Remark—Note that this result, while not stated exactly as here, was used before in the mathematics literature, see e.g. [Gui04a, Men17]. Further, these works relate the distribution $\rho(x, t)$ that appears in the general formula of equation (40) of $I_\Delta[\rho_{\mathbf{A}}, \rho_{\mathbf{B}}]$ to the one of a *Dyson Brownian bridge*. More precisely, [Gui04a] shows that *there exists* a joint distribution of two random matrices \mathbf{A}' and \mathbf{B}' such that ρ_A is the LSD of \mathbf{A}' , ρ_B the one of \mathbf{B}' , and $\rho(t)$ is the asymptotic eigenvalue distribution of $\mathbf{X}(t) = (1 - t/\Delta)\mathbf{A}' + t\mathbf{B}'/\Delta + \sqrt{t(1 - t/\Delta)}\mathbf{W}/\sqrt{m}$, with \mathbf{W} a Gaussian Wigner matrix independent of \mathbf{A}' and \mathbf{B}' . However, the joint distribution of $(\mathbf{A}', \mathbf{B}')$ is unknown in general: result 3.1 shows that when $\mathbf{B} = \mathbf{A} + \sqrt{\Delta}\mathbf{Z}/\sqrt{m}$, one can take $\mathbf{A}' = \mathbf{A}$ and $\mathbf{B}' = \mathbf{B}$.¹⁴ To the best of our knowledge, the only other case in which this joint distribution is known is when \mathbf{A} and \mathbf{B} are independent Gaussian Wigner matrices with arbitrary variances [BBMP14].

¹³For more details on the Stieltjes transform, see appendix A.

¹⁴Indeed, one has then $\mathbf{X}(t) = \mathbf{A} + [t/\sqrt{\Delta}\mathbf{Z} + \sqrt{t(1 - t/\Delta)}\mathbf{W}]/\sqrt{m}$. The last term of this equation is the sum of two independent Gaussian matrices, therefore it is itself Gaussian, with variance $t^2/\Delta + t(1 - t/\Delta) = t$. In distribution one has thus $\mathbf{X}(t) \stackrel{d}{=} \mathbf{A} + \sqrt{t}\mathbf{W}/\sqrt{m}$, i.e. $\mathbf{X}(t)$ is the Dyson Brownian motion starting in \mathbf{A} .

Our second result allows for a direct simplification of equation (46). It can be stated as follows:

Result 3.2

For any well-behaved density ρ_A (corresponding to a random matrix \mathbf{A}), we have with $I = I_{\Delta=1}$ the function of equation (40):

$$\begin{aligned} & \sup_{\rho_L} \left\{ \frac{1}{2} \int \rho_L^{\otimes 2}(dx, dy) \ln|x-y| - \frac{1}{4} \int \rho_L(dx) x^2 + \frac{1}{2} I[\rho_L, \rho_A] \right\} \\ & = -\frac{3}{8} + \frac{1}{4} \int \rho_A(dx) x^2. \end{aligned} \tag{48}$$

Equivalently, if $\rho(x, t)$ solves the Euler–Matytsin equation (41) between ρ_A and ρ_L , we have:

$$\begin{aligned} & \sup_{\rho_L} \left\{ \int \rho_L^{\otimes 2}(dx, dy) \ln|x-y| - \int_0^1 dt \int dx \rho(x, t) \left[\frac{\pi^2}{3} \rho(x, t)^2 + v(x, t)^2 \right] \right\} \\ & = \int \rho_A^{\otimes 2}(dx, dy) \ln|x-y|. \end{aligned} \tag{49}$$

Moreover, in both previous equations the supremum is reached by the *additive free convolution* $\rho_L^* = \rho_A \boxplus \sigma_{s.c.}$ [Voi86, AGZ10], i.e. the eigenvalue density of $\mathbf{A} + \mathbf{Z}/\sqrt{m}$, which is the Dyson Brownian motion at time $t = 1$.

In appendix C we give a complete justification of results 3.1 and 3.2. Using result 3.2, and the fact that $\mathbf{Y}/\sqrt{m} = \mathbf{S} + \sqrt{\Delta}\mathbf{Z}/\sqrt{m}$, equation (46) takes the final form:

$$\begin{aligned} \Psi[\rho_Y, \rho_S] &= \frac{1}{8} \left\{ \frac{2(\alpha - 1)^2 \ln(1 - \alpha)}{\alpha} - \alpha(3 + \ln 4) + 2 \ln \alpha - 2\alpha \ln \pi\alpha + 6 \right\} \\ & - \frac{\sqrt{\alpha}}{2} \int \rho_S(dx) x - \frac{\alpha - 1}{2} \int \rho_S(dx) \ln x \\ & + \frac{\alpha}{2} \int \rho_S(dx) \rho_S(dy) \ln|x-y| - \frac{\alpha}{2} \int \rho_Y(dx) \rho_Y(dy) \ln|x-y|. \end{aligned} \tag{50}$$

Remark—Note that the derivation of equation (48) (presented in appendix C) does not rely on the expression of $I_{\Delta}[\rho_L, \rho_A]$ derived in [Mat94, GZ02] that involved Burgers’ equation, but is a simple consequence of the relation of the HCIZ integral to the large deviations of the Dyson Brownian motion [GZ02, Gui04a, BBMP14, Men17]. One can therefore recover equation (50) without appealing to the hydrodynamical formalism of Matytsin.

3.1.2. *The case $\alpha > 1$.* When $\alpha > 1$, the calculation is very similar, but one considers the limit spectral density of $\mathbf{S}' \equiv \mathbf{X}^T \mathbf{X} / \sqrt{nm}$, which is now a full-rank matrix, rather than the singular matrix $\mathbf{S} = \mathbf{X} \mathbf{X}^T / \sqrt{nm}$. Denoting the asymptotic spectral densities of these matrices by ρ_S and $\rho_{S'}$, we have:

$$\rho_S(x) = \alpha^{-1} \rho_{S'}'(x) + (1 - \alpha^{-1}) \delta(x).$$

The issue that arises in the calculation is that $\rho_S(x)$ has a singular component around $x = 0$. In order to regularize this singularity, we consider an arbitrarily small $\epsilon > 0$, and

we replace in the asymptotic result:

$$\rho_{\mathbf{S}} = \alpha^{-1} \rho'_{\mathbf{S}}(x) + (1 - \alpha^{-1}) \delta(x) \rightarrow \rho_{\mathbf{S}}^{(\epsilon)} = (\alpha^{-1} \rho'_{\mathbf{S}}(x) + (1 - \alpha^{-1}) \delta(x)) \boxplus \sigma_{\text{s.c.}}^{\epsilon}. \quad (51)$$

Here, $\sigma_{\text{s.c.}}^{\epsilon}$ is the asymptotic eigenvalue distribution of $\sqrt{\epsilon/m} \mathbf{Z}$ with \mathbf{Z} a Gaussian Wigner matrix¹⁵, and \boxplus is the free convolution [Voi86, AGZ10]. Assuming that the limit eigenvalue of \mathbf{S} is given by $\rho_{\mathbf{S}}^{(\epsilon)}$, we perform calculations completely similar to what we did for $\alpha \leq 1$, and we reach, at leading order in ϵ :

$$\begin{aligned} \Psi[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}] &= C(\alpha) + \frac{1}{2\alpha} \int \rho'_{\mathbf{S}}(dx) \rho'_{\mathbf{S}}(dy) \ln |x - y| - \frac{\alpha}{4} \int \left(\rho_{\mathbf{S}}^{(\epsilon)}(dx) \rho_{\mathbf{S}}^{(\epsilon)}(dy) \right. \\ &\quad \left. + \rho_{\mathbf{Y}}(dx) \rho_{\mathbf{Y}}(dy) \right) \ln |x - y| + \frac{1 - \alpha^{-1}}{2} \int \rho'_{\mathbf{S}}(dx) \ln x \\ &\quad - \frac{1}{2\sqrt{\alpha}} \int \rho'_{\mathbf{S}}(dx) x - \frac{\alpha}{4} \int_0^{\Delta} dt \int dx \rho^{(\epsilon)}(x, t) \left[\frac{\pi^2}{3} \rho^{(\epsilon)}(x, t)^2 + v^{(\epsilon)}(x, t)^2 \right], \end{aligned} \quad (52)$$

with the constant $C(\alpha)$ given by:

$$C(\alpha) = \frac{2(\alpha - 1)^2 \ln(\alpha - 1) + \alpha(-3\alpha - 2\alpha \ln(2\pi\alpha) + 2 \ln \alpha + 6)}{8\alpha}. \quad (53)$$

We also denoted $v^{(\epsilon)}(x, t)$ and $\rho^{(\epsilon)}(x, t)$ to make explicit the fact that the boundary condition at $t = 0$ is determined by $\rho_{\mathbf{S}}^{(\epsilon)}$. Exactly as in the case $\alpha \leq 1$ we make use of result 3.2, which allows to reach the simplification:

$$\begin{aligned} \Psi[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}] &= C(\alpha) + \frac{1}{2\alpha} \int \rho'_{\mathbf{S}}(dx) \rho'_{\mathbf{S}}(dy) \ln |x - y| + \frac{1 - \alpha^{-1}}{2} \int \rho'_{\mathbf{S}}(dx) \ln x \\ &\quad - \frac{1}{2\sqrt{\alpha}} \int \rho'_{\mathbf{S}}(dx) x - \frac{\alpha}{2} \int \rho_{\mathbf{Y}}(dx) \rho_{\mathbf{Y}}(dy) \ln |x - y|. \end{aligned} \quad (54)$$

3.2. The MMSE of denoising

The information-theoretic MMSE for denoising is defined as:

$$\text{MMSE}(\Delta) \equiv \mathbb{E} \left[\frac{1}{m} \sum_{\mu, \nu} (S_{\mu\nu}^* - \langle S_{\mu\nu} \rangle)^2 \right], \quad (55)$$

where $\mathbf{S}^* = \mathbf{X}^* (\mathbf{X}^*)^{\top} / \sqrt{nm}$ is the original ‘ground-truth’ matrix, and the observation is $\mathbf{Y} = \sqrt{m} \mathbf{S}^* + \sqrt{\Delta} \mathbf{Z}$. Note that adding the diagonal terms $\mu = \nu$ in this definition does not change the asymptotic MMSE, cf appendix G.2. This definition satisfies, for any α :

$$\lim_{\Delta \rightarrow \infty} \text{MMSE}(\Delta) = 1. \quad (56)$$

¹⁵In simpler terms, $\sigma_{\text{s.c.}}^{\epsilon}$ is a Wigner semi-circle law extending from $-2\sqrt{\epsilon}$ to $2\sqrt{\epsilon}$.

As is usually the case for Gaussian channels [GSV05], the MMSE is simply given by a derivative of the asymptotic free entropy:

$$\text{MMSE}(\Delta) = \Delta - \frac{4}{\alpha} \frac{\partial \Psi[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}]}{\partial \Delta^{-1}}. \tag{57}$$

A proof of equation (57) is given in appendix G.3. We can now compute the MMSE explicitly using equations (57), (46) and (54). Therefore the MMSE takes the form:

$$\text{MMSE}(\Delta) = \Delta + \frac{4}{\alpha} \Delta^2 \frac{\partial}{\partial \Delta} \{\Psi[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}]\} = \Delta - 2\Delta^2 \frac{\partial}{\partial \Delta} \int \rho_{\mathbf{Y}}(dx) \rho_{\mathbf{Y}}(dy) \ln |x - y|. \tag{58}$$

Recall that $\rho_{\mathbf{Y}}$ is the asymptotic eigenvalue density of \mathbf{Y}/\sqrt{m} . As we show below, this can be computed using free probability techniques.

A ‘miraculously simple’ solution—Equation (58) is particularly simple, as it only depends on the eigenvalue distribution of the *observations* \mathbf{Y} . This was noticed in [BABP16] for an important class of RIE. Here, this arises as a consequence of the fact that the only dependency in Δ in the free entropy of equations (46) and (54) is through the eigenvalue density of \mathbf{Y} .

Numerical evaluation of the MMSE—In order to compute the Stieltjes transform $g_{\mathbf{Y}}(z)$ (and from it $\rho_{\mathbf{Y}}$ and $v_{\mathbf{Y}}$ by equation (47)), we can use free probability techniques. Indeed, we know the \mathcal{R} -transform of \mathbf{Y}/\sqrt{m} , since we have $\mathbf{Y}/\sqrt{m} = \mathbf{X}\mathbf{X}^\top/\sqrt{nm} + (\Delta/m)^{1/2}\mathbf{Z}$:

$$\mathcal{R}_{\mathbf{Y}}(s) = g_{\mathbf{Y}}^{-1}(-s) - \frac{1}{s} = \frac{1}{\sqrt{\alpha}(1 - \sqrt{\alpha}s)} + \Delta s. \tag{59}$$

This implies that $s = -g_{\mathbf{Y}}(z)$ can be found as the solution to the algebraic equation:

$$s^3 - \frac{\sqrt{\alpha}z + \Delta}{\sqrt{\alpha}\Delta} s^2 + \frac{1}{\sqrt{\alpha}\Delta} (z + \sqrt{\alpha} - \alpha^{-1/2})s - \frac{1}{\sqrt{\alpha}\Delta} = 0. \tag{60}$$

Equation (60) can be solved numerically. We then compute the derivative of the logarithmic potential in equation (58) using finite differences, which allows us to efficiently access the MMSE numerically.

3.3. Bayes-optimal estimator for denoising

3.3.1. From asymptotic results to estimation. While the derivations presented in sections 3.1 and 3.2 allow to characterize the performance of the Bayes-optimal estimator, they did not provide an actual estimation algorithm matching this performance. In this section, we work out such an estimator using the asymptotic limits described above. From Bayesian statistics, the Bayes-optimal estimator of $\mathbf{S} \equiv \mathbf{X}\mathbf{X}^\top/\sqrt{nm}$ is $\hat{\mathbf{S}} \equiv \langle \mathbf{S} \rangle$, with $\langle \cdot \rangle$ the average under the posterior distribution of equation (36).

Starting again from equation (36), we have:

$$\begin{cases} \langle S_{\mu\nu} \rangle = \frac{Y_{\mu\nu}}{\sqrt{m}} - \Delta n \frac{\partial \Phi_{\mathbf{Y},n}}{\partial (Y_{\mu\nu}/\sqrt{m})}, \\ \langle S_{\mu\mu} \rangle = \frac{Y_{\mu\mu}}{\sqrt{m}} - 2\Delta n \frac{\partial \Phi_{\mathbf{Y},n}}{\partial (Y_{\mu\mu}/\sqrt{m})}. \end{cases} \quad (61)$$

Importantly, our results of equations (46) and (54) imply that, as $n \rightarrow \infty$, $\Phi_{\mathbf{Y},n}$ depends on \mathbf{Y} solely via the spectrum of \mathbf{Y}/\sqrt{m} . This means that we can write equation (61) at leading order as $n, m \rightarrow \infty$ as:

$$\begin{cases} \langle S_{\mu\nu} \rangle = \frac{Y_{\mu\nu}}{\sqrt{m}} - \Delta n \sum_{\rho=1}^m \frac{\partial \Phi_{\mathbf{Y},n}}{\partial y_\rho} \frac{\partial y_\rho}{\partial (Y_{\mu\nu}/\sqrt{m})}, \\ \langle S_{\mu\mu} \rangle = \frac{Y_{\mu\mu}}{\sqrt{m}} - 2\Delta n \sum_{\rho=1}^m \frac{\partial \Phi_{\mathbf{Y},n}}{\partial y_\rho} \frac{\partial y_\rho}{\partial (Y_{\mu\mu}/\sqrt{m})}, \end{cases} \quad (62)$$

in which $\{y_\rho\}_{\rho=1}^m$ are the eigenvalues of \mathbf{Y}/\sqrt{m} . Let us denote $\{\mathbf{v}^\rho\}$ the corresponding eigenvectors, and fix $\mu_0 < \nu_0$. By the Hellman–Feynman theorem, if we perturb the matrix \mathbf{Y}/\sqrt{m} as:

$$Y_{\mu\nu} \rightarrow Y_{\mu\nu}^\gamma = Y_{\mu\nu} + \gamma\sqrt{m}\delta_{\mu\mu_0}\delta_{\nu\nu_0}, \quad (63)$$

for a small $\gamma \in \mathbb{R}$, then the corresponding transformation of an eigenvalue y_ρ is:

$$\frac{dy_\rho}{d\gamma} = \frac{\partial y_\rho}{\partial (Y_{\mu_0\nu_0}/\sqrt{m})} = \sum_{\mu\nu} (\delta_{\mu\mu_0}\delta_{\nu\nu_0} + \delta_{\mu\nu_0}\delta_{\nu\mu_0}) v_\mu^\rho v_\nu^\rho = 2v_{\mu_0}^\rho v_{\nu_0}^\rho. \quad (64)$$

In the same way, we obtain for $\mu_0 = \nu_0$:

$$\frac{\partial y_\rho}{\partial (Y_{\mu_0\mu_0}/\sqrt{m})} = \sum_{\mu\nu} \delta_{\mu\mu_0}\delta_{\nu\mu_0} v_\mu^\rho v_\nu^\rho = (v_{\mu_0}^\rho)^2. \quad (65)$$

Plugging equations (64) and (65) in equation (62) yields, for any μ, ν :

$$\langle S_{\mu\nu} \rangle = \frac{Y_{\mu\nu}}{\sqrt{m}} - 2\Delta n \sum_{\rho=1}^m \frac{\partial \Phi_{\mathbf{Y},n}}{\partial y_\rho} v_\mu^\rho v_\nu^\rho = \sum_{\rho=1}^m \left[y_\rho - 2\Delta n \frac{\partial \Phi_{\mathbf{Y},n}}{\partial y_\rho} \right] v_\mu^\rho v_\nu^\rho. \quad (66)$$

Note that equations (50) and (54) have the following form at finite n :

$$\Phi_{\mathbf{Y},n} = F_n(\rho\mathbf{s}) - \frac{1}{2mn} \sum_{\mu \neq \nu} \ln |y_\mu - y_\nu| + \mathcal{O}_n(1), \quad (67)$$

so that

$$n \frac{\partial \Phi_{\mathbf{Y},n}}{\partial y_\mu} = -\frac{1}{m} \sum_{\nu(\neq\mu)} \frac{1}{y_\mu - y_\nu} + \mathcal{O}_n(1) = \text{P.V.} \int dy \rho_{\mathbf{Y}}(y) \frac{1}{y - y_\mu} + \mathcal{O}_n(1), \quad (68)$$

with P.V. the principal value of the integral. Stated differently, the Bayes-optimal estimator has the form:

$$\begin{cases} \hat{\mathbf{S}} = \sum_{\mu=1}^m \hat{\xi}_\mu \mathbf{v}^\mu (\mathbf{v}^\mu)^\top, \\ \hat{\xi}_\mu = y_\mu - 2\Delta v_{\mathbf{Y}}(y_\mu, \Delta). \end{cases} \quad (69)$$

In equation (69) we introduced the function:

$$v_{\mathbf{Y}}(x, \Delta) \equiv -\lim_{\epsilon \downarrow 0} \text{Re}[g_{\mathbf{Y}}(x + i\epsilon)], \quad (70)$$

with $g_{\mathbf{Y}}(z)$ the Stieltjes transform of \mathbf{Y}/\sqrt{m} , cf appendix A.

An important consequence of equation (69) is that the optimal estimator $\hat{\mathbf{S}}$ is *diagonal in the eigenbasis of the observation \mathbf{Y}* . Such estimators are usually called RIE. Note that one can argue that the optimal estimator is a RIE in this problem without going through the derivation above, but simply as a consequence of the rotation-invariance of the posterior distribution of equation (8), which was noticed in [BBP17]. The optimal RIE (in the L^2 sense) for the denoising of an extensive-rank matrix corrupted by Gaussian noise has been worked out in [BABP16], in which the authors exactly obtained the expression of equation (69). We therefore found back their main result using our asymptotic derivation¹⁶.

As we mentioned at the beginning of section 3, one can easily see that none of the results on the MMSE or the optimal denoiser are specific to the denoising of Wishart matrices, but can be easily generalized to any rotationally-invariant matrix, so that our derivation can be compared to the one of [BABP16] based on a direct minimization of the L^2 error, in the case of Gaussian additive noise. As we mentioned above, we leave for future work an extension of our derivation to completely generic additive rotationally-invariant noise (which is analyzed in [BABP16]). Note that the mean squared error of the optimal estimator should be the MMSE of equation (58). We have checked this numerically (see section 3.5) for two types of signals. Figure 4 shows that the two results match, as expected.

3.3.2. Small- α expansion of the optimal denoiser. We shall use the optimal estimator of equation (69) in order to test our PGY expansion for denoising in the Gaussian setting. Recall that we are denoising a $m \times m$ Wishart matrix $\mathbf{S} = \mathbf{X}\mathbf{X}^\top/\sqrt{nm}$, where \mathbf{X} is $m \times n$, and $\alpha = n/m$. We shall consider here an expansion of the denoiser in the overcomplete regime of small $\alpha = m/n$ (i.e. close to the limit where \mathbf{S} becomes Gaussian), and work out in the next section the small- α expansion within the PGY formalism.

