

Bayes-optimal limits in structured PCA, and how to reach them

Jean Barbier^{†,*}, Francesco Camilli^{‡,*},
Marco Mondelli[◇], Manuel Sáenz[†]

[†] *International Center for Theoretical Physics, Italy*

[‡] *Bologna University, Italy*

[‡] *École Normale Supérieure de Paris, France*

[◇] *Institute of Science and Technology, Austria*

Abstract

We study the paradigmatic spiked matrix model of principal components analysis, where the rank-one signal is corrupted by additive noise. While the noise is typically taken from a Wigner matrix with independent entries, here the potential acting on the eigenvalues has a quadratic plus a quartic component. The quartic term induces strong correlations between the matrix elements, which makes the setting relevant for applications but analytically challenging. Our work provides the first characterization of the Bayes-optimal limits for inference in this model with structured noise. If the signal prior is rotational-invariant, then we show that a spectral estimator is optimal. In contrast, for more general priors, the existing approximate message passing algorithm (AMP) falls short of achieving the information-theoretic limits, and we provide a justification for this sub-optimality. Finally, by generalizing the theory of Thouless-Anderson-Palmer equations, we cure the issue by proposing a novel AMP which matches the theoretical limits. Our information-theoretic analysis is based on the replica method, a powerful heuristic from statistical mechanics; instead, the novel AMP comes with a rigorous state evolution analysis tracking its performance in the high-dimensional limit. Even if we focus on a specific noise distribution, our methodology can be generalized to a wide class of trace ensembles, at the cost of more involved expressions.

* Corresponding authors: jbarbier@ictp.it, francesco.camilli2@unibo.it

Contents

1	Introduction, problem setting and main results	3
1.1	Introduction and related works	3
1.2	Probabilistic model of PCA with structured noise	6
1.3	A concrete example: the quartic ensemble	9
1.4	Main results	11
2	The inhomogeneous spherical integral	17
2.1	Definition and variational characterization	18
2.2	Special cases	19
2.2.1	Low-rank HCIZ integral	19
2.2.2	Low-rank spherical integral with external field and diagonal replica coupling	19
2.2.3	Low-rank spherical integral with non-diagonal replica coupling and replica symmetric overlap	20
2.3	Derivation of the variational formula	21
3	Information-theoretic analysis by the replica method	23
3.1	An equivalent quadratic model	24
3.2	Replica symmetric free entropy using the inhomogeneous spherical integral	27
3.3	Replica saddle point equations	30
3.4	Spectral PCA is optimal for rotation-invariant signals	34
4	Sub-optimality of the previously proposed AMP	37
4.1	Analysis of the one-step AMP fixed point performance	38
4.2	Analysis of the replica Bayes-optimal fixed point	40
4.3	What is actually doing this sub-optimal AMP? Mismatched estimation with Gaussian likelihood	42
5	Towards an optimal AMP: AdaTAP formalism	47
5.1	The AdaTAP single-instance free entropy	48
5.2	Saddle point: reduction to an Ising model, AdaTAP equations and optimal pre-processing of the data	50
5.3	Simplifying the AdaTAP equations by self-averaging of the Onsager reaction term	51

6	Approximate message passing, optimally	55
6.1	BAMP: Bayes-optimal AMP	56
6.2	Onsager coefficients and state evolution recursion	57
7	Numerics	61
7.1	Spectral properties of the pre-processed matrix $\mathbf{J}(\mathbf{Y})$	62
7.2	BAMP improves over the existing AMP and matches the replica prediction for the MMSE	62
A	Learning the optimal pre-processing $\mathbf{J}(\mathbf{Y})$ by expectation maximization	66
B	Proofs for BAMP	68
B.1	Auxiliary AMP	68
B.2	State evolution of auxiliary AMP	69
B.3	Proof of Theorem 1	71

1 Introduction, problem setting and main results

1.1 Introduction and related works

Given their ubiquitous appearance in the statistics literature, spiked matrix models, which were originally formulated as probabilistic models for principal component analysis (PCA) [58], are now a paradigm in high dimensional inference. Thanks to their universality features they, and their generalizations, find numerous applications in other central problems such as community detection [1, 63], group synchronization [91, 90], sub-matrix localization or high-dimensional clustering [66]; see [67, 92] for more applications.

In this paper we focus on the following estimation problem: a statistician needs to extract a rank-one matrix (the spike) $\mathbf{X}^*\mathbf{X}^{*\top}$, $\mathbf{X}^* \in \mathbb{R}^N$, from the data

$$\mathbf{Y} = \frac{\lambda}{N}\mathbf{X}^*\mathbf{X}^{*\top} + \mathbf{Z} \in \mathbb{R}^{N \times N} \quad (1)$$

with some additive noise \mathbf{Z} . The positive parameter λ , referred to as signal-to-noise ratio (SNR), sets the strength of the signal with respect to that of the noise.

The spectral properties of finite rank perturbations of large random matrices like (1) were intensively investigated in random matrix theory [7, 8, 94, 45, 30, 81, 24, 25, 6], showing the presence of a spectral transition often called BBP transition (in reference to the authors of [7]): when λ is large enough, the top eigenvalue of \mathbf{Y}

detaches from the bulk of the eigenvalue distribution. Its corresponding eigenvector has then a non trivial projection onto the sought ground truth \mathbf{X}^* , and can be used as its estimator.

The problem has also been approached from the angle of Bayesian inference. In particular, besides the previous spectral estimator, there exists a whole family of iterative algorithms, known as approximate message passing (AMP), that can be tailored to take further advantage of prior structural information known about the signal. AMP algorithms were first proposed for estimation in linear models [59, 23, 22, 35, 64, 72], but have since been applied to a range of statistical estimation problems, including generalized linear models [14, 68, 71, 76, 95, 96, 97] and low-rank matrix estimation [34, 46, 61, 67, 79, 17]. An attractive feature of AMP is that under suitable model assumptions, its performance in the high-dimensional limit is precisely characterized by a succinct deterministic recursion called state evolution [22, 27, 57]. Using the state evolution analysis, it has been proved that AMP achieves Bayes-optimal performance for some models [34, 36, 79, 14], and a conjecture from statistical physics posits that for a wide range of estimation problems, AMP is optimal among polynomial-time algorithms.

The references mentioned above rely on the assumption of Gaussian identically and independently distributed (i.i.d.) noise $Z_{ij} \sim \mathcal{N}(0, 1)$, under which the model identified by (1) is the well-known Wigner spiked model [40, 16, 3, 58]. This independence, or “absence of structure”, in the noise has many advantages from the theoretical point of view due to the numerous simplifications it generates.

In order to relax this property, we can seek inspiration from the statistical physics literature on disordered systems. An idea that was first brought forth in [21, 87] for the Sherrington-Kirkpatrick model, and later imported also in high dimensional inference [4, 5, 53], is that of giving an inhomogeneous variance profile to the noise matrix elements (we mention that this idea in inference is similar to the earlier definition of “spatially coupled systems” [43, 65] in coding theory, see [11, 12] for its use in the present context). This procedure makes the (Z_{ij}) no longer identically distributed, but it leaves them independent. This an important step towards more structure in the noise (and therefore the data). Yet, the independence assumption is a rather strong one. Actually, [53] showed that for a broad class of observation models, as long as the independence assumption holds, the model is information-theoretically equivalent to one with independent Gaussian (possibly inhomogeneous) noise.

One way to go beyond this last assumption is to consider noises that belong to the wider class of *rotationally invariant matrices*. Since the appearance of the seminal works [73, 74, 89], there has been a remarkable development in this direction, as evidenced by the rapidly growing number of papers on spin glasses [26, 84, 86, 42, 70, 47] and inference [24, 25, 48, 49, 69, 98, 41, 102, 101] that try to take into account struc-

tured disorder, including the present one. Indeed, we hereby consider a spiked model in which the noise matrix \mathbf{Z} is drawn from an orthogonal matrix ensemble different from the Gaussian orthogonal ensemble (which is the only rotationally invariant ensemble such that the matrix entries are independent). Intuitively, the presence of dependencies in the noise should be exploitable by an algorithm that is sharp enough to see patterns within it and use them to retrieve the sought rank one matrix more efficiently. Going in that direction, in [41] the author proposed a version of AMP designed for rotationally invariant noises (using earlier ideas of [86, 84]) and provided also a rigorous state evolution analysis for it. Furthermore, in a recent work [13], part of the authors performed a rigorous analysis of a Bayes estimator and an AMP, both assuming Gaussian noise, whereas the actual noise in the data was drawn from a generic orthogonal matrix ensemble. However, besides intuition and the mentioned works, to our best knowledge there is little theoretical understanding of the true role played by noise structure in spiked matrix estimation and more generically in inference. In particular, prior to our work there was no theoretical prediction of optimal performance to benchmark practical inference algorithms.

Organization of the paper.

The end of this section properly defines the model and the quartic random matrix ensemble we consider. In Section 2, we define and analyze an integral dubbed *inhomogeneous spherical integral*, that will play an essential role in the analysis. For those interested mainly in the main results this section can be skipped at first reading. Section 3 contains the core information-theoretic analysis based on the replica method. We also show at the end of it that, for rotationally invariant priors, the spectral estimator is Bayes-optimal in the MMSE sense. Next, in Section 4, we analyze both the fixed point performance of the previously proposed AMP for structured PCA [41] and our replica prediction for the MMSE. We deduce that, in general, the AMP in [41] is sub-optimal, and we provide an explanation for why this is the case. Using the theory of adaptive TAP equations [86], Section 5 lays the foundations for defining an optimal AMP: the main outcome is an optimal pre-processing polynomial function that depends on the statistical properties of the noise and has to be applied to the data, in order to achieve Bayes-optimality. Section 6 demonstrates that, by exploiting this pre-processing function, a novel AMP can be written down which *does match* the MMSE predicted by the replica theory. This algorithm comes with a scalar state evolution recursion which rigorously tracks its performance in the limit of large size. The Onsager reaction coefficients of our AMP are different from those in [41] and their calculation, as well as the state evolution analysis, requires new ideas. To highlight these differences and emphasize the match with the replica

MMSE, this new algorithm is dubbed Bayes-optimal AMP, or *BAMP*. In the final Section 7 we provide a numerical confirmation of our theoretical predictions. In Appendix A, we provide expectation-maximization (EM) equations that learn the optimal pre-processing function to be used by BAMP, when noise statistics are not known. In the last technical Appendix B, we gather the proofs of the various results needed to reach the state evolution of our BAMP algorithm.

Notations.

Bold notations are reserved for vectors and matrices. By default a vector \mathbf{x} is a column vector, and its transpose \mathbf{x}^\top is therefore a row vector. Thus the usual L_2 norm $\|\mathbf{x}\|^2 = \mathbf{x}^\top \mathbf{x}$ and $\mathbf{x}\mathbf{x}^\top$ is a rank-one projector. The notation $\mathbf{x} \xrightarrow{W_2} X$ denotes convergence of the empirical distribution of the random vector \mathbf{x} to the random variable X in Wasserstein-2 distance. Symbol \propto means “equality up to a constant” (often, a normalization constant) and $:=$ is an equality by definition. Tr is the usual trace operator. For a vector \mathbf{x} , the matrix $\text{diag}(\mathbf{x})$ is diagonal with \mathbf{x} on its diagonal. For a diagonal matrix \mathbf{A} and a function $F : \mathbb{R} \mapsto \mathbb{R}$ the matrix $F(\mathbf{A})$ is diagonal with F applied componentwise to each diagonal entry of \mathbf{A} . A function F applied to a real symmetric $N \times N$ matrix diagonalizable as $\mathbf{M} = \mathbf{U}\mathbf{A}\mathbf{U}^\top$ acts in the standard way: $F(\mathbf{M}) := \mathbf{U}F(\mathbf{A})\mathbf{U}^\top$. \mathbb{E}_A is an expectation with respect to the random variable A ; \mathbb{E} is an expectation with respect to all random variables entering the ensuing expression. For a function F of one argument we denote F' its derivative. Notations like $i \leq N$ always implicitly assume that the index i starts at 1. Notation $[t] := \{1, 2, \dots, t\} = \{i \leq t\}$. Powers for vectors apply componentwise (this is however *not* the case for matrices). We often compactly write $\mathbb{E}(\dots)^2 = \mathbb{E}[(\dots)^2] \geq (\mathbb{E}(\dots))^2$ and similarly for other functions, we denote equivalently $\mathbb{E}[f(\dots)]$ and $\mathbb{E}f(\dots)$. Matrix I_N is the identity of size N .

1.2 Probabilistic model of PCA with structured noise

Consider a vector $\mathbf{X}^* = (X_i^*)_{i \leq N}$ whose components are drawn i.i.d. from a given distribution P_X with support bounded uniformly in N . Two cases will be considered: the factorized case

$$dP_X(\mathbf{X}^*) = \prod_{i \leq N} dP_X(X_i^*) = \prod_{i \leq N} P_X(X_i^*) dX_i^*,$$

and the case where dP_X is the uniform measure over the N -sphere of radius \sqrt{N} . If not specified the first case is assumed. We will always consider priors with unit

second moment

$$\int dP_X(x)x^2 = 1.$$

This is just a convention as if one wants to consider a different normalization, it can simply be included through a proper rescaling of the SNR λ .

The inference task we are interested in is the retrieval of the rank-one ‘‘spike’’ $\mathbf{P}^* := \mathbf{X}^*\mathbf{X}^{*\top}$ from the following observed matrix

$$\mathbf{Y} = \frac{\lambda}{N}\mathbf{P}^* + \mathbf{Z}, \quad (2)$$

where \mathbf{Z} is a unknown noise matrix, $\lambda \geq 0$ is the SNR. Whenever \mathbf{Z} is a Wigner matrix this model corresponds to the usual Wigner spike model. But here we no longer assume that the noise is *unstructured* (namely, has independent entries). More specifically, we will assume that is drawn from a certain orthogonal rotationally invariant random matrix ensemble defined by a potential $V : \mathbb{R} \mapsto \mathbb{R}$ and a density (with normalization constant C_V)

$$dP_Z(\mathbf{Z}) = C_V d\mathbf{Z} \exp\left(-\frac{N}{2}\text{Tr}V(\mathbf{Z})\right). \quad (3)$$

Rotational invariance means that \mathbf{Z} equals in distribution $\mathbf{U}^\top\mathbf{Z}\mathbf{U}$ for any orthogonal matrix \mathbf{U} (this follows from the trace in the exponent) [93]. More precisely, when changing variables from matrix \mathbf{Z} to eigenvalues \mathbf{D} and eigenbasis \mathbf{O} via $\mathbf{Z} = \mathbf{O}^\top\mathbf{D}\mathbf{O}$ we have

$$dP_Z(\mathbf{D}, \mathbf{O}) = C_V d\mathbf{O} d\mathbf{D} \exp\left(-\frac{N}{2}\text{Tr}V(\mathbf{D})\right) \prod_{i < j} |D_i - D_j|. \quad (4)$$

The measure $d\mathbf{O}$ is the Haar measure, i.e., uniform measure over the orthogonal group $\mathbf{O}(N)$, and the last term coupling all eigenvalues in a pairwise long-range fashion is the Vandermonde determinant. Note that only the special case $V(x) = x^2/(2\sigma)$ corresponding to the Gaussian orthogonal ensemble induces independent (Gaussian distributed) matrix entries (up to symmetry). Any other potential generates dependencies among matrix elements and thus *structure*. E.g., if we take $V(x) = x^4/4$,

$$dP_Z(\mathbf{Z}) = C_V d\mathbf{Z} \prod_{i,j,k,l} \exp\left(-\frac{N}{8}Z_{ij}Z_{jk}Z_{kl}Z_{li}\right) \quad (5)$$

which clearly is not factorizable over matrix entries.

We now introduce the Bayesian framework which we are going to analyse. Let the projector $\mathbf{P} := \mathbf{x}\mathbf{x}^\top$. This allows us to write the posterior measure of the inference problem:

$$dP_{X|Y}(\mathbf{x} | \mathbf{Y}) = \frac{C_V}{P_Y(\mathbf{Y})} dP_X(\mathbf{x}) \exp\left(-\frac{N}{2} \text{Tr}V\left(\mathbf{Y} - \frac{\lambda}{N} \mathbf{P}\right)\right). \quad (6)$$

Because both the prior $P_X(\mathbf{x})$ matches the density of the signal and the likelihood $P_{Y|X}$ matches the noise density P_Z and moreover the SNR λ is known, the posterior written above is the “correct” one and we are in the *Bayesian-optimal setting*. Studying the limits of inference in this setting draws a fundamental line between what is information-theoretically possible and what is not in terms of performance of inference. The evidence reads

$$P_Y(\mathbf{Y}) = C_V \int dP_X(\mathbf{x}) \exp\left(-\frac{N}{2} \text{Tr}V\left(\mathbf{Y} - \frac{\lambda}{N} \mathbf{P}\right)\right). \quad (7)$$

One of the main object of interest is the *free entropy* (or minus the *free energy*), which is nothing else than minus the Shannon entropy of the data:

$$F_N(\mathbf{Y}) := -H(\mathbf{Y}) = \mathbb{E} \ln P_Y(\mathbf{Y}). \quad (8)$$

Therefore the free entropy is related to the mutual information by an additive constant corresponding to the entropy of the noise, and is therefore simply computed (while the free entropy is not):

$$\begin{aligned} I(\mathbf{P}^*; \mathbf{Y}) &= -F_N(\mathbf{Y}) - H(\mathbf{Y} | \mathbf{X}^*) \\ &= -F_N(\mathbf{Y}) - H(\mathbf{Z}) \\ &= -F_N(\mathbf{Y}) - \ln C_V + \frac{N}{2} \mathbb{E} \text{Tr}V(\mathbf{Z}). \end{aligned} \quad (9)$$

Using the explicit form of the observation model (2) the free entropy reads

$$\begin{aligned} F_N(\mathbf{Y}) &= \mathbb{E} \ln \int dP_X(\mathbf{x}) \exp\left(-\frac{N}{2} \text{Tr}\left[V\left(\mathbf{Z} + \frac{\lambda}{N}(\mathbf{P}^* - \mathbf{P})\right) - V(\mathbf{Z})\right]\right) \\ &\quad + \ln C_V - \frac{N}{2} \mathbb{E} \text{Tr}V(\mathbf{Z}). \end{aligned} \quad (10)$$

We extracted the noise entropy in the second line so that we can isolate the mutual information and to make the argument of the integrated exponential of order N . In this way the problem is naturally mapped onto a statistical mechanics model with extensive Hamiltonian given by minus the log-likelihood:

$$H_N(\mathbf{x}; \mathbf{Z}, \mathbf{X}^*) := \frac{N}{2} \text{Tr}\left[V\left(\mathbf{Z} + \frac{\lambda}{N}(\mathbf{P}^* - \mathbf{P})\right) - V(\mathbf{Z})\right]. \quad (11)$$

Indeed, our Hamiltonian can be rewritten as

$$\frac{1}{2}\text{Tr}(\mathbf{P}^* - \mathbf{P}) \int_0^\lambda dt V' \left(\mathbf{Z} + \frac{t}{N} (\mathbf{P}^* - \mathbf{P}) \right).$$

The difference between the two projectors has only two eigenvalues of order N and the matrix inside the potential derivative has $O(1)$ eigenvalues, hence the previous is of $O(N)$ too. The free entropy is thus directly linked to the expected log-partition function associated to this Hamiltonian:

$$\mathbb{E} \ln \mathcal{Z}(\mathbf{Y}) := \mathbb{E} \ln \int dP_X(\mathbf{x}) \exp \left(- H_N(\mathbf{x}; \mathbf{Z}, \mathbf{X}^*) \right). \quad (12)$$

The notation $;$ in $H_N(\mathbf{x}; \mathbf{Z}, \mathbf{X}^*)$ emphasizes that \mathbf{Z}, \mathbf{X}^* are quenched variables while \mathbf{x} fluctuates according the Gibbs-Boltzmann distribution associated to this Hamiltonian (i.e., the posterior). The same notation with same meaning for Hamiltonians will be used later on.

1.3 A concrete example: the quartic ensemble

Analysing this model for a generic potential V is *possible* through the novel methodology presented in this paper. But as it will become apparent, if we take a generic polynomial potential V , the higher the order of this polynomial, the more technical and cumbersome it becomes. So for the sake of pedagogy we focus in the present contribution on a very concrete example of non trivial correction to the i.i.d. noise hypothesis. As a matter of fact, the simplest inference problem with correlated noise elements is that with the quartic matrix potential: for two positive real numbers μ and γ we restrict our analysis to the potential

$$V(x) = \frac{\mu}{2}x^2 + \frac{\gamma}{4}x^4. \quad (13)$$

This was first studied by Brézin et al in [28] to study the planar approximation of quantum field theories with large internal symmetry groups. We could have also considered a non-symmetric potential with a cubic term too, but for simplicity we restrict ourselves to that case as symmetry will slightly simplify the computations (but there is no barrier to applying our methods to that a more general, possibly non-even, potential).

The matrix ensemble defined by (13) has a known Stieltjes transform \mathcal{S} and asymptotic eigenvalue density ρ , see, e.g., [93]: if \mathbf{Z} is a sequence of matrices of

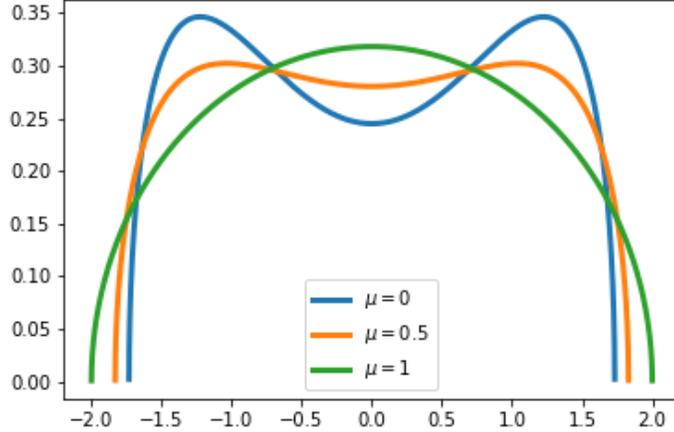


Figure 1: Asymptotic spectral density (14) of the random noise ensemble defined by the potential (13) from less structured (with independent entries) at $(\mu = 1, \gamma = 0)$, corresponding to the standard semi-circle law, to the more structured $(\mu = 0, \gamma = 16/27)$ (recall relation (17)).

increasing size N drawn from (5) with the above quartic potential and whose sequence of eigenvalues is $(D_i)_{i \leq N}$, then

$$\frac{1}{N} \sum_{i \leq N} \delta_{D_i, x} \xrightarrow{N \rightarrow \infty} \rho(x) = \frac{1}{2\pi} (\mu + 2a^2\gamma + \gamma x^2) \sqrt{4a^2 - x^2}, \quad (14)$$

$$\mathcal{S}(z) = \int \frac{d\rho(x)}{z - x} = \frac{1}{2} (\mu z + \gamma z^3 - (\mu + 2a^2\gamma + \gamma z^2) \sqrt{z^2 - 4a^2}), \quad (15)$$

for a z lying outside of the support of ρ , and where

$$a^2 := \frac{\sqrt{\mu^2 + 12\gamma} - \mu}{6\gamma}. \quad (16)$$

It is evident that when $\gamma \rightarrow 0^+$ one has $a^2 \rightarrow 1/\mu$ and consequently $\rho(x) \rightarrow \rho_{\text{sc}}(x)$ the standard semi-circle law, see Figure 1. In principle the choice of γ and μ is totally free, as long as¹ $\gamma > 0$. However, we are interested in a noise with unit variance in

¹We use implicitly the convexity of the potential, which requires $\mu, \gamma > 0$, to obtain the density of eigenvalues [93]. But we believe that this condition can be relaxed if one can get an associated well-defined asymptotic spectral density and that our analysis would still hold.

order to be able to make a meaningful comparison with models with unstructured noise. By enforcing this unitarity constraint one finds a relation between γ and μ :

$$\gamma = \gamma(\mu) = \frac{1}{27}(8 - 9\mu + \sqrt{64 - 144\mu + 108\mu^2 - 27\mu^3}). \quad (17)$$

With this choice one can check that

$$\int d\rho(x)x^2 = 1 \quad \text{for any } \mu \in [0, 1].$$

When $(\mu = 1, \gamma(1) = 0)$ we recover the pure Wigner case already analyzed in great details. On the contrary $(\mu = 0, \gamma(0) = 16/27)$ corresponds to a purely quartic case with unit variance, and to the “most structured” ensemble in our restricted class of noise ensembles. Therefore, μ can be thought of as a parameter allowing to interpolate between unstructured and structured noise ensembles. Even for this simple family of potentials, as soon as $\mu < 1$, neither the Bayes-optimal nor the algorithmic limits of inference are known (except for those of a simple spectral algorithm, see [24]).

1.4 Main results

Our main contributions can be divided in two categories: those on the fundamental, information-theoretic, limitations of inference in structured PCA and, complementary to that, novel algorithmic ideas allowing to match these Bayes-optimal limits efficiently. Both require conceptual insights and technical advances that we emphasize. We gather here these results and state them informally; we refer to the main sections for precise statements.

Information-theoretic results

- Our analysis of the information-theoretic (Bayes-optimal) performance based on the non-rigorous replica method yields first a low-dimensional variational formulation for the free entropy (log-partition function) of the model when P_X is factorized:

Result 1 (Free entropy). The free entropy (i.e., minus Shannon entropy of the data) verifies in the limit of large size the following characterization:

$$\frac{1}{N}F_N(\mathbf{Y}) = -\frac{1}{N}H(\mathbf{Y}) \xrightarrow{N \rightarrow \infty} \text{extr } f_\rho(\boldsymbol{\tau})$$

where $\boldsymbol{\tau} \in \mathbb{R}^{13}$ and for an explicit real-valued function $f_\rho : \mathbb{R}^{13} \mapsto \mathbb{R}$ depending on the noise asymptotic spectral density ρ . See (61) for the complete statement. Here

and everywhere in the paper extr stands for the following “extremization” procedure: if $f : \mathbb{R}^k \mapsto \mathbb{R}$ then

$$\text{extr } f(\boldsymbol{\tau}) := f(\boldsymbol{\tau}_*) \text{ where } \boldsymbol{\tau}_* := \underset{\{\boldsymbol{\tau} \in \mathbb{R}^k : \nabla f(\boldsymbol{\tau}) = \mathbf{0}\}}{\text{argmax}} f(\boldsymbol{\tau}).$$

We will see that despite the apparent mess, the 13-dimensional system of equations defining $\boldsymbol{\tau}_*$ will reduce to a much simpler 3-dimensional one thanks to special symmetries inherent to the Bayes-optimal nature of our analysis, and known as Nishimori identity in physics, which is a simple consequence of Bayes rule [18]:

Nishimori identity. For any bounded function f of the signal \mathbf{X}^* , the data \mathbf{Y} and of conditionally i.i.d. samples from the posterior $\mathbf{x}^j \sim P_{X|Y}(\cdot | \mathbf{Y})$, $j = 1, 2, \dots, n$, we have that

$$\mathbb{E}\langle f(\mathbf{Y}, \mathbf{X}^*, \mathbf{x}^2, \dots, \mathbf{x}^n) \rangle = \mathbb{E}\langle f(\mathbf{Y}, \mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^n) \rangle \quad (18)$$

where the bracket notation $\langle \cdot \rangle$ is used for the joint expectation over the posterior samples $(\mathbf{x}^j)_{j \leq n}$, \mathbb{E} is over the signal \mathbf{X}^* and data \mathbf{Y} .

The reduction of the replica saddle point equations thanks to this identity is done in Section 3.3. As a consequence only three scalar quantities will remain after reduction, one denoted m and called “magnetization” quantifying the overlap between the minimum mean-square error (MMSE) estimator and the signal.

- From the solution of this variational problem we deduce our second main result, namely, an asymptotically exact expression for the minimum mean-square error of inference of the hidden spike with factorized prior:

Result 2 (Minimum mean-square error). The minimum mean-square error verifies

$$\lim_{N \rightarrow \infty} \frac{1}{2N^2} \mathbb{E} \|\mathbf{X}^* \mathbf{X}^{*\top} - \mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} | \mathbf{Y}]\|_{\text{F}}^2 = \frac{1}{2}(1 - m^2)$$

where m is one component of the solution $\boldsymbol{\tau}_*$ to the variational problem for the free entropy, studied in Section 3.3.

The main technical and conceptual novelties which lead to these formulas are:

- To the best of our knowledge, we provide the first adaptation of the replica method to the analysis of the fundamental limits of inference in a model with a *noise* having strongly dependent random entries (instead of a measurement operator, or matrix of covariates, in a regression setting). See Section 3.

- If the structure of the noise (i.e., its statistical properties) is encoded by a polynomial potential V of order $K + 1$, then this induces in the posterior distribution k -wise interactions between the signal’s estimator entries, for all $k \leq K + 1$. Said differently, the underlying factor graph is an hypergraph with hyperedges of degrees $K + 1, K, \dots, 1$. However, we discovered that by exploiting the low-rank structure of the signal, all these interactions can be reduced to effective pair-wise interactions. This allows to reduce the model to an Ising model more convenient for theoretical analysis (a similar reduction is useful for algorithmic approaches too, see next section). The reduction we propose is general and systematic for low-rank signals corrupted by rotational invariant noise matrices. See Section 3.1.
- Our analysis can be mainstreamed once we have identified a key integral that we refer to as the *inhomogeneous spherical integral*. This exactly solvable integral is a generalization of the standard low-rank spherical integral appearing in random matrix theory (as it is related to the R-transform) [93], in spin-glasses [89, 86, 84, 42, 20], the theory of large-deviations for matrix-valued stochastic processes [54, 55] and matrix models in high-energy physics [56, 62, 52]. Given the breadth of applications of this integral, we foresee that the generalization we propose and analyze in Section 2 may have applications well beyond the present setting, for the study of models where rotationally invariant matrices with non-independent matrices appear.
- Another important conclusion from our analysis is the fact that for signals \mathbf{X}^* whose law is rotation-invariant (such as Gaussian or uniformly spherically distributed), the simple spectral PCA procedure of [24] is Bayes-optimal:

Result 3 (Optimality of spectral PCA for rotation-invariant priors). Let \mathbf{X}^* be a standard Gaussian vector or uniformly sampled on the sphere of radius \sqrt{N} . Then its inference from \mathbf{Y} can be optimally achieved from the naive spectral algorithm that constructs an estimator $C\boldsymbol{\nu}\boldsymbol{\nu}^\top$ of \mathbf{P}^* from the eigenvector $\boldsymbol{\nu} = \boldsymbol{\nu}(\mathbf{Y})$ of \mathbf{Y} with leading eigenvalue λ_{\max} and that is then properly rescaled by a certain factor $C = C(\lambda, \rho)$, see [24].

This is verified both by the replica method and an exact computation based on Gaussian integration and a saddle point method, see Section 3.4. We remark that this statement is incorrect for other priors P_X .

Algorithmic results

On the algorithmic side our contributions are the following:

- We analytically show that the existing approximate message passing algorithms [41, 102], whose iterates are based on the data matrix \mathbf{Y} , do not saturate the Bayes-optimal performance predicted by our replica theory. See Section 4.
- We employ in Section (5) the AdaTAP formalism of Opper et al [86] to analyze the model from the algorithmic perspective. What the analysis shows is that, like in the replica method, one can reduce the model with interactions of order higher than two to a pure quadratic Ising model with an effective interaction matrix $\mathbf{J}(\mathbf{Y})$ which is a non-trivial matrix polynomial of the data \mathbf{Y} . This explains the reason why the previously proposed AMP algorithms are sub-optimal: the data \mathbf{Y} is *not* the best choice of matrix to use in the AMP iterates, despite being the most natural one. The Bayes-optimal choice is instead $\mathbf{J}(\mathbf{Y})$ obtained from our theory, which cannot be guessed a-priori. We informally state this fact as one of our main results:

Result 4 (Bayesian-optimal processing of data and optimal AMP). Consider the matrix estimation model under structured noise (2). Given the observed matrix of data \mathbf{Y} , the optimal choice of matrix to use in a Bayesian inference algorithm such as AMP is *not* \mathbf{Y} but instead a proper polynomial of it, i.e., $\mathbf{J}(\mathbf{Y}) = \sum_{k \leq K} c_k \mathbf{Y}^k$, with coefficients $(c_k)_{k \in [K]}$ depending on V . For example, when the potential V is given by (13) we show in Sections 5.1 and 5.2 that the optimal choice is

$$\mathbf{J}(\mathbf{Y}) = \mu\lambda\mathbf{Y} - \gamma\lambda^2\mathbf{Y}^2 + \gamma\lambda\mathbf{Y}^3.$$

Employing this matrix in the AMP iterates leads to a Bayesian-optimal inference algorithm whose complexity scales as the dimension N , see the result below.

- After having defined the Bayesian-optimal AMP recursion, we provide a rigorous state evolution recursion to track its asymptotic performance. We highlight that, since the data matrix \mathbf{Y} is replaced by the polynomial $\mathbf{J}(\mathbf{Y})$, we cannot apply the state evolution result of [41]. More specifically, the Onsager correction terms will have a different form than the ones of [41], and their derivation requires a novel analysis.

Result 5 (State evolution of the Bayes-optimal AMP (BAMP)). Consider the Bayesian-optimal Approximate Message Passing (BAMP) algorithm defined by the recursion

$$\mathbf{f}^t = \mathbf{J}(\mathbf{Y})\mathbf{u}^t - \sum_{i \leq t} c_{t,i} \mathbf{u}^i, \quad \mathbf{u}^{t+1} = g_{t+1}(\mathbf{f}^t), \quad t \geq 1. \quad (19)$$

When a proper choice of coefficients $\{\mathbf{c}_{t,j}\}_{j \in [t]}$ is considered, for a large family of functions $(g_t)_{t \geq 1}$ and ψ , the following holds almost surely:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i \leq N} \psi(u_i^1, \dots, u_i^{t+1}, f_i^1, \dots, f_i^t, X_i^*) = \mathbb{E} \psi(U_1, \dots, U_{t+1}, F_1, \dots, F_t, X^*).$$

Equivalently the joint empirical distribution over the N rows of the $N \times (2t+2)$ matrix $(\mathbf{u}^1, \dots, \mathbf{u}^{t+1}, \mathbf{f}^1, \dots, \mathbf{f}^t, \mathbf{X}^*)$ converges in a certain sense to the $(2t+2)$ -dimensional random vector $(U_1, \dots, U_{t+1}, F_1, \dots, F_t, X^*)$ when N increases. Here

$$U_{i+1} = g_{i+1}(F_t) \quad \text{and} \quad (F_1, \dots, F_t) = (\mu_1, \dots, \mu_t)X^* + (W_1, \dots, W_t)$$

with $(W_i)_{i \leq t}$ a multivariate Gaussian vector whose covariance as well as $(\mu_i)_{i \leq t}$ can be computed via a deterministic state evolution recursion.

The precise rigorous statement can be found in Section 6. The idea of the argument is to construct an auxiliary AMP which tracks the quantities $(\mathbf{Y}^{j-1} \mathbf{u}^t)_{t \geq 1, j \leq K-1}$. By decomposing the iterates of this auxiliary AMP into a component aligned with previous iterates, a component in the direction of the signal and independent Gaussian noise, we obtain the form of the Onsager correction and the state evolution. From this result we can rigorously predict the performance of the novel AMP algorithm we propose. The optimality of the pre-processed matrix $\mathbf{J}(\mathbf{Y})$ and associated AMP is then confirmed by the perfect matching of the fixed point of the state evolution recursion tracking the AMP mean-square error and our replica prediction for the MMSE.

Two important remarks are in order. First, we emphasize that the BAMP algorithm (19) we propose is *not* the usual AMP of [41] where the data matrix \mathbf{Y} is just replaced by the pre-processed matrix $\mathbf{J}(\mathbf{Y})$. Indeed, the correct Onsager coefficients $\{\mathbf{c}_{t,i}\}$ entering BAMP require a novel type of “multi-stage” state evolution recursion which is completely different from the one in [41], see Section 6. The novel acronym we introduce emphasizes that crucial distinction.