¹⁶ As noticed in [BABP16], equation (69) is in a sense a ‘miracle’ similar to the one we described for the MMSE in section 3.2: the eigenvalues $\{\hat{\xi}_\mu\}$ of the optimal estimator are expressed solely as a function of the spectrum of the observations \mathbf{Y} .

Note that in the PGY approach, we will not consider diagonal observations: we use $Y_{\mu\mu} = 0$. Equivalently, by rotation-invariance, we assume here that \mathbf{Y} has asymptotically zero trace, which we can do without any loss of generality. Simply, we shift the observations by a deterministic term:

$$\mathbf{Y} = \frac{1}{\sqrt{n}}\mathbf{X}\mathbf{X}^\top - \frac{1}{\sqrt{\alpha}}\mathbf{I}_m + \sqrt{\Delta}\mathbf{Z}. \quad (71)$$

The computation detailed in appendix D expands the Stieltjes transform of \mathbf{Y}/\sqrt{m} and gives

$$v_{\mathbf{Y}}(y) = \frac{y}{2(\Delta + 1)} + \frac{(\Delta + 1 - y^2)}{2(\Delta + 1)^3}\sqrt{\alpha} + \mathcal{O}(\alpha). \quad (72)$$

Using this expansion in the optimal estimator of equation (69), we get the result:

$$\hat{\xi}_\mu = \frac{1}{\Delta + 1}y_\mu - \frac{\Delta}{(\Delta + 1)^2}\left[1 - \frac{y_\mu^2}{\Delta + 1}\right]\sqrt{\alpha} + \mathcal{O}(\alpha). \quad (73)$$

This expansion describes how every eigenvalue y of \sqrt{m} should be denoised, in the limit of small α . The leading term, $\hat{\xi}_\mu = y_\mu/(\Delta + 1)$, is the standard expression for denoising a Wigner matrix, as expected. The term of order $\sqrt{\alpha}$ is the first non-trivial correction, we shall compare it to PGY-based denoising in the next section.

3.4. Denoising using the PGY expansion

3.4.1. PGY expansion for the denoising problem. In this section, we apply the PGY expansion formalism to the symmetric matrix denoising problem, for a Wishart matrix $\mathbf{S} \equiv \mathbf{X}\mathbf{X}^\top/\sqrt{nm}$. Recall that we observe $Y_{\mu\nu} \sim P_{\text{out}}(\cdot|\sqrt{m}S_{\mu\nu})$. We consider a general channel P_{out} (where the noise is applied independently to every component of the matrix), and a general prior, factorized over the components of the input matrix: $P_X(\mathbf{X}) = \prod_{\mu,i} P_X(X_{\mu i})$. Furthermore, we shall assume that the mean of P_X is 0. Without loss of generality (as it can be absorbed in the channel), we also assume that the variance of the prior is $\mathbb{E}_{P_X}[x^2] = 1$.

The PGY expansion is very close to the computation that we did for the matrix factorization problem, except that, in the denoising problem, we do not fix the first and second moments of the field \mathbf{X} . The free entropy thus takes a form which is simpler than the one of equation (14). We get:

$$nm\Phi_{\mathbf{Y},n} = \sum_{\mu<\nu} \left[-\omega_{\mu\nu}g_{\mu\nu} - \frac{b_{\mu\nu}}{2}(-r_{\mu\nu} + g_{\mu\nu}^2) \right] + \ln \int P_H(d\mathbf{H}) P_X(d\mathbf{x}) e^{-S_{\text{eff}}[\mathbf{x},\mathbf{H}]}, \quad (74)$$

in which we introduced an *effective action*:

$$S_{\text{eff}}[\mathbf{x}, \mathbf{H}] = \sum_{\mu<\nu} \left[\omega_{\mu\nu}(iH)_{\mu\nu} - \frac{b_{\mu\nu}}{2}(iH)_{\mu\nu}^2 \right] + H_{\text{eff}}[\mathbf{x}, \mathbf{H}], \quad (75)$$

and the (un-normalized) distribution and effective Hamiltonian:

$$\begin{cases} P_H[d\mathbf{H}] \equiv d\mathbf{H} \prod_{\mu < \nu} \int \frac{d\tilde{H}}{2\pi} e^{iH_{\mu\nu}\tilde{H}} P_{\text{out}}(Y_{\mu\nu}|\tilde{H}) = \prod_{\mu < \nu} \frac{1}{2\pi} e^{-\frac{\Delta H_{\mu\nu}^2}{2} + iY_{\mu\nu}H_{\mu\nu}} dH_{\mu\nu}, \\ H_{\text{eff}}[\mathbf{x}, \mathbf{H}] \equiv \frac{1}{\sqrt{n}} \sum_{\mu < \nu} \sum_i (iH)_{\mu\nu} x_{\mu i} x_{\nu i}. \end{cases} \quad (76)$$

We only impose the constraints $g_{\mu\nu} = -\langle (iH)_{\mu\nu} \rangle$ and $r_{\mu\nu} = -\langle (iH)_{\mu\nu}^2 \rangle + g_{\mu\nu}^2$, using the Lagrange multipliers $\omega_{\mu\nu}$ and $b_{\mu\nu}$.

The PGY expansion proceeds exactly as in section 2.1: for pedagogical reasons, we detail the PGY expansion specifically for the denoising problem in appendix E. As shown there, the result for denoising can be obtained by using the PGY expansion for matrix factorization and imposing $m_{\mu i} = 0$, $v_{\mu i} = 1$, $\lambda_{\mu i} = 0$, $\gamma_{\mu i} = 0$. The PGY free entropy for denoising at order η^3 takes the form:

$$\begin{aligned} nm\Phi_{\mathbf{Y},n} = & \sum_{\mu < \nu} \left[-\omega_{\mu\nu} g_{\mu\nu} - \frac{b_{\mu\nu}}{2} (-r_{\mu\nu} + g_{\mu\nu}^2) + \ln \int dz \frac{e^{-\frac{1}{2b_{\mu\nu}}(z-\omega_{\mu\nu})^2}}{\sqrt{2\pi b_{\mu\nu}}} P_{\text{out}}(Y_{\mu\nu}|z) \right] \\ & + \frac{\eta^2}{2} \sum_{\mu < \nu} [g_{\mu\nu}^2 - r_{\mu\nu}] + \frac{\eta^3}{6n^{1/2}} \sum_{\substack{\mu_1, \mu_2, \mu_3 \\ \text{pairwise distinct}}} g_{\mu_1\mu_2} g_{\mu_2\mu_3} g_{\mu_3\mu_1}. \end{aligned} \quad (77)$$

3.4.2. *Small- α expansion of the PGY denoiser.* We can now expand the denoising estimator predicted by the PGY expansion at order 3 in α as we did for the optimal denoiser in section 3.3.2. We also remind what we noted above (see section 2.1.3) that the PGY expansion at order η^3 corresponds to an expansion of the optimal denoiser in α , up to order $\sqrt{\alpha}$.

We start from the expression of the free entropy in equation (77), evaluated at $\eta = 1$. Using the fact that the diagonal elements of \mathbf{g} are zero, the last term is just proportional to $\text{Tr}[\mathbf{g}^3]$. Using that $\partial_{g_{\mu\nu}} \text{Tr}[\mathbf{g}^k] = 2k(\mathbf{g}^{k-1})_{\mu\nu}$, the TAP equations corresponding to equation (77) are:

$$\begin{cases} b_{\mu\nu} = 1 & ; & g_{\mu\nu} = \frac{Y_{\mu\nu} - \omega_{\mu\nu}}{\Delta + b_{\mu\nu}}, \\ \omega_{\mu\nu} = -b_{\mu\nu} g_{\mu\nu} + g_{\mu\nu} + \frac{1}{\sqrt{n}} (\mathbf{g}^2)_{\mu\nu} & ; & r_{\mu\nu} = \frac{1}{\Delta + b_{\mu\nu}}. \end{cases} \quad (78)$$

Therefore this can be reduced to a single equation on \mathbf{g} :

$$\frac{\mathbf{Y}}{\sqrt{m}} = (\Delta + 1) \frac{\mathbf{g}}{\sqrt{m}} + \sqrt{\alpha} \left(\frac{\mathbf{g}}{\sqrt{m}} \right)^2 - \frac{\sqrt{\alpha}}{m} \text{Tr} \left[\frac{\mathbf{g}^2}{m} \right] \mathbf{I}_m. \quad (79)$$

Note that in equation (79) we added a last term to impose the zero-trace condition on \mathbf{g} and \mathbf{Y} , since equation (78) only describes the off-diagonal terms¹⁷. We solve

¹⁷Since these are rotationally-invariant matrices, having a zero trace is asymptotically equivalent to having a zero diagonal.

equation (79) perturbatively in α .¹⁸ We denote by y a generic eigenvalue of \mathbf{Y}/\sqrt{m} , and g_y the corresponding eigenvalue of \mathbf{g}/\sqrt{m} (as is clear from equation (79), the two matrices will be diagonal in the same eigenbasis). The first order is trivial, and given by:

$$g_y = \frac{y}{\Delta + 1} + \mathcal{O}(\sqrt{\alpha}). \tag{80}$$

Plugging this back into equation (79) and considering the order $\sqrt{\alpha}$, we reach (imposing the zero-trace condition):

$$g_y = \frac{y}{\Delta + 1} + \frac{(\Delta + 1 - y^2)}{(\Delta + 1)^3} \sqrt{\alpha} + \mathcal{O}(\alpha). \tag{81}$$

This is the same result as the one obtained with the RIE of equation (73), and this is the expression plotted in figure 1 in orange. Therefore the PGY expansion does capture the asymptotic optimal estimator of the problem (in the overcomplete regime), as it should.

Effect of the prior on the denoising problem—Note that the PGY expansion suggests that the Bayes-optimal estimator for denoising is independent of the particularities of the prior distribution P_X , and only depends on its first two moments, assumed here to be fixed. While this observation is only made here up to order η^3 , understanding if (and how) an i.i.d. prior on \mathbf{X} might enhance the denoising of the matrix $\mathbf{X}\mathbf{X}^\top$ by going to higher orders in the PGY expansion is beyond our results. Note finally that the universality suggested by the PGY expansion at order 3 might be better understood in terms of random matrix theory, in which many such universality results are known in the asymptotic limit, e.g. for the bulk spectral statistics of $\mathbf{X}\mathbf{X}^\top$ [MP67].

3.5. Numerical simulations

In this section, we present numerical results illustrating our contributions. The code is available in a public GitHub repository [MKMZ22].

Denoising MMSE for Wishart denoising—In figure 2, we compare the denoising errors reached by different procedures:

- The MMSE, that is the analytical value of the Bayes-optimal error predicted in equation (58), is shown in the full red line. We discretize both parameters (t, x) , to obtain precise estimates of the involved integrals via Romberg’s method [Rom55], and a polynomial solver to compute the Stieltjes transform of $\mathbf{Y}(t)$ above the real axis by using equation (60).
- The denoising MSE reached by the optimal RIE of equation (69). For each value of Δ (shown as green points), we compute the MSE achieved on three different instances of \mathbf{Y} , and error bars are too small to be visible. We use $m = 3000$ to generate the matrices.

¹⁸Note that one could also solve equation (78) iteratively, as what would usually be done for TAP equations. Here, we rather leverage an analytical expression for the spectrum of the solution. In the more general matrix factorization case, there is no such expression for the solution to the TAP equations (18) and (26), and we found ‘naive’ iterations to be numerically very ill-behaved, see our remark at the end of section 2.1.

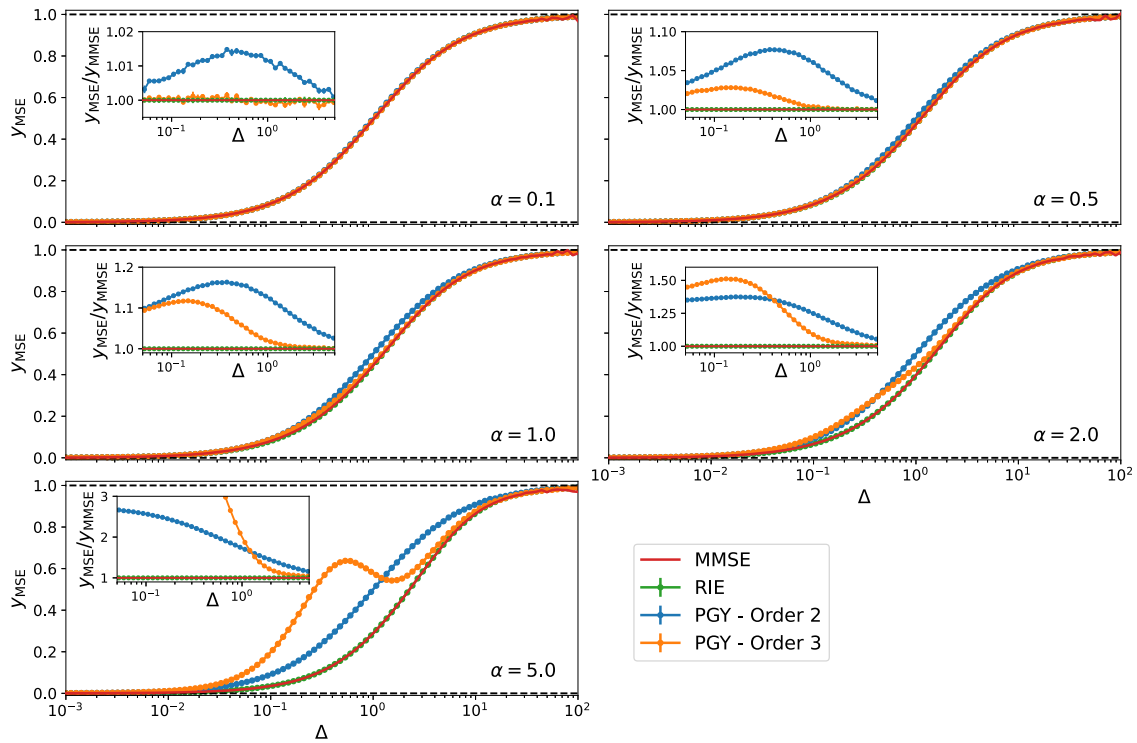


Figure 2. (Extended version of figure 1) Denoising of a Wishart matrix, for different values of the aspect ratio α as a function of the additive noise variance Δ : comparison of the analytic MMSE with the RIE estimator of [BABP16], and the TAP-based estimator obtained from the PGY expansion truncated at orders 2 and 3.

- In blue points, we show the error reached by the PGY equations at order η^2 , that is simply via the denoiser of equation (80). As for the RIE, we use $m = 3000$ and average over three instances, while error bars are in practice invisible. Note that the truncation at order 2 corresponds to assuming that the underlying matrix is a Wigner matrix, as clearly shown in equation (80).
- Finally, orange points correspond to the PGY equations truncated at order η^3 , that is equation (81). We use the same physical parameters as for the PGY at order 2.

In the inset of all figures shown in figure 2, we show the ratio of the error with the analytical MMSE. It is clear that the optimal RIE achieves the Bayes-optimal error within numerical accuracy. Moreover, the PGY at order 3 significantly outperforms the order-2 method when α is not too large (recall that the order of the PGY expansion can be understood as an order of perturbation in α). This strengthens our claim that the PGY expansion captures the optimal estimator, although one would have to take into account all orders of perturbation to turn it into an estimator for the factorization problem, something which is beyond the scope of this paper. Interestingly, for large values of α (e.g. $\alpha = 5$ in figure 2) and small values of Δ , the order-3 method can perform worse than the order-2. Such a peculiar behavior of the PGY expansion series was already noticed for a class of Ising models in [RT12]. We remark that our evidence

Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising




Figure 3. Denoising of a Wishart matrix: comparison of the analytical free entropy computed using Matytsin’s equations (red curve), the one computed with the TAP approach, truncated at orders 2 (blue points) and 3 (orange points), the annealed free energy (dashed pink line) computed in appendix F, the denoising solution of [KKM+16] in dashed cyan, and the (incorrect) solution for the factorization problem proposed in [KKM+16] in purple.

combined with the one of [RT12] suggests that such a phenomenon might also be present in the general PGY expansion of many rotationally-invariant inference problems which was derived in [MFC+19]: this interesting question is however not in the scope of the present paper and would deserve further studies.

Estimating the asymptotic free entropy—Similar observations can be made by computing the free entropy predicted by the solution to the Matytsin problem and comparing it to the predictions of the truncated PGY expansion. We present our results in figure 3, in which we show the following curves for different values of α :

- The red curve is obtained by analytically estimating equations (46) and (54), depending on the regime of α . The discretization technique we use was detailed in the description of figure 2.
- The blue and oranges points correspond to truncating the PGY expansion at order 2 and 3, as in figure 2, and the physical parameters used in the simulations are the same. Note that truncating the PGY series at order 3 does not necessarily yield a better approximation (to the free entropy or the MSE, see figure 2) than order 2 for low values of Δ : indeed, this truncation does not correspond to any physical

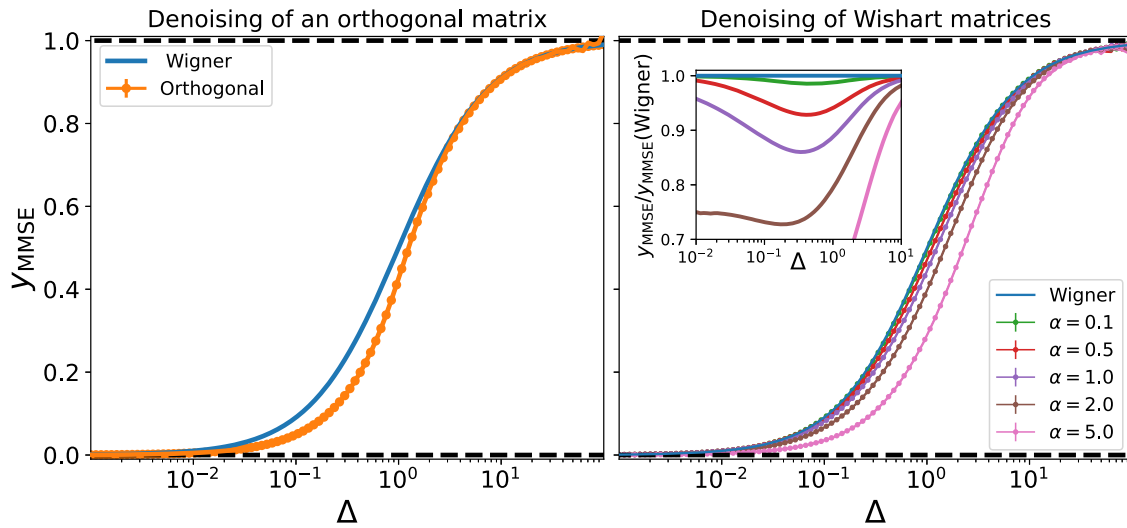


Figure 4. Comparison of the denoising MMSE computed from the optimal RIE of [BABP16], i.e. equation (69), for different ground-truth matrices \mathbf{S}^* . We use $m = 3000$ and average over three instances, with no visible error bars. (Left) Denoising of a uniformly-sampled symmetric orthogonal matrix. (Right) Denoising of a Wishart matrix with various values of $m/n = \alpha$. The well known denoising of a Wigner random matrix (which is the overcomplete limit of a Wishart matrix when $\alpha \rightarrow 0$) is shown for reference.

approximation, while we saw that the truncation at order 2 is equivalent to assuming that the matrix to denoise is Wigner.

- In dashed cyan, we show the free entropy prediction of [KKM⁺16] for the matrix denoising problem, matching exactly the PGY expansion at order 2. Note that [KKM⁺16] also tackled the more involved matrix factorization problem. Since the free energies of factorization and denoising are the same, we show as well in solid purple their (incorrect) prediction for the free entropy when analyzing the matrix factorization problem. Recall that we analytically disproved their approach in section 2.2. The computation of these two curves is elementary, and described in appendix G.1.
- For completeness, we compare all these curves to an annealed bound (in dashed pink) that can be computed analytically, using a calculation presented in appendix F. In particular, the PGY expansion at order 3 clearly violates the annealed bound for small enough Δ and large enough α , indicating that it does not correspond to a physical free entropy and that the PGY expansion would have to be carried to further orders to correct this behavior.

Denoising of other rotationally-invariant matrices—In figure 4, we illustrate the influence of the structure of the signal on the performance of the optimal denoiser. We show how the structure of a symmetric orthogonal matrix (i.e. $\mathbf{S}^* = \mathbf{O}\mathbf{O}^T$, with a uniformly-sampled orthogonal matrix \mathbf{O} , and $D_\mu \stackrel{\text{i.i.d.}}{=} \pm 1$ with probability 1/2), or of a Wishart matrix with different values of α , allows to significantly improve the denoising performance over simple Wigner (scalar) denoising.




Figure 5. Denoising MMSE for Wishart denoising at large α , and different values of Δ . We show both analytical predictions using equation (58) (dashed lines) and numerical applications of the RIE of equation (69) (using $m = 20\,000$ and averaging over three instances, error bars are invisible, dashed lines are covered by the corresponding data points). Numerical errors arise in the MMSE calculation, due to the difficulty of estimating the derivative in equation (69) for very large values of the noise. We compare it with the performance of vanilla PCA, which is optimal in the low-rank limit $n = \mathcal{O}(1)$.

Large- α behavior and the BBP transition—We conclude this presentation of numerical results by figure 5, in which we investigate the denoising MMSE in the large α (or undercomplete) regime. While this regime is not reachable by our PGY expansion, we can evaluate the MMSE with the exact formula of equation (58) (dashed lines), or by sampling large matrices and using the optimal RIE of equation (69). The agreement between these two methods is again extremely good, showing in particular that the RIE calculation is not affected strongly by finite-size effects (e.g. at $\alpha = 200$, we have $n = m/\alpha = 100$, so that finite-size effects could have been present).

Interestingly, we can compare the performance in this regime to the low-rank limit in which $n = \mathcal{O}(1)$. In the extreme case in which $n = 1$ (so that $\mathbf{X} = \mathbf{x}$ is a vector), it is known that the optimal procedure is PCA, which estimates \mathbf{x} as the leading eigenvector of $\mathbf{Y}/\sqrt{\alpha} = \mathbf{x}\mathbf{x}^\top/\sqrt{m} + \sqrt{\Delta/\alpha}\mathbf{Z}$. Its performance has been described very precisely, and exhibits a transition, known as ‘BBP’ [BBAP05] at $\Delta/\alpha = 1$. More precisely, if \mathbf{v} is the leading eigenvector of \mathbf{Y} , we have

$$\lim_{m \rightarrow \infty} \frac{1}{m} |\mathbf{v}^\top \mathbf{x}|^2 = \begin{cases} 0 & \text{if } \Delta \geq \alpha \\ 1 - \frac{\Delta}{\alpha} & \text{if } \Delta \leq \alpha. \end{cases} \quad (82)$$

We show the corresponding denoising MMSE as a pink line in figure 5. It is clear that the optimal denoiser approaches the PCA performance at large α . We note that in the extensive-rank regime we observe no phase transition in the MMSE. We will further comment on the description of the transition between the low-rank and extensive-rank regimes in the coming conclusion.

4. Conclusion and openings

In this paper, we have shown how the PGY formalism of high-temperature expansions at fixed order parameters can be applied to the problems of extensive-rank matrix factorization and denoising. The previous approaches to this problem [SK13, KMZ13, KKM⁺16, PSC14a, PSC14b, ZZY21] turn out to be relying on a hidden assumption that amounts to truncating the PGY series at second order. Both in the matrix factorization and in the denoising problem, we have shown, by an explicit computation of the third order contribution, that the higher-order corrections are not negligible.

The denoising of rotationally-invariant matrices provides a good test of our approach. On the one hand it can be solved analytically using matrix integration techniques. On the other hand there exists an optimal denoising algorithm that relies only on the observed matrix, and denoises each eigenvalue. We have shown that these two approaches give the same mean squared error, and that our PGY-based expansion, truncated at order three, also agrees with these other approaches for denoising Wishart matrices in an overcomplete (or small- α) regime close to the Wigner limit.

The interest of the formalism that we have presented here is that it opens the way to an optimality-matching algorithm for matrix factorization in the extensive rank case accounting for general prior and output channel. Indeed, the optimization of the PGY free entropy leads to TAP equations which, properly iterated in time, develop into AMP-type algorithms. An important step in that direction would of course be to work out the PGY expansion to all orders, and re-sum it. On the other hand, as a first step, it will be interesting to find out if the TAP equations that we have derived at order 3 can be transformed into a useful algorithm for overcomplete matrix factorization.

It would also be interesting to perform a perturbative calculation of the free entropy and MMSE in the opposite regime of large α , to analytically understand the limit behavior observed in figure 5. Such a calculation is related to an important open problem in random matrix theory, that is a sharp description of the transition between the low-rank and extensive-rank spherical HCIZ integrals. Solving this later problem would help to describe the transition between low-rank and extensive-rank results in the PGY expansion.

Another interesting open problem is to relate the PGY-expansion to other approaches. The recent conjecture proposed in [BM21] for solving matrix factorization with replicas should be compared to the present approach. In general it is well known that the two main approaches to disordered systems, the replica approach and the cavity approach (that derives TAP equations) have the same physical content, the second one having the advantage that it can be transformed into an algorithm. The challenge that we face here is the extension of these two methods to the case where the order parameter is a large matrix, and the contact that can be made between the two. In particular, in both of these approaches, we have found that the effective order parameter turns out to be a distribution of eigenvalues. Understanding the generality of this statement is an important open challenge. On the other hand, it has been shown in [MFC⁺19] in finite-rank problems that these expansions are equivalent to other techniques, e.g. the EC or adaTAP approaches. This equivalence does not seem to easily

transfer to the extensive-rank case: understanding how to apply these approaches here is another interesting direction of research.

Acknowledgments

We are grateful to Laura Foini, Christian Schmidt, Alice Guionnet, Jean-Philippe Bouchaud, and Yoshiyuki Kabashima for many insightful discussions throughout different stages of this work. A M is particularly thankful to Laura Foini for collaboration on the annealed calculations and discussions on the PGY expansion. We acknowledge funding from the ERC under the European Union’s Horizon 2020 Research and Innovation Program Grant Agreement 714608-SMiLe. Additional funding is acknowledged by AM from ‘Chaire de recherche sur les modèles et sciences des données’, Fondation CFM pour la Recherche-ENS.

Appendix A. Some definitions of probability theory

We introduce here a few notations and quantities particularly useful to study the asymptotic spectra of random matrices. For more complete mathematical or physical introductions to random matrix theory, the reader can refer to [Meh04, AGZ10, LNV18]. Let us consider a symmetric random matrix $M \in \mathbb{R}^{n \times n}$, with eigenvalues $(\lambda_1, \dots, \lambda_n)$. In general, we will consider random matrices that admit an *asymptotic spectral density*, here denoted $\mu(x)$, such that

$$\mu(x) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \delta(x - \lambda_i). \tag{83}$$

Here the limit is to be understood as the weak limit of probability measures. We denote $\lambda_{\max} \equiv \max \text{supp}(\mu)$ and $\lambda_{\min} \equiv \min \text{supp}(\mu)$, assuming that the support of μ is bounded from above and below. Letting $\mathbb{C}_+ \equiv \{z \in \mathbb{C} \text{ s.t. } \text{Im}(z) > 0\}$, one can introduce the *Stieltjes* transform of μ as:

$$g_\mu(z) \equiv \mathbb{E} \left[\frac{1}{X - z} \right] = \int_{\mathbb{R}} d\lambda \mu(\lambda) \frac{1}{\lambda - z} = \lim_{n \rightarrow \infty} \frac{1}{\lambda_i - z}, \tag{84}$$

for all $z \in \mathbb{C}_+$. Note that then we also have $g_\mu(z) \in \mathbb{C}_+$. Moreover, on $(\lambda_{\max}, +\infty)$, g_μ induces a strictly increasing \mathcal{C}^∞ diffeomorphism $\mathcal{S}_\mu : (\lambda_{\max}, \infty) \leftrightarrow (-\infty, 0)$, and we denote its inverse g_μ^{-1} . One can then introduce the *\mathcal{R} -transform* of μ as:

$$\forall z > 0, \quad \mathcal{R}_\mu(z) \equiv g_\mu^{-1}(-z) - \frac{1}{z}. \tag{85}$$

$\mathcal{R}_\mu(z)$ is *a priori* defined for $-z \in g_\mu[(\lambda_{\min}, \lambda_{\max})^c]$ and admits an analytical expansion around $z = 0$. We can write this expansion as:

$$\mathcal{R}_\mu(z) = \sum_{k=0}^{\infty} c_{k+1}(\mu) z^k. \tag{86}$$

J. Stat. Mech. (2022) 083301

The elements of the sequence $\{c_k(\mu)\}_{k \in \mathbb{N}^*}$ are called the *free cumulants* of μ . In particular, one can show that $c_1(\mu) = \mathbb{E}_\mu(X)$ and $c_2(\mu) = \mathbb{E}_\mu(X^2) - (\mathbb{E}_\mu X)^2$. The free cumulants can be recursively computed from the moments of the measure using the so-called *free cumulant equation*:

$$\forall k \in \mathbb{N}^*, \quad \mathbb{E}_\mu X^k = \sum_{m=1}^k c_m(\mu) \sum_{\substack{\{k_i\}_{i \in [1, m]} \\ \text{s.t. } \sum_i k_i = k}} \prod_{i=1}^m \mathbb{E}_\mu X^{k_i-1}. \quad (87)$$

Appendix B. Technicalities of the PGY expansion

B.1. Setting of the PGY expansion

We describe here in more details the formalism we used to derive result 2.1: we follow the formalism of [GY91] to perform a Plefka expansion [Ple82]. Note that some parts of the derivation are very similar to what is done in [MFC⁺19]. We start from equation (14):

$$\begin{cases} nm\Phi_{\mathbf{Y},n} &= \sum_{\mu,i} \left[\lambda_{\mu i} m_{\mu i} + \frac{\gamma_{\mu i}}{2} (v_{\mu i} + (m_{\mu i})^2) \right] + \sum_{\mu < \nu} \left[-\omega_{\mu\nu} g_{\mu\nu} - \frac{b_{\mu\nu}}{2} (-r_{\mu\nu} + g_{\mu\nu}^2) \right] \\ &+ \ln \int P_H(d\mathbf{H}) P_X(d\mathbf{X}) e^{-S_{\text{eff}}[\mathbf{X}, \mathbf{H}]}, \\ S_{\text{eff}}[\mathbf{X}, \mathbf{H}] &\equiv \sum_{\mu,i} \left[\lambda_{\mu i} X_{\mu i} + \frac{\gamma_{\mu i}}{2} X_{\mu i}^2 \right] + \sum_{\mu < \nu} \left[\omega_{\mu\nu} (iH)_{\mu\nu} - \frac{b_{\mu\nu}}{2} (iH)_{\mu\nu}^2 \right] + H_{\text{eff}}[\mathbf{X}, \mathbf{H}], \\ H_{\text{eff}}[\mathbf{X}, \mathbf{H}] &\equiv \frac{1}{\sqrt{n}} \sum_{\mu < \nu} \sum_i (iH)_{\mu\nu} X_{\mu i} X_{\nu i}. \end{cases}$$

For clarity, we will keep the dependency of all the Lagrange parameters on η explicit when needed. For a given η and a given \mathbf{Y} , one defines the operator U of Georges–Yedidia [GY91]:

$$\begin{aligned} U_{\mathbf{Y},\eta} &\equiv H_{\text{eff}} - \langle H_{\text{eff}} \rangle_{\mathbf{Y},\eta} + \sum_{\mu,i} \partial_\eta \lambda_{\mu i} (X_{\mu i} - m_{\mu i}) + \frac{1}{2} \sum_{\mu,i} \partial_\eta \gamma_{\mu i} [X_{\mu i}^2 - v_{\mu i} - (m_{\mu i})^2] \\ &+ \sum_{\mu < \nu} \partial_\eta \omega_{\mu\nu} (iH)_{\mu\nu} + g_{\mu\nu} - \frac{1}{2} \sum_{\mu < \nu} \partial_\eta b_{\mu\nu} [(iH)_{\mu\nu}^2 + r_{\mu\nu} - g_{\mu\nu}^2]. \end{aligned} \quad (88)$$

B.2. First orders of perturbation

Order 1 in η —At order 1, we have directly:

$$\left(\frac{\partial \Phi_{\mathbf{Y},n}}{\partial \eta} \right)_{\eta=0} = \frac{1}{nm} \langle H_{\text{eff}} \rangle_0 = \frac{1}{n^{3/2}m} \sum_i \sum_{\mu < \nu} g_{\mu\nu} m_{\mu i} m_{\nu i}. \quad (89)$$

We can then use ‘Maxwell-like’ relations [GY91] to compute the derivatives of the Lagrange parameters at $\eta = 0$. For instance, for $\lambda_{\mu i}$ and $\gamma_{\mu i}$ they can be read from equation (14):

$$\begin{cases} \lambda_{\mu i} + m_{\mu i} \gamma_{\mu i} = nm \frac{\partial \Phi_{\mathbf{Y},n}}{\partial m_{\mu i}}, & (90a) \\ \gamma_{\mu i} = 2nm \frac{\partial \Phi_{\mathbf{Y},n}}{\partial v_{\mu i}}, & (90b) \end{cases}$$

so that we can compute:

$$\begin{cases} \partial_{\eta} \gamma_{\mu i}(\eta = 0) = 2nm \frac{\partial}{\partial v_{\mu i}} \left[\frac{\partial \Phi_{\mathbf{Y},n}}{\partial \eta}(\eta = 0) \right] = 0, \\ \partial_{\eta} \lambda_{\mu i}(\eta = 0) = nm \frac{\partial}{\partial m_{\mu i}} \left[\frac{\partial \Phi_{\mathbf{Y},n}}{\partial \eta}(\eta = 0) \right] = \frac{1}{\sqrt{n}} \sum_{\nu} g_{\mu\nu} m_{\nu i}. \end{cases} \quad (91)$$

In a similar fashion, we reach

$$\begin{cases} \partial_{\eta} b_{\mu\nu}(\eta = 0) = 2nm \frac{\partial}{\partial r_{\mu\nu}} \left[\frac{\partial \Phi_{\mathbf{Y},n}}{\partial \eta}(\eta = 0) \right] = 0, \\ \partial_{\eta} \omega_{\mu\nu}(\eta = 0) = -nm \frac{\partial}{\partial g_{\mu\nu}} \left[\frac{\partial \Phi_{\mathbf{Y},n}}{\partial \eta}(\eta = 0) \right] = -\frac{1}{\sqrt{n}} \sum_i m_{\mu i} m_{\nu i}. \end{cases} \quad (92)$$

Using equations (91) and (92), we can compute $U_{\mathbf{Y},0}$ from equation (88). For clarity of the notation, we will denote by lowercase letters *centered variables*, i.e. $x_{\mu i} \equiv X_{\mu i} - m_{\mu i}$ and $(ih)_{\mu\nu} \equiv (iH)_{\mu\nu} + g_{\mu\nu}$. We obtain after a straightforward calculation:

$$U_{\mathbf{Y},\eta=0} = \frac{1}{\sqrt{n}} \sum_i \sum_{\mu < \nu} [(ih)_{\mu\nu} x_{\mu i} x_{\nu i} - g_{\mu\nu} x_{\mu i} x_{\nu i} + (ih)_{\mu\nu} m_{\mu i} x_{\nu i} + (ih)_{\mu\nu} x_{\mu i} m_{\nu i}]. \quad (93)$$

Order 2 in η —Relying on the formulas of [GY91], we have:

$$\frac{1}{2} \left(\frac{\partial^2 \Phi_{\mathbf{Y},n}}{\partial \eta^2} \right)_{\eta=0} = \frac{1}{2nm} \langle U_{\mathbf{Y}}^2 \rangle_0. \quad (94)$$

Recall that at $\eta = 0$, all the variables $\{h_{\mu\nu}, x_{\rho i}\}$ are independent and have zero mean. Using this fact alongside with the expression of $U_{\mathbf{Y},\eta=0}$ given in equation (93) yields:

$$\begin{aligned}
 \frac{1}{2} \left(\frac{\partial^2 \Phi_{\mathbf{Y},n}}{\partial \eta^2} \right)_{\eta=0} &= \frac{1}{2n^2 m} \sum_{i,j} \sum_{\substack{\mu < \nu \\ \mu' < \nu'}} \langle [(ih)_{\mu\nu} x_{\mu i} x_{\nu i} - g_{\mu\nu} x_{\mu i} x_{\nu i} + (ih)_{\mu\nu} m_{\mu i} x_{\nu i} \\
 &\quad + (ih)_{\mu\nu} x_{\mu i} m_{\nu i}] \times \langle (ih)_{\mu'\nu'} x_{\mu' j} x_{\nu' j} - g_{\mu'\nu'} x_{\mu' j} x_{\nu' j} \\
 &\quad + (ih)_{\mu'\nu'} m_{\mu' j} x_{\nu' j} + (ih)_{\mu'\nu'} x_{\mu' j} m_{\nu' j} \rangle_0, \\
 &= \frac{1}{2n^2 m} \sum_i \sum_{\mu < \nu} [-r_{\mu\nu} v_{\mu i} v_{\nu i} + g_{\mu\nu}^2 v_{\mu i} v_{\nu i} \\
 &\quad - r_{\mu\nu} m_{\mu i}^2 v_{\nu i} - r_{\mu\nu} v_{\mu i} m_{\nu i}^2]. \tag{95}
 \end{aligned}$$

Order 3 in η —At order 3, the formula of appendix A of [GY91] yields:

$$\frac{1}{3!} \left(\frac{\partial^3 \Phi_{\mathbf{Y},n}}{\partial \eta^3} \right)_{\eta=0} = -\frac{1}{6nm} \langle U_{\mathbf{Y}}^3 \rangle_0. \tag{96}$$

In order to compute the right-hand side of equation (96), we decompose the operator U of equation (93) as $U_{\mathbf{Y}} = U_c + U_g + U_m$, with

$$\begin{cases}
 U_c \equiv \frac{1}{\sqrt{n}} \sum_i \sum_{\mu < \nu} (ih)_{\mu\nu} x_{\mu i} x_{\nu i}, \\
 U_g \equiv -\frac{1}{\sqrt{n}} \sum_i \sum_{\mu < \nu} g_{\mu\nu} x_{\mu i} x_{\nu i}, \\
 U_m \equiv \frac{1}{\sqrt{n}} \sum_i \sum_{\mu < \nu} [(ih)_{\mu\nu} m_{\mu i} x_{\nu i} + (ih)_{\mu\nu} x_{\mu i} m_{\nu i}].
 \end{cases} \tag{97}$$

Therefore, we obtain (dropping the 0 subscript in the averages to lighten the notations):

$$\begin{aligned}
 \langle U_{\mathbf{Y}}^3 \rangle_0 &= \langle U_c^3 \rangle + \langle U_g^3 \rangle + \langle U_m^3 \rangle + 3\langle U_c^2 U_g \rangle + 3\langle U_c^2 U_m \rangle + 3\langle U_g^2 U_c \rangle + 3\langle U_g^2 U_m \rangle \\
 &\quad + 3\langle U_m^2 U_c \rangle + 3\langle U_m^2 U_g \rangle + 6\langle U_c U_g U_m \rangle. \tag{98}
 \end{aligned}$$

In equation (98), the terms $\langle U_g^2 U_c \rangle$ and $\langle U_g^2 U_m \rangle$ are trivially zero since $\langle \mathbf{h} \rangle = 0$. Let us now argue that, in equation (98), all the terms except $\langle U_g^3 \rangle$ are negligible in the thermodynamic limit. We can for instance consider $\langle U_c^3 \rangle$, which can easily be written from equation (97) as (using again $\langle \mathbf{h} \rangle = 0$):

$$\langle U_c^3 \rangle = \frac{1}{n^{3/2}} \sum_{\mu < \nu} \langle (ih)_{\mu\nu}^3 \rangle \sum_{i,j,k} \langle x_{\mu i} x_{\nu i} x_{\mu j} x_{\nu j} x_{\mu k} x_{\nu k} \rangle, \tag{99}$$

$$= \frac{1}{n^{3/2}} \sum_{\mu < \nu} \langle (ih)_{\mu\nu}^3 \rangle \sum_{i=1}^n \langle x_{\mu i}^3 \rangle \langle x_{\nu i}^3 \rangle. \tag{100}$$

In particular, this directly implies that $\langle U_c^3 \rangle = \mathcal{O}(n^{3/2})$, and therefore that it will not contribute to the asymptotic free energy, which is in the scale $\Theta(n^2)$. In the same way, one can show:

$$\left\{ \begin{aligned} \langle U_m^3 \rangle &= \frac{1}{n^{3/2}} \sum_{\mu < \nu} \langle (ih)_{\mu\nu}^3 \rangle \sum_{i=1}^n [m_{\mu i}^3 \langle x_{\nu i}^3 \rangle + m_{\nu i}^3 \langle x_{\mu i}^3 \rangle] = \mathcal{O}(n^{3/2}), & (101a) \\ \langle U_c^2 U_g \rangle &= \frac{1}{n^{3/2}} \sum_{\mu < \nu} g_{\mu\nu} r_{\mu\nu} \sum_{i=1}^n \langle x_{\mu i}^3 \rangle \langle x_{\nu i}^3 \rangle = \mathcal{O}(n^{3/2}), & (101b) \\ \langle U_c^2 U_m \rangle &= \frac{1}{n^{3/2}} \sum_{\mu < \nu} \langle (ih)_{\mu\nu}^3 \rangle \sum_{i=1}^n [\langle x_{\mu i}^3 \rangle m_{\nu i} v_{\nu i} + \langle x_{\nu i}^3 \rangle m_{\mu i} v_{\mu i}] = \mathcal{O}(n^{3/2}), & (101c) \\ \langle U_m^2 U_c \rangle &= \frac{2}{n^{3/2}} \sum_{\mu < \nu} \langle (ih)_{\mu\nu}^3 \rangle \sum_{i=1}^n m_{\mu i} v_{\mu i} m_{\nu i} v_{\nu i} = \mathcal{O}(n^{3/2}), & (101d) \\ \langle U_m^2 U_g \rangle &= \frac{2}{n^{3/2}} \sum_{\mu < \nu} g_{\mu\nu} r_{\mu\nu} \sum_{i=1}^n m_{\mu i} v_{\mu i} m_{\nu i} v_{\nu i} = \mathcal{O}(n^{3/2}), & (101e) \\ \langle U_c U_g U_m \rangle &= \frac{1}{n^{3/2}} \sum_{\mu < \nu} g_{\mu\nu} r_{\mu\nu} \sum_{i=1}^n [\langle x_{\mu i}^3 \rangle m_{\nu i} v_{\nu i} + \langle x_{\nu i}^3 \rangle m_{\mu i} v_{\mu i}] = \mathcal{O}(n^{3/2}). & (101f) \end{aligned} \right.$$