Secondly, it is true that our replica prediction for the MMSE is non-rigorous. However, our state evolution analysis of BAMP is fully rigorous (just like the analysis of the AMP in [41]). By comparing their asymptotic fixed point performance by state evolution in Section 7, we show that BAMP improves over the AMP in [41]. This improvement is thus a rigorous conclusion, while the conjecture is that, thanks to this improvement, BAMP saturates the Bayes-optimal performance.

Comments on the potential universality of our results

We comment the hypotheses under which our results are conjectured valid, and then extrapolate on the more general settings in which the results may still hold.

We start with a remark concerning the insensitivity of our results to the “statistical details” of the noise eigenvalues. Let us precise the hypotheses on the distribution of the noise, in particular on its eigenvalues, under which our results are conjectured valid. As seen from (4) the eigenvalues of the noise are strongly dependent due to the Vandermonde determinant. However, we conjecture that all our results still hold if one considers instead a simpler ensemble where the N eigenvalues are drawn i.i.d. from $\rho(x)dx$, see (14). The reason is that all the analysis and results rely only on the weak convergence of the empirical density of eigenvalues of the ensemble under consideration towards ρ . Hence, as long as this is the case, our results must hold, even if we do not rigorously prove it. To formally show it, from now on we consider that the diagonal matrix \mathbf{D} of eigenvalues of the noise is *deterministic* with the sole constraint that the empirical density of its diagonal entries converges towards $\rho(x)dx$. This of course includes as special cases the two aforementioned settings (i.i.d. and coupled by Vandermonde determinant). We therefore work in this paper under the following hypothesis.

Hypothesis 1 (Distribution of the noise). The noise $\mathbb{R}^{N \times N} \ni \mathbf{Z} = \mathbf{O}^\top \mathbf{D} \mathbf{O}$ in model (2) is a symmetric rotationally invariant matrix, namely, it is equal in law to $\mathbf{U}^\top \mathbf{Z} \mathbf{U}$ for any orthogonal matrix $\mathbf{U} \in \mathbb{O}(N)$ (the group of $N \times N$ orthogonal matrices). Equivalently, \mathbf{O} is drawn from the Haar (uniform) measure over $\mathbb{O}(N)$. Moreover, we only require for its (possibly deterministic) eigenvalues $(D_i)_{i \leq N}$ that their empirical law $N^{-1} \sum_{i \leq N} \delta_{D_i, x}$ is tending weakly as $N \rightarrow \infty$ to a probability measure with support bounded uniformly in N and with density ρ with respect to the Lebesgue measure. As mentioned earlier, for the purpose of having a uniform measure of SNR when tuning $(\mu, \gamma(\mu))$ we will consider cases where $\int d\rho(x)x^2 = 1$ despite this is not necessary for the analysis to hold.

A second remark concerns the rotational invariance of the noise. We believe that our results may extend beyond this hypothesis to cases where the noise eigenbasis may be invariant under more restrictive transformations (such as permutation invariant), or even “almost deterministic”. This intuition comes from a very recent line of work concerning linear regression and phase retrieval with structured matrices of covariates. Indeed, the authors of [37, 38, 39] show that in this context, the class of rotationally invariant matrices leads to the same performance as a much broader class of almost deterministic matrices (with the same spectral density), also when AMP or its linearized version are used as inference algorithm. This is a different setting from the one we consider, since in our setup the structured matrix is the noise, but it nevertheless suggests that our predictions should remain true more generically. The confirmation of this universality is left for future work.

What is conjectured exact, and what is rigorous

We end this section with a remark concerning the level of rigor of our derivations. Most of our results are based on non-rigorous but well established methods from the statistical mechanics of mean-field disordered systems, in particular the replica method at the replica symmetric level, and the theory of Anderson-Thouless-Palmer equations. For a general background on these techniques we refer to [75, 85, 78]. It is important to keep in mind that despite being non-rigorous, the results obtained from these techniques are conjectured *exact* in the present setting of *Bayesian-optimal inference* (or equivalently, statistical mechanical models living on their *Nishimori line* [82]), in the asymptotic large size limit $N \rightarrow \infty$.

This widely admitted asymptotic exactness, first proved for the Sherrington-Kirkpatrick model [50, 99, 88], spreads in numerous fields and in particular in the analysis of high-dimensional inference. In this context a plethora of rigorous results confirm the validity of replica predictions [10, 40, 16, 18, 9, 2, 19]. In particular, replica symmetric formulas for the free entropy, mutual information and minimum mean-square error have been systematically proved thanks to a combination of concentration techniques specifically adapted to the context of inference [9, 18] together with rigorous versions of the cavity method [100, 88, 33], (adaptive) interpolation techniques [51, 15, 16] or Hamilton-Jacobi approaches [80, 31, 32]. From this fastly growing literature, we conjecture that it is only a matter of time before our replica-based predictions are proven.

Concerning our algorithmic results on the novel approximate message passing we propose (BAMP), the results are completely rigorous; full proofs are provided as appendix. They are based on the theory of message passing algorithms and associated state evolution recursions [22], in particular the most recent results for structured matrices as considered here [41, 102].

2 The inhomogeneous spherical integral

In this section we derive the expression of a useful general integral that will play a crucial role along the whole analysis, and that we believe may have an interest on its own. For the reader interested in the information-theoretic and algorithmic analyses directly, this section can be skipped at first reading as only its main results (22), (23) and (33) will be used in the rest.

Indices $\ell, \ell' \leq n$ will always indicate the “replica dimension” (with n which always remains finite), while $i, j, k \leq N$ index the “spin dimension” (where N will diverge).

2.1 Definition and variational characterization

Let $\mathbf{O} \sim \text{Haar}(\mathbb{O}(N))$ be drawn from the Haar measure over the orthogonal group of $N \times N$ matrices. Consider a fixed matrix $\mathbf{x} \in \mathbb{R}^{N \times n}$ with rows $\mathbf{x}_i \in \mathbb{R}^n$, $i \leq N$, and columns $\mathbf{x}_\ell \in \mathbb{R}^N$, $\ell \leq n$. Assume it has the column-wise overlap structure

$$\mathbf{x}_\ell^\top \mathbf{x}_{\ell'} = N q_{\ell\ell'}, \quad \ell, \ell' \leq n. \quad (20)$$

We let $\mathbf{q} = (q_{\ell\ell'})_{\ell, \ell' \leq n} := N^{-1} \mathbf{x}^\top \mathbf{x}$. Every vector is considered a column vector, so, e.g., $(\mathbf{O}\mathbf{x})_i$ is a n -dimensional column-vector corresponding to the transpose of the i th row of the $N \times n$ matrix $\mathbf{O}\mathbf{x}$, while $(\mathbf{O}\mathbf{x})_i^\top$ is a row-vector.

Let the matrices $\mathbf{C}_{\ell\ell'} = \text{diag}((C_{i,\ell\ell'})_{i \leq N})$, $\mathbf{C}_i = (C_{i,\ell\ell'})_{\ell, \ell' \leq n}$, and the “external fields” $\mathbf{h}_\ell = (h_{i,\ell})_{i \leq N}$, $\mathbf{h}_i = (h_{i,\ell})_{\ell \leq n}$ all having entries bounded uniformly in N . The sequence $(\mathbf{h}_i \in \mathbb{R}^n, \mathbf{C}_i \in \mathbb{R}^{n \times n})_{i \leq N}$ is assumed to have an empirical law tending to that of the random $(\mathbf{h} \in \mathbb{R}^n, \mathbf{C} \in \mathbb{R}^{n \times n})$: for any continuous bounded function $f : \mathbb{R}^{n \times n} \times \mathbb{R}^n \mapsto \mathbb{R}^k$ with k independent of N ,

$$\frac{1}{N} \sum_{i \leq N} f(\mathbf{C}_i, \mathbf{h}_i) \xrightarrow{N \rightarrow \infty} \mathbb{E} f(\mathbf{C}, \mathbf{h}).$$

We denote by $\mathbb{R} \ni I_N = I_N(\mathbf{q}, (\mathbf{C}_{\ell\ell'})_{\ell, \ell' \leq n}, (\mathbf{h}_\ell)_{\ell \leq n}) = I_N(\mathbf{q}, (\mathbf{C}_i, \mathbf{h}_i)_{i \leq N})$ the generalized low-rank spherical integral, which is defined as

$$\begin{aligned} I_N &:= \frac{1}{N} \ln \mathbb{E}_{\mathbf{O}} \exp \sum_{i \leq N} ((\mathbf{O}\mathbf{x})_i^\top \mathbf{C}_i (\mathbf{O}\mathbf{x})_i + (\mathbf{O}\mathbf{x})_i^\top \mathbf{h}_i) \\ &= \frac{1}{N} \ln \mathbb{E}_{\mathbf{O}} \exp \left(\sum_{\ell, \ell' \leq n} (\mathbf{O}\mathbf{x}_\ell)^\top \mathbf{C}_{\ell\ell'} \mathbf{O}\mathbf{x}_{\ell'} + \sum_{\ell \leq n} (\mathbf{O}\mathbf{x}_\ell)^\top \mathbf{h}_\ell \right) \\ &= \frac{1}{N} \ln \mathbb{E}_{\mathbf{O}} \exp \left(\sum_{i, j, k \leq N} \sum_{\ell, \ell' \leq n} O_{ij} O_{ik} x_{j,\ell'} x_{k,\ell} C_{i,\ell\ell'} + \sum_{i, j \leq N} \sum_{\ell \leq n} O_{ij} x_{j,\ell} h_{i,\ell} \right). \end{aligned} \quad (21)$$

Calling the columns $(\mathbf{x}_\ell)_{\ell \leq n}$ “replicas”, the matrices $(\mathbf{C}_i)_{i \leq N}$, $(\mathbf{C}_{\ell\ell'})_{\ell, \ell' \leq n}$ are coupling them (after the replicas have been jointly rotated by the random \mathbf{O}). Therefore we call them “replica coupling matrices”.

As $N \rightarrow \infty$ with n fixed this integral is given by

$$I_N \xrightarrow{N \rightarrow \infty} I_{\mathbf{C}, \mathbf{h}}(\mathbf{q}), \quad (22)$$

with variational formula

$$\begin{aligned} I_{\mathbf{C}, \mathbf{h}}(\mathbf{q}) &:= \frac{1}{2} \text{extr}_{\tilde{\mathbf{q}}} (\text{Tr} \mathbf{q} \tilde{\mathbf{q}} + \mathbb{E} \mathbf{h}^\top (\tilde{\mathbf{q}} - 2\mathbf{C})^{-1} \mathbf{h} - \mathbb{E} \ln \det(\tilde{\mathbf{q}} - 2\mathbf{C})) \\ &\quad - \frac{1}{2} (n + \ln \det \mathbf{q}). \end{aligned} \quad (23)$$

The extremum is over symmetric matrices such $\tilde{\mathbf{q}} - 2\mathbf{C}$ is positive definite for all \mathbf{C} living on its domain.

We remark that it may be the case that the extremum over $\tilde{\mathbf{q}}$ is actually attained on the boundary of the optimization domain, in which case the optimization requires more care than what is done in (35) to solve it (as (35) assumes the extremum to lie inside the optimization domain). This is however not expected in the settings of the present paper. When this phenomenon happens, in the standard low-rank spherical integral this leads to a “sticking phenomenon” where the solution of the optimization is dependent on the maximum eigenvalue of the full-rank random matrix entering the integral’s definition, see [54].

2.2 Special cases

2.2.1 Low-rank HCIZ integral

The special case $\mathbf{h}_i = \mathbf{0}$ and replica coupling matrices $\mathbf{C}_i = \mathbf{C}D_i$ for $i \leq N$ corresponds to the standard rank- n spherical (or HCIZ) integral:

$$\begin{aligned} I_N &= \frac{1}{N} \ln \mathbb{E}_{\mathbf{O}} \exp \sum_{i,j,k \leq N} \sum_{\ell, \ell' \leq n} O_{ij} O_{ik} x_{j,\ell'} x_{k,\ell} C_{\ell\ell'} D_i \\ &= \frac{1}{N} \ln \mathbb{E}_{\mathbf{O}} \exp \sum_{\ell, \ell' \leq n} C_{\ell\ell'} (\mathbf{O}\mathbf{x}_\ell)^\top \mathbf{D} \mathbf{O}\mathbf{x}_{\ell'} \\ &= \frac{1}{N} \ln \mathbb{E}_{\mathbf{O}} \exp \text{Tr } \mathbf{O}^\top \mathbf{D} \mathbf{O} (\mathbf{x}\mathbf{C}\mathbf{x}^\top), \end{aligned}$$

where $\mathbf{D} = \text{diag}((D_i)_{i \leq N})$ and $\mathbf{x}\mathbf{C}\mathbf{x}^\top$ is an arbitrary rank- n symmetric matrix (arbitrary given that \mathbf{x} and \mathbf{C} are so). Its asymptotic expression can also be obtained from the results of [54] after diagonalizing $\mathbf{x}\mathbf{C}\mathbf{x}^\top$ and depends only on the limit of the empirical distribution of (D_i) and on the n non-zero eigenvalues of $\mathbf{x}\mathbf{C}\mathbf{x}^\top$.

2.2.2 Low-rank spherical integral with external field and diagonal replica coupling

Taking diagonal replica coupling matrices $\mathbf{C}_i = I_n D_i / 2$ gives (a generalization of) the spherical integral with external field found in [Prop. 2.7, [42]]:

$$I_N = \frac{1}{N} \ln \mathbb{E}_{\mathbf{O}} \exp \sum_{\ell \leq n} \left(\frac{1}{2} (\mathbf{O}\mathbf{x}_\ell)^\top \mathbf{D} \mathbf{O}\mathbf{x}_\ell + (\mathbf{O}\mathbf{x}_\ell)^\top \mathbf{h}_\ell \right). \quad (24)$$

2.2.3 Low-rank spherical integral with non-diagonal replica coupling and replica symmetric overlap

Let the $N \times N$ diagonal matrices $\mathbf{A} = \text{diag}((A_i)_{i \leq N})$ and similarly for \mathbf{B} . The empirical law of $(A_i, B_i)_{i \leq N}$ tends to that of (A, B) . Of particular interest to us corresponds to taking $\ell \in \{0, \dots, n\}$, $\mathbf{h}_i = \mathbf{0}$ and replica coupling matrices with only non-zero entries being

$$(\mathbf{C}_i)_{\ell 0} = (\mathbf{C}_i)_{0\ell} = \frac{A_i}{2} \text{ for } 1 \leq \ell \leq n, \quad (\mathbf{C}_i)_{\ell\ell} = \frac{B_i}{2}(1 - \delta_{\ell,0}), \quad (25)$$

or equivalently,

$$\mathbf{C}_{0\ell} = \mathbf{C}_{\ell 0} = \frac{\mathbf{A}}{2} \text{ and } \mathbf{C}_{\ell\ell} = \frac{\mathbf{B}}{2} \text{ for } 1 \leq \ell \leq n, \quad \mathbf{C}_{\ell\ell'} = \mathbf{0} \text{ else.} \quad (26)$$

Note that this is *not* a special case of the standard rank- n spherical integral of the first example: here \mathbf{C}_i cannot be written as \mathbf{C} times a function of i ; instead different entries of \mathbf{C}_i vary with i differently. In this case the generalized spherical integral reads (the sum over ℓ below starts at $\ell = 1$)

$$I_N = \frac{1}{N} \ln \mathbb{E}_{\mathbf{O}} \exp \sum_{\ell \leq n} \left((\mathbf{O}\mathbf{x}_0)^\top \mathbf{A} \mathbf{O}\mathbf{x}_\ell + \frac{1}{2} (\mathbf{O}\mathbf{x}_\ell)^\top \mathbf{B} \mathbf{O}\mathbf{x}_\ell \right). \quad (27)$$

So the 0th replica plays here a special role (it corresponds to the planted signal).

We consider a “replica symmetric structure” for the overlap matrix parametrized by the vector $(v_0, v, m, q) \in \mathbb{R}^4$:

$$\mathbf{q} = \begin{pmatrix} v_0 & m & m & m & \dots & m \\ m & v & q & q & \dots & q \\ m & q & v & q & \dots & q \\ m & q & q & v & \dots & q \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ m & q & q & q & \dots & v \end{pmatrix} \in \mathbb{R}^{(n+1) \times (n+1)}, \quad (28)$$

and, coherently, we assume that the extremum over $\tilde{\mathbf{q}}$ is attained for a matrix having the same structure with different constants $(\tilde{v}_0, \tilde{v}, \tilde{m}, \tilde{q})$. Its determinant can be easily computed via Gauss’ reduction:

$$\ln \det \mathbf{q} = \ln v_0 + n \ln(v - q) + \ln \left(1 + n \frac{v_0 q - m^2}{v_0(v - q)} \right). \quad (29)$$

We also need to compute

$$\text{Tr} \mathbf{q} \tilde{\mathbf{q}} = v_0 \tilde{v}_0 + n(2m\tilde{m} + v\tilde{v} + (n-1)q\tilde{q}). \quad (30)$$

Letting \mathbf{C} be defined as (25) but with the random variables A, B replacing A_i, B_i , the last missing term is obtained similarly as (29): under the replica symmetric structure for $\tilde{\mathbf{q}}$,

$$\mathbb{E} \ln \det(\tilde{\mathbf{q}} - 2\mathbf{C}) = \ln \tilde{v}_0 + n \ln(\tilde{v} - B - \tilde{q}) + \mathbb{E} \ln \left(1 + n \frac{\tilde{v}_0 \tilde{q} - (\tilde{m} - A)^2}{\tilde{v}_0(\tilde{v} - B - \tilde{q})} \right).$$

Combining everything in the variational formula (23), and taking into account that \mathbf{q} here is a $(n+1) \times (n+1)$ matrix, we obtain the following expression for the generalized spherical integral with replica coupling (25), and under a replica symmetric structure for the overlap and conjugate matrices (thus the upperscript):

$$\begin{aligned} I_N \rightarrow I_{A,B}^{\text{RS}}(\mathbf{q}) &:= \frac{1}{2} \text{extr}_{(\tilde{v}_0, \tilde{v}, \tilde{m}, \tilde{q})} \left\{ v_0 \tilde{v}_0 - \ln \tilde{v}_0 + n(2m\tilde{m} + v\tilde{v} + (n-1)q\tilde{q}) \right. \\ &\quad \left. - n \mathbb{E} \ln(\tilde{v} - B - \tilde{q}) - \mathbb{E} \ln \left(1 + n \frac{\tilde{v}_0 \tilde{q} - (\tilde{m} - A)^2}{\tilde{v}_0(\tilde{v} - B - \tilde{q})} \right) \right\} \\ &\quad - \frac{1 + \ln v_0}{2} - \frac{n}{2} (1 + \ln(v - q)) - \frac{1}{2} \ln \left(1 + n \frac{v_0 q - m^2}{v_0(v - q)} \right). \end{aligned} \quad (31)$$

By definition (27) of I_N this formula has to cancel when $n = 0$. Thus

$$\text{extr}_{\tilde{v}_0} \{ v_0 \tilde{v}_0 - \ln \tilde{v}_0 \} - 1 - \ln v_0 = 0. \quad (32)$$

The saddle point equation over \tilde{v}_0 then yields $\tilde{v}_0 = 1/v_0$, in which case this latter formula indeed cancels. So the simplified formula reads

$$\begin{aligned} I_{A,B}^{\text{RS}}(\mathbf{q}) &= \frac{1}{2} \text{extr}_{(\tilde{v}, \tilde{m}, \tilde{q})} \left\{ n(2m\tilde{m} + v\tilde{v} + (n-1)q\tilde{q}) - n \mathbb{E} \ln(\tilde{v} - B - \tilde{q}) \right. \\ &\quad \left. - \mathbb{E} \ln \left(1 + n \frac{\tilde{q} - v_0(\tilde{m} - A)^2}{\tilde{v} - B - \tilde{q}} \right) \right\} - \frac{n}{2} (1 + \ln(v - q)) \\ &\quad - \frac{1}{2} \ln \left(1 + n \frac{v_0 q - m^2}{v_0(v - q)} \right). \end{aligned} \quad (33)$$

2.3 Derivation of the variational formula

Let $\tilde{\mathbf{x}}_\ell := \mathbf{O} \mathbf{x}_\ell$ the columns of $\tilde{\mathbf{x}} = \mathbf{O} \mathbf{x}$. Under the law of \mathbf{O} at fixed \mathbf{x} , these random vectors are uniform among all vectors having the overlap structure of (\mathbf{x}_ℓ) . Thus

their law conditional on \mathbf{x} is just a function of the symmetric overlap $\mathbf{q} = (q_{\ell\ell'})$:

$$P(\tilde{\mathbf{x}} | \mathbf{x}) = P(\tilde{\mathbf{x}} | \mathbf{q}) = \frac{1}{\mathcal{Z}(\mathbf{q})} \prod_{\ell \geq \ell'}^{1,n} \delta(Nq_{\ell\ell'} - \tilde{\mathbf{x}}_\ell^\top \tilde{\mathbf{x}}_{\ell'}) = \frac{1}{\mathcal{Z}(\mathbf{q})} \delta(N\mathbf{q} - \tilde{\mathbf{x}}^\top \tilde{\mathbf{x}})$$

with normalization

$$\mathcal{Z}(\mathbf{q}) = \int d\tilde{\mathbf{x}} \delta(N\mathbf{q} - \tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}).$$

Using the Fourier representation of the Delta function, the integral to compute reads (below $\tilde{\mathbf{q}}$ is a $n \times n$ symmetric matrix with complex entries)

$$\begin{aligned} \exp(NI_N) &= \frac{1}{\mathcal{Z}(\mathbf{q})} \int d\tilde{\mathbf{x}} \delta(N\mathbf{q} - \tilde{\mathbf{x}}^\top \tilde{\mathbf{x}}) \exp \sum_{i \leq N} (\tilde{\mathbf{x}}_i^\top \mathbf{C}_i \tilde{\mathbf{x}}_i + \tilde{\mathbf{x}}_i^\top \mathbf{h}_i) \\ &= \frac{1}{\mathcal{Z}(\mathbf{q})} \int d\tilde{\mathbf{x}} d\tilde{\mathbf{q}} \exp \left(\frac{N}{2} \text{Tr} \mathbf{q} \tilde{\mathbf{q}} - \frac{1}{2} \text{Tr} \tilde{\mathbf{x}}^\top \tilde{\mathbf{x}} \tilde{\mathbf{q}} + \sum_{i \leq N} (\tilde{\mathbf{x}}_i^\top \mathbf{C}_i \tilde{\mathbf{x}}_i + \tilde{\mathbf{x}}_i^\top \mathbf{h}_i) \right) \\ &= \frac{1}{\mathcal{Z}(\mathbf{q})} \int d\tilde{\mathbf{q}} \exp \left(\frac{N}{2} \text{Tr} \mathbf{q} \tilde{\mathbf{q}} \right) \prod_{i \leq N} \int d\tilde{\mathbf{x}}_i \exp \left(-\frac{1}{2} \tilde{\mathbf{x}}_i^\top (\tilde{\mathbf{q}} - 2\mathbf{C}_i) \tilde{\mathbf{x}}_i + \tilde{\mathbf{x}}_i^\top \mathbf{h}_i \right). \end{aligned}$$

We will soon evaluate the $\tilde{\mathbf{q}}$ -integral by saddle-point approximation. We now assume that the dominating saddle-point belongs to a set

$$D_\epsilon := \{ \tilde{\mathbf{q}} \in \mathbb{R}^{n \times n} : \tilde{\mathbf{q}} - 2\mathbf{C} \succ \epsilon I_n \text{ for all } \mathbf{C} \text{ living on its domain} \},$$

for some arbitrarily small $\epsilon > 0$ but independent of N . Thus restricting the integral to this domain yields a sub-leading correction $\exp o(N)$. For $\tilde{\mathbf{q}} \in D_\epsilon$ a Gaussian integration over $\tilde{\mathbf{x}}$ is possible: $\exp(NI_N)$ equals

$$\begin{aligned} &\frac{(2\pi)^{Nn/2} e^{o(N)}}{\mathcal{Z}(\mathbf{q})} \int_{D_\epsilon} d\tilde{\mathbf{q}} \exp \frac{N}{2} \frac{1}{N} \sum_{i \leq N} (\text{Tr} \mathbf{q} \tilde{\mathbf{q}} + \mathbf{h}_i^\top (\tilde{\mathbf{q}} - 2\mathbf{C}_i)^{-1} \mathbf{h}_i - \ln \det(\tilde{\mathbf{q}} - 2\mathbf{C}_i)) \\ &= \frac{(2\pi)^{Nn/2} e^{o(N)}}{\mathcal{Z}(\mathbf{q})} \int_{D_\epsilon} d\tilde{\mathbf{q}} \exp \left\{ \frac{N}{2} \mathbb{E}(\text{Tr} \mathbf{q} \tilde{\mathbf{q}} + \mathbf{h}^\top (\tilde{\mathbf{q}} - 2\mathbf{C})^{-1} \mathbf{h} - \ln \det(\tilde{\mathbf{q}} - 2\mathbf{C})) \right\}. \end{aligned}$$

We used the convergence of the empirical law of the sequence $(\mathbf{C}_i, \mathbf{h}_i)_i$ to turn the above empirical mean into a statistical expectation over (\mathbf{C}, \mathbf{h}) , including the correction in the $\exp o(N)$; this is possible because over D_ϵ the summand is a bounded

continuous function of $(\mathbf{C}_i, \mathbf{h}_i)$. As N diverges at fixed n we can estimate the integral by saddle-point and reach that the generalized spherical integral is

$$I_N \rightarrow \frac{1}{2} \text{extr}_{\tilde{\mathbf{q}}} (\text{Tr} \mathbf{q} \tilde{\mathbf{q}} + \mathbb{E} \mathbf{h}^\top (\tilde{\mathbf{q}} - 2\mathbf{C})^{-1} \mathbf{h} - \mathbb{E} \ln \det(\tilde{\mathbf{q}} - 2\mathbf{C})) - \frac{1}{2} \text{extr}_{\tilde{\mathbf{q}}} (\text{Tr} \mathbf{q} \tilde{\mathbf{q}} - \ln \det \tilde{\mathbf{q}}) \quad (34)$$

where the term $-\ln \mathcal{Z}(\mathbf{q})/N$ from the normalization has been obtained by simply setting \mathbf{C} and \mathbf{h} to all-zeros in the first optimization problem. The extremum is over $n \times n$ symmetric matrices $\tilde{\mathbf{q}}$ such $\tilde{\mathbf{q}} - 2\mathbf{C}$ is positive definite for all \mathbf{C} on its domain.

Assuming that the extremum is attained inside the optimization domain we can perform the extremization using $\ln \det \mathbf{A} = \text{Tr} \ln \mathbf{A}$. The extremum is solution of the matrix equation

$$\mathbf{q} = \mathbb{E} \mathbf{h}^\top (\tilde{\mathbf{q}} - 2\mathbf{C})^{-2} \mathbf{h} + \mathbb{E} (\tilde{\mathbf{q}} - 2\mathbf{C})^{-1}. \quad (35)$$

The second extremization leads instead to $\tilde{\mathbf{q}} = \mathbf{q}^{-1}$. Thus the result.

3 Information-theoretic analysis by the replica method

Let us start with a remark. Express the noise $\mathbf{Z} = \mathbf{O}^\top \mathbf{D} \mathbf{O}$ in terms of its random Haar distributed basis \mathbf{O} and eigenvalues \mathbf{D} , so that the observation model becomes

$$\mathbf{Y} = \frac{\lambda}{N} \mathbf{P}^* + \mathbf{O}^\top \mathbf{D} \mathbf{O}. \quad (36)$$

When the signal is rotationally invariant we can consider the noise diagonal right away by absorbing \mathbf{O} into \mathbf{x}, \mathbf{X}^* . If the law P_X is uniform on the sphere, then this joint rotation does not change the distribution of \mathbf{x}, \mathbf{X}^* which greatly simplifies the analysis. In this simpler case, the replica method is not needed as the computation of the free entropy can be carried out simply using a saddle point method. We provide this analysis in Section 3.4. The rotational invariance of the Gaussian law implies that also that case could be treated similarly by direct computation. On the contrary, for other priors than spherical or Gaussian this is no longer possible and the replica method is needed.

In order to deal with such non-rotational invariant priors we are going to adapt an approach developed by Kabashima in [60, 98] to study certain inference models where rotational invariant random matrices appear as quenched disorder. The main

difference compared to the works is the fact that because they consider (generalized) linear regression, the structured matrix plays the role of covariates/data and therefore does not influence the form of the likelihood when writing the posterior. A novelty of the present setting is the fact that because the structured matrix is now the noise itself, the likelihood is a function of its statistics which in turn complicates the analysis.

The goal here is to compute the log-partition function (12) using the replica trick

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \ln \mathcal{Z} = \lim_{N \rightarrow \infty} \frac{1}{N} \lim_{n \rightarrow 0} \partial_n \ln \mathbb{E} \mathcal{Z}^n = \lim_{n \rightarrow 0} \partial_n \lim_{N \rightarrow \infty} \frac{1}{N} \ln \mathbb{E} \mathcal{Z}^n. \quad (37)$$

The expectation is with respect to \mathbf{Y} or equivalently the independent \mathbf{O}, \mathbf{x}_0 (recall \mathbf{D} is deterministic). The last equality *assumes* the commutation of the two limits. Another key assumption of the method is that we are going to make the computation considering $n \in \mathbb{N}$ and then assume an analytic continuation to n in a small neighborhood of 0. Before doing all that we are going to first re-express our model in a form more convenient for analysis.

3.1 An equivalent quadratic model

The Hamiltonian (11) of the model can be written in a more convenient way by introducing the following shorthand notations for order parameters. Despite at the moment only vector \mathbf{x} has been introduced, soon a family of vectors (\mathbf{x}_ℓ) will be introduced when “replicating” the system. So we directly introduce the order parameters for these:

$$v_\ell = v(\mathbf{x}_\ell) := \frac{1}{N} \|\mathbf{x}_\ell\|^2, \quad (38)$$

$$M_{(k)\ell} = M_{(k)}(\mathbf{x}_\ell, \mathbf{Z}) := \frac{1}{N} \mathbf{x}_\ell^\top \mathbf{Z}^k \mathbf{x}_\ell, \quad (39)$$

$$\kappa_\ell = \kappa(\mathbf{x}_\ell, \mathbf{x}_0, \mathbf{Z}) := \frac{1}{N} \mathbf{x}_\ell^\top \mathbf{Z} \mathbf{x}_0, \quad (40)$$

$$q_{\ell\ell'} = q(\mathbf{x}_\ell, \mathbf{x}_{\ell'}) := \frac{1}{N} \mathbf{x}_\ell^\top \mathbf{x}_{\ell'}, \quad (41)$$

where the replica indices $0 \leq \ell, \ell' \leq n$ with the identification $\mathbf{x}_0 := \mathbf{X}^*$.

We now treat the quadratic and quartic part of the matrix potential separately. Let us denote

$$\Delta := \frac{1}{N} (\mathbf{P}^* - \mathbf{P}), \quad M_\ell := M_{(1)\ell}.$$

The quadratic part yields a contribution:

$$\begin{aligned} \frac{N}{4} \text{Tr}[(\mathbf{Z} + \lambda\Delta)^2 - \mathbf{Z}^2] &= \frac{1}{2} \left[\lambda(\mathbf{x}_0^\top \mathbf{Z} \mathbf{x}_0 - \mathbf{x}^\top \mathbf{Z} \mathbf{x}) + N\lambda^2 \left(\frac{1}{2N^2} (\|\mathbf{x}_0\|^4 + \|\mathbf{x}\|^4) - q_{01}^2 \right) \right] \\ &= -\frac{N\lambda}{2} M_1 + \frac{N\lambda^2}{2} \left(\frac{1}{2} (v_0^2 + v_1^2) - q_{01}^2 \right) + o(N). \end{aligned} \quad (42)$$

The subscript 1 indicates that only one replica $\mathbf{x}_1 := \mathbf{x}$ is involved yet, and by convention it is replica number one. We used that by the law of large numbers, and thanks to the symmetry of the chosen matrix potential, we can assert that

$$M_{(2k+1)0} = o_N(1), \quad M_{(2)0} = 1 + o_N(1)$$

due to our choice of normalization, so in particular $M_0 = M_{(1)0} = o_N(1)$. Again by the law of large numbers we have

$$v_0 = \mathbb{E}(X_1^*)^2 + o_N(1) = 1 + o_N(1).$$

The quartic contribution is more complicated due to the non-commutativity of matrices:

$$\begin{aligned} &\frac{N}{8} \text{Tr}[(\mathbf{Z} + \lambda\Delta)^4 - \mathbf{Z}^4] \\ &= \frac{N}{8} \text{Tr}[\lambda^4 \Delta^4 + 4\lambda^3 \mathbf{Z} \Delta^3 + 4\lambda^2 \mathbf{Z}^2 \Delta^2 + 4\lambda \mathbf{Z}^3 \Delta + 2\lambda^2 \mathbf{Z} \Delta \mathbf{Z} \Delta] \\ &= \frac{N}{8} \left[\lambda^4 (2q_{01}^4 + v_1^4 + 1 - 4q_{01}^2 (v_1^2 + 1 - v_1)) \right. \\ &\quad + 4\lambda^3 (M_0(1 - q_{01}^2) - M_1(v_1^2 - q_{01}^2) + 2q_{01}(v_1 - 1)\kappa_1) \\ &\quad + 4\lambda^2 \left(M_{(2)0} + v_1 \frac{1}{N} \mathbf{x}^\top \mathbf{Z}^2 \mathbf{x} - 2q_{01} \frac{1}{N} \mathbf{x}^\top \mathbf{Z}^2 \mathbf{x}_0 \right) + 4\lambda \left(M_{(3)0} - \frac{1}{N} \mathbf{x}^\top \mathbf{Z}^3 \mathbf{x} \right) \\ &\quad \left. + 2\lambda^2 (M_0^2 + M_1^2 - 2\kappa_1^2) \right]. \end{aligned} \quad (43)$$

Note that the only three terms which we did not write in a compact form using order parameters are linear and quadratic forms in \mathbf{x} that do *not* appear elsewhere to a power greater than 1. This is because introducing order parameters for these would add useless redundancy in the final equations (but it is necessary for the other order parameters due to powers of them appearing in the Hamiltonian). Let

$$\begin{aligned} f_\ell = f(q_{0\ell}, v_\ell, M_\ell, \kappa_\ell) &:= \gamma \frac{\lambda^4}{8} \left(2q_{0\ell}^4 + v_\ell^4 - 4q_{0\ell}^2 (v_\ell^2 + 1 - v_\ell) \right) - \gamma \frac{\lambda^3}{2} M_\ell (v_\ell^2 - q_{0\ell}^2) \\ &\quad + \gamma \lambda^3 q_{0\ell} (v_\ell - 1) \kappa_\ell + \gamma \frac{\lambda^2}{4} M_\ell^2 - \gamma \frac{\lambda^2}{2} \kappa_\ell^2 + \mu \frac{\lambda^2}{2} \left(\frac{1}{2} v_\ell^2 - q_{0\ell}^2 \right) - \mu \frac{\lambda}{2} M_\ell. \end{aligned} \quad (44)$$

Plugging the contributions we computed into (11) shows that the Hamiltonian is equivalently written as

$$H_N(\mathbf{x}; \mathbf{Z}, \mathbf{x}_0) = Nf_1 + \gamma \frac{\lambda}{2} \mathbf{x}^\top (\lambda v_1 \mathbf{Z}^2 - \mathbf{Z}^3) \mathbf{x} - \gamma q_{01} \lambda^2 \mathbf{x}^\top \mathbf{Z}^2 \mathbf{x}_0 + C + o(N), \quad (45)$$

where we have put all irrelevant constants inside C . We will neglect the $o(N)$ contribution in the following as it yields a subleading correction to the free entropy. Also the constant C is irrelevant, so we simply forget about it. Keep in mind that at the moment f_1 is still a function of \mathbf{x} . This model is thus *not* (yet) quadratic in \mathbf{x} due to terms such as $M_1(\mathbf{x}, \mathbf{Z})^2$ appearing in f_1 .