All in all, equations (96) and (101) imply that:

$$\frac{1}{3!} \left(\frac{\partial^3 \Phi_{\mathbf{Y},n}}{\partial \eta^3} \right)_{\eta=0} = -\frac{1}{6nm} \langle U_g^3 \rangle + \mathcal{O}_n(1). \quad (102)$$

This remaining term can be computed again by expanding the sums and using the independence of all variables $\{h_{\mu\nu}, x_{\rho i}\}$ at $\eta = 0$. Elementary combinatorics allow to count the number of terms appearing in this expansion, and we reach:

$$\begin{aligned} \langle U_g^3 \rangle &= -\frac{1}{8n^{3/2}} \left\{ \prod_{a=1}^3 \sum_{\mu_a \neq \nu_a} g_{\mu_a \nu_a} \right\} \sum_{i_1, i_2, i_3} \left\langle \prod_{a=1}^3 x_{\mu_a i_a} x_{\nu_a i_a} \right\rangle, \\ &= -\frac{1}{n^{3/2}} \left\{ \prod_{a=1}^3 \sum_{\mu_a \neq \nu_a} g_{\mu_a \nu_a} \right\} \delta_{\mu_1 \nu_2} \delta_{\mu_2 \nu_3} \delta_{\mu_3 \nu_1} \sum_{i=1}^n \prod_{a=1}^3 v_{\mu_a i} - \frac{1}{2n^{3/2}} \sum_{\mu \neq \nu} g_{\mu\nu}^3 \sum_{i=1}^n \langle x_{\mu i}^3 \rangle \langle x_{\nu i}^3 \rangle. \end{aligned} \quad (103)$$

Note that the factor 8 in equation (103) disappeared since all possible pairings of indices $\{\mu_a, \nu_a\}_{a=1}^3$ are equivalent, and one can easily count that there are eight such pairings. The second term in equation (103) is again $\mathcal{O}(n^{3/2})$, and we reach:

$$\frac{1}{3!} \left(\frac{\partial^3 \Phi_{\mathbf{Y},n}}{\partial \eta^3} \right)_{\eta=0} = \frac{1}{6n^{5/2}m} \sum_i \sum_{\substack{\mu_1, \mu_2, \mu_3 \\ \text{pairwise distinct}}} g_{\mu_1 \mu_2} g_{\mu_2 \mu_3} g_{\mu_3 \mu_1} \prod_{a=1}^3 v_{\mu_a i} + \mathcal{O}_n(1). \quad (104)$$

Note that the constraint of having pairwise distinct indices is a simple consequence of the constraint $\mu_a \neq \nu_a$ in equation (103), associated with the form of the pairing.

Remark on higher-order cumulants—An important remark that one can already conjecture by generalizing from equations (95) and (104) is that at any given order of perturbation in η , *only the first two moments of the fields \mathbf{H}, \mathbf{X} will appear at dominant order*. This conjecture arises as a consequence of a simple scaling argument: the higher-order moments constraint too much indices on which we can sum, and thus

the terms involving them can be neglected. This was already formulated for similar Plefka expansions in symmetric and bipartite finite-rank models in [MFC⁺19].

B.3. PGY expansion in the non-symmetric model **FX**

In this section, we detail the derivation of result 2.2 from the PGY expansion formalism. As many things are similar to the derivation of result 2.1, we will shorten some arguments that can be easily transposed from the symmetric setting.

B.3.1. The method. As in the symmetric case, we start from the original expression of the TAP free entropy, very similar to equation (14):

$$\left\{ \begin{aligned} n(m+p)\Phi_{\mathbf{Y},n}(\eta) &= \sum_{\mu,i} \left[\lambda_{\mu i}^F m_{\mu i}^F + \frac{\gamma_{\mu i}^F}{2} (v_{\mu i}^F + (m_{\mu i}^F)^2) \right] + \sum_{i,l} \left[\lambda_{i l}^X m_{i l}^X + \frac{\gamma_{i l}^X}{2} (v_{i l}^X + (m_{i l}^X)^2) \right], \\ &+ \sum_{\mu,l} \left[-\omega_{\mu l} g_{\mu l} - \frac{b_{\mu l}}{2} (-r_{\mu l} + g_{\mu l}^2) \right] + \ln \int P_H(d\mathbf{H}) P_F(d\mathbf{F}) P_X(d\mathbf{X}) e^{-S_{\text{eff},\eta}[\mathbf{F},\mathbf{X},\mathbf{H}]}, \\ S_{\text{eff},\eta}[\mathbf{F},\mathbf{X},\mathbf{H}] &\equiv \sum_{\mu,i} \left[\lambda_{\mu i}^F F_{\mu i} + \frac{\gamma_{\mu i}^F}{2} F_{\mu i}^2 \right] + \sum_{i,l} \left[\lambda_{i l}^X X_{i l} + \frac{\gamma_{i l}^X}{2} X_{i l}^2 \right] \\ &+ \sum_{\mu,l} \left[\omega_{\mu l} (iH)_{\mu l} - \frac{b_{\mu l}}{2} (iH)_{\mu l}^2 \right] + \frac{\eta}{\sqrt{n}} \sum_{\mu,i,l} (iH)_{\mu l} F_{\mu i} X_{i l}, \\ H_{\text{eff}}[\mathbf{F},\mathbf{X},\mathbf{H}] &\equiv \frac{1}{\sqrt{n}} \sum_{\mu,i,l} (iH)_{\mu l} F_{\mu i} X_{i l}. \end{aligned} \right. \tag{105}$$

The operator U of Georges–Yedidia is defined similarly as the one of equation (88):

$$\begin{aligned} U_{\mathbf{Y},\eta} &\equiv H_{\text{eff}} - \langle H_{\text{eff}} \rangle_{\mathbf{Y},\eta} + \sum_{\mu,i} \partial_{\eta} \lambda_{\mu i}^F (F_{\mu i} - m_{\mu i}^F) + \frac{1}{2} \sum_{\mu,i} \partial_{\eta} \gamma_{\mu i}^F [F_{\mu i}^2 - v_{\mu i}^F - (m_{\mu i}^F)^2] \\ &+ \sum_{i,l} \partial_{\eta} \lambda_{i l}^X (X_{i l} - m_{i l}^X) + \frac{1}{2} \sum_{i,l} \partial_{\eta} \gamma_{i l}^X [X_{i l}^2 - v_{i l}^X - (m_{i l}^X)^2] \\ &+ \sum_{\mu,l} \partial_{\eta} \omega_{\mu l} (iH)_{\mu l} + g_{\mu l} - \frac{1}{2} \sum_{\mu,l} \partial_{\eta} b_{\mu l} [(iH)_{\mu l}^2 + r_{\mu l} - g_{\mu l}^2]. \end{aligned} \tag{106}$$

Order 1 in η —At order 1, we have directly:

$$\left(\frac{\partial \Phi_{\mathbf{Y},n}}{\partial \eta} \right)_{\eta=0} = -\frac{1}{n(m+p)} \langle H_{\text{eff}} \rangle_0 = \frac{1}{n^{3/2}(m+p)} \sum_{\mu,i,l} g_{\mu l} m_{\mu i}^F m_{i l}^X. \tag{107}$$

We can then use the Maxwell relations we already described in equation (90), to compute, e.g.

$$\begin{cases} \partial_\eta \gamma_{\mu i}^F(\eta = 0) = 2n(m + p) \frac{\partial}{\partial v_{\mu i}^F} \left[\frac{\partial \Phi_{\mathbf{Y},n}}{\partial \eta}(\eta = 0) \right] = 0, \\ \partial_\eta \lambda_{\mu i}^F(\eta = 0) = n(m + p) \frac{\partial}{\partial m_{\mu i}^F} \left[\frac{\partial \Phi_{\mathbf{Y},n}}{\partial \eta}(\eta = 0) \right] = \frac{1}{\sqrt{n}} \sum_l g_{\mu l} m_{il}^X. \end{cases} \quad (108)$$

Applying this technique to all Lagrange multipliers, we compute $U_{\mathbf{Y},0}$ from equation (106). Again we will denote by lowercase letters *centered variables*, e.g. $x_{\mu i} \equiv X_{\mu i} - m_{\mu i}^X$. We obtain:

$$U_{\mathbf{Y},0} = \frac{1}{\sqrt{n}} \sum_{\mu,i,l} [(ih)_{\mu l} f_{\mu i} x_{il} - g_{\mu l} f_{\mu i} x_{il} + (ih)_{\mu l} m_{\mu i}^F x_{il} + (ih)_{\mu l} f_{\mu i} m_{il}^X]. \quad (109)$$

Order 2 in η —One can then compute, using the formulas of [GY91], and very similarly to model XX[†]:

$$\begin{aligned} \frac{1}{2} \left(\frac{\partial^2 \Phi_{\mathbf{Y},n}}{\partial \eta^2} \right)_{\eta=0} &= \frac{1}{2n(m + p)} \langle U^2 \rangle_0, \\ &= \frac{1}{2n^2(m + p)} \sum_{\mu,i,l} [-r_{\mu l} v_{\mu i}^F v_{il}^X + g_{\mu l}^2 v_{\mu i}^F v_{il}^X - r_{\mu l} (m_{\mu i}^F)^2 v_{il}^X - r_{\mu l} v_{\mu i}^F (m_{il}^X)^2]. \end{aligned} \quad (110)$$

Orders 3 and 4 are more tedious to compute, and we detail them in separate paragraphs.

B.3.2. Order 3 of the expansion. In order to compute this order of the PGY expansion, we use again the formulas of appendix A of [GY91]:

$$\frac{1}{3!} \left(\frac{\partial^3 \Phi_{\mathbf{Y},n}}{\partial \eta^3} \right)_{\eta=0} = -\frac{1}{6n(m + p)} \langle U^3 \rangle_0. \quad (111)$$

We will need to introduce the (centered) third moments of the distributions of the independent variables $\{iH_{\mu l}, F_{\mu i}, X_{il}\}$ at $\eta = 0$. These moments are denoted $\{c_{\mu l}^{(3,H)}, c_{\mu i}^{(3,F)}, c_{il}^{(3,X)}\}$. We then decompose the operator of equation (109) as follows:

$$U_{\mathbf{Y},0} = \frac{1}{\sqrt{n}} \sum_{\mu,i,l} \left[\underbrace{(ih)_{\mu l} f_{\mu i} x_{il}}_A + \underbrace{(-g_{\mu l} f_{\mu i} x_{il})}_{B_H} + \underbrace{(ih)_{\mu l} m_{\mu i}^F x_{il}}_{B_F} + \underbrace{(ih)_{\mu l} f_{\mu i} m_{il}^X}_{B_X} \right]. \quad (112)$$

Since all the variables in the equation above are *centered*, we get easily:

$$\begin{aligned} \langle A^3 + B_H^3 + B_F^3 + B_X^3 \rangle_0 &= \frac{1}{n^{3/2}} \sum_{\mu,i,l} \left[c_{\mu l}^{(3,H)} c_{\mu i}^{(3,F)} c_{il}^{(3,X)} - g_{\mu l}^3 c_{\mu i}^{(3,F)} c_{il}^{(3,X)} \right. \\ &\quad \left. + c_{\mu l}^{(3,H)} (m_{\mu i}^F)^3 c_{il}^{(3,X)} + c_{\mu l}^{(3,H)} c_{\mu i}^{(3,F)} (m_{il}^X)^3 \right]. \end{aligned} \quad (113)$$

Using again the centering of the variables and the decomposition above, we get that the only non-zero terms of the type $\langle X^2 Y \rangle_0$ with $X, Y \in \{A, B_H, B_F, B_X\}$ yields the contribution:

$$3\langle A^2(B_H + B_F + B_X) \rangle_0 \tag{114}$$

$$= \frac{3}{n^{3/2}} \sum_{\mu,i,l} [g_{\mu l} r_{\mu l} c_{\mu i}^{(3,F)} c_{i l}^{(3,X)} + c_{\mu l}^{(3,H)} m_{\mu i}^F v_{\mu i}^F c_{i l}^{(3,X)} + c_{\mu l}^{(3,H)} c_{\mu i}^{(3,F)} m_{i l}^X v_{i l}^X].$$

Finally, the last contribution to $\langle U^3 \rangle_0$ comes from the term:

$$6\langle AB_H B_F + AB_H B_X + AB_F B_X + B_H B_F B_X \rangle_0 \tag{115}$$

$$= \frac{6}{n^{3/2}} \sum_{\mu,i,l} [g_{\mu l} r_{\mu l} m_{\mu i}^F v_{\mu i}^F c_{i l}^{(3,X)} + g_{\mu l} r_{\mu l} c_{\mu i}^{(3,F)} m_{i l}^X v_{i l}^X + c_{\mu l}^{(3,H)} m_{\mu i}^F v_{\mu i}^F m_{i l}^X v_{i l}^X + g_{\mu l} r_{\mu l} m_{\mu i}^F v_{\mu i}^F m_{i l}^X v_{i l}^X].$$

Summing the contributions from equations (113)–(115) yields $\langle U^3 \rangle_0$, which then yields:

$$\frac{1}{3!} \left(\frac{\partial^3 \Phi_{\mathbf{Y},n}}{\partial \eta^3} \right)_{\eta=0} = \frac{-1}{6n^{5/2}(m+p)} \sum_{\mu,i,l} [c_{\mu l}^{(3,H)} c_{\mu i}^{(3,F)} c_{i l}^{(3,X)} - g_{\mu l}^3 c_{\mu i}^{(3,F)} c_{i l}^{(3,X)} + c_{\mu l}^{(3,H)} c_{\mu i}^{(3,F)} (m_{i l}^X)^2 + c_{\mu l}^{(3,H)} (m_{\mu i}^F)^3 c_{i l}^{(3,X)} + 3g_{\mu l} r_{\mu l} c_{\mu i}^{(3,F)} c_{i l}^{(3,X)} + 3c_{\mu l}^{(3,H)} m_{\mu i}^F v_{\mu i}^F c_{i l}^{(3,X)} + 3c_{\mu l}^{(3,H)} c_{\mu i}^{(3,F)} m_{i l}^X v_{i l}^X + 6g_{\mu l} r_{\mu l} m_{\mu i}^F v_{\mu i}^F c_{i l}^{(3,X)} + 6g_{\mu l} r_{\mu l} c_{\mu i}^{(3,F)} m_{i l}^X v_{i l}^X + 6c_{\mu l}^{(3,H)} m_{\mu i}^F v_{\mu i}^F m_{i l}^X v_{i l}^X + 6g_{\mu l} r_{\mu l} m_{\mu i}^F v_{\mu i}^F m_{i l}^X v_{i l}^X]. \tag{116}$$

From equation (116), since all involved terms inside the sum are of order $\mathcal{O}_N(1)$, it is clear that the third order is subdominant:

$$\frac{1}{3!} \left(\frac{\partial^3 \Phi_{\mathbf{Y},n}}{\partial \eta^3} \right)_{\eta=0} = \mathcal{O}_n(1). \tag{117}$$

Higher-order moments—An important remark that one can already conjecture by generalizing from equations (116) and (117) is that at any given order of perturbation in η , *only the first two moments of the fields $\mathbf{H}, \mathbf{F}, \mathbf{X}$ will appear at dominant order.* This conjecture arises as a consequence of a simple scaling argument: the higher-order moments constraint too much indices on which we can sum, and thus the terms involving them can be neglected. Note that we noticed already a completely similar behavior in the symmetric case in appendix B.2.

B.3.3. Order 4 of the expansion. At order 4, one can again use appendix A of [GY91]. The general (and quite heavy) formula is:

$$\begin{aligned}
 & n(m+p) \left(\frac{\partial^4 \Phi_{\mathbf{Y},n}}{\partial \eta^4} \right)_{\eta=0} \\
 &= \langle U^4 \rangle_0 - 3 \langle U^2 \rangle_0^2 - 3 \sum_{\mu,i} \left[\partial_\eta^2 \lambda_{\mu i}^F \langle U^2 f_{\mu i} \rangle_0 + \frac{\partial_\eta^2 \gamma_{\mu i}^F}{2} \langle U^2 (f_{\mu i}^2 + 2m_{\mu i}^F f_{\mu i} - v_{\mu i}^F) \rangle_0 \right] \\
 &\quad - 3 \sum_{i,l} \left[\partial_\eta^2 \lambda_{i l}^X \langle U^2 x_{i l} \rangle_0 + \frac{\partial_\eta^2 \gamma_{i l}^X}{2} \langle U^2 (x_{i l}^2 + 2m_{i l}^X x_{i l} - v_{i l}^X) \rangle_0 \right] \\
 &\quad - 3 \sum_{\mu,l} \left[\partial_\eta^2 \omega_{\mu l} \langle U^2 (ih)_{\mu l} \rangle_0 + \frac{\partial_\eta^2 b_{\mu l}}{2} \langle U^2 ((ih)_{\mu l}^2 - 2g_{\mu l} (ih)_{\mu l} + r_{\mu l}) \rangle_0 \right]. \tag{118}
 \end{aligned}$$

This section describes the calculation of the order 4 perturbation of the free entropy for model **FX**. It is particularly tedious and lengthy, but the techniques involved are not conceptually complicated. For simplicity, we will not consider terms involving cumulants of order 3 and 4 of the variables $ih_{\mu l}$, $f_{\mu i}$ and $x_{i l}$. One can check that the terms involving these moments cancel out, and in the end only yield sub-dominant contributions in the thermodynamic limit. From equation (110) and the Maxwell relations of equation (90), we obtain the derivatives of the Lagrange multipliers at leading order and at $\eta = 0$:

$$\left\{ \begin{aligned}
 \partial_\eta^2 \omega_{\mu l} &= \frac{2}{n} g_{\mu l} \sum_i [(m_{\mu i}^F)^2 v_{i l}^X + v_{\mu i}^F (m_{i l}^X)^2], \\
 \partial_\eta^2 b_{\mu l} &= -\frac{2}{n} \sum_i [v_{\mu i}^F v_{i l}^X + (m_{\mu i}^F)^2 v_{i l}^X + v_{\mu i}^F (m_{i l}^X)^2], \\
 \partial_\eta^2 \lambda_{\mu i}^F &= \frac{2}{n} m_{\mu i}^F \sum_l [r_{\mu l} (m_{i l}^X)^2 - g_{\mu l}^2 v_{i l}^X], \\
 \partial_\eta^2 \gamma_{\mu i}^F &= -\frac{2}{n} \sum_l [r_{\mu l} v_{i l}^X + r_{\mu l} (m_{i l}^X)^2 - g_{\mu l}^2 v_{i l}^X], \\
 \partial_\eta^2 \lambda_{i l}^X &= \frac{2}{n} m_{i l}^X \sum_\mu [r_{\mu l} (m_{\mu i}^F)^2 - g_{\mu l}^2 v_{\mu i}^F], \\
 \partial_\eta^2 \gamma_{i l}^X &= -\frac{2}{n} \sum_\mu [r_{\mu l} v_{\mu i}^F + r_{\mu l} (m_{\mu i}^F)^2 - g_{\mu l}^2 v_{\mu i}^F].
 \end{aligned} \right. \tag{119}$$

Discarding as we mentioned the cumulants of order greater than 3, we can compute:

$$\left\{ \begin{aligned}
 \langle U^2 x_{i l} \rangle_0 &= -\frac{2}{n} m_{i l}^X v_{i l}^X \sum_\mu r_{\mu l} v_{\mu i}^F, \\
 \langle U^2 (x_{i l}^2 + 2m_{i l}^X x_{i l} - v_{i l}^X) \rangle_0 &= \frac{2(v_{i l}^X)^2}{n} \sum_\mu [-r_{\mu l} v_{\mu i}^F + g_{\mu l}^2 v_{\mu i}^F - r_{\mu l} (m_{\mu i}^F)^2] - \frac{4}{n} v_{i l}^X (m_{i l}^X)^2 \sum_\mu r_{\mu l} v_{\mu i}^F.
 \end{aligned} \right.$$

From this and equation (119), one can obtain the term involving the derivatives of the Lagrange parameters λ^X and γ^X in equation (118):

$$\begin{aligned}
 & -3 \sum_{i,l} \left[\partial_\eta^2 \lambda_{il}^X \langle U^2 x_{il} \rangle_0 + \frac{\partial_\eta^2 \gamma_{il}^X}{2} \langle U^2 (x_{il}^2 + 2m_{il}^X x_{il} - v_{il}^X) \rangle_0 \right] \\
 & = -\frac{12}{n^2} \sum_{i,l} v_{il}^X (m_{il}^X)^2 \left[\sum_\mu r_{\mu l} v_{\mu i}^F \right]^2 \\
 & \quad - \frac{6}{n^2} \sum_{i,l} (v_{il}^X)^2 \left[\sum_\mu [r_{\mu l} v_{\mu i}^F + r_{\mu l} (m_{\mu i}^F)^2 - g_{\mu l}^2 v_{\mu i}^F] \right]^2.
 \end{aligned} \tag{120}$$

Doing similarly for $(ih)_{\mu l}$ and $f_{\mu i}$, we obtain all the terms involving derivatives of the Lagrange multipliers in equation (118):

$$\begin{aligned}
 & -3 \sum_{\mu,i} \left[\partial_\eta^2 \lambda_{\mu i}^F \langle U^2 f_{\mu i} \rangle_0 + \frac{\partial_\eta^2 \gamma_{\mu i}^F}{2} \langle U^2 (f_{\mu i}^2 + 2m_{\mu i}^F f_{\mu i} - v_{\mu i}^F) \rangle_0 \right] \\
 & = -\frac{12}{n^2} \sum_{\mu,i} v_{\mu i}^F (m_{\mu i}^F)^2 \left[\sum_l r_{\mu l} v_{il}^X \right]^2
 \end{aligned} \tag{121}$$

$$\begin{aligned}
 & \quad - \frac{6}{n^2} \sum_{\mu,i} (v_{\mu i}^F)^2 \left[\sum_l [r_{\mu l} v_{il}^X + r_{\mu l} (m_{il}^X)^2 - g_{\mu l}^2 v_{il}^X] \right]^2, \\
 & -3 \sum_{\mu,l} \left[\partial_\eta^2 \omega_{\mu l} \langle U^2 (ih)_{\mu l} \rangle_0 + \frac{\partial_\eta^2 b_{\mu l}}{2} \langle U^2 ((ih)_{\mu l}^2 - 2g_{\mu l} (ih)_{\mu l} + r_{\mu l}) \rangle_0 \right]
 \end{aligned} \tag{122}$$

$$\begin{aligned}
 & = \frac{12}{n^2} \sum_{\mu,l} r_{\mu l} g_{\mu l}^2 \left[\sum_i v_{\mu i}^F v_{il}^X \right]^2 \\
 & \quad - \frac{6}{n^2} \sum_{\mu,l} r_{\mu l}^2 \left[\sum_i [v_{\mu i}^F v_{il}^X + (m_{\mu i}^F)^2 v_{il}^X + v_{\mu i}^F (m_{il}^X)^2] \right]^2.
 \end{aligned}$$

The calculation at order 2 already gave (cf equation (110)):

$$\begin{aligned}
 -3 \langle U^2 \rangle_0^2 & = -\frac{3}{n^2} \sum_{\substack{\mu,i,l \\ \mu',i'}} [-r_{\mu l} v_{\mu i}^F v_{il}^X + g_{\mu l}^2 v_{\mu i}^F v_{il}^X - r_{\mu l} (m_{\mu i}^F)^2 v_{il}^X - r_{\mu l} v_{\mu i}^F (m_{il}^X)^2] \\
 & \quad \times [-r_{\mu' l'} v_{\mu' i'}^F v_{i'l'}^X + g_{\mu' l'}^2 v_{\mu' i'}^F v_{i'l'}^X - r_{\mu' l'} (m_{\mu' i'}^F)^2 v_{i'l'}^X - r_{\mu' l'} v_{\mu' i'}^F (m_{i'l'}^X)^2].
 \end{aligned} \tag{123}$$

We finally have to compute $\langle U^4 \rangle_0$, whose calculation is very tedious, but not conceptually difficult. We again make use of the decomposition of equation (112). A first simplification arises when using that the variables are centered and that we can neglect their moments of odd order. This implies:

$$\langle U^4 \rangle_0 = \underbrace{\langle A^4 \rangle_0}_{I_1} + \underbrace{6\langle A^2(B_H^2 + B_F^2 + B_X^2) \rangle_0}_{I_2} + \underbrace{\langle (B_H + B_F + B_X)^4 \rangle_0}_{I_3} + \mathcal{O}_n(1). \tag{124}$$

We now compute the three terms I_1, I_2, I_3 independently. An explicit calculation gives

$$I_1 = \frac{6}{n^2} \sum_{\mu,i,l} r_{\mu l}^2 (v_{\mu i}^F)^2 (v_{i l}^X)^2 + \frac{3}{n^2} \sum_{\substack{\mu,i,l \\ \mu',i',l'}} r_{\mu l} r_{\mu' l'} v_{\mu i}^F v_{\mu' i'}^F v_{i l}^X v_{i' l'}^X \tag{125}$$

$$+ \frac{6}{n^2} \sum_{\mu,i,l} r_{\mu l} v_{\mu i}^F v_{i l}^X \left(\sum_{\mu'} r_{\mu' l} v_{\mu' i}^F v_{i l}^X + \sum_{i'} r_{\mu l} v_{\mu i}^F v_{i' l}^X + \sum_{l'} r_{\mu l} v_{\mu i}^F v_{i l}^{X'} \right) + \mathcal{O}_n(1).$$

Importantly, the indices are not supposed to be pairwise distinct unless explicitly stated so. The terms I_2, I_3 can be explicitly computed as well, and are very lengthy:

$$I_2 = \frac{6}{n^2} \sum_{\substack{\mu,i,l \\ \mu',i',l'}} r_{\mu l} v_{\mu i}^F v_{i l}^X [-g_{\mu' l'}^2 v_{\mu' i'}^F v_{i' l'}^X + r_{\mu' l'} v_{\mu' i'}^F (m_{i' l'}^X)^2 + r_{\mu' l'} (m_{\mu' i'}^F)^2 v_{i l}^X] \tag{126}$$

$$+ \frac{12}{n^2} \sum_{\mu,i,l} r_{\mu l} v_{\mu i}^F v_{i l}^X [-g_{\mu l}^2 v_{\mu i}^F v_{i l}^X + r_{\mu l} v_{\mu i}^F (m_{i l}^X)^2 + r_{\mu l} (m_{\mu i}^F)^2 v_{i l}^X]$$

$$+ \frac{12}{n^2} \sum_{\mu,i,l} \sum_{\mu'} r_{\mu l} v_{\mu i}^F v_{i l}^X [-g_{\mu' l}^2 v_{\mu' i}^F v_{i l}^X + r_{\mu' l} v_{\mu' i}^F (m_{i l}^X)^2 + r_{\mu' l} (m_{\mu' i}^F)^2 v_{i l}^X]$$

$$+ \frac{12}{n^2} \sum_{\mu,i,l} \sum_{i'} r_{\mu l} v_{\mu i}^F v_{i l}^X [-g_{\mu l}^2 v_{\mu i}^F v_{i' l}^X + r_{\mu l} v_{\mu i}^F (m_{i' l}^X)^2 + r_{\mu l} (m_{\mu i}^F)^2 v_{i' l}^X]$$

$$+ \frac{12}{n^2} \sum_{\mu,i,l} \sum_{l'} r_{\mu l} v_{\mu i}^F v_{i l}^X [-g_{\mu l}^2 v_{\mu i}^F v_{i l}^{X'} + r_{\mu l} v_{\mu i}^F (m_{i l}^{X'})^2 + r_{\mu l} (m_{\mu i}^F)^2 v_{i l}^{X'}] + \mathcal{O}_n(1),$$

$$\begin{aligned}
 I_3 = & \frac{3}{n^2} \sum_{\substack{\mu, i, l \\ \mu', i', l'}} [g_{\mu l}^2 g_{\mu' l'}^2 v_{\mu i}^F v_{\mu' i'}^F v_{i l}^X v_{i' l'}^X + r_{\mu l} r_{\mu' l'} (m_{\mu i}^F)^2 (m_{\mu' i'}^F)^2 v_{i l}^X v_{i' l'}^X + r_{\mu l} r_{\mu' l'} v_{\mu i}^F v_{\mu' i'}^F (m_{i l}^X)^2 (m_{i' l'}^X)^2] \\
 & - \frac{6}{n^2} \sum_{\substack{\mu, i, l \\ \mu', i', l'}} [-r_{\mu l} r_{\mu' l'} (m_{\mu' i'}^F)^2 v_{\mu i}^F (m_{i l}^X)^2 v_{i' l'}^X + g_{\mu l}^2 r_{\mu' l'} (m_{\mu' i'}^F)^2 v_{\mu i}^F v_{i l}^X v_{i' l'}^X + g_{\mu l}^2 r_{\mu' l'} v_{\mu' i'}^F v_{\mu i}^F (m_{i' l'}^X)^2 v_{i l}^X] \\
 & - \frac{12}{n^2} \sum_{\mu, i, l} \left[\sum_{\mu'} g_{\mu l}^2 r_{\mu' l} (m_{\mu i}^F)^2 v_{\mu i}^F (v_{i l}^X)^2 - \sum_{i'} r_{\mu l}^2 (m_{\mu i'}^F)^2 v_{\mu i}^F (m_{i l}^X)^2 v_{i' l}^X + \sum_{l'} g_{\mu l}^2 r_{\mu l'} (v_{\mu i}^F)^2 (m_{i l'}^X)^2 v_{i l}^X \right] \\
 & + \frac{6}{n^2} \sum_{\mu, i, l} \left[\sum_{\mu', i'} m_{\mu i}^F m_{\mu' i'}^F m_{\mu' i}^F m_{\mu i'}^F r_{\mu l} r_{\mu' l'} v_{i l}^X v_{i' l}^X + \sum_{i', l'} m_{i l}^X m_{i' l'}^X m_{i' l}^X m_{i l'}^X v_{\mu i}^F v_{\mu i'}^F r_{\mu l} r_{\mu' l'} \right], \quad (127) \\
 & + \frac{6}{n^2} \sum_{\mu, i, l} \left[\sum_{\mu', l'} g_{\mu l} g_{\mu' l'} g_{\mu l'} g_{\mu' l} v_{\mu i}^F v_{\mu' i}^F v_{i l}^X v_{i' l'}^X \right] + \mathcal{O}_n(1).
 \end{aligned}$$

Many simplifications occur in the terms of equations (120) to (123) and (125) to (127). Two type of terms are for instance negligible:

- Terms of the type $n^{-4} \sum_{\mu i l} A_{\mu i l}$, with $A_{\mu i l}$ typically of order 1. These terms are negligible by a simple scaling argument.
- Terms involving m^F or m^X . For instance, the term:

$$\sum_{\mu} \sum_{i \neq i'} \sum_{l \neq l'} m_{i l}^X m_{i' l'}^X m_{i' l}^X m_{i l'}^X r_{\mu l} r_{\mu l'} v_{\mu i}^F v_{\mu i'}^F.$$

By hypothesis H.1, the variables m^F, m^X behave like uncorrelated variables, so that all these terms will be negligible. A more detailed explanation of how uncorrelated variables leads to all these terms being negligible can be found, e.g. in [MFC⁺19]. This is precisely the sort of terms that are not negligible when involving $g_{\mu l}$, because of the structure described in H.2.

We can now sum all equations (120) to (123) and (125) to (127), simplifying the terms that are negligible by the arguments above, and checking that almost all non-negligible terms are canceling each other. This is a lengthy but straightforward calculation, and we reach the result:

$$\begin{aligned}
 \frac{1}{4!} \left(\frac{\partial^4 \Phi_{\mathbf{Y}, n}}{\partial \eta^4} \right)_{\eta=0} &= \frac{1}{4n^3(m+p)} \sum_i \sum_{\mu_1 \neq \mu_2} \sum_{l_1 \neq l_2} g_{\mu_1 l_1} g_{\mu_1 l_2} g_{\mu_2 l_2} g_{\mu_2 l_1} v_{\mu_1 i}^F v_{\mu_2 i}^F v_{i l_1}^X v_{i l_2}^X \\
 &+ \mathcal{O}_n(1). \quad (128)
 \end{aligned}$$

This ends the derivation of result 2.2.

Appendix C. On the solutions to Matytsin’s equations

C.1. Derivation of result 3.1

Note first that this result is shown in the very particular case in which \mathbf{S} is a Wigner matrix in [BBMP14]. Let us recall the complete form of Burgers’ equation:

$$\begin{cases} \partial_t \rho(x, t) + \partial_x [\rho(x, t)v(x, t)] = 0, \\ \partial_t v(x, t) + v(x, t)\partial_x v(x, t) = \pi^2 \rho(x, t)\partial_x \rho(x, t). \end{cases} \quad (129)$$

Recall that $g_{\mathbf{Y}(t)}(z)$ is the Stieltjes transform of the asymptotic eigenvalue distribution of $\mathbf{Y}(t)/\sqrt{m}$, and similarly $g_{\mathbf{S}}$ is the Stieltjes transform of $\rho_{\mathbf{S}}$. It is known by simple free probability that for every z such that $\text{Im}[z] > 0$, we have that $\text{Im}[g_{\mathbf{Y}(t)}(z)] > 0$ and that $g_{\mathbf{Y}(t)}$ satisfies the free convolution identity:

$$g_{\mathbf{Y}(t)}(z) = g_{\mathbf{S}}[z + tg_{\mathbf{Y}(t)}(z)]. \quad (130)$$

Note that defining ρ and v by equation (47), the boundary conditions of equation (41) are directly satisfied by the Stieltjes–Perron theorem. In the following, we denote by $(\rho_\epsilon, v_\epsilon)$ the functions defined by equation (47) without taking the $\epsilon \downarrow 0$ limit. By equation (130), we have for any x , with $f_\epsilon(x, t) \equiv -\overline{g_{\mathbf{Y}(t)}}(x + i\epsilon) = -g_{\mathbf{Y}(t)}(x - i\epsilon) = v_\epsilon(x, t) + i\pi\rho_\epsilon(x, t)$:

$$\begin{aligned} \partial_t f_\epsilon(x, t) &= -\{-f_\epsilon(x, t) - t\partial_t f_\epsilon(x, t)\}g'_{\mathbf{S}}[x - i\epsilon - tf_\epsilon(x, t)], \\ \partial_x f_\epsilon(x, t) &= -\{1 - t\partial_x f_\epsilon(x, t)\}g'_{\mathbf{S}}[x - i\epsilon - tf_\epsilon(x, t)]. \end{aligned}$$

This implies

$$\begin{aligned} \partial_t f_\epsilon + f_\epsilon \partial_x f_\epsilon &= -\{-f_\epsilon(x, t) - t\partial_t f_\epsilon(x, t) + f_\epsilon(x, t) - tf_\epsilon(x, t)\partial_x f_\epsilon(x, t)\} \\ &\quad \times g'_{\mathbf{S}}[x - i\epsilon - tf_\epsilon(x, t)], \\ &= t\{\partial_t f_\epsilon(x, t) + f_\epsilon(x, t)\partial_x f_\epsilon(x, t)\}g'_{\mathbf{S}}[x - i\epsilon - tf_\epsilon(x, t)]. \end{aligned}$$

This implies, that for any $\epsilon > 0$, $f_\epsilon(x, t)$ satisfies the complex Burgers’ equation for all $t \in (0, \Delta)$. Moreover, as $\epsilon \downarrow 0$, this solution satisfies the proper boundary conditions: this ends our justification of result 3.1.

Remark: the derivation of [Sch18]—In the appendix of his PhD thesis [Sch18], Schmidt states a result very similar to result 3.1. However, there is an essential issue in his justification: the argument is based on the use of the methods of characteristics, which is in general wrong for complex PDEs. Fortunately, the solution found remains correct, as we show above.

C.2. Derivation of result 3.2

In this part, we ‘shoot birds with cannons’ (in the words of [Sch18]), in order to derive result 3.2. The idea is quite simple: consider \mathbf{A} a symmetric $n \times n$ matrix, with a well-defined asymptotic spectral distribution ρ_A . We define the Gaussian integral:

$$I_n(\mathbf{A}) \equiv \frac{1}{n^2} \ln \int \prod_{i \leq j} \frac{dH_{ij}}{\sqrt{2(1 + \delta_{ij})\pi/n}} \exp\left\{-\frac{n}{4} \text{Tr}[\mathbf{H}^2] + \frac{n}{2} \text{Tr}[\mathbf{H}\mathbf{A}]\right\}. \quad (131)$$

It is trivial to compute it, and if $I(\rho_A) \equiv \lim_{n \rightarrow \infty} I_n(\mathbf{A})$, we have directly

$$I(\rho_A) = \lim_{n \rightarrow \infty} \frac{1}{4n} \sum_{i,j} A_{ij}^2 = \frac{1}{4} \int \rho_A(dx) x^2. \quad (132)$$

Now comes the cannon part. Can we find back equation (132) using Matytsin’s formalism and HCIZ integrals? We start from equation (131), introducing the change of variable to the eigenvalues of \mathbf{H} , exactly as we did in equation (38):

$$I_n(\mathbf{A}) = \frac{1}{n^2} \ln \mathcal{C}_n + \frac{1}{n^2} \ln \int_{\mathbb{R}^n} d\mathbf{L} \prod_{i < j} |l_i - l_j| e^{-\frac{n}{4} \sum_i l_i^2} \int_{\mathcal{O}(n)} \mathcal{D}\mathbf{O} e^{\frac{n}{2} \text{Tr}[\mathbf{O}\mathbf{L}\mathbf{O}^T \mathbf{A}]},$$

$$\mathcal{C}_n = \frac{\pi^{n(n+1)/4}}{2^{n/2} \Gamma(n+1) \prod_{i=1}^n \Gamma(i/2)} \frac{1}{(2\pi/n)^{n(n-1)/4}} \frac{1}{(4\pi/n)^{n/2}}. \quad (133)$$

We worked out the asymptotics of \mathcal{C}_n in section 3.1, and we derived:

$$\frac{1}{n^2} \ln \mathcal{C}_n = \frac{3}{8} + \mathcal{O}_n(1). \quad (134)$$

We can now use Laplace’s method on \mathbf{L} , and we reach:

$$I(\rho_A) = \frac{3}{8} + \sup_{\rho_L} \left\{ \frac{1}{2} \int \rho_L^{\otimes 2}(dx, dy) \ln |x - y| - \frac{1}{4} \int \rho_L(dx) x^2 + \frac{1}{2} I[\rho_L, \rho_A] \right\}, \quad (135)$$

with $I[\rho_L, \rho_A] = I_{\Delta=1}[\rho_L, \rho_A]$ the Matytsin function of equation (40). Using this explicit form into equation (135) we reach:

$$I(\rho_A) = \frac{1}{4} \int \rho_A(dx) x^2 + \sup_{\rho_L} \left\{ \frac{1}{4} \int \rho_L^{\otimes 2}(dx, dy) \ln |x - y| \right. \\ \left. - \frac{1}{4} \int \rho_A^{\otimes 2}(dx, dy) \ln |x - y| \right. \\ \left. - \frac{1}{4} \int_0^1 dt \int dx \rho(x, t) \left[\frac{\pi^2}{3} \rho(x, t)^2 + v(x, t)^2 \right] \right\}. \quad (136)$$

Comparing with equation (132) shows the equations presented in result 3.2.

The last part of result 3.2 can be justified using known results on high-dimensional HCIZ integrals and their links with large deviations theory [Gui04b, BBMP14]. Indeed, these works show (see e.g. equations (14) and (15) of [BBMP14]) that the following function:

$$S[\rho_L] \equiv -\frac{3}{8} + \frac{1}{4} \int \rho_A(dx) x^2 + \frac{1}{4} \int \rho_L(dx) x^2 - \frac{1}{2} \int \rho_L^{\otimes 2}(dx, dy) \ln |x - y| - \frac{1}{2} I[\rho_L, \rho_A] \tag{137}$$

is the rate function for the large deviations of the Dyson Brownian motion starting from ρ_A , at time $t = 1$. In particular, the minimum of this rate function is reached in the expected eigenvalue density of the Dyson Brownian motion, i.e. $\rho_L^* = \rho_A \boxplus \sigma_{s.c.}$, which shows the last part of result 3.2.