We now use delta functions to fix various order parameters. We are going to use repeatedly the Fourier representation of the delta function, namely

$$\delta(x) = \frac{1}{2\pi} \int d\hat{x} \exp(i\hat{x}x). \quad (46)$$

Because the integrals we will end-up with will always be at some point evaluated by saddle point, implying a deformation of the integration contour in the complex plane, tracking the imaginary unit i in the delta functions will be irrelevant. Similarly, the normalization $1/(2\pi)$ will always contribute to sub-exponential corrections in the integrals at hand. Therefore, we will allow ourselves to formally write

$$\delta(x) = \int d\hat{x} \exp(r\hat{x}x) \quad (47)$$

for a convenient constant r , keeping in mind these considerations (again, as we evaluate the final integrals by saddle point, the choice of r ends-up being irrelevant).

We denote jointly $\boldsymbol{\tau} := (v_1, M_1, \kappa_1, q_{01})$ and $\hat{\boldsymbol{\tau}}$ for their Fourier conjugates. Coming back to the the partition function for this equivalent model (45), it can be re-expressed using delta functions as

$$\begin{aligned} & \int dP_X(\mathbf{x}) d\boldsymbol{\tau} \exp(-H_N(\mathbf{x}; \mathbf{Z}, \mathbf{x}_0)) \\ & \quad \times \delta(Nq_{01} - \mathbf{x}^\top \mathbf{x}_0) \delta(Nv_1 - \|\mathbf{x}\|^2) \delta(NM_1 - \mathbf{x}^\top \mathbf{Z} \mathbf{x}) \delta(N\kappa_1 - \mathbf{x}^\top \mathbf{Z} \mathbf{x}_0) \\ & = \int dP_X(\mathbf{x}) d\boldsymbol{\tau} d\hat{\boldsymbol{\tau}} \exp(-H_N(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{x}; \mathbf{x}_0, \mathbf{Z})), \end{aligned} \quad (48)$$

where

$$H_N(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{x}; \mathbf{x}_0, \mathbf{Z}) := Nh(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}) + \mathbf{x}^\top \mathbf{J}_1(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{Z}) \mathbf{x} + \mathbf{x}^\top \mathbf{J}_0(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{Z}) \mathbf{x}_0 \quad (49)$$

and

$$h(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}) := f_1 - \hat{q}_{01}q_{01} - \frac{\hat{v}_1v_1}{2} - \frac{\hat{M}_1M_1}{2} - \hat{\kappa}_1\kappa_1, \quad (50)$$

$$\mathbf{J}_1(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{Z}) := \frac{\hat{v}_1}{2}I_N + \frac{\hat{M}_1}{2}\mathbf{Z} + \gamma\frac{\lambda^2}{2}v_1\mathbf{Z}^2 - \gamma\frac{\lambda}{2}\mathbf{Z}^3, \quad (51)$$

$$\mathbf{J}_0(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{Z}) := \hat{q}_{01}I_N + \hat{\kappa}_1\mathbf{Z} - \gamma q_{01}\lambda^2\mathbf{Z}^2. \quad (52)$$

So what this shows is that by introducing new variables (order parameters and conjugate Fourier parameters), the original model turns out being equivalent to an extended system with Hamiltonian (49). The key point of all this analysis is that by introducing the new variables $\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}$ we have turned the interactions between the $(x_i)_{i \leq N}$ into purely quadratic ones. This form is now more appropriate to be solved using (generalizations of) known techniques. We emphasize that despite the algebraic manipulations leading from (11) to (49) are cumbersome, given a more complicated polynomial potential V the very same strategy could be applied but would require the introduction of more order parameters. Yet, the equivalent model would still collapse into a quadratic one of the above form but with a more complicated function h and matrices $\mathbf{J}_1, \mathbf{J}_0$ (still being polynomials of the noise \mathbf{Z} of order one less than the order of V). The reason is that the key mechanisms behind these simplifications when expanding the original Hamiltonian (11) are stemming from the low-rank structure of the spike.

3.2 Replica symmetric free entropy using the inhomogeneous spherical integral

Having reduced the model to a quadratic one, we are now ready to replicate the system to compute the free entropy. The partition function \mathcal{Z} is now computed using the equivalent model (49). The expected replicated partition function is

$$\mathbb{E}\mathcal{Z}(\mathbf{x}_0, \mathbf{Z})^n = \int \prod_{\ell=0}^n dP_X(\mathbf{x}_\ell) \prod_{\ell \leq n} d\boldsymbol{\tau}_\ell d\hat{\boldsymbol{\tau}}_\ell \mathbb{E}_{\mathbf{Z}} \exp\left(-\sum_{\ell \leq n} H_N(\boldsymbol{\tau}_\ell, \hat{\boldsymbol{\tau}}_\ell, \mathbf{x}_\ell; \mathbf{x}_0, \mathbf{Z})\right), \quad (53)$$

with replicas $(\mathbf{x}_\ell, \boldsymbol{\tau}_\ell, \hat{\boldsymbol{\tau}}_\ell)_{\ell \leq n}$ and shared quenched disorder \mathbf{x}_0, \mathbf{Z} . What we do next is to replace \mathbf{Z} by $\mathbf{O}^\top \mathbf{D} \mathbf{O}$ and fix the overlap structure between replicas

$$\mathbf{x}_\ell^\top \mathbf{x}_{\ell'} = Nq_{\ell\ell'}, \quad \ell, \ell' \leq n \quad (54)$$

by introducing further variables and their Fourier conjugates (this is already taken care of for the overlaps $\mathbf{x}_\ell^\top \mathbf{x}_0$ with the planted signal). The purpose will become clear

soon. Redefining $\boldsymbol{\tau}_\ell := (v_\ell, M_\ell, \kappa_\ell)$ and similarly for $\hat{\boldsymbol{\tau}}_\ell$, and defining the overlaps $\mathbf{q} = (q_{\ell\ell'})_{0 \leq \ell < \ell' \leq n}$ and similarly for $\hat{\mathbf{q}}$, the log-partition function can be recast as

$$\begin{aligned} \mathbb{E}Z^n &= \int d\mathbf{q} d\hat{\mathbf{q}} \prod_{\ell \leq n} d\boldsymbol{\tau}_\ell d\hat{\boldsymbol{\tau}}_\ell \exp N \left(\sum_{\ell \leq n} \left(\frac{\hat{v}_\ell v_\ell}{2} + \frac{\hat{M}_\ell M_\ell}{2} + \hat{\kappa}_\ell \kappa_\ell - f_\ell \right) + \sum_{0 \leq \ell < \ell' \leq n} \hat{q}_{\ell\ell'} q_{\ell\ell'} \right) \\ &\quad \times \int \prod_{\ell=0}^n dP_X(\mathbf{x}_\ell) \exp \left(- \sum_{0 \leq \ell < \ell' \leq n} \hat{q}_{\ell\ell'} \mathbf{x}_\ell^\top \mathbf{x}_{\ell'} - \frac{1}{2} \sum_{\ell \leq n} \hat{v}_\ell \|\mathbf{x}_\ell\|^2 \right) \\ &\quad \times \mathbb{E}_{\mathbf{O}} \exp \sum_{\ell \leq n} \left((\mathbf{O}\mathbf{x}_0)^\top \mathbf{A}_\ell \mathbf{O}\mathbf{x}_\ell + \frac{1}{2} (\mathbf{O}\mathbf{x}_\ell)^\top \mathbf{B}_\ell \mathbf{O}\mathbf{x}_\ell \right) \end{aligned} \quad (55)$$

where the $N \times N$ ‘‘replica coupling matrices’’ are

$$\mathbf{A}_\ell := -\hat{\kappa}_\ell \mathbf{D} + \gamma q_{0\ell} \lambda^2 \mathbf{D}^2, \quad (56)$$

$$\mathbf{B}_\ell := -\hat{M}_\ell \mathbf{D} - \gamma \lambda^2 v_\ell \mathbf{D}^2 + \gamma \lambda \mathbf{D}^3. \quad (57)$$

We now assume a replica-symmetric ansatz which should lead to the correct solution due to the strong concentration-of-measure effects taking place in the Bayes-optimal setting as well as the Nishimori identities [82, 18]. It means that we assume that the saddle point over the order parameters dominating the partition function as $N \rightarrow \infty$, which are finitely many, lies in the subset verifying the following (note the minus sign introduced for $-\hat{q}$ and $-\hat{m}$ for convenience): for all $\ell \neq \ell' = 1, \dots, n$

$$\text{Replica Symmetry Ansatz: } \begin{cases} M_\ell = M, & \hat{M}_\ell = \hat{M}, \\ \kappa_\ell = \kappa, & \hat{\kappa}_\ell = \hat{\kappa}, \\ v_\ell = v, & \hat{v}_\ell = \hat{v}, \\ q_{\ell\ell'} = q, & \hat{q}_{\ell\ell'} = -\hat{q}, \\ q_{0\ell} = m, & \hat{q}_{0\ell} = -\hat{m}. \end{cases} \quad (58)$$

Using this ansatz, the matrices $(\mathbf{A}_\ell, \mathbf{B}_\ell)_{\ell \leq n}$ become independent of ℓ . We thus call their common value \mathbf{A}, \mathbf{B} . As a consequence the term $\mathbb{E}_{\mathbf{O}}(\cdot)$ at the third line in (55) is recognized to be what we call an *inhomogeneous spherical integral* defined and analyzed in a devoted Section 2.2.3. From Section 2 we know that the result of such integral depends only on the overlap structure; this is the reason why we fixed it earlier. We will thus replace it by $\exp NI_{A,B}^{\text{RS}}(n, v, m, q)$ whose formula is (33) and which is parametrized by the random variables (below $D \sim \rho$)

$$A = -\hat{\kappa}D + \gamma m \lambda^2 D^2, \quad (59)$$

$$B = -\hat{M}D - \gamma \lambda^2 v D^2 + \gamma \lambda D^3. \quad (60)$$

Notice that at this point the only \mathbf{x} -integrals remaining (second line of (55)) are completely factorized over the spin indices i . Hence after taking the saddle point the log-replicated free entropy becomes in the limit $N \rightarrow \infty$

$$\begin{aligned} \frac{1}{N} \ln \mathbb{E} \mathcal{Z}^n &\rightarrow \text{extr} \left\{ n \left(\frac{\hat{v}v}{2} + \frac{\hat{M}M}{2} + \hat{\kappa}\kappa - \hat{m}m + \frac{1-n}{2} \hat{q}q - f(m, v, M, \kappa) \right) \right. \\ &\left. + I_{A,B}^{\text{RS}}(n, v, m, q) + \ln \int \prod_{\ell=0}^n dP_X(x_\ell) e^{\hat{q} \sum_{\ell < \ell' \leq n} x_\ell x_{\ell'} + \hat{m} \sum_{\ell \leq n} x_0 x_\ell - \frac{\hat{v}}{2} \sum_{\ell \leq n} x_\ell^2} \right\} \end{aligned}$$

where the extremum is over all scalars in (58). The last line can be treated by a Hubbard-Stratonovič transform (i.e., Gaussian integral formula) to decouple the integral over the replica indices. Doing so it becomes

$$\mathbb{E} \left(\int dP_X(x) \exp \left(\sqrt{\hat{q}} Z x - \frac{\hat{q} + \hat{v}}{2} x^2 + \hat{m} X_0 x \right) \right)^n,$$

with $Z \sim \mathcal{N}(0, 1)$, $X_0 \sim P_X$.

We now consider the limit of number of replicas going to 0 assuming the analytic continuation of our formulas from integer n to real. To expand the latter term we use $\ln \mathbb{E} X^n = n \mathbb{E} \ln X + O(n^2)$. The inhomogeneous spherical integral given by (33) (with $v_0 = 1$) also has to be expanded in n . We get

$$\begin{aligned} I_{A,B}^{\text{RS}}(n, v, m, q) &= \frac{n}{2} \text{extr}_{(\tilde{v}, \tilde{m}, \tilde{q})} \left\{ 2m\tilde{m} + v\tilde{v} - q\tilde{q} - \mathbb{E} \ln(\tilde{v} - B - \tilde{q}) \right. \\ &\left. - \mathbb{E} \frac{\tilde{q} - (\tilde{m} - A)^2}{\tilde{v} - B - \tilde{q}} \right\} - \frac{n}{2} (1 + \ln(v - q)) - \frac{nq - m^2}{2(v - q)} + O(n^2) \end{aligned}$$

with an expectation over $D \sim \rho$ entering A, B . Now we plug the previous expressions in the log-replicated partition function and expand up to $O(n)$ the resulting expression:

$$\begin{aligned} \frac{1}{N} \ln \mathbb{E} \mathcal{Z}^n &\rightarrow \text{extr} \left\{ n \left(\frac{\hat{v}v}{2} + \frac{\hat{M}M}{2} + \hat{\kappa}\kappa - \hat{m}m + \frac{1-n}{2} \hat{q}q - f(m, v, M, \kappa) \right) \right. \\ &\left. + I_{A,B}^{\text{RS}}(n, v, m, q) + n \mathbb{E} \ln \int dP_X(x) \exp \left(\sqrt{\hat{q}} Z x - \frac{\hat{q} + \hat{v}}{2} x^2 + \hat{m} X_0 x \right) \right\} + O(n^2). \end{aligned}$$

One can check that as it should $\lim_{N \rightarrow \infty} N^{-1} \ln \mathbb{E} \mathcal{Z}^n$ vanishes when $n \rightarrow 0$. Taking the n -derivative (recall (37)) and then sending $n \rightarrow 0$ the final formula for the free entropy is obtained (and recalling that we dropped irrelevant constants along the

computation):

$$\begin{aligned}
\frac{1}{N} \mathbb{E} \ln \mathcal{Z} \rightarrow \text{extr} & \left\{ \frac{\hat{v}v}{2} + \frac{\hat{M}M}{2} + \hat{\kappa}\kappa - \hat{m}m + \frac{\hat{q}q}{2} + m\tilde{m} + \frac{v\tilde{v}}{2} - \frac{q\tilde{q}}{2} \right. \\
& - \gamma \frac{\lambda^4}{8} \left(2m^4 + v^4 - 4m^2(v^2 + 1 - v) \right) + \gamma \frac{\lambda^3}{2} M(v^2 - m^2) \\
& - \gamma \lambda^3 m(v-1)\kappa - \gamma \frac{\lambda^2}{4} M^2 + \gamma \frac{\lambda^2}{2} \kappa^2 - \mu \frac{\lambda^2}{2} \left(\frac{1}{2}v^2 - m^2 \right) + \mu \frac{\lambda}{2} M \\
& + \mathbb{E} \ln \int dP_X(x) \exp \left(\sqrt{\hat{q}}Zx - \frac{\hat{q} + \hat{v}}{2}x^2 + \hat{m}X_0x \right) \\
& - \frac{1}{2} \mathbb{E} \ln (\tilde{v} - \tilde{q} + \hat{M}D + \gamma\lambda^2vD^2 - \gamma\lambda D^3) - \frac{1}{2} \ln(v - q) - \frac{q - m^2}{2(v - q)} \\
& \left. + \frac{1}{2} \mathbb{E} \frac{(\tilde{m} + \hat{\kappa}D - \gamma m\lambda^2 D^2)^2 - \tilde{q}}{\tilde{v} - \tilde{q} + \hat{M}D + \gamma\lambda^2vD^2 - \gamma\lambda D^3} \right\} + \text{constant}. \tag{61}
\end{aligned}$$

The extremization is intended over the set of 13 variational parameters $v, \hat{v}, \tilde{v}, m, \hat{m}, \tilde{m}, q, \hat{q}, \tilde{q}, M, \hat{M}, \kappa, \hat{\kappa}$. However, as we shall see later the saddle point equations will reduce only to two, because thanks to the Nishimori identities the saddle point values of many order parameters can be found right away. This is a specific and rather convenient feature of the Bayesian-optimal setting.

3.3 Replica saddle point equations

Define the following random local measure

$$\langle \cdot \rangle_{\hat{m}, \hat{q}, \hat{v}} = \frac{\int dP_X(x) e^{\sqrt{\hat{q}}Zx + \hat{m}xX_0 - \frac{\hat{q} + \hat{v}}{2}x^2}(\cdot)}{\int dP_X(x) e^{\sqrt{\hat{q}}Zx + \hat{m}xX_0 - \frac{\hat{q} + \hat{v}}{2}x^2}}, \tag{62}$$

the randomness being $Z \sim \mathcal{N}(0, 1)$ and $X_0 \sim P_X$, and the random functions (random in $D \sim \rho$)

$$H = (\tilde{v} - \tilde{q} + \hat{M}D + \gamma\lambda^2vD^2 - \gamma\lambda D^3)^{-1}, \tag{63}$$

$$Q = \gamma m\lambda^2 D^2 - \hat{\kappa}D - \tilde{m}. \tag{64}$$

Below follow the saddle point equations obtained by equating to 0 the gradient w.r.t. the variational parameters of the variational free entropy in (61). The parameter

associated to each equation are reported in the round parenthesis:

$$\begin{aligned}
(m) \quad & \mu\lambda^2 m + \gamma\lambda^4 m(v^2 + 1 - v - m^2) - \gamma\lambda^3 Mm - \hat{m} + \tilde{m} + \frac{m}{v - q} \\
& - \gamma\lambda^3(v - 1)\kappa + \gamma\lambda^2 \mathbb{E}QH D^2 = 0 \\
(\hat{m}) \quad & m = \mathbb{E}X_0 \langle X \rangle_{\hat{m}, \hat{q}, \hat{v}} \\
(\tilde{m}) \quad & m = \mathbb{E}QH \\
(q) \quad & \hat{q} - \tilde{q} = \frac{q - m^2}{(v - q)^2} \\
(\hat{q}) \quad & q = \mathbb{E} \langle X \rangle_{\hat{m}, \hat{q}, \hat{v}}^2 \\
(\tilde{q}) \quad & q = \mathbb{E}(Q^2 - \tilde{q})H^2 \\
(v) \quad & -\mu\lambda^2 v - \gamma\lambda^4(v^3 - m^2(2v - 1)) + 2\gamma\lambda^3 Mv + \hat{v} + \tilde{v} - \frac{1}{v - q} - \frac{m^2 - q}{(v - q)^2} \\
& - \gamma\lambda^2 \mathbb{E}H D^2 - 2\gamma m \lambda^3 \kappa + \gamma\lambda^2 \mathbb{E}D^2(\tilde{q} - Q^2)H^2 = 0 \\
(\hat{v}) \quad & v = \mathbb{E} \langle X^2 \rangle_{\hat{m}, \hat{q}, \hat{v}} \\
(\tilde{v}) \quad & v = \mathbb{E}[H + H^2(Q^2 - \tilde{q})] \\
(M) \quad & \mu\lambda + \gamma\lambda^3(v^2 - m^2) - \gamma\lambda^2 M + \hat{M} = 0 \\
(\hat{M}) \quad & M = \mathbb{E}D[H + H^2(Q^2 - \tilde{q})] \\
(\kappa) \quad & \hat{\kappa} = \gamma\lambda^3 m(v - 1) - \gamma\lambda^2 \kappa \\
(\hat{\kappa}) \quad & \kappa = \mathbb{E}DQH
\end{aligned}$$

As in any replica symmetric mean-field theory, the physical meaning of some order parameters makes it possible to fix their values to their expectation, obtainable using the Nishimori identities and, as a consequence, to drastically reduce this 13-dimensional system. To begin with, recall that we fixed v to be the squared norm of a sample from the posterior re-scaled by the number of components. Assuming concentration effects take place as they should in this optimal setting, and denoting the posterior mean by $\langle \cdot \rangle$, using the Nishimori identity we have that

$$v = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \langle \|\mathbf{x}\|^2 \rangle = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \|\mathbf{X}^*\|^2 = 1. \quad (65)$$

We have $\hat{v} = 0$ because the constraint $v = 1$ is enforced by the prior without the need of a delta constraint. The (κ) -equation can then be used to directly eliminate $\hat{\kappa}$ by inserting $\hat{\kappa} = -\gamma\lambda^2 \kappa$ into Q . The Nishimori identity also imposes

$$m = \mathbb{E}X_0 \langle X \rangle_{\hat{m}, \hat{q}, 0} = q = \mathbb{E} \langle X \rangle_{\hat{m}, \hat{q}, 0}^2. \quad (66)$$

It is not difficult to realize that for this to be true one also needs necessarily $\hat{m} = \hat{q}$. So we have 8 variables left. The most tricky parameter is M , that we introduced to decouple the four body interactions in the Hamiltonian. Notice first that (recall definitions (39) and (41))

$$\begin{aligned} \frac{1}{N}\mathbb{E}\langle \mathbf{x}^\top \mathbf{Z} \mathbf{x} \rangle &= \frac{1}{N}\mathbb{E}\langle \mathbf{x}^\top \left(\mathbf{Y} - \frac{\lambda}{N} \mathbf{P}^* \right) \mathbf{x} \rangle \\ &= \frac{1}{N}\mathbb{E} \mathbf{X}^{*\top} \mathbf{Y} \mathbf{X}^* - \lambda \mathbb{E} \left\langle \left(\frac{1}{N} \mathbf{x}^\top \mathbf{X}^* \right)^2 \right\rangle \\ &= \lambda \left(1 - \mathbb{E} \left\langle \left(\frac{1}{N} \mathbf{x}^\top \mathbf{X}^* \right)^2 \right\rangle \right) + O\left(\frac{1}{N}\right). \end{aligned}$$

We used that by the Nishimori identity

$$\frac{\mathbb{E}\langle \mathbf{x}^\top \mathbf{Y} \mathbf{x} \rangle}{N} = \frac{\mathbb{E} \mathbf{X}^{*\top} \mathbf{Y} \mathbf{X}^*}{N} = \frac{1}{N} \mathbb{E} \mathbf{X}^{*\top} \left(\frac{\lambda}{N} \mathbf{X}^* \mathbf{X}^{*\top} + \mathbf{Z} \right) \mathbf{X}^* = (\mathbb{E}(X_1^*)^2)^2 \lambda = \lambda. \quad (67)$$

Indeed, by diagonalizing the noise,

$$\mathbb{E} \mathbf{X}^{*\top} \mathbf{Z} \mathbf{X}^* = \mathbb{E} \sum_{i \leq N} s_i^2 D_i = \mathbb{E} \|\mathbf{X}^*\|^2 \mathbb{E} D_1 = 0,$$

where \mathbf{s} is a uniform spherical vector of same norm as \mathbf{X}^* , and $\mathbb{E} D_1 = 0$ by symmetry. By concentration happening on the Nishimori line [18] we have

$$\mathbb{E} \left\langle \left(\frac{1}{N} \mathbf{x}^\top \mathbf{X}^* \right)^2 \right\rangle = \left(\mathbb{E} \left\langle \frac{1}{N} \mathbf{x}^\top \mathbf{X}^* \right\rangle \right)^2 + o_N(1) = m^2 + o_N(1).$$

Hence

$$M = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \langle \mathbf{x}^\top \mathbf{Z} \mathbf{x} \rangle = \lambda(1 - m^2). \quad (68)$$

The (M)-equation together with the other identities implies $\hat{M} = -\mu\lambda$. To summarize the Nishimori identities and concentration properties enforce five constraints:

$$v = 1, \quad \hat{v} = 0, \quad m = q, \quad \hat{m} = \hat{q}, \quad M = \lambda(1 - m^2) \quad (69)$$

and we have 6 variables left. Our updated definitions of Q and H are

$$Q = \gamma m \lambda^2 D^2 + \gamma \lambda^2 \kappa D - \tilde{m}, \quad (70)$$

$$H = (\tilde{v} - \tilde{q} - \mu \lambda D + \gamma \lambda^2 D^2 - \gamma \lambda D^3)^{-1}. \quad (71)$$

Using the Nishimori identities we see from the (\tilde{v}) and (\tilde{q}) -equations that

$$m = \mathbb{E}H^2(Q^2 - \tilde{q}) \quad \Rightarrow \quad \mathbb{E}H = 1 - m. \quad (72)$$

The latter has to be interpreted as an equation for the quantity $\tilde{V} := \tilde{v} - \tilde{q}$ as a function of m . Furthermore, one can now express \tilde{m} as a function of κ and m . In fact from equation (\tilde{m}) , unfolding Q and then solving for \tilde{m} , one gets

$$\tilde{m} = \frac{\gamma\lambda^2}{1-m} \mathbb{E}D(mD + \kappa)H - \frac{m}{1-m}. \quad (73)$$

Plugging this back into the (m) -equation we get \hat{m} , equation (79). We stress that inside H there is still an m dependency through \tilde{V} .

With all these simplifications we can close the equations on (m, κ) only:

$$(\hat{m}) \quad m = \mathbb{E}X_0 \langle X \rangle_{\hat{m}, \hat{m}, 0} \quad (74)$$

$$(\tilde{\kappa}) \quad \kappa = \mathbb{E}DQH, \quad (75)$$

where the random variables $Q = Q(m, \kappa, D)$ and $H = H(m, D)$ are

$$Q = \gamma m \lambda^2 D^2 + \gamma \lambda^2 \kappa D - \frac{\gamma \lambda^2}{1-m} \mathbb{E}D(mD + \kappa)H + \frac{m}{1-m}, \quad (76)$$

$$H = (\tilde{V} - \mu \lambda D + \gamma \lambda^2 D^2 - \gamma \lambda D^3)^{-1}, \quad (77)$$

with $\tilde{V} = \tilde{V}(m)$ and $\hat{m} = \hat{m}(m, \kappa)$ being determined respectively by

$$\mathbb{E}H = 1 - m, \quad (78)$$

$$\hat{m} = \gamma \lambda^2 \mathbb{E}HD \left(\frac{mD + \kappa}{1-m} + DQ \right) + \mu \lambda^2 m. \quad (79)$$

Then the replica prediction for the MMSE is

$$\lim_{N \rightarrow \infty} \frac{1}{2N^2} \mathbb{E} \|\mathbf{X}^* \mathbf{X}^{*\top} - \mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} \mid \mathbf{Y}]\|_{\text{F}}^2 = \frac{1}{2}(1 - m^2). \quad (80)$$

From (79) it is evident that when $\gamma = 0$ and $\mu = 1$ (to preserve unit variance of the noise), κ and \hat{m} decouple, $\hat{m} = \lambda^2 m$, and the equation (74) reduces to the standard replica saddle point equation for the Wigner spike model.

There would be also an equation for \tilde{q} , that is decoupled though, meaning that \tilde{q} is a simple function of m and κ in the end:

$$(q) \quad \tilde{q} = \hat{m}(m, \kappa) - \frac{m}{1-m}. \quad (81)$$

3.4 Spectral PCA is optimal for rotation-invariant signals

In this section we show that spectral PCA [24] is optimal for inferring \mathbf{X}^* such that \mathbf{X}^* equals in law $\mathbf{O}\mathbf{X}^*$ for any orthogonal matrix \mathbf{O} . This is the case for Gaussian and spherically uniformly distributed \mathbf{X}^* .

To do so, we first show that the previous computations can be straightforwardly modified to accommodate the case of spherical prior. Let us assume that the signal \mathbf{X}^* is uniformly distributed on a sphere of radius \sqrt{N} . We denote the uniform measure on this sphere by ω . Thanks to the invariance property of the measure on the sphere under rotations we know that \mathbf{x} equals in law $\mathbf{O}\mathbf{x}$ for $\mathbf{x} \sim \omega$ and any orthogonal matrix \mathbf{O} . Therefore, we can directly diagonalize the noise without loss of generality and work with the equivalent model

$$\mathbf{Y} = \frac{\lambda}{N}\mathbf{P}^* + \mathbf{D}. \quad (82)$$

In this way we can get rid of \mathbf{O} and as a consequence replicating the system and the inhomogeneous spherical integral becomes useless. Only Gaussian integrations and a saddle point estimation are needed.

The partition function is (49)–(52) but with the diagonal matrix \mathbf{D} replacing \mathbf{Z} (the constraint $\|\mathbf{x}\|^2 = N$ is taken care of by the Hamiltonian):

$$\int d\mathbf{x}d\boldsymbol{\tau}d\hat{\boldsymbol{\tau}} \exp \left(-Nh(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}) - \mathbf{x}^\top \mathbf{J}_1(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{D})\mathbf{x} - \mathbf{x}^\top \mathbf{J}_0(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{D})\mathbf{x}_0 \right). \quad (83)$$

Because now \mathbf{J}_1 and \mathbf{J}_0 are diagonal matrices, the \mathbf{x} -integral in the partition function is just a Gaussian integral: it is (up to an irrelevant multiplicative constant)

$$\int d\boldsymbol{\tau}d\hat{\boldsymbol{\tau}} \exp N \left(-h(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}) - \frac{1}{2N} \sum_{i \leq N} \ln J_{1,i} + \frac{1}{4N} \sum_{i \leq N} x_{0,i}^2 \frac{J_{0,i}^2}{J_{1,i}} \right) \quad (84)$$

with $v_1 = 1$ (appearing in h). Because \mathbf{x}_0 is a uniform spherical vector combined with the convergence of the empirical law of (D_i) we have

$$-\frac{1}{2N} \sum_{i \leq N} \ln J_{1,i} + \frac{1}{4N} \sum_{i \leq N} x_{0,i}^2 \frac{J_{0,i}^2}{J_{1,i}} = -\frac{1}{2} \mathbb{E} \ln J_{1,1} + \frac{1}{4} \mathbb{E} \frac{J_{0,1}^2}{J_{1,1}} + o_N(1).$$

Thus saddle point estimation of (83) yields

$$\begin{aligned} \frac{1}{N} \ln \mathcal{Z} \rightarrow \text{const} + \text{extr} \left\{ -f(m, 1, M, \kappa) + \hat{m}m + \frac{\hat{v}v}{2} + \frac{\hat{M}M}{2} + \hat{\kappa}\kappa \right. \\ \left. - \frac{1}{2} \mathbb{E} \ln (\hat{v} + \hat{M}D + \gamma\lambda^2 D^2 - \gamma\lambda D^3) + \frac{1}{2} \mathbb{E} \frac{(\hat{m} + \hat{\kappa}D - \gamma\lambda^2 m D^2)^2}{\hat{v} + \hat{M}D + \gamma\lambda^2 D^2 - \gamma\lambda D^3} \right\}, \quad (85) \end{aligned}$$

where recall that f is defined by (44). Note that this strategy does not require the replica method, and it could also be applied in the case of Gaussian prior $P_X = \mathcal{N}(0, 1)$, due to its rotational invariance.

At this point, the saddle point equations can be written and simplified similarly as in the previous section. After doing so and from the numerical solution of the saddle point equations, one can deduce that: (i) in the case of spherical and Gaussian priors the MMSE is the same; and (ii) this MMSE matches the performance of the spectral PCA algorithm studied in [24]. Additionally, (iii) the MMSE obtained from this exact approach matches the replica prediction of the previous section in the case of Gaussian prior (a special case of factorized P_X tackled by our replica theory). This further confirms the validity and consistency of our methodology. Therefore we conclude that spectral PCA is Bayes-optimal in the special case of rotationally invariant priors and noise.

Let us provide a further argument in support of Bayes-optimality of PCA in the present setting. In this argument we consider the noise eigenvalues as quenched random variables, and we are going to average over them. We first notice that the MMSE estimator is diagonal in the basis of the matrix of data \mathbf{Y} . Indeed, letting \mathbf{Y} be diagonalized as $\mathbf{Y} = \mathbf{U}^\top \mathbf{S} \mathbf{U}$ then using the posterior (6),

$$\begin{aligned} \mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} | \mathbf{Y}] &= \frac{C_V}{P_Y(\mathbf{Y})} \int dP_X(\mathbf{x}) \exp\left(-\frac{N}{2} \text{Tr}V\left(\mathbf{S} - \frac{\lambda}{N}(\mathbf{U}\mathbf{x})(\mathbf{U}\mathbf{x})^\top\right)\right) \mathbf{x}\mathbf{x}^\top \\ &= \frac{C_V}{P_Y(\mathbf{Y})} \mathbf{U}^\top \left(\int dP_X(\mathbf{x}) \exp\left(-\frac{N}{2} \text{Tr}V\left(\mathbf{S} - \frac{\lambda}{N}\mathbf{x}\mathbf{x}^\top\right)\right) \mathbf{x}\mathbf{x}^\top \right) \mathbf{U} \end{aligned} \quad (86)$$

where we changed $\mathbf{U}\mathbf{x}$ to \mathbf{x} , which leaves the prior invariant by rotational invariance. We would then like to see that the matrix

$$\mathbf{L} = \frac{C_V}{P_Y(\mathbf{Y})} \int dP_X(\mathbf{x}) \exp\left(-\frac{N}{2} \text{Tr}V\left(\mathbf{S} - \frac{\lambda}{N}\mathbf{x}\mathbf{x}^\top\right)\right) \mathbf{x}\mathbf{x}^\top$$

is a diagonal. Indeed, because $\mathbf{S} = \text{diag}(s_1, \dots, s_N)$ is diagonal, $\text{Tr}V(\mathbf{S} - (\lambda/N)\mathbf{x}\mathbf{x}^\top)$ can be easily seen (see, e.g., the steps leading to (152)) to be a polynomial of degree k of the k variables

$$\left(\sum_{i \leq N} x_i^2, \sum_{i \leq N} s_i x_i^2, \dots, \sum_{i \leq N} s_i^{k-1} x_i^2 \right).$$

Then, for every $1 \leq j \leq N$, the integrand that defines \mathbf{L} takes the same value for \mathbf{x} and the point \mathbf{x}' which results from changing the sign of the j -th coordinate of \mathbf{x} . We thus have that \mathbf{L} is a diagonal matrix.

For $1 \leq k \leq N$, let \mathbf{u}_k be the eigenvector of the k -largest eigenvalue of \mathbf{Y} . Then we can express $\mathbf{L}(\mathbf{Y})$ as $\text{diag}(\gamma_1(\mathbf{Y}), \dots, \gamma_N(\mathbf{Y}))$, where by definition we have that

$$\mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} | \mathbf{Y}] = \sum_{k \leq N} \gamma_k \mathbf{u}_k \mathbf{u}_k^\top, \quad (87)$$

i.e., $\gamma_k = \mathbf{u}_k^\top \mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} | \mathbf{Y}] \mathbf{u}_k$ with the ordering $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_N$. This therefore means that the ‘‘matrix magnetization’’ may be written according to

$$\frac{1}{N^2} \mathbb{E} \text{Tr}(\mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} | \mathbf{Y}] \mathbf{X}^* \mathbf{X}^{*\top}) = \frac{1}{N^2} \sum_{k \leq N} \mathbb{E}[(\mathbf{u}_k^\top \mathbf{X}^*)^2 \gamma_k].$$

We would like now to compute the asymptotic magnetization of the Bayes estimator. For this we will use Nishimori identities and a bound over the projections of \mathbf{X}^* onto the eigenvectors of \mathbf{Y} that we verify numerically. More specifically, we will assume that there is some constant $K > 0$ such that for all $k \geq 2$ it holds that

$$(\mathbf{u}_k^\top \mathbf{X}^*)^2 \leq K. \quad (88)$$

As mentioned before, inequality (88), which is an explicit rate of convergence for the limit in [24, Theorem 2], has been verified through many numerical experiments for different noise potentials and SNRs. In every case, a bound of this type is observed, although for experiments close to the corresponding phase transition, the constant K takes larger values and the quantity bounded exhibits a larger variance (this type of behavior is expected to hold very close to the transition point).