A long remark: an alternative derivation of result 3.2—Finally, we note that we can show equation (48) (or equivalently equation (49)) in an alternative manner that does not appeal to the Gaussian integral calculation we described. To see this, recall that we know that $\rho_L^* = \rho_A \boxplus \sigma_{s.c.}$, and we denote $\rho(x, t)$ the density of the Dyson Brownian motion (that solves the Euler–Matytsin equations between ρ_A and ρ_L^* , cf result 3.1). Let us define

$$G(t) \equiv \int dx dy \rho(x, t) \rho(y, t) \ln |x - y| - \int_0^t du \int dx \rho(x, u) \left[\frac{\pi^2}{3} \rho(x, u)^2 + v(x, u)^2 \right]. \tag{138}$$

We will obtain that $G(t)$ is constant by showing $G'(t) = 0$. From the definition of G , we have:

$$G'(t) = 2 \int \partial_t \rho(dx, t) \rho(dy, t) \ln |x - y| - \int dx \rho(x, t) \left[\frac{\pi^2}{3} \rho(x, t)^2 + v(x, t)^2 \right]. \tag{139}$$

Using equation (41) and integration by parts, we have, with P.V. the principal value of the integral:

$$G'(t) = 2 \int dx \rho(x, t) v(x, t) \text{P.V.} \left\{ \int dy \rho(y, t) \frac{1}{x - y} \right\} - \int dx \rho(x, t) \left[\frac{\pi^2}{3} \rho(x, t)^2 + v(x, t)^2 \right]. \tag{140}$$

However, we know by result 3.1 that here we have:

$$v(x, t) = \text{P.V.} \int \frac{\rho(y, t)}{x - y} dy, \tag{141}$$

Let us emphasize that this equality, shown in result 3.1, is really specific to the fact that we are considering the Matytsin equations between ρ_A and $\rho_L^* = \rho_A \boxplus \sigma_{s.c.}$. In general,

there is a remainder term in $v(x, t)$ beyond the Hilbert transform of $\rho(x, t)$ [GZ02, Gui04a, Men17]. This identity implies that equation (140) becomes:

$$G'(t) = \int dx \rho(x, t)v(x, t)^2 - \frac{\pi^2}{3} \int dx \rho(x, t)^3. \tag{142}$$

The result then follows by a simple consequence of the theory of Hilbert transformation, which we give here as a lemma. It can be found in different forms, e.g. in [Tri85] or [GZ02, Gui04a], and we give a short proof here for completeness.

Lemma C.1 (Properties of Hilbert transformation)

Let ρ be a (well-behaved) probability density, and denote $\mathcal{H}[\rho]$ its Hilbert transform:

$$\mathcal{H}[\rho](x) \equiv \frac{1}{\pi} \text{P.V.} \int \frac{\rho(y)}{x-y} dy = -\frac{1}{\pi} \lim_{\epsilon \downarrow 0} \text{Re}[g_\rho(x + i\epsilon)]. \tag{143}$$

Then one has the identity:

$$\frac{1}{3} \int dx \rho(x)^3 = \int dx \rho(x) \mathcal{H}[\rho](x)^2. \tag{144}$$

Applying lemma C.1 with $v(t) = \pi\mathcal{H}[\rho(t)]$ ends our alternative derivation of result 3.2.

Proof of lemma C.1. This simple fact can be shown in (at least) two ways. The more pedestrian way is to use the Fourier transform of the Hilbert transformation, which reads

$$\hat{\mathcal{H}}[\rho](u) \equiv \int e^{-2i\pi xu} \mathcal{H}[\rho](x) dx = -i \text{sign}(u) \hat{\rho}(u). \tag{145}$$

Here we report a more clever proof that uses the general identity:

$$\mathcal{H}[\rho]^2 = \rho^2 + 2\mathcal{H}[\rho\mathcal{H}[\rho]]. \tag{146}$$

This can be shown by an argument of complex analysis, based on the following theorem: let $f(z)$ be an analytic function in the upper-half of the complex plane. Then if $a(x) \equiv \lim_{\epsilon \downarrow 0} \text{Re}[f(x + i\epsilon)]$ and $b(x) \equiv \lim_{\epsilon \downarrow 0} \text{Im}[f(x + i\epsilon)]$, we have $b = \mathcal{H}[a]$. Using this result for f^2 (with $f = \rho + i\mathcal{H}[\rho]$) we reach equation (146). Therefore we have:

$$\int \mathcal{H}[\rho](x)^2 \rho(x) dx = \int \rho(x)^3 dx + 2 \int dx \rho(x) \mathcal{H}[\rho\mathcal{H}[\rho]](x). \tag{147}$$

It is also a common property of the Hilbert transform (for all the mentioned properties see e.g. [Tao04]) that $\int g\mathcal{H}[f] = -\int f\mathcal{H}[g]$. Applying it in the last equation yields:

$$\int dx \rho(x) \mathcal{H}[\rho](x)^2 = \int \rho(x)^3 dx - 2 \int dx \rho(x) \mathcal{H}[\rho](x)^2, \quad (148)$$

which is the sought result. \square

Appendix D. Small- α expansion of the optimal RIE

Here we study the behavior of the optimal rotationally-invariant denoiser of [BABP16] in the overcomplete limit of small α . Starting from the optimal denoiser found in equation (69), we need to understand the behavior of the Stieltjes transform of \mathbf{Y}/\sqrt{m} , with \mathbf{Y} the shifted observation matrix of equation (71).

At leading order as $\alpha \rightarrow 0$, by the random matrix equivalent of the central limit theorem (and somewhat informally), \mathbf{Y}/\sqrt{m} should behave like $\mathbf{W}/\sqrt{m} + \sqrt{\Delta}\mathbf{Z}/\sqrt{m}$, with \mathbf{W} and \mathbf{Z} independent Gaussian Wigner matrices. The Stieltjes transform of a Wigner matrix is well-known, and we get for any y eigenvalue of \mathbf{Y}/\sqrt{m} that:

$$v_{\mathbf{Y}}(y) = \frac{y}{2(\Delta + 1)} + \mathcal{O}_{\alpha}(1). \quad (149)$$

This yields:

$$\hat{\xi}_{\mu} = \frac{1}{\Delta + 1} y_{\mu} + \mathcal{O}_{\alpha}(1), \quad (150)$$

which corresponds to the limit of scalar denoising [BABP16]. Let us now go to higher order in α . We need to consider the corrections of the spectrum of \mathbf{Y} at first non-trivial order in α . This can be done by recalling the \mathcal{R} -transform of the shifted Wishart matrix $\mathbf{R}/\sqrt{m} = \mathbf{X}\mathbf{X}^{\top}/\sqrt{nm} - \alpha^{-1/2}\mathbf{I}_m$:

$$\mathcal{R}_{\mathbf{R}}(s) = \frac{1}{\sqrt{\alpha}(1 - s\sqrt{\alpha})} - \frac{1}{\sqrt{\alpha}} = s + \sqrt{\alpha}s^2 + \mathcal{O}(\alpha). \quad (151)$$

Therefore by free addition the \mathcal{R} -transform of \mathbf{Y}/\sqrt{m} is:

$$\mathcal{R}_{\mathbf{Y}}(s) = s(1 + \Delta) + \sqrt{\alpha}s^2 + \mathcal{O}(\alpha) = g_{\mathbf{Y}}^{-1}(-s) - \frac{1}{s}.$$

Letting $s = -g_{\mathbf{Y}}(z)$, we reach the expansion:

$$g_{\mathbf{Y}}(z) = \frac{\overbrace{\sqrt{-4\Delta - 4 + z^2} - z}^{\text{Wigner}}}{2(\Delta + 1)} - \left[\frac{(z - \sqrt{z^2 - 4(\Delta + 1)})^4}{8(\Delta + 1)^3 (4(\Delta + 1) + z(\sqrt{z^2 - 4(\Delta + 1)} - z))} \right] \sqrt{\alpha} + \mathcal{O}(\alpha). \quad (152)$$

For $|y| \leq 2\sqrt{\Delta + 1}$, this gives the expansion of $v_Y(y)$ as:

$$v_Y(y) = \frac{y}{2(\Delta + 1)} + \frac{(\Delta + 1 - y^2)}{2(\Delta + 1)^3} \sqrt{\alpha} + \mathcal{O}(\alpha). \quad (153)$$

In the end, we reach the first non-trivial order:

$$\hat{\xi}_\mu = \frac{1}{\Delta + 1} y_\mu - \frac{\Delta}{(\Delta + 1)^2} \left[1 - \frac{y_\mu^2}{\Delta + 1} \right] \sqrt{\alpha} + \mathcal{O}(\alpha). \quad (154)$$

Appendix E. PGY expansion for denoising

As for the factorization problem (cf section 2), we replace H_{eff} by ηH_{eff} , and we do the PGY expansion in η , and then take $\eta = 1$. The order 1 and the U operator of [GY91] at $\eta = 0$ are very simple:

$$\begin{cases} nm \frac{\partial \Phi}{\partial \eta} &= -\langle H_{\text{eff}} \rangle = 0, \\ U_{\eta=0} &= \frac{1}{\sqrt{n}} \sum_{\mu < \nu} \sum_i (iH_{\mu\nu}) x_{\mu i} x_{\nu i} = \frac{1}{2\sqrt{n}} \sum_{\mu \neq \nu} \sum_i (iH_{\mu\nu}) x_{\mu i} x_{\nu i}. \end{cases} \quad (155)$$

Recall that here we consider the denoising problem, so $x_{\mu i} \stackrel{\text{i.i.d.}}{\sim} P_X$, with $\mathbb{E}_{P_X}[x] = 0$ and $\mathbb{E}_{P_X}[x^2] = 1$. For generality, we will not assume that P_{out} is necessarily Gaussian so that we can consider any noise applied component-wise, and not only additive Gaussian noise.

Order 2—From the Georges–Yedidia formalism, we can obtain the order 2:

$$\begin{aligned} \frac{1}{2} \frac{\partial^2 \Phi}{\partial \eta^2} &= \frac{1}{2nm} \langle U^2 \rangle, \\ &= \frac{1}{8n^2m} \sum_{\mu_1 \neq \nu_1} \sum_{\mu_2 \neq \nu_2} \sum_{i,j} \langle (iH)_{\mu_1 \nu_1} (iH)_{\mu_2 \nu_2} \rangle \langle x_{\mu_1 i} x_{\nu_1 i} x_{\mu_2 j} x_{\nu_2 j} \rangle, \\ &= \frac{1}{8n^2m} \sum_{\substack{\mu_1 \neq \nu_1 \\ \mu_2 \neq \nu_2}} \sum_{i,j} (g_{\mu_1 \nu_1} g_{\mu_2 \nu_2} - r_{\mu_1 \nu_1} [\delta_{\mu_1 \mu_2} \delta_{\nu_1 \nu_2} + \delta_{\mu_1 \nu_2} \delta_{\nu_1 \mu_2}]) \\ &\quad \times (\delta_{\mu_1 \mu_2} \delta_{\nu_1 \nu_2} \delta_{ij} + \delta_{\mu_1 \nu_2} \delta_{\nu_1 \mu_2} \delta_{ij}), = \frac{1}{4nm} \sum_{\mu \neq \nu} (g_{\mu\nu}^2 - r_{\mu\nu}) = \frac{1}{2nm} \sum_{\mu < \nu} [g_{\mu\nu}^2 - r_{\mu\nu}]. \end{aligned}$$

Order 3—The order 3 can also be computed using [GY91]. We can actually use the calculation that we described in appendix B.2, and decompose the operator U_Y as $U_Y = U_c + U_g$, with terms given in equation (97). As we described very precisely in appendix B.2, we reach:

$$\langle U^3 \rangle = -\frac{1}{\sqrt{n}} \sum_{\substack{\mu_1, \mu_2, \mu_3 \\ \text{pairwise distinct}}} g_{\mu_1 \mu_2} g_{\mu_2 \mu_3} g_{\mu_3 \mu_1} + \mathcal{O}(n^{3/2}), \quad (156)$$

in which the term $\mathcal{O}(n^{3/2})$ contains all the dependency on the third-order moments of the i.i.d. variables $\{x_{\mu i}\}$, and of the independent variables $\{(ih)_{\mu\nu}\}$. All in all, we reach the expression of equation (77):

$$\frac{1}{3!} \frac{\partial^3 \Phi}{\partial \eta^3} = -\frac{1}{3! n m} \langle U^3 \rangle = \frac{1}{6 n^{3/2} m} \sum_{\substack{\sigma_1, \sigma_2, \sigma_3 \\ \text{pairwise distincts}}} g_{\sigma_1 \sigma_2} g_{\sigma_2 \sigma_3} g_{\sigma_3 \sigma_1} + \mathcal{O}(n^{-1/2}). \quad (157)$$

Appendix F. Annealed calculations in the Gaussian setting

F.1. The symmetric case

Outline of the calculation—We consider the inference problem of model \mathbf{XX}^\top , in the Gaussian setting. Equivalently, it can be written as (with $\lambda \equiv \Delta^{-1}$ and rescaling the observation \mathbf{Y}):

$$\mathbf{Y} = \sqrt{\frac{\lambda}{n}} \mathbf{X} \mathbf{X}^\top + \mathbf{Z}, \quad (158)$$

in which $\mathbf{X} \in \mathbb{R}^{m \times n}$, with a signal-to-noise ratio $\lambda > 0$, and $m/n \rightarrow \alpha > 0$. Without loss of generality, we can assume that $X_{\mu i} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$. Note that the scaling of equation (158) implies that the empirical spectral distribution of \mathbf{Y}/\sqrt{n} converges (as $n \rightarrow \infty$) to a limit measure μ_Y . Since we assumed a standard Gaussian prior on \mathbf{X} , one can write the partition function for this model as:

$$\mathcal{Z}_{\mathbf{Y}, n} \equiv \int_{\mathbb{R}^{m \times n}} \frac{d\mathbf{x}}{(2\pi)^{mn/2}} \exp\left\{-\frac{1}{2} \text{Tr}(\mathbf{x} \mathbf{x}^\top)\right\} \frac{\exp\left\{-\frac{1}{4} \text{Tr}\left(\mathbf{Y} - \sqrt{\frac{\lambda}{n}} \mathbf{x} \mathbf{x}^\top\right)^2\right\}}{(2\pi)^{m(m+1)/4}},$$

Note that compared to equation (4a) we added the diagonal terms $\mu = \nu$. As there are $\mathcal{O}(n)$ such terms, and the free energy is of order $\mathcal{O}(n^2)$, this addition does not affect the limit free entropy. We compute here the annealed free entropy $\Phi_{\text{an}}(\alpha) \equiv \lim_{n \rightarrow \infty} \{(nm)^{-1} \ln \mathbb{E} \mathcal{Z}_{\mathbf{Y}, n}\}$.

Averaging the first moment—Using the Bayes-optimality assumption, we have:

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}}[\mathcal{Z}_{\mathbf{Y}, n}] &= \int \prod_{\mu \leq \nu} dy_{\mu\nu} \prod_{a=0,1} \\ &\times \left[\int_{\mathbb{R}^{m \times n}} \frac{d\mathbf{X}^a}{(2\pi)^{\frac{nm}{2}}} e^{-\frac{1}{2} \sum_{\mu, i} (x_{\mu i}^a)^2} e^{-\frac{1}{4} \sum_{\mu, \nu} \left(y_{\mu\nu} - \sqrt{\frac{\lambda}{n}} \sum_i x_{\mu i}^a x_{\nu i}^a\right)^2} \right]. \end{aligned}$$

We can now integrate over \mathbf{Y} . We get:

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}}[\mathcal{Z}_{\mathbf{Y},n}] &= \prod_{a \in \{0,1\}} \left[\int_{\mathbb{R}^{m \times n}} \frac{d\mathbf{X}^a}{(2\pi)^{\frac{nm}{2}}} e^{-\frac{1}{2} \sum_{\mu,i} (x_{\mu i}^a)^2} \right] \\ &\times \prod_{\mu \leq \nu} \int \frac{dy_{\mu\nu}}{2\pi} e^{-\frac{1}{2(1+\delta_{\mu\nu})} \sum_a \left(y_{\mu\nu} - \sqrt{\frac{\lambda}{n}} \sum_i x_{\mu i}^a x_{\nu i}^a \right)^2}, \\ &= D_n(\alpha) \prod_{a \in \{0,1\}} \left\{ \int_{\mathbb{R}^{m \times n}} \frac{d\mathbf{X}^a}{(2\pi)^{\frac{nm}{2}}} e^{-\frac{1}{2} \sum_{\mu,i} (x_{\mu i}^a)^2} \right\} \\ &\times e^{-\frac{\lambda}{4n} \sum_{\mu,\nu} \sum_a \left(\sum_i x_{\mu i}^a x_{\nu i}^a \right)^2 + \frac{\lambda}{8n} \sum_{\mu,\nu} \left(\sum_a \sum_i x_{\mu i}^a x_{\nu i}^a \right)^2}, \end{aligned}$$

with a constant $D_n(\alpha) > 0$. Note that the above equation can be rewritten, using the identity:

$$\begin{aligned} &e^{-\frac{\lambda}{4n} \sum_{\mu,\nu} \sum_{a \in \{0,1\}} \left(\sum_i x_{\mu i}^a x_{\nu i}^a \right)^2 + \frac{\lambda}{8n} \sum_{\mu,\nu} \left(\sum_{a \in \{0,1\}} \sum_i x_{\mu i}^a x_{\nu i}^a \right)^2} \\ &= \exp \left\{ -\frac{\lambda}{8n} \sum_{\mu,\nu} \left[\sum_i x_{\mu i}^0 x_{\nu i}^0 - \sum_i x_{\mu i}^1 x_{\nu i}^1 \right]^2 \right\}, \end{aligned}$$

so that we decouple the replicas (up to a constant that only depends on n and α):

$$\begin{aligned} &\exp \left\{ -\frac{\lambda}{8n} \sum_{\mu,\nu} \left[\sum_i x_{\mu i}^0 x_{\nu i}^0 - \sum_i x_{\mu i}^1 x_{\nu i}^1 \right]^2 \right\} \\ &\simeq \int \prod_{\mu \leq \nu} dQ_{\mu\nu} e^{-\frac{n}{4} \sum_{\mu,\nu} Q_{\mu\nu}^2} \exp \left\{ \frac{i\sqrt{\lambda}}{2\sqrt{2}} \sum_{\mu,\nu} Q_{\mu\nu} \left[\sum_i x_{\mu i}^0 x_{\nu i}^0 - \sum_i x_{\mu i}^1 x_{\nu i}^1 \right] \right\}. \end{aligned}$$

We can integrate over the prior distribution on $\{\mathbf{X}^0, \mathbf{X}^1\}$ and reach, up to a constant $B_n(\alpha)$ independent of λ :

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}}[\mathcal{Z}_{\mathbf{Y},n}] &= B_n(\alpha) \int \prod_{\mu \leq \nu} dQ_{\mu\nu} e^{-\frac{n}{4} \sum_{\mu,\nu} Q_{\mu\nu}^2} e^{-\frac{n}{2} \ln \det [I_m + i\sqrt{\frac{\lambda}{2}} \mathbf{Q}]} e^{-\frac{n}{2} \ln \det [I_m - i\sqrt{\frac{\lambda}{2}} \mathbf{Q}]}, \\ &= B_n(\alpha) \int \prod_{\mu \leq \nu} dQ_{\mu\nu} e^{-\frac{n}{4} \sum_{\mu,\nu} Q_{\mu\nu}^2} e^{-\frac{n}{2} \ln \det [I_m + \frac{\lambda}{2} \mathbf{Q}^2]}. \end{aligned} \tag{159}$$

The saddle point—We note that in equation (159) the integrand only depends on the *spectrum* of the matrix \mathbf{Q} . Since \mathbf{Q} is integrated over the set of symmetric matrices, we can write $\mathbf{Q} = \mathbf{O}\mathbf{D}\mathbf{O}^\top$ and use the classical change of variables to this representation (see e.g. proposition 4.1.1 of [AGZ10]):

$$\mathbb{E}_{\mathbf{Y}}[\mathcal{Z}_{\mathbf{Y},n}] = C_n(\alpha) \int_{\mathbb{R}^m} d\mathbf{L} \times \exp \left\{ -\frac{n}{4} \sum_{\mu=1}^m l_{\mu}^2 + \frac{1}{2} \sum_{\mu \neq \nu} \ln |l_{\mu} - l_{\nu}| - \frac{n}{2} \sum_{\mu=1}^m \ln \left(1 + \frac{\lambda}{2} l_{\mu}^2 \right) \right\}, \quad (160)$$

with a constant $C_n(\alpha)$ that we will compute at the end by taking the $\lambda \downarrow 0$ limit. In the end we obtain the following variational principle for the annealed free entropy:

$$\Phi_{\text{an}}(\alpha) = c(\alpha) + \sup_{\nu \in \mathcal{M}_1^+(\mathbb{R})} \left\{ -\frac{1}{4} \mathbb{E}_{\nu}[X^2] + \frac{\alpha}{2} \mathbb{E}_{X,Y \sim \nu} \ln |X - Y| - \frac{1}{2} \mathbb{E}_{\nu} \left[\ln \left(1 + \frac{\lambda}{2} x^2 \right) \right] \right\}, \quad (161)$$

in which $c(\alpha) \equiv \lim_{n \rightarrow \infty} \{ (nm)^{-1} \ln C_n(\alpha) \}$, and $\mathcal{M}_1^+(\mathbb{R})$ is the set of real probability measures.

Computing $c(\alpha)$ —When $\lambda = 0$ the solution to the variational problem is easily known, as it is related to the celebrated Wigner semi-circle law. More precisely, the solution to the variational problem of equation (161) at $\lambda = 0$ is:

$$\nu_{\alpha}^*(dx) = \frac{1}{2\pi\alpha} \mathbb{1}_{|x| \leq 2\sqrt{\alpha}} \sqrt{4\alpha - x^2} dx.$$

And it satisfies:

$$-\frac{1}{4} \int \nu_{\alpha}^*(dx) x^2 + \frac{\alpha}{2} \int \nu_{\alpha}^*(dx) \nu_{\alpha}^*(dy) \ln |x - y| = -\frac{3\alpha}{8} + \frac{\alpha}{4} \ln \alpha.$$

Moreover, a direct calculation at $\lambda = 0$ gives $\Phi_{\text{an}}(\alpha)|_{\lambda=0} = -(\alpha/4) \ln 4\pi$. This yields

$$c(\alpha) = \frac{3\alpha}{8} - \frac{\alpha}{4} \ln 4\pi - \frac{\alpha}{4} \ln \alpha. \quad (162)$$

Exact resolution of the variational principle—We now detail an exact resolution of the variational principle of equation (161). This calculation is very much inspired by a similar derivation, for the Wigner semi-circle law, done in [LNV18]. We define the potential

$$V_{\alpha}(x) \equiv \frac{1}{4\alpha} x^2 + \frac{1}{2\alpha} \ln \left(1 + \frac{\lambda}{2} x^2 \right),$$

Going back to equation (160), prior to taking the $n \rightarrow \infty$ limit, we can apply the saddle point method to the set of eigenvalues $\{l_{\mu}\}_{\mu=1}^m$. The saddle-point equation for the eigenvalues reads:

$$\frac{1}{m} \sum_{\nu(\neq\mu)} \frac{1}{l_\mu - l_\nu} = \frac{1}{2\alpha} \left(l_\mu + \frac{\lambda l_\mu}{1 + \frac{\lambda}{2} l_\mu^2} \right) = V'_\alpha(l_\mu). \quad (163)$$

A classical way to solve equations of the type of equation (163) is to use Tricomi's theorem [Tri85], as explained for instance in [LNV18]. Here we follow an alternative (equivalent) way by writing a self-consistent equation on the Stieltjes transform. We will multiply both sides of equation (163) by $(z - l_\mu)^{-1}$, before summing the result over μ . Recall the definition of the Stieltjes transform for any $z \in \mathbb{C}_+$:

$$\mathcal{S}(z) \equiv \lim_{n \rightarrow \infty} \frac{1}{m} \sum_{\mu=1}^m \frac{1}{l_\mu - z}.$$

As can be found, e.g. in [LNV18] (or very quickly derived), we have at large n, m :

$$\frac{1}{m^2} \sum_{\mu \neq \nu} \frac{1}{l_\mu - l_\nu} \frac{1}{z - l_\mu} = \frac{1}{2} \mathcal{S}(z)^2 + \mathcal{O}(1/n).$$

We now focus on the right-hand side of equation (163), multiplied by $(z - l_\mu)^{-1}$, and summed over μ . The first term can be simplified using the relation:

$$\frac{1}{2\alpha} \frac{1}{m} \sum_{\mu=1}^m \frac{l_\mu}{z - l_\mu} = -\frac{1 + z\mathcal{S}(z)}{2\alpha} + \mathcal{O}(1/n).$$

The other term can also be simplified:

$$\begin{aligned} \Gamma(z) &\equiv \frac{1}{2\alpha} \frac{1}{m} \sum_{\mu=1}^m \frac{\lambda l_\mu}{1 + \frac{\lambda}{2} l_\mu^2} \frac{1}{z - l_\mu}, \\ &= \frac{1}{2\alpha z} \left(\frac{1}{m} \sum_{\mu=1}^m \frac{\lambda l_\mu}{1 + \frac{\lambda}{2} l_\mu^2} + \frac{1}{m} \sum_{\mu=1}^m \frac{2}{z - l_\mu} \right) - \frac{1}{2\alpha z} \frac{1}{m} \sum_{\mu=1}^m \frac{2}{(1 + \frac{\lambda}{2} l_\mu^2)(z - l_\mu)}. \end{aligned} \quad (164)$$

Note that by the symmetry $x \rightarrow -x$ in the potential $V_\alpha(x)$, the optimizing measure should be symmetric with respect to 0. This implies that

$$\lim_{n \rightarrow \infty} \frac{1}{m} \sum_{\mu=1}^m \frac{\lambda l_\mu}{1 + \frac{\lambda}{2} l_\mu^2} = 0.$$

Let us denote $L > 0$ the limit:

$$\frac{1}{m} \sum_{\mu=1}^m \frac{1}{1 + \frac{\lambda}{2} l_\mu^2} \equiv L + \mathcal{O}_n(1).$$

Finally, we have the relation

$$\frac{1}{m} \sum_{\mu=1}^m \frac{2}{\left(1 + \frac{\lambda}{2} l_{\mu}^2\right)(z - l_{\mu})} = \frac{2}{\lambda z} \frac{1}{m} \sum_{\mu=1}^m \frac{\lambda l_{\mu}}{\left(1 + \frac{\lambda}{2} l_{\mu}^2\right)(z - l_{\mu})} + \frac{2}{z} \frac{1}{m} \sum_{\mu=1}^m \frac{1}{1 + \frac{\lambda}{2} l_{\mu}^2}.$$

Using these tricks, we reach for the term we were treating:

$$\Gamma(z) = \frac{1}{2\alpha z} (-2\mathcal{S}(z) + \mathcal{O}_n(1)) - \frac{1}{2\alpha z} \left[\frac{2}{\lambda z} 2\alpha \Gamma(z) + \frac{2}{z} L + \mathcal{O}_n(1) \right],$$

so that

$$\Gamma(z) = -\frac{\lambda}{2\alpha} \frac{L + z\mathcal{S}(z)}{1 + \frac{\lambda}{2} z^2} + \mathcal{O}_n(1).$$

Finally going back to equation (163) and taking the $n \rightarrow \infty$ limit, we obtain:

$$\mathcal{S}(z)^2 = -\frac{1}{\alpha} [1 + z\mathcal{S}(z)] - \frac{\lambda}{\alpha} \frac{L + z\mathcal{S}(z)}{1 + \frac{\lambda}{2} z^2}.$$

We can solve this quadratic equation to obtain:

$$\mathcal{S}(z) = -\frac{1}{2} \left[\frac{\lambda z}{\alpha \left(\frac{\lambda z^2}{2} + 1\right)} + \frac{z}{\alpha} - \sqrt{\left(\frac{\lambda z}{\alpha \left(\frac{\lambda z^2}{2} + 1\right)} + \frac{z}{\alpha} \right)^2 - 4 \left(\frac{1}{\alpha} + \frac{\lambda L}{\alpha \left(\frac{\lambda z^2}{2} + 1\right)} \right)} \right].$$

The solution is taken such that $\text{Im}[\mathcal{S}(z)] > 0$ for $z \in \mathbb{C}_+$ and $\mathcal{S}(z) \sim -1/z$ for $z \rightarrow +\infty$. Using the Stieltjes–Perron inversion formula, the density of the maximizing measure ν_{α}^* is thus:

$$\frac{d\nu_{\alpha}^*}{dx} = \lim_{\epsilon \rightarrow 0^+} \frac{1}{\pi} \text{Im}[\mathcal{S}(x + i\epsilon)] = \frac{1}{2\pi\alpha} \sqrt{4 \left(\alpha + \frac{\lambda\alpha L}{1 + \frac{\lambda x^2}{2}} \right) - \left(\frac{\lambda x}{1 + \frac{\lambda x^2}{2}} + x \right)^2}. \tag{165}$$

The constant $L > 0$ has to be chosen in order to ensure the proper normalization of the probability distribution.

F.2. The non-symmetric case

The calculation—We consider now the non-symmetric setting, i.e. model FX, with $m/n \rightarrow \alpha > 0$ and $p/n \rightarrow \psi > 0$, and in the Gaussian setting:

$$\mathbf{Y} = \sqrt{\frac{\lambda}{n}} \mathbf{F}\mathbf{X} + \mathbf{Z},$$

in which $\mathbf{Y} \in \mathbb{R}^{m \times p}$, $\mathbf{F} \in \mathbb{R}^{m \times n}$ and $\mathbf{X} \in \mathbb{R}^{n \times p}$. Without loss of generality, we assume an i.i.d. standard Gaussian prior for both \mathbf{F} , \mathbf{X} , and that the quenched noise \mathbf{Z} is Gaussian

Perturbative construction of mean-field equations in extensive-rank matrix factorization and denoising as well. The partition function of this model is:

$$\mathcal{Z}_{\mathbf{Y},n} = \int d\mathbf{f} d\mathbf{x} \frac{e^{-\frac{1}{2}\sum_{\mu,i} f_{\mu i}^2} e^{-\frac{1}{2}\sum_{i,l} x_{il}^2} e^{-\frac{1}{2}\sum_{\mu,l} \left(y_{\mu l} - \sqrt{\frac{\lambda}{n}} \sum_i f_{\mu i} x_{il}\right)^2}}{(2\pi)^{\frac{mn}{2}} (2\pi)^{\frac{mp}{2}} (2\pi)^{\frac{mp}{2}}}.$$

We can write the first moment of the partition function using Bayes-optimality:

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}}[\mathcal{Z}_{\mathbf{Y},n}] &= \int \prod_{\mu,l} dy_{\mu l} \prod_{a \in \{0,1\}} \left[\int d\mathbf{f}^a d\mathbf{x}^a \frac{e^{-\frac{1}{2} \text{Tr} \mathbf{x}^a (\mathbf{x}^a)^\top - \frac{1}{2} \text{Tr} \mathbf{f}^a (\mathbf{f}^a)^\top}}{(2\pi)^{\frac{mn}{2} + \frac{mp}{2}}} \right] \\ &\times \frac{e^{-\frac{1}{2} \sum_{a \in \{0,1\}} \sum_{\mu,l} \left(y_{\mu l} - \sqrt{\frac{\lambda}{n}} \sum_i f_{\mu i}^a x_{il}^a\right)^2}}{(2\pi)^{mp}}. \end{aligned}$$

Integrating over \mathbf{Y} , we reach:

$$\begin{aligned} \mathbb{E}_{\mathbf{Y}}[\mathcal{Z}_{\mathbf{Y},n}] &= \prod_{a \in \{0,1\}} \left[\int d\mathbf{f}^a d\mathbf{x}^a \frac{e^{-\frac{1}{2} \text{Tr} \mathbf{x}^a (\mathbf{x}^a)^\top - \frac{1}{2} \text{Tr} \mathbf{f}^a (\mathbf{f}^a)^\top}}{(2\pi)^{\frac{mn}{2} + \frac{mp}{2}}} \right] \\ &\times \frac{\exp\left\{-\frac{\lambda}{4n} \sum_{\mu,l} \left[\sum_i f_{\mu i}^0 x_{il}^0 - \sum_i f_{\mu i}^1 x_{il}^1\right]^2\right\}}{2^{\frac{mp}{2}} (2\pi)^{\frac{mp}{2}}}. \end{aligned}$$

We can then use a Gaussian transformation in the last term:

$$\begin{aligned} &\exp\left\{-\frac{\lambda}{4n} \sum_{\mu,l} \left[\sum_i f_{\mu i}^0 x_{il}^0 - \sum_i f_{\mu i}^1 x_{il}^1\right]^2\right\} \\ &= \left(\frac{n}{2\pi}\right)^{\frac{mp}{2}} \int \prod_{\mu,l} dQ_{\mu l} e^{-\frac{n}{2} \sum_{\mu,l} Q_{\mu l}^2} \\ &\times \exp\left\{i\sqrt{\frac{\lambda}{2}} \sum_{\mu,l} Q_{\mu l} \left[\sum_i f_{\mu i}^0 x_{il}^0 - \sum_i f_{\mu i}^1 x_{il}^1\right]\right\}. \end{aligned}$$

One integrates now over the prior distributions on $\mathbf{f}^a, \mathbf{x}^a$. The two replicas yield the same contribution and we reach:

$$\mathbb{E}_{\mathbf{Y}}[\mathcal{Z}_{\mathbf{Y},n}] = D_n(\alpha, \psi) \int \prod_{\mu,l} dQ_{\mu l} e^{-\frac{n}{2} \sum_{\mu,l} Q_{\mu l}^2} e^{-n \ln \det[\mathbf{I}_{p+\frac{1}{2}} \mathbf{Q}^\top \mathbf{Q}]}, \quad (166)$$

with D_n a constant that only depends on n, α, ψ . To compute this integral, we use a Weyl-type formula for integrating a function that only depends on the singular values of the integrated matrix. It is stated for instance in proposition 4.1.3 of [AGZ10]. We need to separate the two possibilities $\alpha \leq \psi$ and $\alpha > \psi$.

The case $\alpha \leq \psi$ —We directly obtain from equation (166) and the Weyl-type formula:

$$\mathbb{E}_{\mathbf{Y}} \mathcal{Z}_n(\mathbf{Y}) = C_n(\alpha, \psi) \int \prod_{\mu=1}^m dl_{\mu} \\ \times e^{-\frac{n}{2} \sum_{\mu} l_{\mu}^2 - n \sum_{\mu} \ln(1 + \frac{\lambda}{2} l_{\mu}^2) + (p-m) \sum_{\mu} \ln l_{\mu} + \frac{1}{2} \sum_{\mu \neq \nu} \ln |l_{\mu}^2 - l_{\nu}^2|},$$

in which $C_n(\alpha, \psi)$ is a constant and we denote $C(\alpha, \psi) \equiv \lim_{n \rightarrow \infty} \{\ln C_n(\alpha, \psi) / (n(m + p))\}$. Performing a saddle point finally yields:

$$\Phi_{\text{an}}(\alpha, \psi) = C(\alpha, \psi) + \sup_{\nu \in \mathcal{M}_1^+(\mathbb{R}_+)} \left\{ -\frac{\alpha}{2(\alpha + \psi)} \mathbb{E}_{\nu}[X] - \frac{\alpha}{\alpha + \psi} \mathbb{E}_{\nu} \left[\ln \left(1 + \frac{\lambda}{2} X \right) \right] \right. \\ \left. + \frac{\alpha(\psi - \alpha)}{2(\alpha + \psi)} \mathbb{E}_{\nu}[\ln X] + \frac{\alpha^2}{2(\alpha + \psi)} \mathbb{E}_{X, Y \sim \nu} \ln |X - Y| \right\}. \quad (167)$$

The case $\alpha > \psi$ —This time, we obtain from equation (166) and the Weyl-type formula:

$$\mathbb{E}_{\mathbf{Y}} \mathcal{Z}_n(\mathbf{Y}) = D_n(\alpha, \psi) \int \prod_{l=1}^p da_l \\ \times e^{-\frac{n}{2} \sum_l a_l^2 - n \sum_l \ln(1 + \frac{\lambda}{2} a_l^2) + (m-p) \sum_l \ln a_l + \frac{1}{2} \sum_{l \neq l'} \ln |a_l^2 - a_{l'}^2|},$$

in which $D_n(\alpha, \psi)$ is a constant and we denote $D(\alpha, \psi) \equiv \lim_{n \rightarrow \infty} [\ln D_n(\alpha, \psi) / (n(m + p))]$. Performing a saddle point finally yields:

$$\Phi_{\text{an}}(\alpha, \psi) = D(\alpha, \psi) + \sup_{\nu \in \mathcal{M}_1^+(\mathbb{R}_+)} \left\{ -\frac{\psi}{2(\alpha + \psi)} \mathbb{E}_{\nu}[X] - \frac{\psi}{\alpha + \psi} \mathbb{E}_{\nu} \left[\ln \left(1 + \frac{\lambda}{2} X \right) \right] \right. \\ \left. + \frac{\psi(\alpha - \psi)}{2(\alpha + \psi)} \mathbb{E}_{\nu}[\ln X] + \frac{\psi^2}{2(\alpha + \psi)} \mathbb{E}_{X, Y \sim \nu} \ln |X - Y| \right\}. \quad (168)$$

Computing the constants—The constants can again be found using the $\lambda \downarrow 0$ limit. If $\alpha \leq \psi$, the extremizing measure $\nu_{\alpha, \psi}^*$ is a scaled Marchenko–Pastur distribution [MP67] with ratio $\alpha/\psi \leq 1$:

$$\nu_{\alpha, \psi}^*(dx) = \frac{1}{2\pi\alpha} \frac{\sqrt{(\lambda_+ - x)(x - \lambda_-)}}{x} \mathbb{1}_{\{\lambda_- \leq x \leq \lambda_+\}} dx, \quad (169)$$

with $\lambda_{\pm} = (\sqrt{\psi} \pm \sqrt{\alpha})^2$. Indeed, $\nu_{\alpha, \psi}^*$ is the asymptotic eigenvalue distribution of $\mathbf{Q}\mathbf{Q}^T$, in which $\{Q_{\mu l}\}$ are independent Gaussian random variables with zero mean and variance $n^{-1} = \psi p^{-1}$. As an easy calculation gives $\Phi_{\text{an}}(\alpha, \psi) = -[\alpha\psi / (2(\alpha + \psi))] \ln 4\pi$ for $\lambda = 0$, this yields after some computation:

$$C(\alpha, \psi) = -\frac{\alpha^2 \ln(\alpha) + \alpha\psi(-3 + 4 \ln(2)) + 2 \ln(\pi) - (\alpha - \psi)^2 \ln(\psi - \alpha) + \psi^2 \ln(\psi)}{4(\alpha + \psi)}.$$

In the same way, one can get $D(\alpha, \psi) = C(\psi, \alpha)$.

Solving for the maximizing density—We can solve the variational principles of equations (167) and (168) in a similar way as what we did in appendix F.1 for the symmetric case. In the limit $n \rightarrow \infty$, the saddle point equation for $x_\mu \equiv l_\mu^2$ gives:

$$-\frac{1}{2} - \frac{\lambda}{2} \frac{1}{1 + \lambda x_\mu/2} + \frac{1}{n} \sum_{\nu(\neq\mu)} \frac{1}{x_\mu - x_\nu} + \frac{(\psi - \alpha)}{2} \frac{1}{x_\mu} = 0.$$

As before we multiply by $(z - x_\mu)^{-1}$ and we sum over $\mu = 1, \dots, m$. In this way we get:

$$\frac{1}{nm} \sum_{\mu=1}^m \sum_{\nu(\neq\mu)=1}^m \frac{1}{x_\mu - x_\nu} \frac{1}{z - x_\mu} = \frac{\alpha \mathcal{S}(z)^2}{2} + \mathcal{O}(1/n).$$

Next, we add an arbitrary parameter $\gamma > 0$ in the calculation, before taking the limit $\gamma \downarrow 0$, as done in [LNV18]. In this way we can compute:

$$\frac{1}{m} \sum_{\mu=1}^m \left(\frac{(\psi - \alpha)}{2} + \frac{\gamma}{2} \right) \frac{1}{(z - x_\mu)} \frac{1}{x_\mu} = \left(\frac{(\psi - \alpha)}{2} + \frac{\gamma}{2} \right) \left(\frac{Q - \mathcal{S}(z)}{z} \right),$$

with the constant

$$Q \equiv \frac{1}{m} \sum_{\mu=1}^m \frac{1}{x_\mu} = \int dx \frac{\nu(x)}{x}.$$

Next we have:

$$\frac{\lambda}{2} \frac{1}{m} \sum_{\mu=1}^m \frac{1}{1 + \lambda x_\mu/2} \frac{1}{z - x_\mu} = \frac{\lambda}{2} \frac{1}{\lambda z/2 + 1} (Y - \mathcal{S}(z)),$$

with

$$Y \equiv \frac{\lambda}{2} \int dx \frac{\mu(x)}{1 + \lambda x/2} = \mathcal{S}(-2/\lambda).$$

Altogether the equation for $\mathcal{S}(z)$ reads:

$$\mathcal{S}(z) - \frac{\lambda}{\lambda z/2 + 1} (Y - \mathcal{S}(z)) + (\psi - \alpha + \gamma) \frac{Q - \mathcal{S}(z)}{z} + \alpha \mathcal{S}(z)^2 = 0,$$

from which we have (after taking the limit $\gamma \downarrow 0$), for $z > 0$:

$$\mathcal{S}(z) = \frac{1}{2\alpha} \left[-1 + \frac{(\psi - \alpha)}{z} - \frac{\lambda}{1 + \lambda z/2} + \frac{1}{z} \sqrt{\left(z - \psi - \alpha + \frac{z\lambda}{1 + \lambda z/2} \right)^2 - 4\alpha z \left((\psi - \alpha)Q - \frac{z\lambda}{1 + \lambda z/2} Y \right)} \right].$$

The behavior at $z \rightarrow \infty$ fixes the constant Y :

$$Y = \frac{1}{2}((\psi - \alpha)Q - 1).$$

From this one can deduce the final density which reads:

$$\frac{d\nu_{\alpha,\psi}^*}{dx} = \frac{1}{2\alpha\pi x} \sqrt{-\left(x - \psi - \alpha + \frac{x\lambda}{1 + \lambda x/2}\right)^2 + 4\alpha x \left((\psi - \alpha)Q - \frac{x\lambda}{1 + \lambda x/2} \frac{1}{2}((\psi - \alpha)Q - 1)\right)}, \quad (170)$$

where the value of Q should be determined by imposing the normalization of $\nu_{\alpha,\psi}^*$.