Now, notice that by Nishimori identities the following holds

$$\mathbb{E} \gamma_k = \mathbb{E}(\mathbf{u}_k^\top \mathbf{X}^*)^2. \quad (89)$$

Also, by [24, Theorem 2] we have that (below R is the R-transform associated with the noise spectral density ρ)

$$\frac{1}{N^2} \mathbb{E}[\gamma_1 (\mathbf{u}_1^\top \mathbf{X}^*)^2] = \frac{1}{N} \left(1 - \frac{R'(1/\lambda)}{\lambda^2} \right) \mathbb{E} \gamma_1 + \frac{1}{N} \mathbb{E} \left[\gamma_1 \left(\frac{(\mathbf{u}_1^\top \mathbf{X}^*)^2}{N} - 1 + \frac{R'(1/\lambda)}{\lambda^2} \right) \right],$$

where the second term on the r.h.s. is a vanishing function of N . If we use (89) and [24, Theorem 2] a second time, we get that

$$\lim_{N \rightarrow \infty} \frac{1}{N^2} \mathbb{E}[\gamma_1 (\mathbf{u}_1^\top \mathbf{X}^*)^2] = \left(1 - \frac{R'(1/\lambda)}{\lambda^2} \right)^2.$$

On the other hand, by inequality (88) and the Nishimori identities (89) we get

$$\frac{1}{N^2} \sum_{2 \leq k \leq N} \mathbb{E}[\gamma_k (\mathbf{u}_k^\top \mathbf{X}^*)^2] \leq \frac{K}{N^2} \sum_{2 \leq k \leq N} \mathbb{E} \gamma_k = \frac{K}{N^2} \sum_{2 \leq k \leq N} \mathbb{E}(\mathbf{u}_k^\top \mathbf{X}^*)^2.$$

that by [24, Theorem 2], vanishes in the limit. We then conclude that

$$\frac{1}{N^2} \mathbb{E} \text{Tr}(\mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} \mid \mathbf{Y}] \mathbf{X}^* \mathbf{X}^{*\top}) = \left(1 - \frac{R'(1/\lambda)}{\lambda^2}\right)^2 + o_N(1).$$

This in turn implies that

$$\lim_{N \rightarrow \infty} \frac{1}{2N^2} \mathbb{E} \|\mathbb{E}[\mathbf{X}^* \mathbf{X}^{*\top} \mid \mathbf{Y}] - \mathbf{X}^* \mathbf{X}^{*\top}\|^2 = 1 - \left(1 - \frac{R'(1/\lambda)}{\lambda^2}\right)^2,$$

which is the MSE of the optimally scaled PCA estimator [24].

4 Sub-optimality of the previously proposed AMP

Consider the following AMP iteration for $t \geq 1$:

$$\mathbf{f}^t = \mathbf{Y} \mathbf{u}^t - \sum_{i=1}^t \mathbf{b}_{t,i} \mathbf{u}^i, \quad \mathbf{u}^{t+1} = h_{t+1}(\mathbf{f}^t). \quad (90)$$

Here, $\mathbf{f}^t = (f_1^t, \dots, f_N^t)$, $\mathbf{u}^{t+1} = (u_1^{t+1}, \dots, u_N^{t+1}) \in \mathbb{R}^N$ and the *denoiser* function $h_{t+1} : \mathbb{R} \rightarrow \mathbb{R}$ is continuously differentiable, Lipschitz and applied component-wise, namely $h_{t+1}(\mathbf{f}^t) = (h_{t+1}(f_1^t), \dots, h_{t+1}(f_N^t))$. The time-dependent AMP estimate of the spike \mathbf{P}^* is $(\mathbf{u}^t)^\top \mathbf{u}^t$.

The Onsager coefficients $\{\mathbf{b}_{t,i}\}_{i \in [t], t \geq 1}$ are carefully chosen so that, conditioned on the signal, the empirical distribution of the components of iterate \mathbf{f}^t is Gaussian. The form of these Onsager coefficients was derived by [84] using non-rigorous dynamic functional theory techniques, and a rigorous state evolution result was recently proved in [41]. More formally, assume that $\mathbf{X}^* \xrightarrow{W_2} X^*$. Then, the state evolution result of [41] gives that

$$(\mathbf{f}^1, \dots, \mathbf{f}^t) \xrightarrow{W_2} (F_1, \dots, F_t) := \boldsymbol{\mu}_t X^* + \mathbf{W}_t, \quad (91)$$

where $\boldsymbol{\mu}_t = (\mu_1, \dots, \mu_t)$ and $\mathbf{W}_t = (W_1, \dots, W_t)$ is a multivariate Gaussian with zero mean and covariance $\boldsymbol{\Sigma}_t = (\sigma_{ij})_{i,j \leq t}$ independent of X^* . Furthermore, the mean vectors $\{\boldsymbol{\mu}_t\}_{t \geq 1}$ and the covariance matrices $\{\boldsymbol{\Sigma}_t\}_{t \geq 1}$ are tracked by a deterministic state evolution recursion. We refer to [41] for more details on this AMP and associated state evolution. Such details won't be crucial for our argument, as we are going to focus directly on the fixed point performance, and not on the dynamics.

For this section, we restrict the analysis to (i) Rademacher prior $P_X = \frac{1}{2}(\delta_1 + \delta_{-1})$, and (ii) a ‘‘large enough’’ signal-to-noise ratio. We remark that our methodology

extends to more generic factorized priors. However, since our goal is to prove sub-optimality of AMP, this setting suffices. Moreover, we will further restrict our proof of sub-optimality to (iii) the “one-step memory” version of the AMP in [41]. This means that the denoiser h_{t+1} in (90) is allowed to depend only on the past iterate \mathbf{f}^t . A more general “multi-step memory AMP” was proposed in [102], where the denoiser h_{t+1} can depend on *all* the past iterates $\mathbf{f}^1, \dots, \mathbf{f}^t$. We remark that the analysis of [84] suggests that the fixed points of both these versions are the same; the longer memory of the latter AMP being only useful to improve its convergence properties. Therefore, despite our analysis below holds under hypotheses (i)–(iii), we conclude more generically that *the existing AMP algorithms for structured PCA in [41, 102] are sub-optimal, and this is the case for most SNR values and prior/signal’s distributions that are not rotationally invariant*². From the findings in the following sections, the reason for the sub-optimality of these AMPs will become clear. Essentially, the data \mathbf{Y} is *not* the best choice of matrix to use in the AMP iterates, despite being the most natural one.

4.1 Analysis of the one-step AMP fixed point performance

In this section we analyse the AMP algorithm (90) for structured PCA proposed in [41], with a posterior mean denoiser with a single-step memory term:

$$h_{t+1}(f_i^t) = \mathbb{E}[X | f_i^t]. \quad (92)$$

In [[41], Section 3] it is shown that the fixed point of this AMP algorithm is, for λ sufficiently large, described by the following system:

$$1 - \Delta_* = \text{mmse}\left(\frac{\lambda^2 \Delta_*^2}{\Sigma_*}\right), \quad \Sigma_* = \Delta_* R'\left(\frac{\lambda \Delta_* (1 - \Delta_*)}{\Sigma_*}\right). \quad (93)$$

Here, $R'(\cdot)$ denotes the derivative of the R transform of the (limiting) distribution of the noise eigenvalues D . For details about the R -transform, the interested reader is referred to [83]. The above is related to the asymptotic overlap of the AMP estimator through

$$\lim_{t \rightarrow \infty} \lim_{N \rightarrow \infty} \left| \frac{1}{N} \mathbf{X}^{*T} \hat{\mathbf{x}}^t \right| = \lim_{t \rightarrow \infty} \lim_{N \rightarrow \infty} \frac{1}{N} \|\hat{\mathbf{x}}^t\|^2 = \Delta_* \quad (94)$$

²We do not discard the possibility that for very peculiar choices of SNR regimes and/or priors these generically sub-optimal AMPs end-up being optimal, but that would be for highly specific setting-dependent reasons. One case where the AMPs of [41], and also the spectral PCA algorithm [24], are actually optimal is when the prior is rotationally invariant (spherical or Gaussian prior), see Section 3.4.

and thus the AMP mean-square error is

$$\lim_{t \rightarrow \infty} \lim_{N \rightarrow \infty} \frac{1}{2N^2} \mathbb{E} \|\hat{\mathbf{x}}^t (\hat{\mathbf{x}}^t)^\top - \mathbf{X}^* \mathbf{X}^{*\top}\|^2 = \frac{1}{2} (1 - \Delta_*^2). \quad (95)$$

In the case of Rademacher prior the explicit form of the posterior-mean denoiser is

$$h_{t+1}(f^t) = \tanh \left(\frac{f^t \mu_t}{\sigma_{tt}^2} \right) \quad (96)$$

where (μ_t, σ_{tt}) are the mean and variance of the (empirically) ‘‘Gaussian observation’’ f^t computed from the state evolution of [41]. The associated mmse function is (below $Z \sim \mathcal{N}(0, 1)$ is a standard Gaussian random variable and $X^* \sim \frac{1}{2}(\delta_{-1} + \delta_1)$)

$$\text{mmse}(x) = 1 - \mathbb{E} \left[X^* \frac{\int dP_X(x) x e^{Zx\sqrt{\hat{m}} + \hat{m}xX^* - \frac{\hat{m}}{2}x^2}}{\int dP_X(x) e^{Zx\sqrt{\hat{m}} + \hat{m}xX^* - \frac{\hat{m}}{2}x^2}} \right] \quad (97)$$

$$= 1 - \mathbb{E} \tanh(x + \sqrt{x}Z). \quad (98)$$

We now consider the limit $(\lambda, \Delta_*, \Sigma_*) \rightarrow (\infty, 1, 1)$ which indeed is a fixed point of (93) as we verify at the end of this section. Moreover it is unique, see [[41], Theorem 3.1]. It implies $x := \lambda^2 \Delta_*^2 / \Sigma_* \rightarrow \lambda^2 \rightarrow \infty$. We have in this limit

$$\begin{aligned} \text{mmse}(x) &= 1 - \int dt \frac{e^{-\frac{1}{2x}(t-x)^2}}{\sqrt{2\pi x}} \tanh(t) \\ &= \sqrt{\frac{\pi}{2}} \frac{e^{-\frac{x}{2}}}{\sqrt{x}} (1 + O(1/x)) \\ &= \exp \left(-\frac{x}{2} (1 + o_x(1)) \right). \end{aligned} \quad (99)$$

We plug this in the first equation of (93) which gives at leading order

$$\Delta_* = 1 - \exp \left(-\frac{\lambda^2}{2} (1 + o_\lambda(1)) \right). \quad (100)$$

It just remains to check that $(\lambda, \Delta_*, \Sigma_*) = (\infty, 1, 1)$ is indeed the unique fixed point of (93) in the large SNR regime. From our analysis we already know that this fixed point is consistent with the first equation of (93). So we simply need to verify the second one, namely,

$$R'(\lambda(1 - \Delta_*)(1 + o_\lambda(1))) \rightarrow 1 \quad (101)$$

as $\lambda \rightarrow \infty$. From (100) we have in this limit $\lambda(1 - \Delta_*) \rightarrow 0$ exponentially fast in λ , and it can be readily verified that $R'(0) = 1$, as the noise distribution D has unit second moment. This ends the argument.

4.2 Analysis of the replica Bayes-optimal fixed point

We now analyse in the same large SNR regime the replica fixed point equations that we recall below for convenience: let us rename $\tilde{V} := \tilde{v} - \tilde{q}$ as they always appear together. We consider that all quantities below are at their saddle point values maximizing the replica free entropy (61).

Let us recall the outcome of the Section 3.3 on the saddle point equations. Consider the random variables (random through their dependence in D)

$$Q = \gamma m \lambda^2 D^2 + \gamma \lambda^2 \kappa D - \frac{\gamma \lambda^2}{1-m} \mathbb{E} D(mD + \kappa) H + \frac{m}{1-m}, \quad (102)$$

$$H = (\tilde{V} - \mu \lambda D + \gamma \lambda^2 D^2 - \gamma \lambda D^3)^{-1}. \quad (103)$$

For a given value of the parameter m , the saddle point equations require $\tilde{V} = \tilde{V}(m)$ to be the solution of the implicit equation

$$\mathbb{E} H = 1 - m. \quad (104)$$

Using this implicit solution, H is a function $H(m)$ and $Q = Q(m, \kappa)$. Let $Z \sim \mathcal{N}(0, 1)$ and $X^* \sim P_X$. The saddle point equations over the order parameters (m, κ) read

$$m = 1 - \text{mmse}(\hat{m}), \quad (105)$$

$$\kappa = \mathbb{E} D Q H, \quad (106)$$

where $\text{mmse}(\hat{m})$ is the same function (98) as before and

$$\hat{m} = \hat{m}(m, \kappa) = \gamma \lambda^2 \mathbb{E} H \left(\frac{m D^2 + \kappa D}{1-m} + D^2 Q \right) + \mu \lambda^2 m. \quad (107)$$

Recall that the replica prediction for the MMSE is (80). In the regime $\lambda \rightarrow \infty$ we thus necessarily have $m \rightarrow 1^-$. Since the solution (m, κ) of the replica saddle point equations yields the MMSE (80) which must be at least as good as the AMP MSE (95) then $m \geq \Delta_*$. Thus from (100) we deduce

$$1 - m = O\left(\exp\left(-\frac{\lambda^2}{2}(1 + o_\lambda(1))\right)\right). \quad (108)$$

The support of the density of D is bounded, therefore from (103) it is then clear that for (104) to be verified under the scaling (108) in the large λ limit, the solution \tilde{V} of (104) must verify

$$\frac{\lambda^2}{\tilde{V}} = o_\lambda(1). \quad (109)$$

Thus from (104) we obtain

$$(1 - m)\tilde{V} = \mathbb{E}\left(1 + \frac{\gamma\lambda D^2(\lambda - D) - \mu\lambda D}{\tilde{V}}\right)^{-1} = 1 + o_\lambda(1) \quad (110)$$

from which we deduce using (108) that

$$\tilde{V} = \Theta\left(\frac{1}{1 - m}\right) = \Omega\left(\exp\left(\frac{\lambda^2}{2}(1 + o_\lambda(1))\right)\right). \quad (111)$$

This also implies that in the limit of large SNR, H becomes deterministic:

$$H = \tilde{V}^{-1} + O\left(\frac{\lambda^2}{\tilde{V}^2}\right). \quad (112)$$

This equality means that H can be written as \tilde{V}^{-1} plus a possibly random term dependent of D , which can be bounded by a non-random constant of order $O(\lambda^2/\tilde{V}^2)$. Similarly for Q : using that κ is bounded (recall that it is the limit of the expectation of (40)), (112) and (110), we get the following deterministic scaling in the large SNR regime:

$$Q = \frac{m}{1 - m} + O(\lambda^2). \quad (113)$$

Using all these scalings together with the fact that $\mathbb{E}D = 0$ and κ is bounded (actually it can now be seen from the $(\hat{\kappa})$ -equation of Section 3.3 that $\kappa = o_\lambda(1)$) we reach, using $\mathbb{E}D^2 = 1$ and (109), (110),

$$\begin{aligned} \hat{m} &= \gamma\lambda^2\mathbb{E}H\left(\frac{mD^2 + \kappa D}{1 - m} + D^2Q\right) + \mu\lambda^2m \\ &= \gamma\lambda^2\left(\tilde{V}^{-1} + O\left(\frac{\lambda^2}{\tilde{V}^2}\right)\right)\left(\frac{2m}{1 - m} + O(\lambda^2)\right) + \mu\lambda^2m \\ &= \gamma\lambda^2\left(\frac{2m}{\tilde{V}(1 - m)} + O\left(\frac{\lambda^2}{\tilde{V}}\right) + O\left(\frac{\lambda^2}{\tilde{V}} \times \frac{1}{\tilde{V}(1 - m)}\right) + O\left(\frac{\lambda^4}{\tilde{V}^2}\right)\right) + \mu\lambda^2m \\ &= \gamma\lambda^2(2m + o_\lambda(1)) + \mu\lambda^2m \\ &= \lambda^2(2\gamma + \mu)(1 + o_\lambda(1)) \end{aligned} \quad (114)$$

where also used $m = 1 + o_\lambda(1)$, see (108). Recall $m = 1 - \text{mmse}(\hat{m})$ as well as the scaling (99). So we have

$$m = 1 - \exp\left(-\frac{\lambda^2}{2}(2\gamma + \mu)(1 + o_\lambda(1))\right). \quad (115)$$

By comparing with (100) we see that $m \neq \Delta_*$. Moreover, since m is the Bayes-optimal overlap, it has to be the case that $m \geq \Delta_*$, namely, $2\gamma + \mu \geq 1$. From (17) it can be verified that $2\gamma + \mu > 1$ strictly for $\mu < 1$. Equality holds for the pure Wigner case ($\mu = 1, \gamma = 0$), as expected. This ends the proof that the MMSE (80) is asymptotically in λ strictly exponentially smaller than the MSE of AMP with one-term memory (95) whenever $\mu < 1, \gamma > 0$.

4.3 What is actually doing this sub-optimal AMP? Mismatched estimation with Gaussian likelihood

In the same spirit as [13], we study here a mismatched estimation where the statistician assumes the noise to be Gaussian, thus a wrong likelihood, whereas the noise is drawn from the quartic ensemble with potential (13). In the same way as we did for the quartic potential, the mismatched posterior associated to (2) is written as

$$d\bar{P}_{X|Y}(\mathbf{x} | \mathbf{Y}) = \frac{1}{\bar{\mathcal{Z}}(\mathbf{Y})} dP_X(\mathbf{x}) \exp\left(\frac{\lambda}{2} \text{Tr} \mathbf{Y} \mathbf{x} \mathbf{x}^\top - \frac{\lambda^2}{4N} \|\mathbf{x}\|^4\right) \quad (116)$$

where we have re-absorbed \mathbf{x} -independent terms in the normalization. The corresponding log-partition function is

$$\mathbb{E} \ln \bar{\mathcal{Z}}(\mathbf{Y}). \quad (117)$$

Notice that we have barred some quantities to distinguish them from their Bayes-optimal analogues. We further stress that, with Gaussian likelihood, the spin-glass model that arises already contains only two body interactions.

We aim at *approximating* (117). Indeed, we are going to perform a replica symmetric computation, which has no a-priori reasons to be exact as we are not anymore in the Bayesian-optimal setting [18] (nor the mismatched posterior is log-concave [19], see also [29] as a counter-example). We denote jointly $\boldsymbol{\tau} = (v_1, q_{01})$ and $\hat{\boldsymbol{\tau}}$ their Fourier conjugates. The partition function can then be expressed using deltas to fix the $\boldsymbol{\tau}$ parameters and expanding \mathbf{Y} as in (2). Up to irrelevant constants it reads

$$\bar{\mathcal{Z}}(\mathbf{Y}) = \int dP_X(\mathbf{x}) d\boldsymbol{\tau} d\hat{\boldsymbol{\tau}} \exp\left(-\bar{H}_N(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{x}; \mathbf{x}_0, \mathbf{Z})\right) \quad (118)$$

where

$$\bar{H}_N(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{x}; \mathbf{x}_0, \mathbf{Z}) := N\bar{h}(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}) + \mathbf{x}^\top \bar{\mathbf{J}}(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{Z}) \mathbf{x} + \hat{q}_{01} \mathbf{x}^\top \mathbf{x}_0 \quad (119)$$

and

$$\bar{h}(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}) := \frac{\lambda^2}{4} v_1^2 - \frac{\lambda^2}{2} q_{01}^2 - q_{01} \hat{q}_{01} - \frac{v_1 \hat{v}_1}{2}, \quad (120)$$

$$\bar{\mathbf{J}}(\boldsymbol{\tau}, \hat{\boldsymbol{\tau}}, \mathbf{Z}) := \frac{\hat{v}_1}{2} I_N - \frac{\lambda}{2} \mathbf{Z}. \quad (121)$$

While replicating we will need as before to fix the entire overlap structure (and not only q_{01}), i.e., $(N\mathbf{q})_{\ell\ell'} = Nq_{\ell\ell'} = \mathbf{x}_\ell^\top \mathbf{x}_{\ell'}$, the diagonal elements being denoted as v_ℓ . As usual, we also introduce the corresponding Fourier conjugates $\hat{\mathbf{q}}$. The expected replicated partition function then reads as

$$\begin{aligned} \mathbb{E} \bar{Z}^n &= \int d\mathbf{q} d\hat{\mathbf{q}} \exp N \left(\sum_{\ell \leq n} \left(\frac{\lambda^2}{2} q_{0\ell}^2 - \frac{\lambda^2}{4} v_\ell^2 + \frac{v_\ell \hat{v}_\ell}{2} \right) + \sum_{0 \leq \ell < \ell' \leq n} \hat{q}_{\ell\ell'} q_{\ell\ell'} \right) \\ &\times \int \prod_{\ell=0}^n dP_X(\mathbf{x}_\ell) \exp \left(- \sum_{0 \leq \ell < \ell' \leq n} \hat{q}_{\ell\ell'} \mathbf{x}_\ell^\top \mathbf{x}_{\ell'} - \frac{1}{2} \sum_{\ell \leq n} \hat{v}_\ell \|\mathbf{x}_\ell\|^2 \right) \\ &\times \mathbb{E}_{\mathbf{O}} \exp \left(\frac{\lambda}{2} \text{Tr} \mathbf{O} \mathbf{D} \mathbf{O}^\top \sum_{\ell \leq n} \mathbf{x}_\ell \mathbf{x}_\ell^\top \right). \end{aligned} \quad (122)$$

In the last line we recognize a rank- n (standard) spherical integral, see Section 2.2.1 and [54]. Recall the spectrum is deterministic with empirical law tending weakly to ρ . Hence we can use the results from Section 2.2.3, with the difference that $\mathbf{C} = I_n \frac{\lambda}{2} D$ is virtually a scalar random variable, and thus w.l.o.g. we can also assume \mathbf{q} , and thus $\hat{\mathbf{q}}$ to be diagonal in (23). If we aim for a replica symmetric ansatz

$$\text{Replica Symmetric Ansatz: } \begin{cases} v_\ell = v, & \hat{v}_\ell = \hat{v} \\ q_{0\ell} = m, & \hat{q}_{0\ell} = -\hat{m} \\ q_{\ell\ell'} = q, & \hat{q}_{\ell\ell'} = -\hat{q} \ (\ell \neq \ell') \end{cases} \quad (123)$$

then \mathbf{q} has a non degenerate eigenvalue $v + (n-1)q$ and $n-1$ degenerate eigenvalues $v - q$. Within this ansatz we can thus replace the mentioned spherical integral with

$$\begin{aligned} \mathbb{E}_{\mathbf{O}} \exp \left(\frac{\lambda}{2} \text{Tr} \mathbf{O} \mathbf{D} \mathbf{O}^\top \sum_{\ell \leq n} \mathbf{x}_\ell \mathbf{x}_\ell^\top \right) &= \exp N \left((n-1) I_D(v - q) + I_D(v - q + nq) \right) \\ &= \exp N n \left(I_D(v - q) + I'_D(v - q) q + O(n) \right) \end{aligned} \quad (124)$$

as done in [86], where $I_D(\cdot)$ are rank-one spherical integrals. The rest can be treated

exactly as in Section 3.2, yielding

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \ln \bar{\mathcal{Z}}(\mathbf{Y}) = \text{extr} \left\{ \frac{\lambda^2}{2} m^2 - \frac{\lambda^2}{4} v^2 + \frac{\hat{v}v}{2} - \hat{m}m + \frac{\hat{q}q}{2} + I_D(v - q) \right. \\ \left. + qI'_D(v - q) + \mathbb{E} \ln \int dP_X(x) \exp \left(\sqrt{\hat{q}} Zx - \frac{\hat{q} + \hat{v}}{2} x^2 + \hat{m}X_0x \right) \right\} \quad (125)$$

where extremization is intended over $m, \hat{m}, q, \hat{q}, v, \hat{v}$. With the same notation for the local measure (62), the fixed point equations read

$$(m) \quad \hat{m} = \lambda^2 m \quad (126)$$

$$(\hat{m}) \quad m = \mathbb{E} X_0 \langle X \rangle_{\hat{m}, \hat{q}, \hat{v}} \quad (127)$$

$$(q) \quad \hat{q} = 2qI''_D(v - q) \quad (128)$$

$$(\hat{q}) \quad q = \mathbb{E} \langle X \rangle_{\hat{m}, \hat{q}, \hat{v}}^2 \quad (129)$$

$$(v) \quad \hat{v} = \lambda^2 v - 2I'_D(v - q) - 2qI''_D(v - q) \quad (130)$$

$$(\hat{v}) \quad v = \mathbb{E} \langle X^2 \rangle_{\hat{m}, \hat{q}, \hat{v}}. \quad (131)$$

The computation above follows the same lines as that in [86], with the only difference being the presence of a planted signal. In case of Gaussian likelihood, the term arising from the spike though is easily tractable, as well as the term containing the fourth norm of the estimator (see (116)). This suggests that the AMP algorithm designed in [41], whose aim was to make the results in [86, 84] rigorous, has to match the performance predicted by our replica computation, measured by the MSE

$$\lim_{N \rightarrow \infty} \frac{1}{2N^2} \mathbb{E} \|\mathbf{X}^* \mathbf{X}^{*\top} - \bar{\mathbb{E}}[\mathbf{X}^* \mathbf{X}^{*\top} | \mathbf{Y}]\|_F^2 = \frac{1}{2} (1 - 2m^2 + q^2). \quad (132)$$

in the large N limit, where the $\bar{\mathbb{E}}$ denotes the expectation w.r.t. (116), and m and q solve (126)–(131).

An alternative to (126)–(131), which turns out to be more practical from the numerical point of view, can be obtained by keeping \mathbf{q} as it is, without diagonalizing it. In this case one needs the entire formula (23), with $\tilde{\mathbf{q}}$ having the same RS structure as \mathbf{q} , in a similar fashion as that of Section 2.2.3. The spherical integral then takes the form (up to constants)

$$\mathbb{E}_{\mathbf{O}} \exp \left(\frac{\lambda}{2} \text{Tr} \mathbf{O} \mathbf{D} \mathbf{O}^\top \sum_{\ell \leq n} \mathbf{x}_\ell \mathbf{x}_\ell^\top \right) \propto \exp \left(Nn \text{extr} \left\{ \frac{v\tilde{v} - q\tilde{q}}{2} - \frac{1}{2} \mathbb{E} \ln(\tilde{v} - \tilde{q} - \lambda D) \right. \right. \\ \left. \left. - \frac{\tilde{q}}{2} \mathbb{E}(\tilde{v} - \tilde{q} - \lambda D)^{-1} - \frac{1}{2} \ln(v - q) - \frac{q}{2(v - q)} + O(n) \right\} \right) \quad (133)$$

where extremization is w.r.t. the tilded variables only, for now. Consequently, the free entropy rewrites as follows

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \ln \bar{\mathcal{Z}}(\mathbf{Y}) &= \text{extr} \left\{ \frac{\lambda^2}{2} m^2 - \frac{\lambda^2}{4} v^2 + \frac{(\hat{v} + \tilde{v})v}{2} - \hat{m}m + \frac{(\hat{q} - \tilde{q})q}{2} \right. \\ &\quad - \frac{1}{2} \mathbb{E} \ln(\tilde{v} - \tilde{q} - \lambda D) - \frac{\tilde{q}}{2} \mathbb{E}(\tilde{v} - \tilde{q} - \lambda D)^{-1} - \frac{1}{2} \ln(v - q) - \frac{q}{2(v - q)} \\ &\quad \left. + \mathbb{E} \ln \int dP_X(x) \exp \left(\sqrt{\hat{q}} Z x - \frac{\hat{q} + \hat{v}}{2} x^2 + \hat{m} X_0 x \right) \right\}. \end{aligned} \quad (134)$$

Here instead, extremization is intended over the tilded and hatted variables, together with m, q, v .

The fixed point equations are

$$(m) \quad \hat{m} = \lambda^2 m \quad (135)$$

$$(\hat{m}) \quad m = \mathbb{E} X_0 \langle X \rangle_{\hat{m}, \hat{q}, \hat{v}} \quad (136)$$

$$(q) \quad \hat{q} - \tilde{q} = \frac{q}{(v - q)^2} \quad (137)$$

$$(\hat{q}) \quad q = \mathbb{E} \langle X \rangle_{\hat{m}, \hat{q}, \hat{v}}^2 \quad (138)$$

$$(\tilde{q}) \quad q = -\tilde{q} \mathbb{E}(\tilde{v} - \tilde{q} - \lambda D)^{-2} \quad (139)$$

$$(v) \quad \hat{v} + \tilde{v} - \lambda^2 v - \frac{1}{v - q} + \frac{q}{(v - q)^2} = 0 \quad (140)$$

$$(\hat{v}) \quad v = \mathbb{E} \langle X^2 \rangle_{\hat{m}, \hat{q}, \hat{v}} \quad (141)$$

$$(\tilde{v}) \quad v - \mathbb{E}(\tilde{v} - \tilde{q} - \lambda D)^{-1} + \tilde{q} \mathbb{E}(\tilde{v} - \tilde{q} - \lambda D)^{-2} = 0. \quad (142)$$

Plugging (\tilde{q}) into (\tilde{v}) we readily see that

$$v - q = \mathbb{E}(\tilde{V} - \lambda D)^{-1} \quad (143)$$

that works as an equation for $\tilde{V} := \tilde{v} - \tilde{q}$ as a function of v, q . Analogously, we can plug (q) into (v) obtaining

$$\hat{v} + \hat{q} = \lambda^2 v + \frac{1}{v - q} - \tilde{V} \quad (144)$$

that determines $\hat{v} + \hat{q}$ as a function of v and q , thanks to the above equation for \tilde{V} . Finally, from (\tilde{q}) and (q) we have respectively

$$\tilde{q} = -\frac{q}{\mathbb{E}(\tilde{V} - \lambda D)^{-2}} \quad (145)$$

$$\hat{q} = \frac{q}{(v - q)^2} + \tilde{q}. \quad (146)$$

Notice that, being in a mismatched setting, there cannot be any simplifications due to the Nishimori identities.

It is not difficult to verify a posteriori that the systems (126)–(131) and (135)–(142) are equivalent. The extremization over the tilded variables has indeed the purpose of reproducing I_D and its derivatives. From (143) one can infer

$$\tilde{V} = R_{\lambda\mathbf{D}}(v - q) + \frac{1}{v - q} \quad (147)$$

where $R_{\lambda\mathbf{D}}$ denotes the R-transform of $\lambda\mathbf{D}$, and deriving both sides w.r.t. v one also has

$$\tilde{V}' = -\frac{1}{\mathbb{E}(\tilde{V} - \lambda D)^{-2}} = R'_{\lambda\mathbf{D}}(v - q) - \frac{1}{(v - q)^2}. \quad (148)$$

Therefore, from (145)

$$\tilde{q} = qR'_{\lambda\mathbf{D}}(v - q) - \frac{q}{(v - q)^2} \quad \Rightarrow \quad \hat{q} = qR'_{\lambda\mathbf{D}}(v - q), \quad (149)$$

and from (144)

$$\hat{v} + \hat{q} = \lambda^2 v - R_{\lambda\mathbf{D}}(v - q), \quad (150)$$

both in perfect agreement with (128) and (130), as long as $R_{\lambda\mathbf{D}} = 2I'_D$ [54].

The system of fixed point equations (135)–(142) can be solved numerically as follows: *(i)* initialize $m = m_0$, $q = q_0$, $v = v_0$ (the latter being identically 1 if we use a Rademacher prior); *(ii)* solve (143) for \tilde{V} ; *(iii)* compute \hat{q} , \tilde{q} , \hat{m} and $\hat{v} + \hat{q}$ from (146), (145), (135) and (144) respectively; *(iv)* update the values of m, q, v through (\hat{m}) , (\hat{q}) and (\hat{v}) obtaining m^1, q^1 and v^1 ; *(v)* repeat the steps *(i)*–*(iv)* starting from $m = m^1$, $q = q^1$ and $v = v^1$, thus obtaining m^2, q^2 and v^2 , and so forth.

The numerics arising from this procedure though turns out to be delicate for extreme values of the overlap, namely when $v - q$ is really small, which in turn happens when λ is large (typically > 3 for Rademacher prior). The equation that seems to generate numerical instability is (146), and in particular the two contributions there appearing. With reference to the Rademacher prior, and the related Figure 2, when $\lambda > 3$ the overlap gets close to ~ 0.999 . At this value $1/(v - q)^2 \sim 10^6$. \tilde{q} , that is also contributing to (146), on the contrary becomes really negative, and is such that \hat{q} is typically ~ 10 near $\lambda \sim 3$. The subtraction of these two big numbers apparently dooms the iterations for larger SNRs. This was not the case in the Bayes-optimal setting, thanks to the simplifications introduced by the Nishimori identities. Indeed,

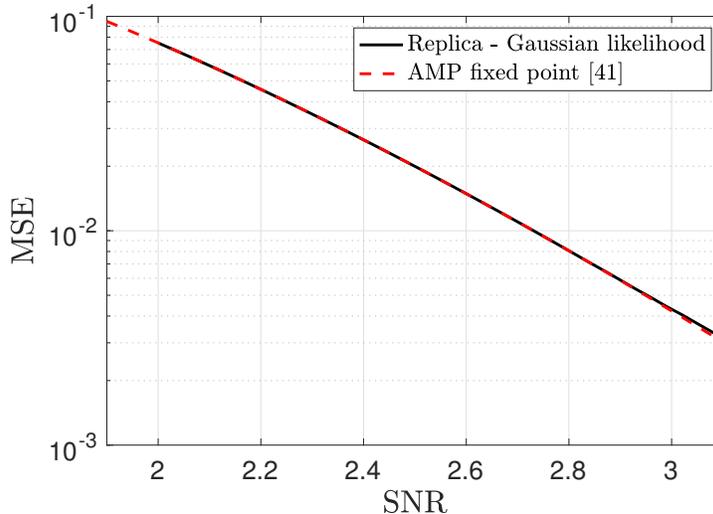


Figure 2: Comparison between the fixed point of the AMP algorithm in [41] and that obtained via the replica computation (cf. (135)–(142)), for i.i.d. Rademacher distributed $(X_i^*)_i$. The agreement between these two fixed points is excellent when the SNR is between 2 and 3.

from (76), (77), (78) and (79) we see that $1 - m$ appears at most at the first power in denominators. The only issue there was that \tilde{V} can grow exponentially fast, and this can be solved by allowing for a wide range of search of the solution of (78).

The fixed point of the MSE arising from (135)–(142) is compared with the fixed point (93), which corresponds to the MSE of the AMP proposed in [41]. The match between these two computations is excellent, as long as the SNR is not too large, because of the aforementioned numerical issues in iterating (135)–(142). The plot of Figure 2 is a compelling numerical confirmation of the arguments put forward in this section. The conclusion is the following: the AMP algorithm of [41] is solving a *replica symmetric approximation* to the TAP equations associated with the mismatched posterior distribution (116).

5 Towards an optimal AMP: AdaTAP formalism

We have previously shown that the AMP found in the literature for structured PCA [41] is sub-optimal. In this section we understand the fundamental reason behind this issue by generalizing the Adaptive Thouless-Anderson-Palmer (AdaTAP) formalism

of [86, 84]. Using our new insights we will then be able in the next section to cure the issue and derive a Bayes-optimal AMP. Like in the replica method and in particular Section 3.1, a key ingredient will be to reduce the model to a quadratic one of the Ising type.

5.1 The AdaTAP single-instance free entropy

Recall that the posterior distribution is given by (6). Denoting $\mathbf{p} := \mathbf{x}\mathbf{x}^\top/N$ and $v := \|\mathbf{x}\|^2/N$ the trace of the matrix potential (13) can be expanded as follows:

$$\begin{aligned} \text{Tr}V(\mathbf{Y} - \lambda\mathbf{p}) &= C + \frac{\mu}{2}\text{Tr}\{\lambda^2v^2 - 2\lambda\mathbf{Y}\mathbf{p}\} \\ &\quad + \frac{\gamma}{4}\text{Tr}\{\lambda^4v^4 - 4\lambda^3v^2\mathbf{Y}\mathbf{p} + 4\lambda^2v\mathbf{Y}^2\mathbf{p} - 4\lambda\mathbf{Y}^3\mathbf{p} + 2\lambda^2\mathbf{Y}\mathbf{p}\mathbf{Y}\mathbf{p}\} \end{aligned}$$

where C is independent of \mathbf{x} . Define the matrix polynomial:

$$\mathbf{R}(v, \mathbf{Y}) := -(\mu\lambda + \gamma\lambda^3v^2)\mathbf{Y} + \gamma\lambda^2v\mathbf{Y}^2 - \gamma\lambda\mathbf{Y}^3. \quad (151)$$

Then

$$\begin{aligned} -\frac{N}{2}\text{Tr}V(\mathbf{Y} - \lambda\mathbf{p}) &\propto \\ &-\frac{N}{4}\lambda^2v^2\left(\mu + \frac{\gamma\lambda^2v^2}{2}\right) - \frac{\mathbf{x}^\top\mathbf{R}(v, \mathbf{Y})\mathbf{x}}{2} - \frac{N}{4}\gamma\lambda^2\left(\frac{\mathbf{x}^\top\mathbf{Y}\mathbf{x}}{N}\right)^2. \end{aligned} \quad (152)$$

The partition function of the model defined by (12) can then be written in the form

$$\begin{aligned} \mathcal{Z} &\propto \int dP_X(\mathbf{x})dvdf\delta(Nv - \|\mathbf{x}\|^2)\delta(Nf - \mathbf{x}^\top\mathbf{Y}\mathbf{x}) \\ &\quad \times \exp\left(-\frac{N}{4}\lambda^2v^2\left(\mu + \frac{\gamma\lambda^2v^2}{2}\right) - \frac{1}{2}\mathbf{x}^\top\mathbf{R}(v, \mathbf{Y})\mathbf{x} - \frac{N}{4}\gamma\lambda^2f^2\right) \\ &= \int dvd\hat{v}df\hat{d}f \exp\left(N\hat{v}v + N\hat{f}f - \frac{N}{4}\lambda^2v^2\left(\mu + \frac{\gamma\lambda^2v^2}{2}\right) - \frac{N}{4}\gamma\lambda^2f^2\right) \\ &\quad \times \int dP_X(\mathbf{x}) \exp\left(-\hat{v}\|\mathbf{x}\|^2 - \hat{f}\mathbf{x}^\top\mathbf{Y}\mathbf{x} - \frac{1}{2}\mathbf{x}^\top\mathbf{R}(v, \mathbf{Y})\mathbf{x}\right) \\ &= \int dvd\hat{v}df\hat{d}f \exp\left(N\hat{v}v + N\hat{f}f - \frac{N}{4}\lambda^2v^2\left(\mu + \frac{\gamma\lambda^2v^2}{2}\right) - \frac{N}{4}\gamma\lambda^2f^2\right) \\ &\quad \times \int dP_X(\mathbf{x}) \exp\left(\frac{1}{2}\mathbf{x}^\top\mathbf{J}(v, \hat{v}, \hat{f}, \mathbf{Y})\mathbf{x}\right), \end{aligned} \quad (153)$$

where the overall symmetric interaction matrix of this “Ising model” is

$$\mathbf{J}(v, \hat{v}, \hat{f}, \mathbf{Y}) := -\mathbf{R}(v, \mathbf{Y}) - 2\hat{v}\mathbf{I}_N - 2\hat{f}\mathbf{Y}. \quad (154)$$

Now, defining the free entropy at fixed (v, \hat{v}, \hat{f})

$$\Phi_N(v, \hat{v}, \hat{f}, \mathbf{Y}) := \ln \int dP_X(\mathbf{x}) \exp\left(\frac{1}{2}\mathbf{x}^\top \mathbf{J}(v, \hat{v}, \hat{f}, \mathbf{Y})\mathbf{x}\right), \quad (155)$$

because the prior is factorized and we have an Ising-type of model, we can directly use the AdaTAP result [86]: it tells us that

$$\begin{aligned} \Phi_N(v, \hat{v}, \hat{f}, \mathbf{Y}) = & -\text{extr}_{\mathbf{m}, \boldsymbol{\tau}, \mathbf{V}} \left\{ \frac{1}{2}\mathbf{m}^\top \mathbf{J}(v, \hat{v}, \hat{f}, \mathbf{Y})\mathbf{m} \right. \\ & + \frac{1}{2} \ln \det (\boldsymbol{\Omega} - \mathbf{J}(v, \hat{v}, \hat{f}, \mathbf{Y})) - \frac{1}{2}\mathbf{V}^\top \mathbf{m}^2 + \frac{1}{2} \sum_{i \leq N} \ln(\tau_i - m_i^2) \\ & \left. - \sum_{i \leq N} \ln \int dP_X(x) \exp\left(\frac{1}{2}V_i x^2 + ((\mathbf{J}(v, \hat{v}, \hat{f}, \mathbf{Y})\mathbf{m})_i - V_i m_i)x\right)\right\} + o_N(1). \end{aligned} \quad (156)$$

The extremization is over $(\mathbf{m}, \boldsymbol{\tau}, \mathbf{V}) \in \mathbb{R}^N \times (\mathbb{R}_{\geq 0}^N)^2$, $\mathbf{m}^2 = (m_i^2)_{i \leq N}$, and the diagonal matrix

$$\boldsymbol{\Omega} := \text{diag}(\mathbf{V} + (\boldsymbol{\tau} - \mathbf{m}^2)^{-1}). \quad (157)$$

Let the bracket notation $\langle \cdot \rangle$ be used as expectation with respect to the posterior (6), while $\langle \cdot \rangle_{\setminus i}$ is the mean with respect to the Gibbs measure of the “cavity graph” where $(J_{ij})_j$ are set to 0. Define also the cavity fields

$$h_i := (\mathbf{J}\mathbf{x})_i.$$

The various variables at their extremum values are (asymptotically exact approximations to) the marginals means, second moments and variances of the cavity fields

$$m_i = \langle x_i \rangle, \quad \tau_i = \langle x_i^2 \rangle, \quad V_i = \langle h_i^2 \rangle_{\setminus i} - \langle h_i \rangle_{\setminus i}^2.$$

From the AdaTAP free entropy at fixed (v, \hat{v}, \hat{f}) we can compute the total log-partition function by saddle-point and get

$$\begin{aligned} \frac{1}{N} \ln \mathcal{Z}(\mathbf{Y}) \propto & o_N(1) \\ & + \text{extr} \left\{ \hat{v}v + \hat{f}f - \frac{1}{4}\lambda^2 v^2 \left(\mu + \frac{\gamma\lambda^2 v^2}{2} \right) - \frac{1}{4}\gamma\lambda^2 f^2 + \Phi_N(v, \hat{v}, \hat{f}, \mathbf{Y}) \right\} \end{aligned} \quad (158)$$

where the extremization is over (v, \hat{v}, f, \hat{f}) .

5.2 Saddle point: reduction to an Ising model, AdaTAP equations and optimal pre-processing of the data

By extremization of the AdaTAP single-instance free entropy (158) we derive the AdaTAP equations. We start with the intensive parameters. The extremization with respect to f is trivial and gives

$$\hat{f} = \frac{1}{2}\gamma\lambda^2 f.$$

So the leading order of the AdaTAP free entropy simplifies to

$$\text{extr}_{v, \hat{v}, f} \left\{ \hat{v}v + \frac{1}{4}\gamma\lambda^2 f^2 - \frac{1}{4}\lambda^2 v^2 \left(\mu + \frac{\gamma\lambda^2 v^2}{2} \right) + \Phi_N \left(v, \hat{v}, \frac{1}{2}\gamma\lambda^2 f, \mathbf{Y} \right) \right\}. \quad (159)$$

The remaining saddle point equations can simply be written down. But this is not necessary as the solution of the three remaining intensive order parameters at the saddle point is simply deduced from their physical meaning, concentration properties, and the Nishimori identity: in the large size limit,

$$v \rightarrow \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \langle \|\mathbf{x}\|^2 \rangle = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \|\mathbf{X}^*\|^2 = 1,$$

as well as (recall (67))

$$f \rightarrow \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \langle \mathbf{x}^\top \mathbf{Y} \mathbf{x} \rangle = \lambda.$$

Moreover we know that

$$\hat{v} \rightarrow 0$$

because the prior is already enforcing the constraint that $v = \|\mathbf{x}\|^2/N \rightarrow 1$ in (153) without the need to introducing a further, redundant, delta constraint; note that for Rademacher or spherical prior this is simply true as no delta function is needed. Therefore the AdaTAP free entropy becomes

$$\frac{1}{8}\gamma\lambda^4 - \frac{1}{4}\mu\lambda^2 + \Phi_N \left(1, 0, \frac{1}{2}\gamma\lambda^3, \mathbf{Y} \right) + o_N(1). \quad (160)$$

From this AdaTAP free entropy we see that the values of the marginal means and variances correspond to the solution of the variational problem (156) with interaction matrix

$$\mathbf{J} \left(1, 0, \frac{1}{2}\gamma\lambda^3, \mathbf{Y} \right) = \mu\lambda\mathbf{Y} - \gamma\lambda^2\mathbf{Y}^2 + \gamma\lambda\mathbf{Y}^3 =: \mathbf{J}(\mathbf{Y}). \quad (161)$$

So we end-up with the following effective partition function of an Ising-like model:

$$\int dP_X(\mathbf{x}) \exp\left(\frac{1}{2}\mathbf{x}^\top \mathbf{J}(\mathbf{Y})\mathbf{x}\right). \quad (162)$$

This shows that the original model is equivalent to an Ising model with interaction matrix $\mathbf{J}(\mathbf{Y})$, which can thus be interpreted as a *Bayes-optimal pre-processing of the data*. This will be verified in Section 6, as the use of $\mathbf{J}(\mathbf{Y})$ instead of \mathbf{Y} will turn AMP into an optimal algorithm. Ising models like this are precisely studied in [86] and we can therefore again exploit directly the AdaTAP formalism. Let

$$\eta_i(\mathbf{J}, \mathbf{m}, V_i) := \frac{\int dP_X(x) x e^{\frac{1}{2}V_i x^2 + ((\mathbf{J}\mathbf{m})_i - V_i m_i)x}}{\int dP_X(x) e^{\frac{1}{2}V_i x^2 + ((\mathbf{J}\mathbf{m})_i - V_i m_i)x}}, \quad (163)$$

$$g_i(\mathbf{J}, \mathbf{m}, V_i) := \frac{\int dP_X(x) x^2 e^{\frac{1}{2}V_i x^2 + ((\mathbf{J}\mathbf{m})_i - V_i m_i)x}}{\int dP_X(x) e^{\frac{1}{2}V_i x^2 + ((\mathbf{J}\mathbf{m})_i - V_i m_i)x}}. \quad (164)$$

The associated AdaTAP equations over $(\mathbf{m}, \boldsymbol{\tau}, \mathbf{V})$, namely the saddle point equations associated with the AdaTAP free entropy (156) with $\mathbf{J}(v, \hat{v}, \hat{f}, \mathbf{Y})$ replaced by $\mathbf{J} = \mathbf{J}(\mathbf{Y})$, read

$$m_i = \eta_i(\mathbf{J}, \mathbf{m}, V_i), \quad (165)$$

$$\tau_i = g_i(\mathbf{J}, \mathbf{m}, V_i), \quad (166)$$

$$\tau_i - m_i^2 = ([\text{diag}(\mathbf{V} + (\boldsymbol{\tau} - \mathbf{m}^2)^{-1}) - \mathbf{J}]^{-1})_{ii}, \quad (167)$$

where the last equation is understood as an implicit equation for \mathbf{V} .

5.3 Simplifying the AdaTAP equations by self-averaging of the Onsager reaction term

The variances V_i are expected to be self-averaging with respect to the interaction matrix, i.e., in the large size limit $V_i = \bar{V} := \lim_{N \rightarrow \infty} \mathbb{E}_{\mathbf{J}} V_i$. The computation we are going to carry out now could be performed in various ways leading to different but equivalent expressions. For pedagogical reasons we take a path that remains as close as possible to the approach of [86]. Following this reference we compute the expectation of the AdaTAP equation for \mathbf{V} . In this section, all quantities \mathbf{V} , \mathbf{m} and $\boldsymbol{\tau}$ are fixed to a solution of the AdaTAP equations (165)–(167).

We start from the convenient identity

$$([\boldsymbol{\Omega} - \mathbf{J}]^{-1})_{ii} = \partial_{\Omega_{ii}} \ln \det(\boldsymbol{\Omega} - \mathbf{J}). \quad (168)$$

We are going to average the right-hand side. As for a Gaussian model there is no spin glass phase and strong concentrations take place, the quenched and annealed averages match [86]: we can thus simply compute the logarithm of the average of the determinant. A Gaussian identity then gives

$$\mathbb{E} \det(\mathbf{\Omega} - \mathbf{J})^{-1/2} = \int \frac{d\mathbf{z}}{(2\pi)^{N/2}} \exp\left(-\frac{1}{2}\mathbf{z}^\top \mathbf{\Omega} \mathbf{z}\right) \mathbb{E} \exp\left(\frac{1}{2}\mathbf{z}^\top \mathbf{J} \mathbf{z}\right). \quad (169)$$

We denote $\mathbf{J} = \sum_{k \leq 3} c_k \mathbf{Y}^k$ where $\mathbf{c} = (\mu\lambda, -\gamma\lambda^2, \gamma\lambda)$. The term we need to compute therefore reads

$$\mathbb{E} \exp\left(\frac{1}{2}\mathbf{z}^\top \mathbf{J} \mathbf{z}\right) = \mathbb{E} \exp\left(\frac{1}{2}\mathbf{z}^\top (c_1 \mathbf{Y} + c_2 \mathbf{Y}^2 + c_3 \mathbf{Y}^3) \mathbf{z}\right) \quad (170)$$

Define the order parameters

$$p := \frac{1}{N} \mathbf{z}^\top \mathbf{X}^*, \quad v := \frac{1}{N} \|\mathbf{z}\|^2, \quad p_D := \frac{1}{N} (\mathbf{O}\mathbf{z})^\top \mathbf{D} \mathbf{O} \mathbf{X}^*. \quad (171)$$

We also have $\|\mathbf{X}^*\|^2/N = 1 + o_N(1)$. Our goal is to identify the generalized spherical integral (21). Replacing \mathbf{Y} by $\lambda \mathbf{p}^* + \mathbf{O}^\top \mathbf{D} \mathbf{O}$ (with $\mathbf{p}^* := \mathbf{X}^* \mathbf{X}^{*\top}/N$) we expand the various terms. The first term is then simply

$$c_1 \mathbf{z}^\top (\lambda \mathbf{p}^* + \mathbf{O}^\top \mathbf{D} \mathbf{O}) \mathbf{z} = c_1 (\lambda N p^2 + (\mathbf{O}\mathbf{z})^\top \mathbf{D} \mathbf{O} \mathbf{z}). \quad (172)$$

The second term is

$$\begin{aligned} & c_2 \mathbf{z}^\top (\lambda^2 (\|\mathbf{X}^*\|^2/N) \mathbf{p}^* + \lambda \mathbf{p}^* \mathbf{O}^\top \mathbf{D} \mathbf{O} + \lambda \mathbf{O}^\top \mathbf{D} \mathbf{O} \mathbf{p}^* + \mathbf{O}^\top \mathbf{D}^2 \mathbf{O}) \mathbf{z} \\ & = c_2 (N \lambda^2 p^2 + 2N \lambda p p_D + (\mathbf{O}\mathbf{z})^\top \mathbf{D}^2 \mathbf{O} \mathbf{z}) + o(N). \end{aligned} \quad (173)$$

Finally the last term is a bit more cumbersome:

$$\begin{aligned} & c_3 \mathbf{z}^\top (\lambda^3 (\|\mathbf{X}^*\|^4/N^2) \mathbf{p}^* + \lambda^2 \mathbf{p}^* \mathbf{O}^\top \mathbf{D} \mathbf{O} \mathbf{p}^* + \lambda^2 (\|\mathbf{X}^*\|^2/N) \mathbf{O}^\top \mathbf{D} \mathbf{O} \mathbf{p}^* + \lambda \mathbf{O}^\top \mathbf{D}^2 \mathbf{O} \mathbf{p}^* \\ & \quad + \lambda^2 (\|\mathbf{X}^*\|^2/N) \mathbf{p}^* \mathbf{O}^\top \mathbf{D} \mathbf{O} + \lambda \mathbf{p}^* \mathbf{O}^\top \mathbf{D}^2 \mathbf{O} + \lambda \mathbf{O}^\top \mathbf{D} \mathbf{O} \mathbf{p}^* \mathbf{O}^\top \mathbf{D} \mathbf{O} + \mathbf{O}^\top \mathbf{D}^3 \mathbf{O}) \mathbf{z} \\ & = c_3 (N \lambda^3 p^2 + \lambda^2 p^2 (\mathbf{O}\mathbf{X}^*)^\top \mathbf{D} \mathbf{O} \mathbf{X}^* + 2\lambda p (\mathbf{O}\mathbf{z})^\top \mathbf{D}^2 \mathbf{O} \mathbf{X}^* \\ & \quad + 2N \lambda^2 p p_D + \lambda N p_D^2 + (\mathbf{O}\mathbf{z})^\top \mathbf{D}^3 \mathbf{O} \mathbf{z}) + o(N). \end{aligned} \quad (174)$$

Combining all we reach

$$\begin{aligned} \mathbb{E} \exp\left(\frac{1}{2}\mathbf{z}^\top \mathbf{J} \mathbf{z}\right) & = \int d\tau d\hat{\tau} \exp\left(NK + \frac{1}{2}\hat{v} \|\mathbf{z}\|^2 - \frac{N}{2}\hat{v}v + o(N)\right) \\ & \quad \times \mathbb{E}_{\mathbf{O}} \exp\left((\mathbf{O}\mathbf{z})^\top \mathbf{C}_{\mathbf{z},\mathbf{z}} \mathbf{O} \mathbf{z} + (\mathbf{O}\mathbf{X}^*)^\top \mathbf{C}_{*,*} \mathbf{O} \mathbf{X}^* + (\mathbf{O}\mathbf{X}^*)^\top \mathbf{C}_{\mathbf{z},*} \mathbf{O} \mathbf{z}\right) \end{aligned} \quad (175)$$

with $d\boldsymbol{\tau} := (dp, dv, dp_D)$ and $d\hat{\boldsymbol{\tau}} := (d\hat{p}, d\hat{v}, d\hat{p}_D)$, and (all coupling matrices below are $N \times N$ and symmetric)

$$\begin{aligned} K &:= \frac{1}{2} (\mu\lambda^2 p^2 - \gamma\lambda^2 (\lambda^2 p^2 + 2\lambda p p_D) + \gamma\lambda (\lambda^3 p^2 + 2\lambda^2 p p_D + \lambda p_D^2) + \hat{p}p + \hat{p}_D p_D), \\ \mathbf{C}_{*,*} &:= \frac{1}{2} \gamma \lambda^3 p^2 \mathbf{D}, \\ \mathbf{C}_{\mathbf{z},\mathbf{z}} &:= \frac{1}{2} (\mu\lambda \mathbf{D} - \gamma\lambda^2 \mathbf{D}^2 + \gamma\lambda \mathbf{D}^3), \\ \mathbf{C}_{\mathbf{z},*} &:= \frac{1}{2} (-\hat{p}I_N - \hat{p}_D \mathbf{D} + 2\gamma\lambda^2 p \mathbf{D}^2). \end{aligned}$$

Note the asymmetry for the variable \hat{v} compared to the other hat-variables, which has not been injected in the definition of the coupling matrices as the others, but instead leads to a term appearing explicitly in (175) (both choices are equivalently valid ones). The term averaged over \mathbf{O} is an inhomogeneous spherical integral as studied in Section 2. In particular, we are in the case of Section 2.2.3 with $\ell \in \{0, 1\}$ with the exception that \mathbf{X}^* also (playing the role of the 0th replica) has a non-zero self-coupling. So this trivial modification of the computation of Section 2.2.3 yields

$$\mathbb{E} e^{\frac{1}{2} \mathbf{z}^\top \mathbf{J} \mathbf{z}} = \int d\boldsymbol{\tau} d\hat{\boldsymbol{\tau}} \exp \left(NK + \frac{1}{2} \hat{v} \|\mathbf{z}\|^2 - \frac{N}{2} \hat{v} v + N I_{\mathbf{C}}(p, v, \hat{p}, \hat{p}_D) + o(N) \right)$$

where the 2×2 random coupling matrix \mathbf{C} has entries

$$2C_{00} = \gamma\lambda^3 p^2 D, \quad (176)$$

$$2C_{11} = \mu\lambda D - \gamma\lambda^2 D^2 + \gamma\lambda D^3, \quad (177)$$

$$2C_{01} = 2C_{10} = \frac{1}{2} (-\hat{p} - \hat{p}_D D + 2\gamma\lambda^2 p D^2), \quad (178)$$

with $D \sim \rho$ drawn from the noise asymptotic spectral density, and

$$\begin{aligned} I_{\mathbf{C}}(p, v, \hat{p}, \hat{p}_D) &= \frac{1}{2} \text{extr}_{(\tilde{v}_0, \tilde{v}, \tilde{p})} \left\{ \tilde{v}_0 + 2\tilde{p}p + \tilde{v}v \right. \\ &\quad \left. - \mathbb{E} \ln \left((\tilde{v}_0 - 2C_{00})(\tilde{v} - 2C_{11}) - (\tilde{p} - 2C_{01})^2 \right) \right\} - \frac{1}{2} \ln(v - p^2) - 1. \end{aligned} \quad (179)$$

One can check that $I_{\mathbf{C}}$ is null when $C_{00} = C_{11} = C_{01}$ as it should. Therefore equation

(169) becomes at leading exponential order

$$\begin{aligned}
& \ln \mathbb{E} \det(\boldsymbol{\Omega} - \mathbf{J})^{-1/2} \\
&= \ln \int \frac{d\mathbf{z}}{(2\pi)^{N/2}} d\boldsymbol{\tau} d\hat{\boldsymbol{\tau}} \exp\left(-\frac{1}{2}\mathbf{z}^\top(\boldsymbol{\Omega} - \hat{v}I_N)\mathbf{z} + NK - \frac{N}{2}\hat{v}v + NI_{\mathbf{C}} + o(N)\right) \\
&= \ln \int d\boldsymbol{\tau} d\hat{\boldsymbol{\tau}} \exp\left(NK - \frac{N}{2}\hat{v}v + NI_{\mathbf{C}} - \frac{1}{2}\ln \det(\boldsymbol{\Omega} - \hat{v}I_N) + o(N)\right) \\
&= \text{extr}\left\{NK - \frac{N}{2}\hat{v}v + NI_{\mathbf{C}} - \frac{1}{2}\ln \det(\boldsymbol{\Omega} - \hat{v}I_N)\right\} + o(N),
\end{aligned}$$

where we used Gaussian integration followed by a saddle point estimation. By the aforementioned strong concentration properties of the Gaussian model, this is also equal to $-\frac{1}{2}\ln \mathbb{E} \det(\boldsymbol{\Omega} - \mathbf{J}) \approx -\frac{1}{2}\mathbb{E} \ln \det(\boldsymbol{\Omega} - \mathbf{J})$ so we reach at leading order

$$\begin{aligned}
\mathbb{E} \ln \det(\boldsymbol{\Omega} - \mathbf{J}) &\approx \text{extr}\left\{-2NK + N\hat{v}v - 2NI_{\mathbf{C}} + \ln \det(\boldsymbol{\Omega} - \hat{v}I_N)\right\} \\
&= \text{extr}_{(\hat{v}, v)}\left\{N\hat{v}v + \sum_{i \leq N} \ln(\Omega_{ii} - \hat{v}) - 2N\tilde{G}(v)\right\} \quad (180)
\end{aligned}$$

where the extremization is over all variables and

$$\tilde{G}(v) := \text{extr}_{(p, p_D, \hat{p}, \hat{p}_D)}\left\{I_{\mathbf{C}}(p, v, \hat{p}, \hat{p}_D) + K(p, p_D, \hat{p}, \hat{p}_D)\right\}. \quad (181)$$

This is the analogue of the G-function appearing, e.g., in [86]. The extremization over \hat{v} in (180) yields that at the saddle point,

$$v = \frac{1}{N} \sum_{i \leq N} \frac{1}{\Omega_{ii} - \hat{v}}.$$

Moreover, combining the TAP equation (167) with (168) and (180) we have

$$\mathbb{E}(\tau_i - m_i^2) = \partial_{\Omega_{ii}} \mathbb{E} \ln \det(\boldsymbol{\Omega} - \mathbf{J}) = \frac{1}{\Omega_{ii} - \hat{v}} \quad (182)$$

where \hat{v} is evaluated at its saddle point value. Therefore, summing over i the last identity and recalling the definition of Ω_{ii} we reach

$$\bar{\chi} := \frac{1}{N} \mathbb{E} \sum_{i \leq N} (\tau_i - m_i^2) = v = \frac{1}{N} \mathbb{E} \sum_{i \leq N} \frac{1}{V_i + (\tau_i - m_i^2)^{-1} - \hat{v}}. \quad (183)$$

Under the concentration assumption $V_i = \bar{V}$ for all $i \leq N$, this identity implies

$$V_i = \hat{v}. \quad (184)$$

Additionally the saddle point equation for v extracted from (180) yields

$$\hat{v} = 2\partial_v \tilde{G}(v)|_{v=\bar{\chi}} \quad \Rightarrow \quad V_i = \bar{V} := 2\partial_v \tilde{G}(v)|_{v=\bar{\chi}}. \quad (185)$$

The variable $\bar{\chi}$ is instance-independent and can be deduced from our replica theory: it is equal to twice the MMSE (80), namely,

$$\bar{\chi} = 1 - m^2 \quad (186)$$

where m is solution to the replica fixed point equations (74)–(79). Computing \bar{V} from (185) is then easy, as taking a derivative w.r.t. v of $\tilde{G}(v)$ is straightforward: all the quantities appearing on the right-hand side of (181) are at the saddle point, so it simply amounts to a partial derivative of (179). It gives

$$\bar{V} = \tilde{v} - \frac{1}{1 - m^2 - p^2} \quad (187)$$

where $\tilde{v} = \tilde{v}(p, v)$ takes its saddle point value from (179) while $p = p(v)$ from (181) with $v = \bar{\chi}$ fixed.

Thanks to these simplifications the AdaTAP equation reads in the large size limit

$$m_i = \eta_i(\mathbf{J}, \mathbf{m}, \bar{V}). \quad (188)$$

Or, when written in a fashion closer to the form of AMP algorithms, the AdaTAP equations read

$$\mathbf{f} = \mathbf{J}\mathbf{m} - \bar{V}\mathbf{m}, \quad \mathbf{m} = \eta_{\bar{V}}(\mathbf{f}), \quad (189)$$

where the “denoiser”, which is applied component-wise above, is

$$\eta_{\bar{V}}(f) := \frac{\int dP_X(x) x e^{\frac{1}{2}\bar{V}x^2 + fx}}{\int dP_X(x) e^{\frac{1}{2}\bar{V}x^2 + fx}}. \quad (190)$$

6 Approximate message passing, optimally

We will now describe an AMP algorithm that matches the replica prediction for the minimum mean-square error. We therefore conjecture it to be Bayes-optimal. The main difference between this new AMP and the previously proposed one for structured PCA is that it is constructed from iterates based on the pre-processed matrix $\mathbf{J}(\mathbf{Y})$ rather than \mathbf{Y} as in [41]. Consequently, the Onsager reaction terms will have to be adapted.

6.1 BAMP: Bayes-optimal AMP

The AdaTAP approach described in Section 5 suggests that, in order to achieve Bayes-optimal performance, one should consider the BAMP iteration which is of the form

$$\mathbf{f}^t = \mathbf{J}(\mathbf{Y})\mathbf{u}^t - \sum_{i=1}^t \mathbf{c}_{t,i}\mathbf{u}^i, \quad \mathbf{u}^{t+1} = g_{t+1}(\mathbf{f}^t), \quad t \geq 1. \quad (191)$$

As in the AMP iteration (90), the *denoiser* function $g_{t+1} : \mathbb{R} \rightarrow \mathbb{R}$ is continuously differentiable, Lipschitz and applied component-wise. Crucially, the Onsager coefficients $\{\mathbf{c}_{t,i}\}_{i \in [t], t \geq 1}$ need to ensure that, conditioned on the signal, the empirical distribution of the iterate \mathbf{f}^t is Gaussian, namely, the convergence result in (91) holds for some mean vector $\boldsymbol{\mu}_t$ and covariance matrix $\boldsymbol{\Sigma}_t$.

We highlight that the matrix \mathbf{Y} in (90) is replaced by the matrix $\mathbf{J}(\mathbf{Y})$ in (191). This means that the state evolution result of [41] cannot be applied and the Onsager coefficients $\{\mathbf{c}_{t,i}\}_{i \in [t], t \geq 1}$ will have a different form with respect to $\{\mathbf{b}_{t,i}\}_{i \in [t], t \geq 1}$.

In what follows, we will consider the general case in which $\mathbf{J}(\mathbf{Y})$ is an arbitrary polynomial of degree K in \mathbf{Y} , namely,

$$\mathbf{J}(\mathbf{Y}) = \sum_{i \leq K} c_i \mathbf{Y}^i.$$

To compute $\{\mathbf{c}_{t,i}\}_{i \in [t], t \geq 1}$ and obtain a state evolution result for the iteration (191), the key idea is to map the first T iterations of (191) to the first $K \times T$ iterations of an *auxiliary* AMP with iterates $(\tilde{\mathbf{z}}^t, \tilde{\mathbf{u}}^t)_{t \in [KT]}$ and denoisers $\{\tilde{h}_{t+1}\}_{t \in [KT]}$, whose state evolution can be deduced from [41]. The denoisers $\{\tilde{h}_{t+1}\}_{t \in [KT]}$ of this auxiliary AMP are chosen so that, for $t \in [T]$ and $\ell \in [K]$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \|\tilde{\mathbf{u}}^{K(t-1)+\ell} - \mathbf{Y}^{\ell-1} \mathbf{u}^t\|_2^2 = 0. \quad (192)$$

More specifically, for $t \in [T]$ and $\ell \in \{2, \dots, K\}$, the denoiser $\tilde{h}_{K(t-1)+\ell}$ giving $\tilde{\mathbf{u}}^{K(t-1)+\ell}$ is a linear combinations of the past iterates $\tilde{\mathbf{u}}^1, \dots, \tilde{\mathbf{u}}^{K(t-1)+\ell-1}$ and of $\tilde{\mathbf{z}}^{K(t-1)+\ell-1}$; furthermore, the coefficients of these linear combinations are chosen to ensure that $\tilde{\mathbf{u}}^{K(t-1)+\ell} \approx \mathbf{Y}^{\ell-1} \mathbf{u}^t$. Hence, from $\tilde{\mathbf{z}}^{Kt}$ and $(\tilde{\mathbf{u}}^{K(t-1)+\ell})_{\ell \in \{2, \dots, K\}}$, one obtains $(\mathbf{Y}^\ell \mathbf{u}^t)_{\ell \in [K]}$ (up to an $o_N(1)$ error). As a result, $\mathbf{J}(\mathbf{Y})\mathbf{u}^t$ can be expressed as a linear combination of $(\tilde{\mathbf{u}}^1, \dots, \tilde{\mathbf{u}}^{Kt}, \tilde{\mathbf{z}}^{Kt})$, which in turn is a linear combination of (i) the past iterates $\{\mathbf{u}^i\}_{i \in [t]}$, (ii) the signal \mathbf{X}^* , plus (iii) independent Gaussian noise. By inspecting the coefficients of this linear combination, one deduces (i) the values of the Onsager coefficients $\{\mathbf{c}_{t,i}\}_{i \in [t], t \geq 1}$ (as the coefficients multiplying the past iterates $\{\mathbf{u}^i\}_{i \in [t]}$), (ii) the mean $\boldsymbol{\mu}_t$ (as the coefficient multiplying the signal \mathbf{X}^*),

and (iii) the covariance matrix Σ_t (as the covariance matrix of the remaining noise terms). Finally, by making \tilde{h}_{Kt+1} depend on g_{t+1} , we enforce that $\tilde{\mathbf{u}}^{Kt+1} \approx \mathbf{u}^{t+1}$. We highlight that the auxiliary AMP is employed purely as a proof technique. Its formal description is deferred to Appendix B.1, and its state evolution follows in Appendix B.2.

For simplicity, we assume to have access to an initialization $\mathbf{u}^1 \in \mathbb{R}^N$, which is independent of the noise \mathbf{Z} and has a strictly positive correlation with \mathbf{X}^* , i.e.,

$$(\mathbf{X}^*, \mathbf{u}^1) \xrightarrow{W_2} (X^*, U_1), \quad \mathbb{E}[X^* U_1] := \epsilon > 0, \quad \mathbb{E}[U_1^2] = 1. \quad (193)$$

The requirement (193) is rather standard in the analysis of AMP algorithms. However, as having access to such an initialization is often impractical, a recent line of work has designed AMP iterations which are initialized with the eigenvector of the data matrix \mathbf{Y} associated to the largest eigenvalue, see [79, 77, 102]. By following the approach detailed in [77], one can design a Bayes-optimal AMP with spectral initialization. As this would be out of the scope of the current contribution – whose goal is to obtain an algorithm with a Bayes-optimal *fixed point* – we will not pursue this extension here.

6.2 Onsager coefficients and state evolution recursion

We now detail the calculation of the Onsager coefficients $\{\mathbf{c}_{t,i}\}_{i \in [t], t \geq 1}$ and of the state evolution parameters $\boldsymbol{\mu}_t, \Sigma_t$ associated to the AMP algorithm (191). We obtain these quantities from the state evolution recursion of the auxiliary AMP which, up to a $o_N(1)$ error, tracks $(\mathbf{Y}^{\ell-1} \mathbf{u}^t)_{\ell \in [K]}$ and, as such, has a number of iterations K times larger. To express the latter, we define a number of auxiliary quantities: the vector $\tilde{\boldsymbol{\mu}}_{Kt} \in \mathbb{R}^{Kt}$, the matrices $\tilde{\boldsymbol{\Delta}}_{Kt}, \tilde{\boldsymbol{\Phi}}_{Kt}, \tilde{\boldsymbol{\Sigma}}_{Kt}, \tilde{\mathbf{B}}_{Kt} \in \mathbb{R}^{Kt \times Kt}$, and the coefficients $\{\alpha_{i,j}\}_{j \in [i], i \in [Kt]}$, $\{\beta_{i,j}\}_{j \in [(i-1)/K]+1, i \in [Kt]}$, $\{\gamma_i\}_{i \in [Kt]}$, $\{\theta_{i,j}\}_{i \in [t], j \in [Kt]}$. The quantities $\tilde{\boldsymbol{\mu}}_{Kt}, \tilde{\boldsymbol{\Delta}}_{Kt}, \tilde{\boldsymbol{\Phi}}_{Kt}, \tilde{\boldsymbol{\Sigma}}_{Kt}, \tilde{\mathbf{B}}_{Kt}$ are directly connected to the state evolution of the auxiliary AMP (see the remark at the end of Appendix B.2). Furthermore, the coefficients $\{\alpha_{i,j}\}_{j \in [i], i \in [Kt]}$, $\{\beta_{i,j}\}_{j \in [(i-1)/K]+1, i \in [Kt]}$, $\{\gamma_i\}_{i \in [Kt]}$, $\{\theta_{i,j}\}_{i \in [t], j \in [Kt]}$ allow for a useful (approximate) decomposition of the vectors $(\mathbf{Y}^\ell \mathbf{u}^t)_{\ell \in [K-1]}$, see the remark at the end of this section.

We start with the initialization

$$\tilde{U}_1 := U_1, \quad (194)$$

where U_1 satisfies (193), and we set

$$\begin{aligned} \tilde{\mu}_1 &:= \lambda \epsilon, & (\tilde{\boldsymbol{\Delta}}_1)_{1,1} &:= 1, & (\tilde{\boldsymbol{\Phi}}_1)_{1,1} &:= 0, & (\tilde{\mathbf{B}}_1)_{1,1} &:= \bar{\kappa}_1, & (\tilde{\boldsymbol{\Sigma}}_1)_{1,1} &:= \bar{\kappa}_2, \\ \alpha_{1,1} &:= 0, & \beta_{1,1} &:= 1, & \gamma_1 &:= 0. \end{aligned} \quad (195)$$

Here and in what follows, we denote by $\{\bar{\kappa}_k\}_{k \geq 1}$ the sequence of free cumulants associated to D . The free cumulants can be recursively computed from the moments, see e.g. [83, Section 2.5].