Appendix G. Other technicalities

G.1. The free entropy of [KKM+16] in the Gaussian setting

In this symmetric Gaussian setting, with prior of variance 1 and noise $\Delta > 0$ in the channel, the prediction of [KKM+16] for the quenched free entropy is:

$$\Phi_{\text{sec.}} = \text{extr}_{q,\hat{q}} \left\{ \frac{1}{2}q\hat{q} - \frac{\hat{q}}{2} - \frac{1}{2} \ln(1 - \hat{q}) - \frac{\alpha}{4}(1 + \ln 2\pi) - \frac{\alpha}{4} \ln[\Delta + (1 - q^2)] \right\}. \quad (171)$$

We denote it $\Phi_{\text{sec.}}$, since we saw in section 2 that it corresponds to a second-order truncation of the PGY expansion for the factorization problem. The extremum in equation (171) is given by the set of equations:

$$\hat{q} = -\frac{\alpha q}{\Delta + (1 - q^2)}; \quad q = -\frac{\hat{q}}{1 - \hat{q}}.$$

In the paramagnetic phase (which corresponds to the denoising solution) we moreover have $q = \hat{q} = 0$ and thus:

$$\Phi_{\text{Kab.}}(q = 0) = -\frac{\alpha}{4}(1 + \ln 2\pi) - \frac{\alpha}{4} \ln(1 + \Delta).$$

Coming back to the factorization problem, the equation on q is $\hat{q} = -q/(1 - q)$ and (if $q \neq 0$):

$$\alpha(1 - q) = \Delta + 1 - q^2,$$

which easily solves into:

$$q = \frac{\alpha \pm \sqrt{(2 - \alpha)^2 + 4\Delta}}{2}. \quad (172)$$

Computing $\Phi_{\text{sec.}}$ can then easily be done: one must consider three possible solutions: the paramagnetic one, and the two solutions of equation (172). One then checks if said solutions are physical (i.e. if $q \in [0, 1)$). To finish, one must compute the free entropies of all physical solutions, and the actual prediction for the free entropy is given by the largest of those.

G.2. Absence of contribution of diagonal terms to the MMSE

Recall that we consider the denoising model:

$$\mathbf{Y} = \frac{1}{\sqrt{n}} \mathbf{X}^* (\mathbf{X}^*)^\top + \sqrt{\Delta} \mathbf{Z}. \tag{173}$$

We let $\mathbf{S}^* \equiv \mathbf{X}^* (\mathbf{X}^*)^\top / \sqrt{n}$, and we define the associated asymptotic free energy:

$$\Phi_{\mathbf{Y},n} \equiv \frac{1}{nm} \mathbb{E} \ln \int \mathcal{D}\mathbf{X} e^{-\frac{1}{4\Delta} \sum_{\mu,\nu} (Y_{\mu\nu} - \frac{1}{\sqrt{n}} \sum_{i=1}^n X_{\mu i} X_{\nu i})^2}. \tag{174}$$

And the asymptotic MMSE is:

$$\text{MMSE} \equiv \frac{1}{m^2} \mathbb{E} \sum_{\mu,\nu} (S_{\mu\nu}^* - \langle S_{\mu\nu} \rangle)^2. \tag{175}$$

In this paragraph, we also consider a definition of the MMSE ‘not taking into account the diagonal’:

$$\text{MMSE} \equiv \frac{1}{m^2} \mathbb{E} \sum_{\mu \neq \nu} (S_{\mu\nu}^* - \langle S_{\mu\nu} \rangle)^2. \tag{176}$$

In the limit $n \rightarrow \infty$, we now show that equations (175) and (176) are equivalent. Equivalently, we will show:

$$\text{MMSE}_{\text{diag}} \equiv \frac{1}{m^2} \mathbb{E} \left\{ \sum_{\mu=1}^m (S_{\mu\mu}^* - \langle S_{\mu\mu} \rangle)^2 \right\} \rightarrow_{n \rightarrow \infty} 0. \tag{177}$$

Proof of equation (177). We will make use of the *Nishimori identity*, a consequence of Bayes-optimality. Physically speaking, the Nishimori identity shows that the planted solution \mathbf{S}^* behaves like a replica of the system under the posterior distribution, and we refer the reader to [BKM⁺19] for its elementary proof:

Proposition G.1 (Nishimori identity)

Let $(X, Y) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ be a couple of random variables. Let $k \geq 1$ and let $X^{(1)}, \dots, X^{(k)}$ be k i.i.d. samples (given Y) from the conditional distribution $\mathbb{P}(X = \cdot | Y)$. Let us denote $\langle - \rangle$ the expectation operator w.r.t. $\mathbb{P}(X = \cdot | Y)$ and \mathbb{E} the expectation w.r.t. (X, Y) . Then, for all continuous bounded function g we have

$$\mathbb{E} \langle g(Y, X^{(1)}, \dots, X^{(k)}) \rangle = \mathbb{E} \langle g(Y, X^{(1)}, \dots, X^{(k-1)}, X) \rangle.$$

By proposition G.1, we have

$$\text{MMSE}_{\text{diag}} = \frac{1}{m^2} \mathbb{E} \left\{ \sum_{\mu=1}^m (S_{\mu\mu}^*)^2 - \langle S_{\mu\mu} \rangle^2 \right\}. \tag{178}$$

Using the simple variance bound $\mathbb{E}X^2 \geq (\mathbb{E}X)^2$, and again the Nishimori identity, we have:

$$\begin{aligned} \text{MMSE}_{\text{diag}} &\leq \frac{1}{m^2} \sum_{\mu=1}^m \left\{ \mathbb{E} \left[(S_{\mu\mu}^*)^2 \right] - (\mathbb{E} \langle S_{\mu\mu} \rangle)^2 \right\}, \\ &\leq \frac{1}{m^2} \sum_{\mu=1}^m \left\{ \mathbb{E} \left[(S_{\mu\mu}^*)^2 \right] - (\mathbb{E} S_{\mu\mu}^*)^2 \right\}, \\ &\leq \frac{1}{m^2 n} \sum_{\mu=1}^m \left\{ \sum_{i,j} \mathbb{E} [(X_{\mu i}^*)^2 (X_{\mu j}^*)^2] - \sum_{i,j} \mathbb{E} [(X_{\mu i}^*)^2] \mathbb{E} [(X_{\mu j}^*)^2] \right\}, \\ &\leq \frac{1}{m^2 n} \sum_{\mu=1}^m \{n^2 + 2n - n^2\} \leq \frac{2}{m}, \end{aligned}$$

since $\mathbb{E}[(X_{\mu i}^*)^2] = 1$. This ends the proof of equation (177). □

G.3. Proof of the I-MMSE theorem for denoising

We shall derive equation (57), making use several times of the Nishimori identity (proposition G.1). We start with:

$$\frac{\partial \Psi[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}]}{\partial \Delta^{-1}} = \frac{\alpha \Delta}{4} + \frac{1}{4nm} \sum_{\mu, \nu} \left\{ -2m [\mathbb{E}(S_{\mu\nu}^*)^2 - \mathbb{E}(\langle S_{\mu\nu} \rangle^2)] + \sqrt{\Delta m} \mathbb{E}[Z_{\mu\nu} \langle S_{\mu\nu} \rangle] \right\}. \quad (179)$$

Moreover we have (taking into account both the $\mu \leq \nu$ and $\mu > \nu$ terms):

$$\mathbb{E}[Z_{\mu\nu} \langle S_{\mu\nu} \rangle] = \sqrt{\frac{m}{\Delta}} [\mathbb{E}(S_{\mu\nu}^*)^2 - \mathbb{E}(\langle S_{\mu\nu} \rangle^2)]. \quad (180)$$

So in the end we reach:

$$\frac{\partial \Psi[\rho_{\mathbf{Y}}, \rho_{\mathbf{S}}]}{\partial \Delta^{-1}} = \frac{\alpha \Delta}{4} - \frac{1}{4n} \sum_{\mu, \nu} [\mathbb{E}(S_{\mu\nu}^*)^2 - \mathbb{E}(\langle S_{\mu\nu} \rangle^2)] = \frac{\alpha \Delta}{4} - \frac{\alpha}{4} \text{MMSE}(\Delta), \quad (181)$$

which ends the proof.

References

- [AGZ10] Anderson G W, Guionnet A and Zeitouni O 2010 *An Introduction to Random Matrices* vol 118 (Cambridge: Cambridge University Press)
- [BABP16] Bun J, Allez R, Bouchaud J-P and Potters M 2016 Rotational invariant estimator for general noisy matrices *IEEE Trans. Inf. Theory* **62** 7475–90
- [BAMCM97] Belouchrani A, Abed-Meraim K, Cardoso J-F and Moulines E 1997 A blind source separation technique using second-order statistics *IEEE Trans. Signal Process.* **45** 434–44
- [BBAP05] Baik J, Ben Arous G and Pécché S 2005 Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices *Ann. Probab.* **33** 1643–97

- [BBMP14] Bun J, Bouchaud J-P, Majumdar S N and Potters M 2014 Instanton approach to large n Harish–Chandra–Itzykson–Zuber integrals *Phys. Rev. Lett.* **113** 070201
- [BBP17] Bun J, Bouchaud J-P and Potters M 2017 Cleaning large correlation matrices: tools from random matrix theory *Phys. Rep.* **666** 1–109
- [BG11] Benaych-Georges F 2011 Rectangular R -transform as the limit of rectangular spherical integrals *J. Theor. Probab.* **24** 969
- [BKM+19] Barbier J, Krzakala F, Macris N, Miolane L and Zdeborová L 2019 Optimal errors and phase transitions in high-dimensional generalized linear models *Proc. Natl Acad. Sci.* **116** 5451–60
- [BLM15] Bayati M, Lelarge M and Montanari A 2015 Universality in polytope phase transitions and message passing algorithms *Ann. Appl. Probab.* **25** 753–822
- [BM11] Bayati M and Montanari A 2011 The dynamics of message passing on dense graphs, with applications to compressed sensing *IEEE Trans. Inf. Theory* **57** 764–85
- [BM21] Barbier J and Macris N 2021 Statistical limits of dictionary learning: random matrix theory and the spectral replica method (arXiv:2109.06610)
- [Bol14] Bolthausen E 2014 An iterative construction of solutions of the TAP equations for the Sherrington–Kirkpatrick model *Commun. Math. Phys.* **325** 333–66
- [CLMW11] Candès E J, Li X, Ma Y and Wright J 2011 Robust principal component analysis? *J. ACM* **58** 1–37
- [CR09] Candès E J and Recht B 2009 Exact matrix completion via convex optimization *Found. Comput. Math.* **9** 717–72
- [CT10] Candès E J and Tao T 2010 The power of convex relaxation: near-optimal matrix completion *IEEE Trans. Inf. Theory* **56** 2053–80
- [DGJ18] Donoho D L, Gavish M and Johnstone I M 2018 Optimal shrinkage of eigenvalues in the spiked covariance model *Ann. Stat.* **46** 1742
- [DJM13] Donoho D L, Javanmard A and Montanari A 2013 Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing *IEEE Trans. Inf. Theory* **59** 7434–64
- [EAK18] El Alaoui A and Krzakala F 2018 Estimation in the spiked Wigner model: a short proof of the replica formula *2018 IEEE Int. Symp. Information Theory (ISIT)* (Piscataway, NJ: IEEE) pp 1874–8
- [FG16] Forrester P J and Grela J 2016 Hydrodynamical spectral evolution for random matrices *J. Phys. A: Math. Theor.* **49** 085203
- [GB21] Gerbelot C and Berthier R 2021 Graph-based approximate message passing iterations (arXiv:2109.11905)
- [GH21] Guionnet A and Huang J 2021 Large deviations asymptotics of rectangular spherical integral (arXiv:2106.07146)
- [GSV05] Guo D, Shamai S and Verdú S 2005 Mutual information and minimum mean-square error in Gaussian channels *IEEE Trans. Inf. Theory* **51** 1261–82
- [Gui04a] Guionnet A 2004 First order asymptotics of matrix integrals; a rigorous approach towards the understanding of matrix models *Commun. Math. Phys.* **244** 527–69
- [Gui04b] Guionnet A 2004 Large deviations and stochastic calculus for large random matrices *Probab. Surv.* **1** 72–172
- [GY91] Georges A and Yedidia J S 1991 How to expand around mean-field theory using high-temperature expansions *J. Phys. A: Math. Gen.* **24** 2173
- [GZ02] Guionnet A and Zeitouni O 2002 Large deviations asymptotics for spherical integrals *J. Funct. Anal.* **188** 461–515
- [Har57] Harish-Chandra 1957 Differential operators on a semisimple Lie algebra *Am. J. Math.* **79** 87–120
- [IZ80] Itzykson C and Zuber J-B 1980 The planar approximation: II *J. Math. Phys.* **21** 411–21
- [JM13] Javanmard A and Montanari A 2013 State evolution for general approximate message passing algorithms, with applications to spatial coupling *Inf. Inference* **2** 115–44
- [KDMMR+03] Kreutz-Delgado K, Murray J F, Rao B D, Engan K, Lee T-W and Sejnowski T J 2003 Dictionary learning algorithms for sparse representation *Neural Comput.* **15** 349–96
- [KKM+16] Kabashima Y, Krzakala F, Mézard M, Sakata A and Zdeborová L 2016 Phase transitions and sample complexity in Bayes-optimal matrix factorization *IEEE Trans. Inf. Theory* **62** 4228–65
- [KMZ13] Krzakala F, Mézard M and Zdeborová L 2013 Phase diagram and approximate message passing for blind calibration and dictionary learning *2013 IEEE Int. Symp. Information Theory* (Piscataway, NJ: IEEE) pp 659–63
- [LKZ17] Lesieur T, Krzakala F and Zdeborová L 2017 Constrained low-rank matrix estimation: phase transitions, approximate message passing and applications *J. Stat. Mech.* **073403**
- [LML+17] Lesieur T, Miolane L, Lelarge M, Krzakala F and Zdeborová L 2017 Statistical and computational phase transitions in spiked tensor estimation *2017 IEEE Int. Symp. Information Theory (ISIT)* (Piscataway, NJ: IEEE) pp 511–5
- [LNV18] Livan G, Novaes M and Vivo P 2018 *Introduction to Random Matrices: Theory and Practice* vol 26 (Berlin: Springer)
- [LPPZ21] Lucibello C, Pittorino F, Perugini G and Zecchina R 2021 Deep learning via message passing algorithms based on belief propagation (arXiv:2110.14583)

- [Mat94] Matytsin A 1994 On the large- N limit of the Itzykson–Zuber integral *Nucl. Phys. B* **411** 805–20
- [MBPS09] Mairal J, Bach F, Ponce J and Sapiro G 2009 Online dictionary learning for sparse coding *Proc. 26th Annual Int. Conf. Machine Learning* pp 689–96
- [Meh04] Mehta M L 2004 *Random Matrices* vol 142 (Amsterdam: Elsevier)
- [Men17] Menon G 2017 The complex Burgers’ equation, the HCIZ integral and the Calogero–Moser system *Random Matrix Theory at CMSA (Harvard)*
- [Méz89] Mézard M 1989 The space of interactions in neural networks: Gardner’s computation with the cavity method *J. Phys. A: Math. Gen.* **22** 2181
- [MFC+19] Maillard A, Foini L, Lage Castellanos A, Krzakala F, Mézard M and Zdeborová L 2019 High-temperature expansions and message-passing algorithms *J. Stat. Mech.* **113301**
- [Min01] Minka T P 2001 Expectation propagation for approximate Bayesian inference *Proc. 17th Conf. Uncertainty in Artificial Intelligence* (Morgan Kaufmann Publishers Inc.) pp 362–9
- [MKMZ22] Maillard A, Krzakala F, Mézard M and Zdeborová L 2022 Numerical code used to produce the figures https://github.com/sphinxsteam/perturbative_mean_field_matrix_factorization
- [MM09] Mézard M and Montanari A 2009 *Information, Physics, and Computation* (Oxford: Oxford University Press)
- [MP67] Marchenko V A and Pastur L A 1967 Distribution of eigenvalues for some sets of random matrices *Mat. Sb.* **114** 507–36
- [MPV86] Mézard M, Parisi G and Virasoro M A 1986 SK model: the replica solution without replicas *Europhys. Lett.* **1** 77–82
- [MPV87] Mézard M, Parisi G and Virasoro M 1987 *Spin Glass Theory and Beyond: An Introduction to the Replica Method and its Applications* vol 9 (Singapore: World Scientific)
- [Nic14] Nicolaescu L I 2014 Complexity of random smooth functions on compact manifolds *Indiana Univ. Math. J.* **63** 1037–65
- [Nis01] Nishimori H 2001 *Statistical Physics of Spin Glasses and Information Processing: An Introduction* vol 111 (Oxford: Clarendon)
- [OCW16] Opper M, Cakmak B and Winther O 2016 A theory of solving TAP equations for Ising models with general invariant random matrices *J. Phys. A: Math. Theor.* **49** 114002
- [OF96] Olshausen B A and Field D J 1996 Emergence of simple-cell receptive field properties by learning a sparse code for natural images *Nature* **381** 607–9
- [OF97] Olshausen B A and Field D J 1997 Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vis. Res.* **37** 3311–25
- [OW01a] Opper M and Winther O 2001 Adaptive and self-averaging Thouless–Anderson–Palmer mean-field theory for probabilistic modeling *Phys. Rev. E* **64** 056131
- [OW01b] Opper M and Winther O 2001 Tractable approximations for probabilistic models: the adaptive Thouless–Anderson–Palmer mean field approach *Phys. Rev. Lett.* **86** 3695
- [OW05] Opper M and Winther O 2005 Expectation consistent approximate inference *J. Mach. Learn. Res.* **6** 2177–204
- [Ple82] Plefka T 1982 Convergence condition of the TAP equation for the infinite-ranged Ising spin glass model *J. Phys. A: Math. Gen.* **15** 1971
- [PSC14a] Parker J T, Schniter P and Cevher V 2014 Bilinear generalized approximate message passing: I: Derivation *IEEE Trans. Signal Process.* **62** 5839–53
- [PSC14b] Parker J T, Schniter P and Cevher V 2014 Bilinear generalized approximate message passing: II: Applications *IEEE Trans. Signal Process.* **62** 5854–67
- [Ran11] Rangan S 2011 Generalized approximate message passing for estimation with random linear mixing *2011 IEEE Int. Symp. Information Theory Proc.* (Piscataway, NJ: IEEE) pp 2168–72
- [Rom55] Romberg W 1955 Vereinfachte numerische integration *Norske Vid. Selsk. Forhdl.* **28** 30–6
- [RSF17] Rangan S, Schniter P and Fletcher A K 2017 Vector approximate message passing *2017 IEEE Int. Symp. Information Theory (ISIT)* (Piscataway, NJ: IEEE) pp 1588–92
- [RT12] Ricci-Tersenghi F 2012 The Bethe approximation for solving the inverse Ising problem: a comparison with other inference methods *J. Stat. Mech.* **P08015**
- [Sch18] Schmidt H C 2018 Statistical physics of sparse and dense models in optimization and inference *PhD Thesis* Université Paris-Saclay (ComUE) <https://tel.archives-ouvertes.fr/tel-03227132/>
- [SK75] Sherrington D and Kirkpatrick S 1975 Solvable model of a spin-glass *Phys. Rev. Lett.* **35** 1792
- [SK13] Sakata A and Kabashima Y 2013 Statistical mechanics of dictionary learning *Europhys. Lett.* **103** 28008
- [SRF16] Schniter P, Rangan S and Fletcher A K 2016 Vector approximate message passing for the generalized linear model *2016 50th Asilomar Conf. Signals, Systems and Computers* (Piscataway, NJ: IEEE) pp 1525–9
- [Tao04] Tao T 2004 Lecture notes 4 for 247A <http://math.ucla.edu/tao/247a>
- [TAP77] Thouless D J, Anderson P W and Palmer R G 1977 Solution of ‘solvable model of a spin glass’ *Phil. Mag.* **35** 593–601
- [Tri85] Tricomi F G 1985 *Integral Equations* vol 5 (New York: Dover)

- [Voi86] Voiculescu D 1986 Addition of certain non-commuting random variables *J. Funct. Anal.* **66** 323–46
- [ZHT06] Zou H, Hastie T and Tibshirani R 2006 Sparse principal component analysis *J. Comput. Graph. Stat.* **15** 265–86
- [ZK16] Zdeborová L and Krzakala F 2016 Statistical physics of inference: thresholds and algorithms *Adv. Phys.* **65** 453–552
- [ZZY21] Zou Q, Zhang H and Yang H 2021 Multi-layer bilinear generalized approximate message passing *IEEE Trans. Signal Process.* **69** 4529–43