For $t \geq 1$, let us define

$$\tilde{U}_{K(t-1)+1+\ell} := \tilde{Z}_{K(t-1)+\ell} + \tilde{\mu}_{K(t-1)+\ell} X^* + \sum_{j=1}^{K(t-1)+\ell} (\tilde{\mathbf{B}}_{K(t-1)+\ell})_{K(t-1)+\ell, j} \tilde{U}_j, \quad \ell \in [K-1], \quad (196)$$

$$\tilde{U}_{Kt+1} := g_{t+1} \left(\mu_t X^* + \sum_{j=1}^{Kt} \theta_{t,j} \tilde{Z}_j \right), \quad (197)$$

$$(\tilde{Z}_1, \dots, \tilde{Z}_{Kt}) \sim \mathcal{N}(0, \tilde{\Sigma}_{Kt}) \text{ and independent of } X^*, U_1. \quad (198)$$

We note that the function g_{t+1} in (197) is the AMP denoiser in (191). Let us also define

$$\tilde{\mu}_{K(t-1)+1+\ell} = \lambda \mathbb{E}[\tilde{U}_{K(t-1)+1+\ell} X^*], \quad (199)$$

$$(\tilde{\Delta}_{K(t-1)+1+\ell})_{K(t-1)+1+\ell, j} = (\tilde{\Delta}_{K(t-1)+1+\ell})_{j, K(t-1)+1+\ell} = \mathbb{E}[\tilde{U}_{K(t-1)+1+\ell} \tilde{U}_j], \quad (200)$$

$$j \in [K(t-1) + 1 + \ell],$$

$$(\tilde{\Phi}_{K(t-1)+1+\ell})_{K(t-1)+1+\ell, j} = \mathbb{E}[\partial_{\tilde{Z}_j} \tilde{U}_{K(t-1)+1+\ell}], \quad j \in [K(t-1) + \ell], \quad (201)$$

$$\tilde{\mathbf{B}}_{K(t-1)+1+\ell} = \sum_{j=0}^{K(t-1)+\ell} \bar{\kappa}_{j+1} \tilde{\Phi}_{K(t-1)+1+\ell}^j, \quad (202)$$

$$\tilde{\Sigma}_{K(t-1)+1+\ell} = \sum_{j=0}^{2(K(t-1)+\ell)} \bar{\kappa}_{j+2} \sum_{k=0}^j (\tilde{\Phi}_{K(t-1)+1+\ell})^k \tilde{\Delta}_{K(t-1)+1+\ell} (\tilde{\Phi}_{K(t-1)+1+\ell}^\top)^{j-k}. \quad (203)$$

Now, we obtain $\tilde{\mu}_{K(t-1)+1}$, $\tilde{\Delta}_{K(t-1)+1}$, $\tilde{\Phi}_{K(t-1)+1}$, $\tilde{\mathbf{B}}_{K(t-1)+1}$, $\tilde{\Sigma}_{K(t-1)+1}$ by setting $\ell = 0$ in (199)–(203) (and by using the initialization (195) for $t = 1$). This allows us to define $\tilde{U}_{K(t-1)+2}$ by setting $\ell = 1$ in (196). Next, we obtain $\tilde{\mu}_{K(t-1)+2}$, $\tilde{\Delta}_{K(t-1)+2}$, $\tilde{\Phi}_{K(t-1)+2}$, $\tilde{\mathbf{B}}_{K(t-1)+2}$, $\tilde{\Sigma}_{K(t-1)+2}$ by setting $\ell = 1$ in (199)–(203). This allows us to define $\tilde{U}_{K(t-1)+2}$ by setting $\ell = 2$ in (196). We iterate this procedure until we have obtained $(\tilde{\mu}_{K(t-1)+\ell}, \tilde{\Delta}_{K(t-1)+\ell}, \tilde{\Phi}_{K(t-1)+\ell}, \tilde{\mathbf{B}}_{K(t-1)+\ell}, \tilde{\Sigma}_{K(t-1)+\ell})_{\ell \in [K]}$ and $(\tilde{U}_{K(t-1)+1+\ell})_{\ell \in [K-1]}$. We note that, for any $i \geq 1$, $\tilde{\mathbf{B}}_i$ and $\tilde{\Sigma}_i$ are the top left submatrices of $\tilde{\mathbf{B}}_{i+1}$ and $\tilde{\Sigma}_{i+1}$, respectively.

At this point, for $\ell \in [K-1]$, we compute the quantities $\{\alpha_{K(t-1)+1+\ell, j}\}_{j \in [K(t-1)+\ell]}$,

$\{\beta_{K(t-1)+1+\ell,j}\}_{j \in [t]}$, $\gamma_{K(t-1)+1+\ell}$ as

$$\alpha_{K(t-1)+1+\ell,j} = \delta_{K(t-1)+\ell,j} + \sum_{\substack{i=1 \\ i \not\equiv 1 \pmod{K}}}^{K(t-1)+\ell} \alpha_{i,j} (\tilde{\mathbf{B}}_{K(t-1)+\ell})_{K(t-1)+\ell,i}, \quad j \in [K(t-1) + \ell], \quad (204)$$

$$\beta_{K(t-1)+1+\ell,j} = (\tilde{\mathbf{B}}_{K(t-1)+\ell})_{K(t-1)+\ell,K(j-1)+1} + \sum_{\substack{i=1 \\ i \not\equiv 1 \pmod{K}}}^{K(t-1)+\ell} \beta_{i,j} (\tilde{\mathbf{B}}_{K(t-1)+\ell})_{K(t-1)+\ell,i}, \quad j \in [t], \quad (205)$$

$$\gamma_{K(t-1)+1+\ell} = \tilde{\mu}_{K(t-1)+\ell} + \sum_{\substack{i=1 \\ i \not\equiv 1 \pmod{K}}}^{K(t-1)+\ell} (\tilde{\mathbf{B}}_{K(t-1)+\ell})_{K(t-1)+\ell,i} \gamma_i. \quad (206)$$

In (204), $\delta_{i,j}$ denotes the Kronecker symbol ($\delta_{i,j} = 1$ if $i = j$ and 0 otherwise), and $\alpha_{i,j}$ is assumed to be 0 if $j \geq i$; in (205), $\beta_{i,j}$ is assumed to be 0 if $j > \lceil (i-1)/K \rceil$.

Recall that $\{c_i\}_{i=1}^K$ are the coefficients of the polynomial $\mathbf{J}(\mathbf{Y})$ (in \mathbf{Y}), i.e., $\mathbf{J}(\mathbf{Y}) = \sum_{i=1}^K c_i \mathbf{Y}^i$. Finally, we are ready to express μ_t , $\{\theta_{t,j}\}_{j \in [Kt]}$:

$$\mu_t = \sum_{i=1}^K c_i \left(\tilde{\mu}_{K(t-1)+i} + \sum_{k=1}^{K(t-1)+i} \gamma_k (\tilde{\mathbf{B}}_{K(t-1)+i})_{K(t-1)+i,k} \right), \quad (207)$$

$$\theta_{t,j} = \sum_{i=1}^K c_i \left(\delta_{K(t-1)+i,j} + \sum_{k=1}^{K(t-1)+i} \alpha_{k,j} (\tilde{\mathbf{B}}_{K(t-1)+i})_{K(t-1)+i,k} \right), \quad j \in [Kt]. \quad (208)$$

As before, $\alpha_{i,j}$ is assumed to be 0 if $j \geq i$. This allows us to define \tilde{U}_{Kt+1} via (197) and, after setting $\beta_{Kt+1,t} = 1$, $\beta_{Kt+1,j} = 0$ for all $j \in [t-1]$, $\alpha_{Kt+1,j} = 0$ for all $j \in [Kt+1]$ and $\gamma_{Kt+1} = 0$, the definition of the state evolution recursion is complete.

From the state evolution recursion defined above, we can derive the Onsager coefficients $\{\mathbf{c}_{t,j}\}_{j \in [t]}$ as

$$\mathbf{c}_{t,j} = \sum_{i=1}^K c_i \sum_{k=1}^{K(t-1)+i} \beta_{k,j} (\tilde{\mathbf{B}}_{K(t-1)+i})_{K(t-1)+i,k}, \quad j \in [t]. \quad (209)$$

At this point, we are ready to present our result concerning the characterization of the iterates of the AMP algorithm (191), with Onsager coefficients given by (209),

in the high-dimensional limit $N \rightarrow \infty$: we prove that the convergence (91) holds, where μ_t is given by (207) and $W_t = \sum_{j=1}^{Kt} \theta_{t,j} \tilde{Z}_j$, with $\{\theta_{t,j}, \tilde{Z}_j\}_{j \in [Kt]}$ described by the recursion above. Equivalently [44, Corollary 7.21], the convergence can be expressed in terms of pseudo-Lipschitz test functions. A function $\psi: \mathbb{R}^m \rightarrow \mathbb{R}$ is *pseudo-Lipschitz of order 2*, denoted by $\psi \in \text{PL}(2)$, if there exists a constant $C > 0$ such that

$$\|\psi(\mathbf{x}) - \psi(\mathbf{y})\|_2 \leq C \left(1 + \|\mathbf{x}\|_2 + \|\mathbf{y}\|_2\right) \|\mathbf{x} - \mathbf{y}\|_2,$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^m$.

Theorem 1 (State evolution of the BAMP). *Let \mathbf{Y} be given by (2) and which verifies Hypothesis 1, and let $\mathbf{J}(\mathbf{Y}) = \sum_{i=1}^K c_i \mathbf{Y}^i$. Consider the AMP algorithm (191), with initialization (193), Onsager coefficients $\{c_{t,j}\}_{j \in [t]}$ given by (209) and where, for $t \geq 1$, g_{t+1} is continuously differentiable and Lipschitz. Then, the following limit holds almost surely for any $\text{PL}(2)$ function $\psi: \mathbb{R}^{2t+2} \rightarrow \mathbb{R}$, for $t \geq 1$ as $N \rightarrow \infty$:*

$$\frac{1}{N} \sum_{i \leq N} \psi(u_i^1, \dots, u_i^{t+1}, f_i^1, \dots, f_i^t, X_i^*) \rightarrow \mathbb{E} \psi(U_1, \dots, U_{t+1}, F_1, \dots, F_t, X^*). \quad (210)$$

Equivalently, as $N \rightarrow \infty$, the joint empirical distribution of $(\mathbf{u}^1, \dots, \mathbf{u}^{t+1}, \mathbf{f}^1, \dots, \mathbf{f}^t, \mathbf{X}^*)$ converges almost surely in Wasserstein-2 distance to $(U_1, \dots, U_{t+1}, F_1, \dots, F_t, X^*)$. Here, for $i \in [t]$, $U_{i+1} = g_{i+1}(F_t)$ and $(F_1, \dots, F_t) = \boldsymbol{\mu}_t X^* + (W_1, \dots, W_t)$, with $W_t = \sum_{j=1}^{Kt} \theta_{t,j} Z_j$ and where $\boldsymbol{\mu}_t$ can be computed via (207), $\{\theta_{t,j}\}_{j \in [Kt]}$ via (208) and $\{Z_j\}_{j \in [Kt]}$ is given by (198).

The proof of Theorem 1 is deferred to Appendix B.3. A few remarks are now in order. First, we highlight that (210) directly implies a high-dimensional characterization of the performance of the AMP (191). In fact, by taking the pseudo-Lipschitz functions $\psi(U_{t+1}, X^*) = (U_{t+1} - X^*)^2$, $\psi(U_{t+1}, X^*) = U_{t+1} \cdot X^*$ and $\psi(U_{t+1}, X^*) = (U_{t+1})^2$, we obtain the limit mean-square error and overlap of the AMP iterates as

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{2N^2} \mathbb{E} \|\mathbf{X}^* (\mathbf{X}^*)^\top - \mathbf{u}^t (\mathbf{u}^t)^\top\|_{\mathbb{F}}^2 &= \frac{1}{2} \left(1 - 2 \left(\mathbb{E}[U_t \cdot X^*]\right)^2 + (\mathbb{E}[(U_t)^2])^2\right), \\ \lim_{N \rightarrow \infty} \frac{|\langle \mathbf{X}^*, \mathbf{u}^t \rangle|}{\|\mathbf{u}^t\| \cdot \|\mathbf{X}^*\|} &= \frac{|\mathbb{E}[U_t \cdot X^*]|}{\sqrt{\mathbb{E}[(U_t)^2]}}. \end{aligned} \quad (211)$$

Next, note that Theorem 1 holds for any family of denoisers $\{g_{t+1}\}_{t \geq 1}$, subject to some mild regularity requirement. A natural choice is to pick the posterior mean

$$g_{t+1}(f) = \mathbb{E}[U_* \mid F_t = f]. \quad (212)$$

Such a choice requires estimating the state evolution parameters μ_t , $\{\theta_{t,j}\}_{j \in [Kt]}$ and $\tilde{\Sigma}_{Kt}$. These parameters, as well as the Onsager coefficients (209), can be estimated consistently from the data. To do so, first we obtain $\tilde{\Delta}_{Kt}$ and $\tilde{\Phi}_{Kt}$ by replacing expectations with empirical averages in (200) and (201), respectively. Next, we compute $\tilde{\mathbf{B}}_{Kt}$ and $\tilde{\Sigma}_{Kt}$ by plugging in such estimates in (202) and (203), respectively. Having done that, we obtain $\{\alpha_{K(t-1)+1+\ell,j}\}_{j \in [K(t-1)+\ell], \ell \in [K-1]}$, $\{\beta_{K(t-1)+1+\ell,j}\}_{j \in [t], \ell \in [K-1]}$, $\{\gamma_{K(t-1)+1+\ell}\}_{\ell \in [K-1]}$ via (204)–(206). Finally, μ_t , $\{\theta_{t,j}\}_{j \in [Kt]}$ and $\{\mathbf{c}_{t,j}\}_{j \in [t]}$ can be computed from (207), (208) and (209), respectively.

As a final remark, we provide an interpretation of the coefficients $\{\alpha_{i,j}\}$, $\{\beta_{i,j}\}$, $\{\gamma_i\}$. As a by-product of the argument proving Theorem 1, we will show that, for $\ell \in [K-1]$, (cf. (234)–(235))

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{Y}^\ell \mathbf{u}^t - \sum_{j=1}^{K(t-1)+\ell} \alpha_{K(t-1)+1+\ell,j} \tilde{\mathbf{z}}^j - \sum_{j=1}^t \beta_{K(t-1)+1+\ell,j} \mathbf{u}^j - \gamma_{K(t-1)+1+\ell} \mathbf{X}^*\|^2}{N} = 0. \quad (213)$$

This formalizes the fact that $\mathbf{Y}^\ell \mathbf{u}^t$ can be approximately expressed as a linear combination of (i) the past iterates $\{\mathbf{u}^j\}_{j \in [t]}$, (ii) the signal \mathbf{X}^* , plus (iii) independent Gaussian noise (represented by the $\tilde{\mathbf{z}}^j$'s). The quantities $\{\alpha_{i,j}\}$, $\{\beta_{i,j}\}$, $\{\gamma_i\}$ represent the coefficients of this linear combination. The characterization (213) allows to subtract from $\mathbf{J}(\mathbf{Y})\mathbf{u}^k$ just the right Onsager terms, so that this difference equals a component in the direction of the signal (whose size is captured by μ_t) plus independent Gaussian noise (given by the linear combination of the $\tilde{\mathbf{z}}^j$'s via the coefficients $\{\theta_{i,j}\}$).

7 Numerics

For all experiments in this section, random instances of \mathbf{Y} are generated according to the model (2). The signal has Rademacher prior, i.e., i.i.d. entries $X_i^* \sim \frac{1}{2}(\delta_1 + \delta_{-1})$. The noise matrices $\mathbf{Z} = \mathbf{O}^\top \mathbf{D} \mathbf{O}$ are generated by first drawing N i.i.d. eigenvalues $(D_i)_{i \leq N}$ according to the density (14), and then multiplying from left and right the diagonal matrix of eigenvalues \mathbf{D} by a random Haar distributed orthogonal matrix \mathbf{O} sampled independently for each realization. As mentioned at the end of Section 1.4, the results are expected to be the same if we were to draw \mathbf{Z} according to the harder to sample³ measure (5).

³This can be done using the Dyson Brownian motion, see [93].

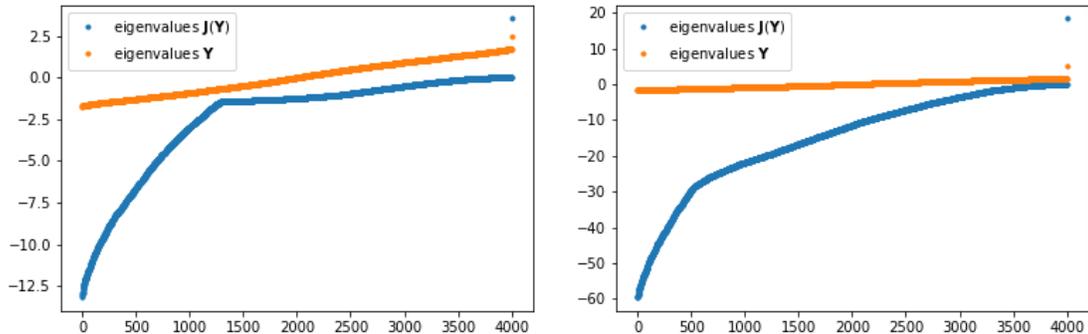


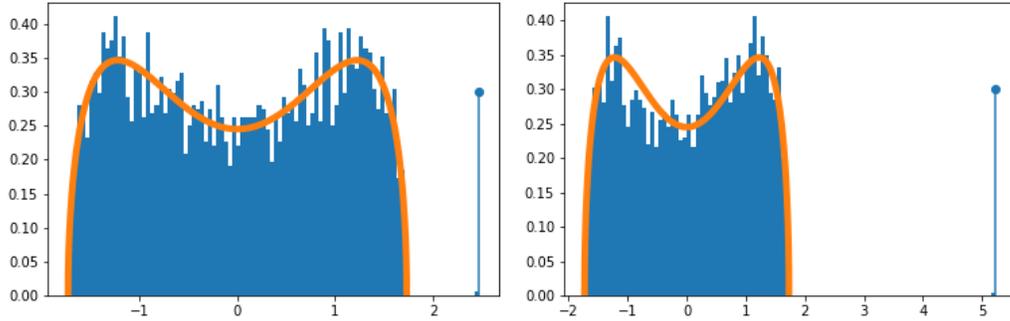
Figure 3: Ranked eigenvalues of the data matrix \mathbf{Y} (orange) and the optimally pre-processed matrix $\mathbf{J}(\mathbf{Y})$ (blue) for $N = 4000$ for (left) $\lambda = 2$ and (right) $\lambda = 5$. The gap between the largest detached eigenvalue on the extreme right and the second highest one is much bigger for the pre-processed matrix. Moreover, all the eigenvalues of $\mathbf{J}(\mathbf{Y})$ in its non-informative bulk are negative.

7.1 Spectral properties of the pre-processed matrix $\mathbf{J}(\mathbf{Y})$

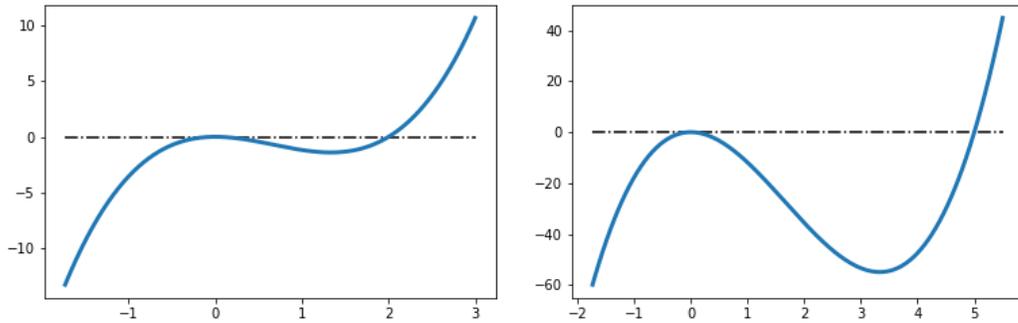
Let us discuss the effect on the spectrum of \mathbf{Y} that has the application of the optimal pre-processing function $J(\cdot)$; clearly, this function does not influence the eigenvectors of \mathbf{Y} which therefore has the same basis as $\mathbf{J}(\mathbf{Y})$. From Figures 3 and 4, the effect is clear: the function J (Figure 4, middle plots (b)) “cleans” the eigenvalues of the data \mathbf{Y} (Figure 4, upper plots (a)) by shifting the non-informative bulk eigenvalues of \mathbf{Y} to negative values, while the largest, informative, eigenvalue is further separated from the bulk. This results in the histograms (Figure 4, lower plots (c)) for the processed data $\mathbf{J}(\mathbf{Y})$. It thus become much easier to distinguish the informative eigenvalue, which may be of interest for smaller instances where the finite-size effects are stronger.

7.2 BAMP improves over the existing AMP and matches the replica prediction for the MMSE

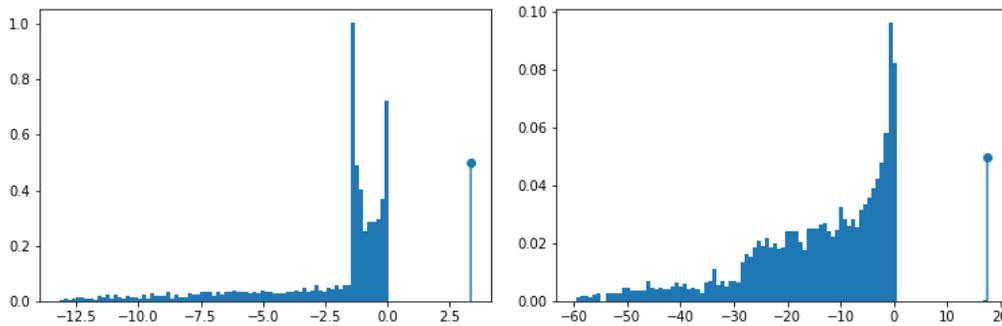
The plots of Figure 5 consider the quartic ensemble discussed in Section 1.3 for three values of the parameter μ , namely, $\mu \in \{0, 0.5, 1\}$ (recall $\gamma = \gamma(\mu)$ is fixed by relation (17)). Again, in all cases the signal \mathbf{X}^* is assumed to have a Rademacher prior. The estimators of the spike $\mathbf{X}^* \mathbf{X}^{*\top}$ are compared in terms of the MSE (y -axis) achieved



(a) Empirical spectral density of \mathbf{Y} . The largest, informative, eigenvalue is emphasized.



(b) The function $J(x) = \mu\lambda x - \gamma\lambda^2 x^2 + \gamma\lambda x^3$ with $(\mu = 0, \gamma(0) = 16/27)$ is used to optimally pre-process the (eigenvalues of the) data \mathbf{Y} and obtain $\mathbf{J}(\mathbf{Y})$. The dashed curve indicates 0. By comparison with the plots (a) above, we understand that the noise bulk will be pushed to negative values, while the spike towards the right, which results in a “cleaning” effect.



(c) Empirical spectral density of the pre-processed matrix $\mathbf{J}(\mathbf{Y})$. The largest eigenvalue is emphasized and well separated from the negative bulk by the application of $J(x)$.

Figure 4: Effect of the optimal pre-processing $J(x)$ on the eigenvalues of \mathbf{Y} . All experiments are for the most structured noise ensemble $(\mu = 0, \gamma(0) = 16/27)$ and $N = 4000$. The left column corresponds to $\lambda = 2$, while the right column to $\lambda = 5$.

at the fixed point, as a function of the SNR λ (x -axis). All algorithms are run for $N = 8000$ and the results are averaged over $n_{\text{trials}} = 50$ independent trials; the state evolution recursions (and the replica prediction as well) correspond to $N \rightarrow \infty$. We compare the following inference procedures:

- In black, we plot the replica prediction (80), obtained as the fixed point of (74)–(79).
- In red, we plot the performance of the BAMP algorithm described in Section 6, where g_{t+1} is the posterior mean denoiser (212). More specifically, the red line corresponds to the fixed point of the MSE given by the state evolution recursion discussed in Section 6.2 (cf. (211)), and the red stars denote the MSE obtained by running the BAMP algorithm (191).
- In blue, we plot the performance of the AMP proposed in [41]. More specifically, the blue line corresponds to the fixed point of the MSE (93) obtained by choosing the posterior mean denoiser with a single-step memory term (92). The blue diamonds denote the MSE obtained by running the AMP (90) with this single-step denoiser.
- Finally, the green squares denote the MSE obtained by running the AMP in [102], which consists in employing the following posterior mean denoiser with a multi-step memory term in the iteration (90):

$$h_{t+1}(f_1, \dots, f_t) = \mathbb{E}[X^* \mid (F_1, \dots, F_t) = (f_1, \dots, f_t)]. \quad (214)$$

We note that all algorithms converge rapidly: 10 iterations are sufficient to reach the corresponding fixed points. A few remarks concerning the numerical results displayed in Figure 5 are now in order:

- In all settings, the fixed point of the BAMP state evolution (in red) matches the replica prediction (in black). This is a strong numerical evidence supporting our conjecture that the proposed BAMP algorithm is Bayes-optimal. These theoretical curves for $N \rightarrow \infty$ are also remarkably close to the MSE achieved by the BAMP algorithm (191) at $N = 8000$.
- When $\mu = 0$, i.e., the noise is sufficiently far from being independent Gaussian, there is a clear performance gap between our proposed BAMP (in red) and the existing AMP algorithms [41, 102] (single-step denoiser in blue, and multi-step in green). As predicted by our theory, this gap is reduced for $\mu = 0.5$, and all curves collapse for $\mu = 1$.

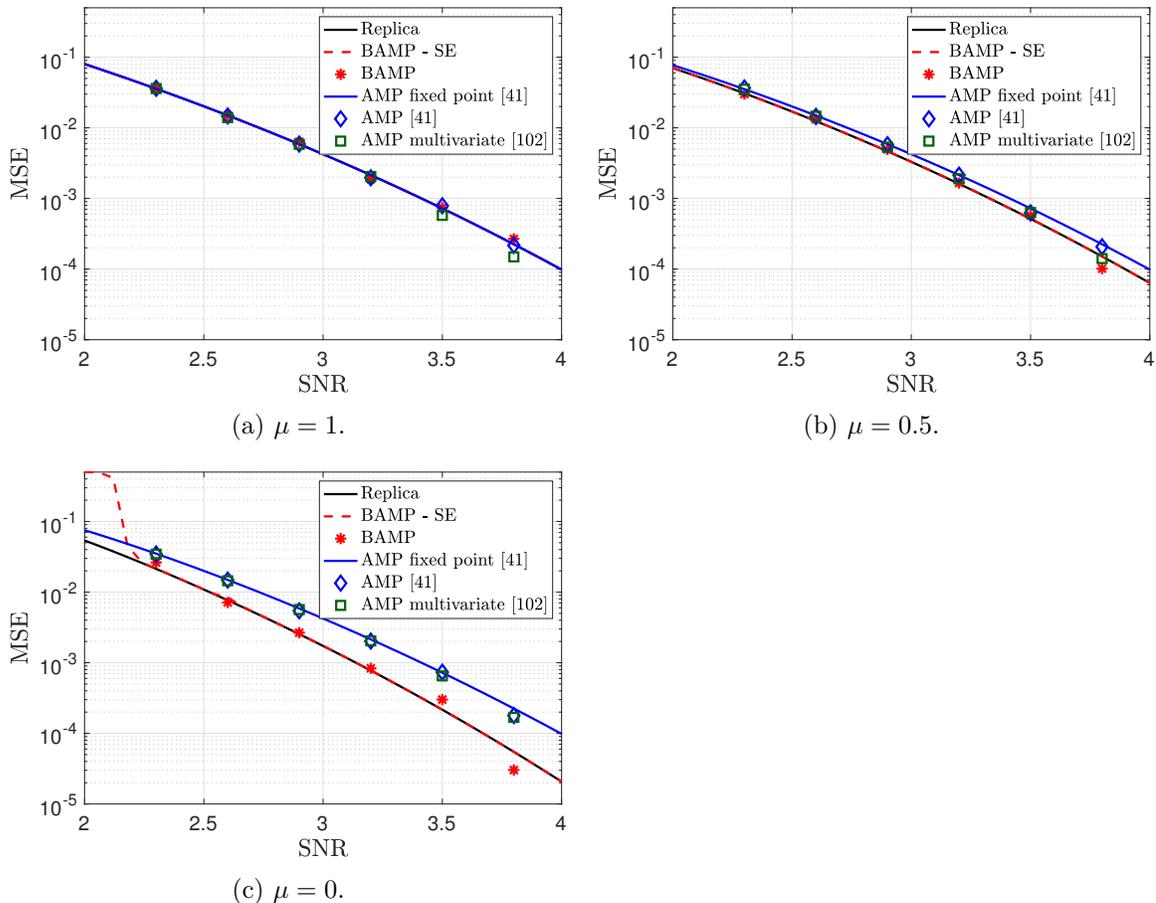


Figure 5: Performance comparison between the replica prediction for the MMSE (in black), the proposed BAMP (in red), and the existing AMP [41, 102] (in blue and green). BAMP matches the Bayes-optimal MSE predicted via the replica method, and it outperforms the existing AMP when the noise is not Gaussian. This improvement is more evident as the noise distribution gets further from a Wigner distribution. Taken all together, these numerical results provide an empirical confirmation of the (Bayes-)optimality of the proposed BAMP algorithm.

- Finally, we note that the BAMP algorithm exhibits a numerical instability for low SNR. More specifically, when $\mu = 0$ and $\lambda = 2.3$, 5 out of the 50 trials of the iteration (191) do not reach the fixed point of state evolution (and are therefore discarded). Furthermore, by inspecting Figure 5c, one notices that the curve representing the BAMP state evolution detaches from the replica prediction as

the SNR get smaller than 2.3. As expected, considering an initialization closer to the fixed point mitigates the issue. This numerical instability is likely due to the state evolution of BAMP corresponding to the recursion of an auxiliary AMP that *triples* the number of iterations. This fact leads to an amplification of the numerical errors.

Let us re-emphasize that all these results hold in the Bayesian-optimal setting where all hyper-parameters of the model are known and optimally used. In practical situations this may not be the case. In particular the statistical properties of the correlated noise \mathbf{Z} may be only partially known, preventing one to obtain the coefficients (c_k) defining the optimal pre-processing of the data $\mathbf{J}(\mathbf{Y}) = \sum_{k \leq K} c_k \mathbf{Y}^k$ as done in Section 5. In Appendix A we provide a learning procedure based on expectation maximization to overcome this issue and which can be of help to practitioners aiming at using BAMP in more realistic situations. Its testing is left for future work.

Acknowledgments

J. Barbier acknowledges funding from the European Research Council for the project “Computational hardness of representation learning” — CHORAL (grant agreement n. 101039794). M. Mondelli was supported by the 2019 Lopez-Loreta prize.

A Learning the optimal pre-processing $\mathbf{J}(\mathbf{Y})$ by expectation maximization

Until now we have assumed that we are in the Bayesian-optimal setting where, in particular, the polynomial potential V defining the noise statistics is completely known and correctly exploited. As seen from section 5.2, given a potential V we could deduce from the AdaTAP formalism an optimal polynomial

$$\mathbf{J} = \mathbf{J}(\mathbf{Y}) = \sum_{k \leq K} c_k \mathbf{Y}^k$$

to pre-process the data \mathbf{Y} before using it in AMP. The Bayes-optimal case corresponds to matrix (161), i.e., $\mathbf{J} = c_1 \mathbf{Y} + c_2 \mathbf{Y}^2 + c_3 \mathbf{Y}^3$ with $\mathbf{c} = (\mu\lambda, -\gamma\lambda^2, \gamma\lambda)$.

We here consider an extension of the previously derived AMP to a case where V is not known and therefore the optimal \mathbf{J} cannot be deduced by the AdaTAP approach as we did in Section 5.2. What is known instead is an upper bound on the order of V . In the base-case model studied in details in the present paper the order

is four. The procedure we propose below will not be tested numerically yet, but we believe it may be of interest to practitioners eager to improve the Bayes-optimal AMP for more practical settings than the specific ones studied here.

To directly learn the coefficients $(c_k)_{k \leq K}$ from the data, we propose to use an expectation maximization (EM) approach, with a routine inside AMP performing the parameter estimation by maximizing the current estimate of the free entropy, i.e., of the log-likelihood of the observed data $\ln P(\mathbf{Y} \mid \mathbf{c})$.

Assume that, at the AMP iterate t , the current estimate of the unknown coefficients $\mathbf{c} = (c_k)_{k \leq K}$ is $\mathbf{c}(t) = (c_k(t))_{k \leq K}$, the AMP estimate of the marginal means is $\mathbf{m}(t)$, and of the Onsager reaction term is $\bar{V}(t)$ (which is related to the set of Onsager coefficients, see Section 6.2). Let also the data matrix polynomial currently used by AMP be

$$\mathbf{J}(t) := \sum_{k \leq K} c_k(t) \mathbf{Y}^k.$$

From the analysis of Section 5.3 we know that at the saddle point we can safely replace the Onsager reaction term V_i by \bar{V} in the AdaTAP equations. When this is plugged back into (182), this identity implies that also the following concentration is consistently valid: $\mathbb{E}(\tau_i - m_i^2) = \tau_i - m_i^2$, which is also equal by exchangeability to $N^{-1} \sum_{i \leq N} \mathbb{E}(\tau_i - m_i^2)$. Let us call $\bar{\chi}(t)$ the AMP estimate of the variance $\mathbb{E}(\tau_i - m_i^2)$. Applying these simplifications to the AMP iterates we get that the matrix $\boldsymbol{\Omega}(t) := \text{diag}(\mathbf{V}(t) + (\boldsymbol{\tau}(t) - \mathbf{m}(t)^2)^{-1})$ can be simplified as

$$\boldsymbol{\Omega}(t) = (\bar{V}(t) + \bar{\chi}(t)^{-1}) I_N.$$

From section 5.1 the AdaTAP approximation to the free entropy at iterate t then reads, using these simplifications, as

$$\begin{aligned} \Phi_N(t, \mathbf{c}(t)) &= \frac{1}{2} \mathbf{m}(t)^\top \mathbf{J}(t) \mathbf{m}(t) + \frac{1}{2} \ln \det (\boldsymbol{\Omega}(t) - \mathbf{J}(t)) - \frac{1}{2} \bar{V}(t) \sum_{i \leq N} m_i(t)^2 + \frac{1}{2} \bar{\chi}(t) \\ &\quad - \sum_{i \leq N} \ln \int dP_X(x) \exp \left(\frac{1}{2} \bar{V}(t) x^2 + ((\mathbf{J}(t) \mathbf{m}(t))_i - \bar{V}(t) m_i(t)) x \right). \end{aligned} \quad (215)$$

The free entropy $\Phi_N(t, \mathbf{c}(t))$ is the current best approximation to the marginal log-likelihood of the data $\ln P(\mathbf{Y} \mid \mathbf{c})$, which we thus aim at maximizing with respect to the unknown parameters, all other quantities being fixed at their current values:

$$\begin{aligned} \partial_{c_k} \Phi_N|_{t, \mathbf{c}(t)} &= \mathbf{m}(t)^\top \mathbf{Y}^k \left(\frac{1}{2} \mathbf{m}(t) - \eta(\mathbf{J}(t), \mathbf{m}(t), \bar{V}(t)) \right) \\ &\quad - \frac{1}{2} \text{Tr}(\mathbf{Y}^k (\boldsymbol{\Omega}(t) - \mathbf{J}(t))^{-1}), \end{aligned} \quad (216)$$

where we used (163) and the notation $\eta(\mathbf{J}(t), \mathbf{m}(t), \bar{V}(t)) = (\eta_i(\mathbf{J}(t), \mathbf{m}(t), \bar{V}(t)))_{i \leq N}$. Because \mathbf{J} is diagonalizable in the same basis as the data \mathbf{Y} , the eigenvalues of which are denoted $\sigma_i = \sigma_i(\mathbf{Y})$, we have

$$\text{Tr}(\mathbf{Y}^k(\boldsymbol{\Omega}(t) - \mathbf{J}(t))^{-1}) = \sum_{i \leq N} \frac{\sigma_i^k}{\bar{V}(t) + \bar{\chi}(t)^{-1} - \sum_{\ell \leq K} c_\ell(t) \sigma_i^\ell}. \quad (217)$$

Then

$$\begin{aligned} \partial_{c_k} \Phi_N|_{t, \mathbf{c}(t)} &= \mathbf{m}(t)^\top \mathbf{Y}^k \left(\frac{1}{2} \mathbf{m}(t) - \eta(\mathbf{J}(t), \mathbf{m}(t), \bar{V}(t)) \right) \\ &\quad - \frac{1}{2} \sum_{i \leq N} \frac{\sigma_i^k}{\bar{V}(t) + \bar{\chi}(t)^{-1} - \sum_{\ell \leq K} c_\ell(t) \sigma_i^\ell}. \end{aligned} \quad (218)$$

We aim at maximizing the free entropy so given a learning rate $\zeta > 0$ the learning rule finally reads

$$c_k(t+1) = c_k(t) + \zeta \partial_{c_k} \Phi_N|_{t, \mathbf{c}(t)}. \quad (219)$$

B Proofs for BAMP

B.1 Auxiliary AMP

The iterates of the auxiliary AMP are denoted by $\tilde{\mathbf{z}}^t, \tilde{\mathbf{u}}^t \in \mathbb{R}^N$, and they are computed as follows, for $t \geq 1$:

$$\tilde{\mathbf{z}}^t = \mathbf{Z} \tilde{\mathbf{u}}^t - \sum_{i=1}^t \bar{\mathbf{b}}_{t,i} \tilde{\mathbf{u}}^i, \quad \tilde{\mathbf{u}}^{t+1} = \tilde{h}_{t+1}(\tilde{\mathbf{z}}^1, \dots, \tilde{\mathbf{z}}^t, \mathbf{u}^1, \mathbf{X}^*). \quad (220)$$

The iteration (220) is initialized with $\tilde{\mathbf{u}}^1 = \mathbf{u}^1$, where \mathbf{u}^1 satisfies (193). For $t \geq 1$, the functions $\tilde{h}_{t+1} : \mathbb{R}^{t+2} \rightarrow \mathbb{R}$ are applied component-wise, and they are recursively defined as

$$\begin{aligned} \tilde{h}_{K(t-1)+1+\ell}(z_1, \dots, z_{K(t-1)+\ell}, u_1, x^*) &= z_{K(t-1)+\ell} + (\tilde{\mathbf{B}}_{K(t-1)+\ell})_{K(t-1)+\ell, 1} u_1 \\ &\quad + \sum_{i=2}^{K(t-1)+\ell} (\tilde{\mathbf{B}}_{K(t-1)+\ell})_{K(t-1)+\ell, i} \tilde{h}_i(z_1, \dots, z_{i-1}, u_1, x^*) + \tilde{\mu}_{K(t-1)+\ell} x^*, \quad \ell \in [K-1], \\ \tilde{h}_{Kt+1}(z_1, \dots, z_{Kt}, u_1, x^*) &= g_{t+1} \left(\mu_t x^* + \sum_{i=1}^{Kt} \theta_{t,i} z_i \right). \end{aligned} \quad (221)$$

The idea is that the choice (221) for the denoisers $\{\tilde{h}_{t+1}\}_{t \geq 1}$ ensures that $\tilde{\mathbf{u}}^{K(t-1)+\ell}$ tracks the quantity $\mathbf{Y}^{\ell-1} \mathbf{u}^t$ for $\ell \in [K]$ and $t \geq 1$, where $\{\mathbf{u}^t\}$ are the iterates of the AMP iteration (191) we are interested in analyzing.

In (221), g_{t+1} is the denoiser of the AMP (191). The parameters $(\tilde{\mathbf{B}}_{K(t-1)+\ell}, \tilde{\mu}_{K(t-1)+\ell}, \mu_t, \theta_{t,i})$ come from the state evolution recursion detailed in Section 6.2: $\tilde{\mathbf{B}}_{K(t-1)+\ell}$ is given by (202), $\tilde{\mu}_{K(t-1)+\ell}$ by (199), μ_t by (207) and $\theta_{t,i}$ by (208). We now discuss how to obtain the coefficients $\{\bar{\mathbf{b}}_{t,i}\}_{i=1}^t$ needed in (220). Let us define the matrix $\bar{\Phi}_t \in \mathbb{R}^{t \times t}$ as

$$(\bar{\Phi}_t)_{i,j} = 0, \quad \text{for } i \leq j, \quad (\bar{\Phi}_t)_{i,j} = \langle \partial_j \tilde{\mathbf{u}}^i \rangle, \quad \text{for } i > j, \quad (222)$$

where, for $j < i$, the vector $\langle \partial_j \tilde{\mathbf{u}}^i \rangle \in \mathbb{R}^N$ denotes the partial derivative of $\tilde{h}_i : \mathbb{R}^{i+1} \rightarrow \mathbb{R}$ with respect to the j -th input (applied component-wise). Then, the vector $(\bar{\mathbf{b}}_{t,1}, \dots, \bar{\mathbf{b}}_{t,t})$ is given by the last row of the matrix $\bar{\mathbf{B}}_t \in \mathbb{R}^{t \times t}$ defined as

$$\bar{\mathbf{B}}_t = \sum_{j=0}^{t-1} \kappa_{j+1} \bar{\Phi}_t^j. \quad (223)$$

where $\{\kappa_k\}_{k \geq 1}$ denotes the sequence of free cumulants associated to the matrix \mathbf{Z} .

B.2 State evolution of auxiliary AMP

Using Theorem 2.3 in [102], we provide a state evolution result for the auxiliary AMP (220). In particular, we show in Proposition 2 that the joint empirical distribution of $(\tilde{\mathbf{z}}^1, \dots, \tilde{\mathbf{z}}^t)$ converges to a t -dimensional Gaussian $\mathcal{N}(\mathbf{0}, \hat{\Sigma}_t)$.

The covariance matrices $\{\hat{\Sigma}_t\}_{t \geq 1}$ are defined recursively, starting with $\hat{\Sigma}_1 = \bar{\kappa}_2 \mathbb{E}[U_1^2]$, where U_1 is defined in (193). Given $\hat{\Sigma}_t$, let

$$\begin{aligned} (\hat{Z}_1, \dots, \hat{Z}_t) &\sim \mathcal{N}(\mathbf{0}, \hat{\Sigma}_t) \text{ and independent of } (X^*, U_1), \\ \hat{U}_s &= \tilde{h}_s(\hat{Z}_1, \dots, \hat{Z}_{s-1}, U_1, X^*), \quad \text{for } s \in \{2, \dots, t+1\}, \end{aligned} \quad (224)$$

where \tilde{h}_s is defined via (221) and we set $\hat{U}_1 = U_1$. Let $\hat{\Phi}_{t+1}, \hat{\Delta}_{t+1} \in \mathbb{R}^{(t+1) \times (t+1)}$ be matrices with entries given by

$$\begin{aligned} (\hat{\Phi}_{t+1})_{i,j} &= 0, \quad \text{for } i \leq j, & (\hat{\Phi}_{t+1})_{i,j} &= \mathbb{E}[\partial_j \hat{U}_i], \quad \text{for } i > j, \\ (\hat{\Delta}_{t+1})_{i,j} &= \mathbb{E}[\hat{U}_i \hat{U}_j], \quad 1 \leq i, j \leq t+1, \end{aligned} \quad (225)$$

where $\partial_j \hat{U}_i$ denotes the partial derivative $\partial_{\hat{Z}_j} \tilde{h}_i(\hat{Z}_1, \dots, \hat{Z}_{i-1}, U_1, X)$. Then, we compute the covariance matrix $\hat{\Sigma}_{t+1}$ as

$$\hat{\Sigma}_{t+1} = \sum_{j=0}^{2t} \bar{\kappa}_{j+2} \sum_{i=0}^j (\hat{\Phi}_{t+1})^i \hat{\Delta}_{t+1} (\hat{\Phi}_{t+1}^\top)^{j-i}. \quad (226)$$

It can be verified that the $t \times t$ top left sub-matrix of $\hat{\Sigma}_{t+1}$ is given by $\hat{\Sigma}_t$.

Proposition 2 (State evolution for auxiliary AMP). *Consider the auxiliary AMP in (220) and the state evolution random variables defined in (224). Let $\tilde{\psi} : \mathbb{R}^{2t+2} \rightarrow \mathbb{R}$ be a PL(2) function. Then, for each $t \geq 1$, we almost surely have*

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \tilde{\psi}(\tilde{z}_i^1, \dots, \tilde{z}_i^t, \tilde{u}_i^1, \dots, \tilde{u}_i^{t+1}, X_i^*) \\ = \mathbb{E}[\tilde{\psi}(\hat{Z}_1, \dots, \hat{Z}_t, \hat{U}_1, \dots, \hat{U}_{t+1}, X^*)]. \end{aligned} \quad (227)$$

Equivalently, as $N \rightarrow \infty$, almost surely:

$$(\tilde{z}^1, \dots, \tilde{z}^t, \tilde{\mathbf{u}}^1, \dots, \tilde{\mathbf{u}}^{t+1}, \mathbf{X}^*) \xrightarrow{W_2} (\hat{Z}_1, \dots, \hat{Z}_t, \hat{U}_1, \dots, \hat{U}_{t+1}, X^*). \quad (228)$$

Furthermore,

$$(\hat{Z}_1, \dots, \hat{Z}_t, \hat{U}_1, \dots, \hat{U}_{t+1}, X^*) \stackrel{d}{=} (\tilde{Z}_1, \dots, \tilde{Z}_t, \tilde{U}_1, \dots, \tilde{U}_{t+1}, X^*), \quad (229)$$

where $(\tilde{Z}_1, \dots, \tilde{Z}_t, \tilde{U}_1, \dots, \tilde{U}_{t+1}, X^*)$ are obtained via (196)–(198).

Proof. The result follows from Theorem 2.3 in [102]. In fact, Assumption 2.1 of [102] holds because of the model assumptions on \mathbf{Z} , Assumption 2.2(a) holds because $(\mathbf{X}^*, \tilde{\mathbf{u}}^1) = (\mathbf{X}^*, \mathbf{u}^1) \xrightarrow{W_2} (X^*, U_1)$ from (193), and Assumption 2.2(b) follows from the definition of \tilde{h}_{t+1} in (221) and the fact that g_{t+1} is continuously differentiable and Lipschitz. As the auxiliary AMP in (220) is of the standard form for which the state evolution result of Theorem 2.3 in [102] holds, we readily obtain (228). The equivalence between (228) and (227) follows from [44, Corollary 7.21]. Finally, by inspecting the state evolution recursions (196)–(198) and (224) giving $(\tilde{Z}_1, \dots, \tilde{Z}_t, \tilde{U}_1, \dots, \tilde{U}_{t+1}, X^*)$ and $(\hat{Z}_1, \dots, \hat{Z}_t, \hat{U}_1, \dots, \hat{U}_{t+1}, X^*)$ respectively, (229) is readily obtained. \square

Proposition 2 gives that the state evolution recursion discussed in Section 6.2 (cf. (196)–(198)) coincides with the state evolution tracking the iterates of the auxiliary AMP algorithm (220). In particular, $\hat{\Delta}_{3t} = \tilde{\Delta}_{3t}$, $\tilde{\Phi}_{3t} = \hat{\Phi}_{3t}$, and $\tilde{\Sigma}_{3t} = \hat{\Sigma}_{3t}$. Furthermore, in the proof of Theorem 1 contained in Appendix B.3, we will show that $\tilde{\mathbf{B}}_{3t} \rightarrow \hat{\mathbf{B}}_{3t}$ as $N \rightarrow \infty$.

B.3 Proof of Theorem 1

We start by presenting a useful technical lemma.

Lemma 3. *Let $F : \mathbb{R}^t \rightarrow \mathbb{R}$ be a Lipschitz function, and let $\partial_k F$ denote its derivative with respect to the k -th argument, for $1 \leq k \leq t$. Assume that $\partial_k F$ is continuous almost everywhere in the k -th argument, for each k . Let $(V_1^{(m)}, \dots, V_t^{(m)})$ be a sequence of random vectors in \mathbb{R}^t converging in distribution to the random vector (V_1, \dots, V_t) as $m \rightarrow \infty$. Furthermore, assume that the distribution of (V_1, \dots, V_t) is absolutely continuous with respect to the Lebesgue measure. Then,*

$$\lim_{m \rightarrow \infty} \mathbb{E}[\partial_k F(V_1^{(m)}, \dots, V_t^{(m)})] = \mathbb{E}[\partial_k F(V_1, \dots, V_t)], \quad 1 \leq k \leq t. \quad (230)$$

The result was proved for $t = 2$ in [22, Lemma 6]. The proof for $t > 2$ is basically the same, see also [44, Lemma 7.14]. At this point, we are ready to give the proof of Theorem 1.

Proof of Theorem 1. We show that, for any PL(2) function $\psi : \mathbb{R}^{2t+2} \rightarrow \mathbb{R}$, the following limit holds almost surely for $t \geq 1$:

$$\begin{aligned} \lim_{N \rightarrow \infty} \left| \frac{1}{N} \sum_{i=1}^N \psi(u_i^1, u_i^2, \dots, u_i^{t+1}, f_i^1, f_i^2, \dots, f_i^t, X_i^*) \right. \\ \left. - \frac{1}{N} \sum_{i=1}^N \psi(\tilde{u}_i^1, \tilde{u}_i^{K+1}, \dots, \tilde{u}_i^{Kt+1}, \tilde{f}_i^1, \tilde{f}_i^2, \dots, \tilde{f}_i^t, X_i^*) \right| = 0, \end{aligned} \quad (231)$$

where we have defined for $s \in \{1, \dots, t\}$,

$$\tilde{\mathbf{f}}^s = \mu_s \mathbf{X}^* + \sum_{i=1}^{Ks} \theta_{s,i} \tilde{\mathbf{z}}^i. \quad (232)$$

From here till the end of the argument, all the limits hold almost surely, and we use C to denote a generic positive constant, which can change from line to line and is

independent of N . By using that ψ is pseudo-Lipschitz, we have that

$$\begin{aligned}
& \left| \frac{1}{N} \sum_{i=1}^N \psi(u_i^1, u_i^2, \dots, u_i^{t+1}, f_i^1, f_i^2, \dots, f_i^t, X_i^*) \right. \\
& \quad \left. - \frac{1}{N} \sum_{i=1}^N \psi(\tilde{u}_i^1, \tilde{u}_i^{K+1}, \dots, \tilde{u}_i^{Kt+1}, \tilde{f}_i^1, \tilde{f}_i^2, \dots, \tilde{f}_i^t, X_i^*) \right| \\
& \leq \frac{C}{N} \sum_{i=1}^N \left(1 + |X_i^*| + 2|u_i^1| + \sum_{k=1}^t \left(|f_i^k| + |\tilde{f}_i^k| + |u_i^{k+1}| + |\tilde{u}_i^{Kk+1}| \right) \right) \\
& \quad \cdot \left(\sum_{k=1}^t \left(|f_i^k - \tilde{f}_i^k|^2 + |u_i^{k+1} - \tilde{u}_i^{Kk+1}|^2 \right) \right)^{1/2} \\
& \leq C(4t+3) \left[1 + \frac{\|\mathbf{X}^*\|^2}{N} + \sum_{k=1}^t \left(\frac{\|\mathbf{f}^k\|^2}{N} + \frac{\|\tilde{\mathbf{f}}^k\|^2}{N} + \frac{\|\mathbf{u}^{k+1}\|^2}{N} + \frac{\|\tilde{\mathbf{u}}^{Kk+1}\|^2}{N} \right) \right]^{1/2} \\
& \quad \cdot \left(\sum_{k=1}^t \left(\frac{\|\mathbf{f}^k - \tilde{\mathbf{f}}^k\|^2}{N} + \frac{\|\mathbf{u}^{k+1} - \tilde{\mathbf{u}}^{Kk+1}\|^2}{N} \right) \right)^{1/2}, \tag{233}
\end{aligned}$$

where the last step uses twice Cauchy-Schwarz inequality. We now inductively show that as $N \rightarrow \infty$: (i) each of the terms in the last line of (233) converges to zero, and (ii) the terms within the square brackets in (233) all converge to finite, deterministic limits. To achieve this goal, we will also show that, for $k \in [t]$ and $\ell \in [K-1]$,

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{Y}^\ell \mathbf{u}^k - \tilde{\mathbf{u}}^{K(k-1)+1+\ell}\|^2}{N} = 0, \tag{234}$$

$$\lim_{N \rightarrow \infty} \frac{\|\tilde{\mathbf{u}}^{K(k-1)+1+\ell} - \sum_{j=1}^{K(k-1)+\ell} \alpha_{K(k-1)+1+\ell, j} \tilde{\mathbf{z}}^j - \sum_{j=1}^k \beta_{K(k-1)+1+\ell, j} \mathbf{u}^j - \gamma_{K(k-1)+1+\ell} \mathbf{X}^*\|^2}{N} = 0. \tag{235}$$

The limit (234) formalizes the idea discussed in Section 6.1 (see (192)) that the iterate $\tilde{\mathbf{u}}^{K(k-1)+1+\ell}$ of the auxiliary AMP tracks the quantity $\mathbf{Y}^\ell \mathbf{u}^k$, where \mathbf{u}^k is the iterate of the AMP we wish to analyze, up to an $o_N(1)$ error. The limit (235) formalizes the interpretation of the coefficients $\{\alpha_{i,j}\}$, $\{\beta_{i,j}\}$, $\{\gamma_i\}$ provided at the end of Section 6.2 (see (213)).

Base case ($t = 1$). We have that

$$\begin{aligned}\mathbf{Y}\mathbf{u}^1 - \tilde{\mathbf{u}}^2 &= \mathbf{Z}\mathbf{u}^1 + \lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^1 \rangle}{N} \mathbf{X}^* - \tilde{\mathbf{z}}^1 - (\tilde{\mathbf{B}}_1)_{1,1} \mathbf{u}^1 - \tilde{\mu}_1 \mathbf{X}^* \\ &= \left(\lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^1 \rangle}{N} - \tilde{\mu}_1 \right) \mathbf{X}^* + (\bar{\mathbf{b}}_{1,1} - (\tilde{\mathbf{B}}_1)_{1,1}) \mathbf{u}^1,\end{aligned}\tag{236}$$

where the first equality uses the definition of \mathbf{Y} and of \tilde{h}_2 (see (221)), and the second equality uses (220) and that $\tilde{\mathbf{u}}^1 = \mathbf{u}^1$. Hence, by triangle inequality,

$$\begin{aligned}\frac{\|\mathbf{Y}\mathbf{u}^1 - \tilde{\mathbf{u}}^2\|^2}{N} &\leq 2 \left(\lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^1 \rangle}{N} - \tilde{\mu}_1 \right)^2 \frac{\|\mathbf{X}^*\|^2}{N} + 2 (\bar{\mathbf{b}}_{1,1} - (\tilde{\mathbf{B}}_1)_{1,1})^2 \frac{\|\mathbf{u}^1\|^2}{N} \\ &\leq C \left(\left(\lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^1 \rangle}{N} - \tilde{\mu}_1 \right)^2 + (\bar{\mathbf{b}}_{1,1} - (\tilde{\mathbf{B}}_1)_{1,1})^2 \right),\end{aligned}\tag{237}$$

where the last inequality uses that $(\mathbf{X}^*, \mathbf{u}^1)$ converges in W_2 to a pair of random variables with finite second moments. As $\tilde{\mu}_1 = \lambda\epsilon$ (cf. (195)), we have

$$\lim_{N \rightarrow \infty} \lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^1 \rangle}{N} = \lambda \mathbb{E}[U_1 X^*] = \lambda\epsilon = \tilde{\mu}_1.\tag{238}$$

Furthermore, note that $(\tilde{\mathbf{B}}_1)_{1,1} = \bar{\kappa}_1$ (cf. (195)) and $\bar{\mathbf{b}}_{1,1} = \kappa_1$ (cf. (223)). Hence, by the model assumptions, as $N \rightarrow \infty$, $\kappa_1 \rightarrow \bar{\kappa}_1$ and, therefore, $\bar{\mathbf{b}}_{1,1} \rightarrow (\tilde{\mathbf{B}}_1)_{1,1}$. By combining this observation with (237) and (238), we obtain that (234) holds for $k = 1$ and $\ell = 1$.

By using (204)–(206), we readily obtain that $\alpha_{2,1} = 1$, $\beta_{2,1} = (\tilde{\mathbf{B}}_1)_{1,1}$ and $\gamma_2 = \tilde{\mu}_1$. Hence, by using the definition (221) of \tilde{h}_2 , we obtain that (235) holds for $k = 1$ and $\ell = 1$.

Next, by using the definitions of \mathbf{Y} , of the auxiliary AMP (220) and of \tilde{h}_3 (cf. (221)), we have

$$\begin{aligned}\mathbf{Y}^2 \mathbf{u}^1 - \tilde{\mathbf{u}}^3 &= \mathbf{Y}(\mathbf{Y}\mathbf{u}^1 - \tilde{\mathbf{u}}^2) + \mathbf{Y}\tilde{\mathbf{u}}^2 - \tilde{\mathbf{z}}^2 - (\tilde{\mathbf{B}}_2)_{2,1} \mathbf{u}^1 - (\tilde{\mathbf{B}}_2)_{2,2} \tilde{\mathbf{u}}^2 - \tilde{\mu}_2 \mathbf{X}^* \\ &= \mathbf{Y}(\mathbf{Y}\mathbf{u}^1 - \tilde{\mathbf{u}}^2) + \mathbf{Z}\tilde{\mathbf{u}}^2 - \tilde{\mathbf{z}}^2 - (\tilde{\mathbf{B}}_2)_{2,1} \mathbf{u}^1 - (\tilde{\mathbf{B}}_2)_{2,2} \tilde{\mathbf{u}}^2 + \left(\lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^2 \rangle}{N} - \tilde{\mu}_2 \right) \mathbf{X}^* \\ &= \mathbf{Y}(\mathbf{Y}\mathbf{u}^1 - \tilde{\mathbf{u}}^2) + (\bar{\mathbf{b}}_{2,1} - (\tilde{\mathbf{B}}_2)_{2,1}) \mathbf{u}^1 + (\bar{\mathbf{b}}_{2,2} - (\tilde{\mathbf{B}}_2)_{2,2}) \tilde{\mathbf{u}}^2 + \left(\lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^2 \rangle}{N} - \tilde{\mu}_2 \right) \mathbf{X}^*.\end{aligned}\tag{239}$$

Hence, by triangle inequality,

$$\begin{aligned}
\frac{\|\mathbf{Y}^2 \mathbf{u}^1 - \tilde{\mathbf{u}}^3\|^2}{N} &\leq C \left(\frac{\|\mathbf{Y}(\mathbf{Y} \mathbf{u}^1 - \tilde{\mathbf{u}}^2)\|^2}{N} + (\bar{\mathbf{b}}_{2,1} - (\tilde{\mathbf{B}}_2)_{2,1})^2 \frac{\|\mathbf{u}^1\|^2}{N} \right. \\
&\quad \left. + (\bar{\mathbf{b}}_{2,2} - (\tilde{\mathbf{B}}_2)_{2,2})^2 \frac{\|\tilde{\mathbf{u}}^2\|^2}{N} + \left(\lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^2 \rangle}{N} - \tilde{\mu}_2 \right)^2 \frac{\|\mathbf{X}^*\|^2}{N} \right) \quad (240) \\
&:= C(T_1 + T_2 + T_3 + T_4).
\end{aligned}$$

Consider the first term. As \mathbf{Y} has bounded operator norm and (234) holds for $k = 1$ and $\ell = 1$, we have that $T_1 \rightarrow 0$ as $N \rightarrow \infty$.

Consider the second and third terms. The following chain of equalities holds

$$\lim_{N \rightarrow \infty} (\tilde{\Phi}_2)_{2,1} = \lim_{N \rightarrow \infty} \langle \partial_1 \tilde{\mathbf{u}}^2 \rangle = \mathbb{E}[\partial_1 \hat{U}_2] = \mathbb{E}[\partial_1 \tilde{U}_2] = (\tilde{\Phi}_2)_{2,1}. \quad (241)$$

Here, the first equality uses the definition (222); the second equality follows from Lemma 3, as $\tilde{\mathbf{u}}^2$ converges in W_2 (and therefore in distribution) to \tilde{U}_2 and $\partial_1 \tilde{U}_2$ is continuous; the third equality uses (229); and the fourth equality uses the definition of $(\tilde{\Phi}_2)_{2,1}$ in (201). By the model assumptions, as $N \rightarrow \infty$, $\kappa_j \rightarrow \bar{\kappa}_j$ for all j . Thus, by combining (241) with the definitions of $\tilde{\mathbf{B}}_2$ and $\tilde{\mathbf{B}}_2$ in (223) and (202), respectively, we conclude that, as $N \rightarrow \infty$, $\bar{\mathbf{b}}_{2,i} \rightarrow (\tilde{\mathbf{B}}_2)_{2,i}$ for $i \in \{1, 2\}$. By Proposition 2, $\|\tilde{\mathbf{u}}^2\|^2/N$ converges to a finite limit, hence we conclude that $T_2, T_3 \rightarrow 0$ as $N \rightarrow \infty$.

Consider the fourth term. Then,

$$\lim_{N \rightarrow \infty} \lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^2 \rangle}{N} = \lambda \mathbb{E}[X^* \tilde{U}_2] = \tilde{\mu}_2.$$

Here, the first equality uses Proposition 2 and the second equality uses the definition of $\tilde{\mu}_2$ in (199). As $\|\mathbf{X}^*\|^2/N = 1$, we conclude that $T_4 \rightarrow 0$ as $N \rightarrow \infty$. This proves that the RHS of (240) vanishes and gives that (234) holds for $k = 1$ and $\ell = 2$.

By using (204)–(206), we readily obtain that $\alpha_{3,1} = (\tilde{\mathbf{B}}_2)_{2,2}$, $\alpha_{3,2} = 1$, $\beta_{3,1} = (\tilde{\mathbf{B}}_2)_{2,1} + (\tilde{\mathbf{B}}_2)_{2,2} (\tilde{\mathbf{B}}_2)_{1,1}$ and $\gamma_3 = \tilde{\mu}_2 + \tilde{\mu}_1 (\tilde{\mathbf{B}}_2)_{2,2}$. Hence, by using the definition (221) of \tilde{h}_3 , we obtain that (235) holds for $k = 1$ and $\ell = 2$.

The proof of (234)–(235) for $k = 1$ and $\ell \in \{3, \dots, K - 1\}$ follows from similar

arguments. In particular, we write

$$\begin{aligned}
\mathbf{Y}^\ell \mathbf{u}^1 - \tilde{\mathbf{u}}^{1+\ell} &= \mathbf{Y}(\mathbf{Y}^{\ell-1} \mathbf{u}^1 - \tilde{\mathbf{u}}^\ell) + \mathbf{Y} \tilde{\mathbf{u}}^\ell - \tilde{\mathbf{z}}^\ell - (\tilde{\mathbf{B}}_\ell)_{\ell,1} \mathbf{u}^1 - \sum_{j=2}^{\ell} (\tilde{\mathbf{B}}_\ell)_{\ell,j} \tilde{\mathbf{u}}^j - \tilde{\mu}_\ell \mathbf{X}^* \\
&= \mathbf{Y}(\mathbf{Y}^{\ell-1} \mathbf{u}^1 - \tilde{\mathbf{u}}^\ell) + \mathbf{Z} \tilde{\mathbf{u}}^\ell - \tilde{\mathbf{z}}^\ell - (\tilde{\mathbf{B}}_\ell)_{\ell,1} \mathbf{u}^1 - \sum_{j=2}^{\ell} (\tilde{\mathbf{B}}_\ell)_{\ell,j} \tilde{\mathbf{u}}^j + \left(\lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^\ell \rangle}{N} - \tilde{\mu}_\ell \right) \mathbf{X}^* \\
&= \mathbf{Y}(\mathbf{Y}^{\ell-1} \mathbf{u}^1 - \tilde{\mathbf{u}}^\ell) + (\bar{\mathbf{b}}_{\ell,1} - (\tilde{\mathbf{B}}_\ell)_{\ell,1}) \mathbf{u}^1 + \sum_{j=2}^{\ell} (\bar{\mathbf{b}}_{\ell,j} - (\tilde{\mathbf{B}}_\ell)_{\ell,j}) \tilde{\mathbf{u}}^j \\
&\quad + \left(\lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^\ell \rangle}{N} - \tilde{\mu}_\ell \right) \mathbf{X}^*,
\end{aligned} \tag{242}$$

which by triangle inequality gives

$$\begin{aligned}
\frac{\|\mathbf{Y}^\ell \mathbf{u}^1 - \tilde{\mathbf{u}}^{1+\ell}\|^2}{N} &\leq C \left(\frac{\|\mathbf{Y}(\mathbf{Y}^{\ell-1} \mathbf{u}^1 - \tilde{\mathbf{u}}^\ell)\|^2}{N} + (\bar{\mathbf{b}}_{\ell,1} - (\tilde{\mathbf{B}}_\ell)_{\ell,1})^2 \frac{\|\mathbf{u}^1\|^2}{N} \right. \\
&\quad \left. + \sum_{j=2}^{\ell} (\bar{\mathbf{b}}_{\ell,j} - (\tilde{\mathbf{B}}_\ell)_{\ell,j})^2 \frac{\|\tilde{\mathbf{u}}^j\|^2}{N} + \left(\lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^\ell \rangle}{N} - \tilde{\mu}_\ell \right)^2 \frac{\|\mathbf{X}^*\|^2}{N} \right).
\end{aligned} \tag{243}$$

As \mathbf{Y} has bounded operator norm and $\|\mathbf{Y}^{\ell-1} \mathbf{u}^1 - \tilde{\mathbf{u}}^\ell\|^2/N \rightarrow 0$ (by the previous step), we have that

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{Y}(\mathbf{Y}^{\ell-1} \mathbf{u}^1 - \tilde{\mathbf{u}}^\ell)\|^2}{N} = 0.$$

Next, by following passages analogous to those in (241), we have that $\lim_{N \rightarrow \infty} \bar{\Phi}_\ell = \tilde{\Phi}_\ell$. As $\kappa_j \rightarrow \bar{\kappa}_j$ for all j , this implies that $\lim_{N \rightarrow \infty} \bar{\mathbf{B}}_\ell = \tilde{\mathbf{B}}_\ell$. Hence, for all $j \in [\ell]$, as $\|\tilde{\mathbf{u}}^j\|/N$ is bounded, we have that

$$\lim_{N \rightarrow \infty} \left((\bar{\mathbf{b}}_{\ell,1} - (\tilde{\mathbf{B}}_\ell)_{\ell,1})^2 \frac{\|\mathbf{u}^1\|^2}{N} + \sum_{j=2}^{\ell} (\bar{\mathbf{b}}_{\ell,j} - (\tilde{\mathbf{B}}_\ell)_{\ell,j})^2 \frac{\|\tilde{\mathbf{u}}^j\|^2}{N} \right) = 0.$$

Finally, as

$$\lim_{N \rightarrow \infty} \lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^\ell \rangle}{N} = \lambda \mathbb{E}[X^* \tilde{U}_\ell] = \tilde{\mu}_\ell,$$

we conclude that the last term in the RHS of (243) vanishes as well, which proves that (234) holds for $k = 1$ and a generic $\ell \in \{3, \dots, K-1\}$. Furthermore, by using

(204)–(206) and the definition (221) of $\tilde{h}_{\ell+1}$, one can readily verify that (235) holds for $k = 1$ and a generic $\ell \in \{3, \dots, K - 1\}$.

By using (220) and the definition of \mathbf{Y} , we have that

$$\mathbf{Y}^K \mathbf{u}^1 - \tilde{\mathbf{z}}^K - \sum_{i=1}^K \bar{\mathbf{b}}_{K,i} \tilde{\mathbf{u}}^i - \tilde{\mu}_K \mathbf{X}^* = \mathbf{Z}(\mathbf{Y}^{K-1} \mathbf{u}^1 - \tilde{\mathbf{u}}^K) + \left(\lambda \frac{\langle \mathbf{X}^*, \mathbf{Y}^{K-1} \mathbf{u}^1 \rangle}{N} - \tilde{\mu}_K \right) \mathbf{X}^*. \quad (244)$$

Hence, by using the definition of $\tilde{\mu}_K$ in (199) and (234) with $k = 1$, $\ell = K - 1$, we obtain

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{Y}^K \mathbf{u}^1 - \tilde{\mathbf{z}}^K - \sum_{i=1}^K \bar{\mathbf{b}}_{K,i} \tilde{\mathbf{u}}^i - \tilde{\mu}_K \mathbf{X}^*\|^2}{N} = 0. \quad (245)$$

Recall that $\mathbf{J}(\mathbf{Y}) = \sum_{j=1}^K c_j \mathbf{Y}^j$. Then, by combining (234) with $k = 1$ and (245), we have

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{J}(\mathbf{Y}) \mathbf{u}^1 - \sum_{j=1}^K c_j (\tilde{\mathbf{z}}^j + \sum_{i=1}^j \bar{\mathbf{b}}_{j,i} \tilde{\mathbf{u}}^i + \tilde{\mu}_j \mathbf{X}^*)\|^2}{N} = 0. \quad (246)$$

By following the same argument as in (241), we have that $\lim_{N \rightarrow \infty} \bar{\Phi}_K = \tilde{\Phi}_K$. As $\kappa_j \rightarrow \bar{\kappa}_j$ for all j , this implies that $\lim_{N \rightarrow \infty} \bar{\mathbf{B}}_K = \tilde{\mathbf{B}}_K$. Therefore,

$$\lim_{N \rightarrow \infty} \frac{\left\| \sum_{j=1}^K c_j (\tilde{\mathbf{z}}^j + \sum_{i=1}^j \bar{\mathbf{b}}_{j,i} \tilde{\mathbf{u}}^i + \tilde{\mu}_j \mathbf{X}^*) - \sum_{j=1}^K c_j (\tilde{\mathbf{z}}^j + \sum_{i=1}^j (\tilde{\mathbf{B}}_j)_{j,i} \tilde{\mathbf{u}}^i + \tilde{\mu}_j \mathbf{X}^*) \right\|^2}{N} = 0. \quad (247)$$

Recall that $\tilde{\mathbf{u}}^1 = \mathbf{u}^1$ and (235) holds for $k = 1$. Hence, by plugging in the formulas for $\mathbf{c}_{1,1}$, μ_1 and $\{\theta_{1,i}\}_{i \in [K]}$ (cf. (209), (207) and (208)), we have

$$\lim_{N \rightarrow \infty} \frac{\left\| \sum_{j=1}^K c_j (\tilde{\mathbf{z}}^j + \sum_{i=1}^j (\tilde{\mathbf{B}}_j)_{j,i} \tilde{\mathbf{u}}^i + \tilde{\mu}_j \mathbf{X}^*) - \mathbf{c}_{1,1} \mathbf{u}^1 - \mu_1 \mathbf{X}^* - \sum_{i=1}^K \theta_{1,i} \tilde{\mathbf{z}}^i \right\|^2}{N} = 0. \quad (248)$$

By combining (246)–(248) with the definitions of \mathbf{f}^1 and $\tilde{\mathbf{f}}^1$ (cf. (191) and (232)), we conclude that

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{f}^1 - \tilde{\mathbf{f}}^1\|^2}{N} = 0. \quad (249)$$

As g_2 is Lipschitz, (249) immediately implies that

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{u}^2 - \tilde{\mathbf{u}}^{K+1}\|^2}{N} = 0. \quad (250)$$

An application of the triangle inequality gives that, for any $i \geq 1$,

$$\begin{aligned} \|\tilde{\mathbf{f}}^i\| - \|\mathbf{f}^i - \tilde{\mathbf{f}}^i\| &\leq \|\mathbf{f}^i\| \leq \|\tilde{\mathbf{f}}^i\| + \|\mathbf{f}^i - \tilde{\mathbf{f}}^i\|, \\ \|\tilde{\mathbf{u}}^{Ki+1}\| - \|\mathbf{u}^{i+1} - \tilde{\mathbf{u}}^{Ki+1}\| &\leq \|\mathbf{u}^{i+1}\| \leq \|\tilde{\mathbf{u}}^{Ki+1}\| + \|\mathbf{u}^{i+1} - \tilde{\mathbf{u}}^{Ki+1}\|. \end{aligned} \quad (251)$$

Thus, by using (251) with $i = 1$ and Proposition 2, we obtain that

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{\|\mathbf{f}^1\|^2}{N} &= \lim_{N \rightarrow \infty} \frac{\|\tilde{\mathbf{f}}^1\|^2}{N} = \mathbb{E}\left[\left(\mu_1 X^* + \sum_{i=1}^K \theta_{1,i} \tilde{Z}_i\right)^2\right], \\ \lim_{N \rightarrow \infty} \frac{\|\mathbf{u}^2\|^2}{N} &= \lim_{N \rightarrow \infty} \frac{\|\tilde{\mathbf{u}}^{K+1}\|^2}{N} = \mathbb{E}[(\tilde{U}_{K+1})^2], \end{aligned} \quad (252)$$

which concludes the base step.

Induction step. Assume towards induction that (234)–(235) hold for $k \in [t]$, $\ell \in [K-1]$ and that, for $k \in [t]$,

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{f}^k - \tilde{\mathbf{f}}^k\|^2}{N} = 0, \quad (253)$$

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{u}^{k+1} - \tilde{\mathbf{u}}^{Kk+1}\|^2}{N} = 0, \quad (254)$$

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{f}^k\|^2}{N} = \lim_{N \rightarrow \infty} \frac{\|\tilde{\mathbf{f}}^k\|^2}{N} = \mathbb{E}\left[\left(\mu_k X^* + \sum_{i=1}^{Kk} \theta_{k,i} \tilde{Z}_i\right)^2\right], \quad (255)$$

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{u}^{k+1}\|^2}{N} = \lim_{N \rightarrow \infty} \frac{\|\tilde{\mathbf{u}}^{Kk+1}\|^2}{N} = \mathbb{E}[\tilde{U}_{Kk+1}^2]. \quad (256)$$

We now show that (253)–(256) hold for $k = t+1$, and that (234)–(235) hold for $k = t+1$, $\ell \in [K-1]$. By doing so, we will have proved also the induction step and consequently that (231) holds.

Using similar passages as in (236), we obtain

$$\begin{aligned} \mathbf{Y}\mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+2} &= \mathbf{Z}\mathbf{u}^{t+1} + \lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^{t+1} \rangle}{N} \mathbf{X}^* - \tilde{\mathbf{z}}^{Kt+1} - \sum_{i=1}^{Kt+1} (\tilde{\mathbf{B}}_{Kt+1})_{Kt+1,i} \tilde{\mathbf{u}}^i - \tilde{\mu}_{Kt+1} \mathbf{X}^* \\ &= \mathbf{Z}(\mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+1}) + \left(\lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^{t+1} \rangle}{N} - \tilde{\mu}_{Kt+1}\right) \mathbf{X}^* + \sum_{i=1}^{Kt+1} (\bar{\mathbf{b}}_{Kt+1,i} - (\tilde{\mathbf{B}}_{Kt+1})_{Kt+1,i}) \tilde{\mathbf{u}}^i. \end{aligned} \quad (257)$$

Hence, by triangle inequality,

$$\begin{aligned}
\frac{\|\mathbf{Y}\mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+2}\|^2}{N} &\leq C \left(\frac{\|\mathbf{Z}(\mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+1})\|^2}{N} \right. \\
&+ \left(\lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^{t+1} \rangle}{N} - \tilde{\mu}_{Kt+1} \right)^2 \frac{\|\mathbf{X}^*\|^2}{N} + \sum_{i=1}^{Kt+1} (\bar{\mathbf{b}}_{Kt+1,i} - (\tilde{\mathbf{B}}_{Kt+1})_{Kt+1,i})^2 \frac{\|\tilde{\mathbf{u}}^i\|^2}{N} \Big) \\
&:= C(\bar{T}_1 + \bar{T}_2 + \bar{T}_3).
\end{aligned} \tag{258}$$

Consider the first term. Since $\|\mathbf{Z}\|_{\text{op}} \leq C$, the induction hypothesis (254) implies that $\bar{T}_1 \rightarrow 0$ as $N \rightarrow \infty$.

Consider the second term. The following chain of equalities holds:

$$\lim_{N \rightarrow \infty} \lambda \frac{\langle \mathbf{X}^*, \mathbf{u}^{t+1} \rangle}{N} = \lim_{N \rightarrow \infty} \lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^{Kt+1} \rangle}{N} = \lambda \mathbb{E}[X \tilde{U}_{Kt+1}] = \tilde{\mu}_{Kt+1}. \tag{259}$$

Here, the first equality uses (254) together with the fact that $\|\mathbf{X}^*\|^2/N = 1$; the second equality follows from Proposition 2; and the third equality uses the definition of $\tilde{\mu}_{Kt+1}$ in (199). Finally, using (259) and again that $\|\mathbf{X}^*\|^2/N = 1$ gives that $\bar{T}_2 \rightarrow 0$ as $N \rightarrow \infty$.

Consider the third term. By following the same argument as in (241), we have that $\lim_{N \rightarrow \infty} \bar{\Phi}_{Kt+1} = \tilde{\Phi}_{Kt+1}$. As $\kappa_j \rightarrow \bar{\kappa}_j$ for all j , this implies that $\lim_{N \rightarrow \infty} \bar{\mathbf{B}}_{Kt+1} = \tilde{\mathbf{B}}_{Kt+1}$. By using the induction hypothesis (256), which shows that $\|\tilde{\mathbf{u}}^i\|^2/N$ converges to a finite limit, we conclude that $\bar{T}_3 \rightarrow 0$ as $N \rightarrow \infty$. This proves that the RHS of (258) vanishes and gives that (234) holds for $k = t + 1$ and $\ell = 1$.

For $\ell \in \{2, \dots, K - 1\}$, by following passages similar to (242), we have

$$\begin{aligned}
\mathbf{Y}^\ell \mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+\ell+1} &= \mathbf{Y}(\mathbf{Y}^{\ell-1} \mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+\ell}) + \sum_{i=1}^{Kt+\ell} (\bar{\mathbf{b}}_{Kt+\ell,i} - (\tilde{\mathbf{B}}_{Kt+\ell})_{Kt+\ell,i}) \tilde{\mathbf{u}}^i \\
&+ \left(\lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^{Kt+\ell} \rangle}{N} - \tilde{\mu}_{Kt+\ell} \right) \mathbf{X}^*,
\end{aligned}$$

which by triangle inequality gives

$$\begin{aligned}
\frac{\|\mathbf{Y}^\ell \mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+\ell+1}\|^2}{N} &\leq C \left(\frac{\|\mathbf{Y}(\mathbf{Y}^{\ell-1} \mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+\ell})\|^2}{N} \right. \\
&+ \sum_{i=1}^{Kt+\ell} (\bar{\mathbf{b}}_{Kt+\ell,i} - (\tilde{\mathbf{B}}_{Kt+\ell})_{Kt+\ell,i})^2 \frac{\|\tilde{\mathbf{u}}^i\|^2}{N} + \left. \left(\lambda \frac{\langle \mathbf{X}^*, \tilde{\mathbf{u}}^{Kt+\ell} \rangle}{N} - \tilde{\mu}_{Kt+\ell} \right)^2 \frac{\|\mathbf{X}^*\|^2}{N} \right).
\end{aligned} \tag{260}$$

The first term on the RHS of (260) vanishes as \mathbf{Y} has bounded operator norm and we have just proved in the previous step that $\|\mathbf{Y}^{\ell-1}\mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+\ell}\|^2/N \rightarrow 0$. To bound the second term, note that, by following the same argument as in (241), we have that $\lim_{N \rightarrow \infty} \bar{\Phi}_{Kt+\ell} = \tilde{\Phi}_{Kt+\ell}$. As $\kappa_j \rightarrow \bar{\kappa}_j$ for all j , this implies that $\lim_{N \rightarrow \infty} \bar{\mathbf{B}}_{Kt+\ell} = \tilde{\mathbf{B}}_{Kt+\ell}$. By using the induction hypothesis (256), we have that $\|\tilde{\mathbf{u}}^i\|^2/N$ converges to a finite limit for $i \in [Kt + \ell - 1]$. Furthermore, as $\|\mathbf{Y}^{\ell-1}\mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{Kt+\ell}\|^2/N \rightarrow 0$, we also have that $\|\tilde{\mathbf{u}}^{Kt+\ell}\|^2/N$ converges to a finite limit. As a result, the second term on the RHS of (260) vanishes. Finally, we can write a chain of equalities analogous to (259) with $Kt + \ell$ in place of $Kt + 1$, from which we deduce that the third term vanishes. This concludes the proof that (234) holds for $k = t + 1$ and $\ell \in [K - 1]$.

For $\ell \in [K - 1]$, by definition (221) of $h_{Kt+1+\ell}$, we have

$$\tilde{\mathbf{u}}^{Kt+1+\ell} = \tilde{\mathbf{z}}^{Kt+\ell} + \tilde{\mu}_{Kt+\ell}\mathbf{X}^* + \sum_{i=1}^{Kt+\ell} (\tilde{\mathbf{B}}_{Kt+\ell})_{Kt+\ell,i} \tilde{\mathbf{u}}^i. \quad (261)$$

Let us define:

$$\begin{aligned} \hat{\mathbf{u}}^{Kt+1+\ell} := & \tilde{\mathbf{z}}^{Kt+\ell} + \tilde{\mu}_{Kt+\ell}\mathbf{X}^* + \sum_{i=1}^{t+1} (\tilde{\mathbf{B}}_{Kt+\ell})_{Kt+\ell, K(i-1)+1} \mathbf{u}^i \\ & + \sum_{\substack{i=1 \\ i \neq 1 \pmod{K}}}^{Kt+\ell} (\tilde{\mathbf{B}}_{Kt+\ell})_{Kt+\ell,i} \left(\sum_{j=1}^{i-1} \alpha_{i,j} \tilde{\mathbf{z}}^j + \sum_{j=1}^{\lceil (i-1)/K \rceil} \beta_{i,j} \mathbf{u}^j + \gamma_i \mathbf{X}^* \right). \end{aligned} \quad (262)$$

Then, by using the recursive definitions (204)–(206), we readily have that the RHS of (262) is equal to

$$\sum_{j=1}^{Kt+\ell} \alpha_{Kt+1+\ell,j} \tilde{\mathbf{z}}^j + \sum_{j=1}^{t+1} \beta_{Kt+1+\ell,j} \mathbf{u}^j + \gamma_{Kt+1+\ell} \mathbf{X}^*. \quad (263)$$

Recall that, by induction hypothesis, (254) holds for $k \in [t]$, and (235) holds for $k \in [t]$ and $\ell \in [K - 1]$. Thus, by using the expressions in (261) and (262) for $\ell = 1$, one readily obtains that

$$\lim_{N \rightarrow \infty} \frac{\|\tilde{\mathbf{u}}^{Kt+2} - \hat{\mathbf{u}}^{Kt+2}\|^2}{N} = 0. \quad (264)$$

Since the RHS of (262) is equal to the expression in (263) for $\ell = 1$, we conclude that (235) holds for $k = t + 1$ and $\ell = 1$. At this point, we have that (235) holds for

$k \in [t]$, $\ell \in [K - 1]$ and also for $k = t + 1$, $\ell = 1$. Hence, by using the expressions in (261) and (262) for $\ell = 2$, we obtain

$$\lim_{N \rightarrow \infty} \frac{\|\tilde{\mathbf{u}}^{Kt+3} - \hat{\mathbf{u}}^{Kt+3}\|^2}{N} = 0. \quad (265)$$

Since the RHS of (262) is equal to the expression in (263) for $\ell = 2$, we conclude that (235) holds for $k = t + 1$, $\ell = 2$. By iterating this procedure for $\ell \in \{3, \dots, K - 1\}$, we obtain that (235) holds for $k = t + 1$, $\ell \in [K - 1]$.

By using (220) and the definition of \mathbf{Y} , we have that

$$\begin{aligned} \mathbf{Y}^K \mathbf{u}^{t+1} - \tilde{\mathbf{z}}^{K(t+1)} - \sum_{i=1}^{K(t+1)} \bar{\mathbf{b}}_{K(t+1),i} \tilde{\mathbf{u}}^i - \tilde{\mu}_{K(t+1)} \mathbf{X}^* \\ = \mathbf{z}(\mathbf{Y}^{K-1} \mathbf{u}^{t+1} - \tilde{\mathbf{u}}^{K(t+1)}) + \left(\lambda \frac{\langle \mathbf{X}^*, \mathbf{Y}^{K-1} \mathbf{u}^{t+1} \rangle}{N} - \tilde{\mu}_{K(t+1)} \right) \mathbf{X}^*. \end{aligned} \quad (266)$$

Hence, by using (234) with $k = t + 1$, $\ell = K - 1$ and the definition of $\tilde{\mu}_{K(t+1)}$ in (199), we obtain

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{Y}^K \mathbf{u}^{t+1} - \tilde{\mathbf{z}}^{K(t+1)} - \sum_{i=1}^{K(t+1)} \bar{\mathbf{b}}_{K(t+1),i} \tilde{\mathbf{u}}^i - \tilde{\mu}_{K(t+1)} \mathbf{X}^*\|^2}{N} = 0. \quad (267)$$

As $\mathbf{J}(\mathbf{Y}) = \sum_{j=1}^K c_j \mathbf{Y}^j$, by combining (267) with (234) with $k = t + 1$, $\ell \in [K - 1]$, we have

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{J}(\mathbf{Y}) \mathbf{u}^{t+1} - \sum_{j=1}^K c_j (\tilde{\mathbf{z}}^{Kt+j} + \sum_{i=1}^{Kt+j} \bar{\mathbf{b}}_{Kt+j,i} \tilde{\mathbf{u}}^i + \tilde{\mu}_{Kt+j} \mathbf{X}^*)\|^2}{N} = 0. \quad (268)$$

By following the same argument as in (241), we have that $\lim_{N \rightarrow \infty} \bar{\Phi}_{Kt+j} = \tilde{\Phi}_{Kt+j}$ for all $j \in [K]$. As $\kappa_j \rightarrow \bar{\kappa}_j$ for all j , this implies that $\lim_{N \rightarrow \infty} \bar{\mathbf{B}}_{Kt+j} = \tilde{\mathbf{B}}_{Kt+j}$ for all $j \in [K]$. Therefore, (268) implies that

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{J}(\mathbf{Y}) \mathbf{u}^{t+1} - \sum_{j=1}^K c_j (\tilde{\mathbf{z}}^{Kt+j} + \sum_{i=1}^{Kt+j} (\tilde{\mathbf{B}}_{Kt+j})_{Kt+j,i} \tilde{\mathbf{u}}^i + \tilde{\mu}_{Kt+j} \mathbf{X}^*)\|^2}{N} = 0. \quad (269)$$

Recall that (254) holds for $k \in [t]$ by the induction hypothesis and (235) holds for $k \in [t + 1]$, $\ell \in [K - 1]$ (thanks to the induction hypothesis and the argument above).

Hence, by plugging in the formulas for $\{\mathbf{c}_{t+1,i}\}_{i \in [t+1]}$, μ_{t+1} and $\{\theta_{t+1,i}\}_{i \in [K(t+1)]}$ (cf. (209), (207) and (208)), we have

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{J}(\mathbf{Y})\mathbf{u}^{t+1} - \sum_{i=1}^{t+1} \mathbf{c}_{t+1,i} \mathbf{u}^i - \mu_t \mathbf{X}^* - \sum_{i=1}^{K(t+1)} \theta_{t+1,i} \tilde{\mathbf{Z}}^i\|^2}{N} = 0. \quad (270)$$

By recalling the definitions of \mathbf{f}^{t+1} and $\tilde{\mathbf{f}}^{t+1}$ (cf. (191) and (232)), (270) implies that

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{f}^{t+1} - \tilde{\mathbf{f}}^{t+1}\|^2}{N} = 0. \quad (271)$$

As g_{t+2} is Lipschitz, (271) also gives that

$$\lim_{N \rightarrow \infty} \frac{\|\mathbf{u}^{t+2} - \tilde{\mathbf{u}}^{K(t+1)+1}\|^2}{N} = 0. \quad (272)$$

Then, by using (251) with $i = t + 1$ and Proposition 2, we obtain that (255) and (256) hold for $k = t + 1$, thus concluding the inductive proof. The result we have just proved by induction, combined with (233), gives that (231) holds.

Another application of Proposition 2, together with (231), gives that

$$\begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \psi(\tilde{u}_i^1, \tilde{u}_i^{K+1}, \dots, \tilde{u}_i^{Kt+1}, \tilde{f}_i^1, \tilde{f}_i^2, \dots, \tilde{f}_i^t, X_i^*) \\ = \mathbb{E}[\psi(\tilde{U}_1, \tilde{U}_{K+1}, \dots, \tilde{U}_{Kt+1}, F_1, \dots, F_t, X^*)], \end{aligned} \quad (273)$$

where we recall that, by the definition in the theorem statement, for $s \in \{1, \dots, t\}$,

$$F_s = \mu_s X^* + \sum_{i=1}^{Ks} \theta_{s,i} \tilde{Z}_i. \quad (274)$$

As $U_{s+1} = g_{s+1}(F_s)$, we have $\tilde{U}_{Ks+1} = U_{s+1}$ for all $s \in [t]$, and the proof is complete. \square

References

- [1] Emmanuel Abbe. Community detection and stochastic block models: Recent developments. *Journal of Machine Learning Research*, 18(177):1–86, 2018.
- [2] Ahmed El Alaoui, Florent Krzakala, and Michael Jordan. Fundamental limits of detection in the spiked Wigner model. *The Annals of Statistics*, 48(2):863–885, 2020.
- [3] Ahmed El Alaoui, Florent Krzakala, and Michael I Jordan. Finite size corrections and likelihood ratio fluctuations in the spiked wigner model. *arXiv preprint arXiv:1710.02903*, 2017.
- [4] Diego Alberici, Francesco Camilli, Pierluigi Contucci, and Emanuele Mingione. The multi-species mean-field spin-glass on the nishimori line. *Journal of Statistical Physics*, 182, 01 2021.
- [5] Diego Alberici, Francesco Camilli, Pierluigi Contucci, and Emanuele Mingione. The solution of the deep boltzmann machine on the nishimori line. *Communications in Mathematical Physics*, 387, 10 2021.
- [6] Zhidong Bai and Jianfeng Yao. On sample eigenvalues in a generalized spiked population model. *Journal of Multivariate Analysis*, 106:167–177, 2012.
- [7] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.
- [8] Jinho Baik and Jack W. Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of multivariate analysis*, 97(6):1382–1408, 2006.
- [9] Jean Barbier. Overlap matrix concentration in optimal Bayesian inference. *Information and Inference: A Journal of the IMA*, 10(2):597–623, 05 2020.
- [10] Jean Barbier, Mohamad Dia, and Nicolas Macris. Proof of threshold saturation for spatially coupled sparse superposition codes. In *IEEE International Symposium on Information Theory (ISIT)*, pages 1173–1177, 2016.
- [11] Jean Barbier, Mohamad Dia, Nicolas Macris, Florent Krzakala, Thibault Lesieur, and Lenka Zdeborová. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. In *Advances in Neural Information Processing Systems*, 2016.

- [12] Jean Barbier, Mohamad Dia, Nicolas Macris, Florent Krzakala, and Lenka Zdeborová. Rank-one matrix estimation: analysis of algorithmic and information theoretic limits by the spatial coupling method. *arXiv preprint, arXiv:1812.02537*, 2018.
- [13] Jean Barbier, TianQi Hou, Marco Mondelli, and Manuel Sáenz. The price of ignorance: how much does it cost to forget noise structure in low-rank matrix estimation? In *Advances in Neural Information Processing Systems*, 2022.
- [14] Jean Barbier, Florent Krzakala, Nicolas Macris, Léo Miolane, and Lenka Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the National Academy of Sciences*, 116(12):5451–5460, 2019.
- [15] Jean Barbier and Nicolas Macris. The adaptive interpolation method: a simple scheme to prove replica formulas in bayesian inference. *Probability Theory and Related Fields*, 174(3-4):1133–1185, 2019.
- [16] Jean Barbier and Nicolas Macris. The adaptive interpolation method for proving replica formulas. Applications to the Curie–Weiss and Wigner spike models. *Journal of Physics A: Mathematical and Theoretical*, 52(29):294002, 2019.
- [17] Jean Barbier, Nicolas Macris, and Cynthia Rush. All-or-nothing statistical and computational phase transitions in sparse spiked matrix estimation. In *Advances in Neural Information Processing Systems*, volume 33, pages 14915–14926. Curran Associates, Inc., 2020.
- [18] Jean Barbier and Dmitry Panchenko. Strong replica symmetry in high-dimensional optimal bayesian inference. *Communications in Mathematical Physics*, 393:1–41, 08 2022.
- [19] Jean Barbier, Dmitry Panchenko, and Manuel Sáenz. Strong replica symmetry for high-dimensional disordered log-concave Gibbs measures. *Information and Inference: A Journal of the IMA*, 11(3):1079–1108, 12 2021.
- [20] Jean Barbier and Manuel Sáenz. Marginals of a spherical spin glass model with correlated disorder. *arXiv preprint arXiv:2112.02066*, 2021.
- [21] Adriano Barra, Pierluigi Contucci, Emanuele Mingione, and Daniele Tantari. Multi-species mean field spin glasses. rigorous results. *Annales Henri Poincaré*, 16:691–708, 2013.

- [22] Mohsen Bayati and Andrea Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785, 2011.
- [23] Mohsen Bayati and Andrea Montanari. The lasso risk for gaussian matrices. *IEEE Transactions on Information Theory*, 58(4):1997–2017, 2012.
- [24] Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.
- [25] Florent Benaych-Georges and Raj Rao Nadakuditi. The singular values and vectors of low rank perturbations of large rectangular random matrices. *Journal of Multivariate Analysis*, 111:120–135, 2012.
- [26] Bhaswar B. Bhattacharya and Subhabrata Sen. High temperature asymptotics of orthogonal mean-field spin glasses. *Journal of Statistical Physics*, 162(1):63–80, 2016.
- [27] Erwin Bolthausen. An iterative construction of solutions of the TAP equations for the Sherrington–Kirkpatrick model. *Communications in Mathematical Physics*, 325(1):333–366, 2014.
- [28] Edouard Brézin, Claude Itzykson, Giorgio Parisi, and Jean-Bernard Zuber. Planar diagrams. *Communications in Mathematical Physics*, 59(1):35 – 51, 1978.
- [29] Francesco Camilli, Pierluigi Contucci, and Emanuele Mingione. An inference problem in a mismatched setting: a spin-glass model with Mattis interaction. *SciPost Phys.*, 12:125, 2022.
- [30] Mireille Capitaine, Catherine Donati-Martin, and Delphine Féral. The largest eigenvalues of finite rank deformation of large Wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability*, 37(1):1–47, 2009.
- [31] Hong-Bin Chen, Jean-Christophe Mourrat, and Jiaming Xia. Statistical inference of finite-rank tensors. *arXiv preprint arXiv:2104.05360*, 2021.
- [32] Hong-Bin Chen and Jiaming Xia. Limiting free energy of multi-layer generalized linear models. *arXiv preprint arXiv:2108.12615*, 2021.

- [33] Amin Coja-Oghlan, Florent Krzakala, Will Perkins, and Lenka Zdeborová. Information-theoretic thresholds from the cavity method. *Advances in Mathematics*, 333:694–795, 2018.
- [34] Yash Deshpande and Andrea Montanari. Information-theoretically optimal sparse PCA. *IEEE International Symposium on Information Theory*, pages 2197–2201, 2014.
- [35] David Donoho, Arian Maleki, and Andrea Montanari. Message passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences of the United States of America*, 106:18914–9, 11 2009.
- [36] David L. Donoho, Adel Javanmard, and Andrea Montanari. Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing. *IEEE Transactions on Information Theory*, 59(11):7434–7464, 2013.
- [37] Rishabh Dudeja and Milad Bakhshizadeh. Universality of linearized message passing for phase retrieval with structured sensing matrices. *IEEE Transactions on Information Theory*, 2022.
- [38] Rishabh Dudeja, Yue M Lu, and Subhabrata Sen. Universality of approximate message passing with semi-random matrices. *arXiv preprint arXiv:2204.04281*, 2022.
- [39] Rishabh Dudeja, Subhabrata Sen, and Yue M Lu. Spectral universality of regularized linear regression with nearly deterministic sensing matrices. *arXiv preprint arXiv:2208.02753*, 2022.
- [40] Ahmed El Alaoui and Florent Krzakala. Estimation in the spiked wigner model: a short proof of the replica formula. In *IEEE International Symposium on Information Theory (ISIT)*, pages 1874–1878. IEEE, 2018.
- [41] Zhou Fan. Approximate message passing algorithms for rotationally invariant matrices. *The Annals of Statistics*, 50(1):197–224, 2022.
- [42] Zhou Fan and Yihong Wu. The replica-symmetric free energy for ising spin glasses with orthogonally invariant couplings. *arXiv preprint arXiv:2105.02797*, 2021.
- [43] Alberto Jimenez Felstrom and Kamil Sh Zigangirov. Time-varying periodic convolutional codes with low-density parity-check matrix. *IEEE Transactions on Information Theory*, 45(6):2181–2191, 1999.

- [44] Oliver Y Feng, Ramji Venkataramanan, Cynthia Rush, Richard J Samworth, et al. A unifying tutorial on approximate message passing. *Foundations and Trends® in Machine Learning*, 15(4):335–536, 2022.
- [45] Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large Wigner matrices. *Communications in mathematical physics*, 272(1):185–228, 2007.
- [46] Alyson Fletcher and Sundeep Rangan. Iterative reconstruction of rank-one matrices in noise. *Information and Inference: A Journal of the IMA*, 7:531–562, 09 2018.
- [47] Laura Foini and Jorge Kurchan. Annealed averages in spin and matrix models. *SciPost Phys.*, 12:080, 2022.
- [48] Marylou Gabrié, André Manoel, Clément Luneau, Jean Barbier, Nicolas Macris, Florent Krzakala, and Lenka Zdeborová. Entropy and mutual information in models of deep neural networks. In *Advances in Neural Information Processing Systems*, pages 1826–1836. NeurIPS, 2018.
- [49] Cedric Gerbelot, Alia Abbara, and Florent Krzakala. Asymptotic errors for teacher-student convex generalized linear models (or: how to prove Kabashima’s replica formula). *arXiv preprint, arXiv:2006.06581*, 2020.
- [50] Francesco Guerra. Broken replica symmetry bounds in the mean field spin glass model. *Communications in mathematical physics*, 233(1):1–12, 2003.
- [51] Francesco Guerra and Fabio Toninelli. The thermodynamic limit in mean field spin glass models. *Communications in Mathematical Physics*, 230(1):71–79, 2002.
- [52] Alice Guionnet. First order asymptotics of matrix integrals: a rigorous approach towards the understanding of matrix models. *Communications in mathematical physics*, 244(3):527–569, 2004.
- [53] Alice Guionnet, Justin Ko, Florent Krzakala, and Lenka Zdeborová. Low-rank matrix estimation with inhomogeneous noise. *arXiv preprint, arXiv:2208.05918*, August 2022.
- [54] Alice Guionnet and Mylène Maida. A Fourier view on the R-transform and related asymptotics of spherical integrals. *Journal of Functional Analysis*, 222(2):435–490, 2005.

- [55] Alice Guionnet and Ofer Zeitouni. Large deviations asymptotics for spherical integrals. *Journal of functional analysis*, 188(2):461–515, 2002.
- [56] Claude Itzykson and Jean-Bernard Zuber. The planar approximation. ii. *Journal of Mathematical Physics*, 21(3):411–421, 1980.
- [57] Adel Javanmard and Andrea Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Information and Inference: A Journal of the IMA*, 2(2):115–144, 2013.
- [58] Iain M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of statistics*, 29(2):295–327, 2001.
- [59] Yoshiyuki Kabashima. A cdma multiuser detection algorithm on the basis of belief propagation. *Journal of Physics A: Mathematical and General*, 2003.
- [60] Yoshiyuki Kabashima. Inference from correlated patterns: a unified theory for perceptron learning and linear vector channels. In *Journal of Physics: Conference Series*, volume 95, page 012001. IOP Publishing, 2008.
- [61] Yoshiyuki Kabashima, Florent Krzakala, Marc Mézard, Ayaka Sakata, and Lenka Zdeborová. Phase transitions and sample complexity in bayes-optimal matrix factorization. *IEEE Transactions on Information Theory*, 62(7):4228–4265, 2016.
- [62] Vladimir Kazakov. Solvable matrix models. *Random matrix models and their applications*, 40:271–283, 2001.
- [63] Chiheon Kim, Afonso S Bandeira, and Michel X Goemans. Community detection in hypergraphs, spiked tensor models, and sum-of-squares. In *2017 International Conference on Sampling Theory and Applications (SampTA)*, pages 124–128. IEEE, 2017.
- [64] Florent Krzakala, Marc Mézard, Francois Sausset, Yifan Sun, and Lenka Zdeborová. Probabilistic reconstruction in compressed sensing: algorithms, phase diagrams, and threshold achieving matrices. *Journal of Statistical Mechanics: Theory and Experiment*, 2012.
- [65] Shrinivas Kudekar, Thomas J Richardson, and Rüdiger L Urbanke. Threshold saturation via spatial coupling: Why convolutional ldpc ensembles perform so well over the bec. *IEEE Transactions on Information Theory*, 57(2):803–834, 2011.

- [66] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Mmse of probabilistic low-rank matrix estimation: Universality with respect to the output channel. In *Annual Allerton Conference*, 2015.
- [67] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(7):073403, 2017.
- [68] Junjie Ma, Ji Xu, and Arian Maleki. Optimization-based amp for phase retrieval: The impact of initialization and ℓ_2 regularization. *IEEE Transactions on Information Theory*, 65(6):3600–3629, 2019.
- [69] Junjie Ma, Ji Xu, and Arian Maleki. Analysis of sensing spectral for signal recovery under a generalized linear model. In *Advances in Neural Information Processing Systems*, volume 34, pages 22601–22613, 2021.
- [70] Antoine Maillard, Laura Foini, Alejandro Lage Castellanos, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. High-temperature expansions and message passing algorithms. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(11):113301, 2019.
- [71] Antoine Maillard, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. Phase retrieval in high dimensions: Statistical and computational phase transitions. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- [72] Arian Maleki, Laura Anitori, Zai Yang, and Richard G Baraniuk. Asymptotic analysis of complex lasso via complex approximate message passing (CAMP). *IEEE Transactions on Information Theory*, 59(7):4290–4308, 2013.
- [73] Enzo Marinari, Giorgio Parisi, and Felix Ritort. Replica field theory for deterministic models: I. binary sequences with low autocorrelation. *Journal of Physics A*, 27:7615–7645, 1994.
- [74] Enzo Marinari, Giorgio Parisi, and Felix Ritort. Replica field theory for deterministic models: II. A non-random spin glass with glassy behaviour. *Journal of Physics A*, 27:7647–7668, 1994.
- [75] Marc Mézard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.

- [76] Marco Mondelli and Ramji Venkataramanan. Approximate message passing with spectral initialization for generalized linear models. In *24th International Conference on Artificial Intelligence and Statistics*, volume 130, pages 397–405, 2021.
- [77] Marco Mondelli and Ramji Venkataramanan. PCA initialization for approximate message passing in rotationally invariant models. In *Advances in Neural Information Processing Systems*, volume 34, pages 29616–29629, 2021.
- [78] Andrea Montanari and Subhabrata Sen. A short tutorial on mean-field spin glass techniques for non-physicists. *arXiv preprint arXiv:2204.02909*, 2022.
- [79] Andrea Montanari and Ramji Venkataramanan. Estimation of low-rank matrices via approximate message passing. *Annals of Statistics*, 45(1):321–345, 2021.
- [80] Jean-Christophe Mourrat. Hamilton–jacobi equations for finite-rank matrix inference. *The Annals of Applied Probability*, 30(5):2234–2260, 2020.
- [81] Raj Rao Nadakuditi and Jack W. Silverstein. Fundamental limit of sample generalized eigenvalue based detection of signals in noise using relatively few signal-bearing and noise-only samples. *IEEE Journal of Selected Topics in Signal Processing*, 4(3):468–480, 2010.
- [82] Hidetoshi Nishimori. *Statistical Physics of Spin Glasses and Information Processing: an Introduction*. Oxford University Press, Oxford; New York, 2001.
- [83] Jonathan Novak. Three lectures on free probability. *Random matrix theory, interacting particle systems, and integrable systems*, 65(309-383):13, 2014.
- [84] Manfred Opper, Burak Cakmak, and Ole Winther. A theory of solving tap equations for ising models with general invariant random matrices. *Journal of Physics A: Mathematical and Theoretical*, 49(11):114002, 2016.
- [85] Manfred Opper and David Saad. *Advanced mean field methods: Theory and practice*. MIT press, 2001.
- [86] Manfred Opper and Ole Winther. Adaptive and self-averaging thouless-anderson-palmer mean-field theory for probabilistic modeling. *Physical Review E*, 64, 2011.
- [87] Dmitry Panchenko. The free energy in a multi-species sherrington-kirkpatrick model. *The Annals of Probability*, 43, 10 2013.

- [88] Dmitry Panchenko. *The Sherrington-Kirkpatrick model*. Springer Science & Business Media, 2013.
- [89] Giorgio Parisi and Marc Potters. Mean-field equations for spin models with orthogonal interaction matrices. *Journal of Physics A: Mathematical and General*, 28:5267, 01 1999.
- [90] Amelia Perry, Alexander S Wein, Afonso S Bandeira, and Ankur Moitra. Optimality and sub-optimality of pca for spiked random matrices and synchronization. *arXiv preprint arXiv:1609.05573*, 2016.
- [91] Amelia Perry, Alexander S. Wein, Afonso S. Bandeira, and Ankur Moitra. Message-passing algorithms for synchronization problems over compact groups. *Communications on Pure and Applied Mathematics*, 71(11):2275–2322, 2018.
- [92] Amelia Perry, Alexander S Wein, Afonso S Bandeira, and Ankur Moitra. Optimality and sub-optimality of PCA I: Spiked random matrix models. *The Annals of Statistics*, 46(5):2416–2451, 2018.
- [93] Marc Potters and Jean-Philippe Bouchaud. *A First Course in Random Matrix Theory: For Physicists, Engineers and Data Scientists*. Cambridge University Press, 2020.
- [94] Sandrine Péché. The largest eigenvalue of small rank perturbations of hermitian random matrices. *probab. theory relat. fields* 134, 127-173. *Probability Theory and Related Fields*, 134:127–173, 01 2006.
- [95] Sundeep Rangan. Generalized approximate message passing for estimation with random linear mixing. In *International Symposium on Information Theory*, pages 2168–2172, 2011.
- [96] Philip Schniter and Sundeep Rangan. Compressive phase retrieval via generalized approximate message passing. *IEEE Transactions on Signal Processing*, 63(4):1043–1055, 2014.
- [97] Pragya Sur and Emmanuel J. Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116(29):14516–14525, 2019.
- [98] Takashi Takahashi and Yoshiyuki Kabashima. Macroscopic analysis of vector approximate message passing in a model mismatch setting. In *International Symposium on Information Theory*, pages 1403–1408, 2020.

- [99] Michel Talagrand. The parisi formula. *Annals of mathematics*, pages 221–263, 2006.
- [100] Michel Talagrand. *Mean Field Models for Spin Glasses. Volume I: Basic Examples*. Springer Verlag, 2011.
- [101] Ramji Venkataramanan, Kevin Kögler, and Marco Mondelli. Estimation in rotationally invariant generalized linear models via approximate message passing. In *International Conference on Machine Learning*, pages 22120–22144, 2022.
- [102] Xinyi Zhong, Tianhao Wang, and Zhou Fan. Approximate message passing for orthogonally invariant ensembles: Multivariate non-linearities and spectral initialization. *arXiv preprint, arXiv:2110.02318*, 2021.