
The price of ignorance: how much does it cost to forget noise structure in low-rank matrix estimation?

Jean Barbier*

International Center for Theoretical Physics
jbarbier@ictp.it

Tianqi Hou

Huawei
thou@connect.ust.hk

Marco Mondelli*

Institute of Science and Technology Austria
marco.mondelli@ist.ac.at

Manuel Sáenz*

International Center for Theoretical Physics
msaenz@ictp.it

Abstract

We consider the problem of estimating a rank-1 signal corrupted by structured rotationally invariant noise, and address the following question: *how well do inference algorithms perform when the noise statistics is unknown and hence Gaussian noise is assumed?* While the matched Bayes-optimal setting with unstructured noise is well understood, the analysis of this mismatched problem is only at its premises. In this paper, we make a step towards understanding the effect of the strong source of mismatch which is the noise statistics. Our main technical contribution is the rigorous analysis of a Bayes estimator and of an approximate message passing (AMP) algorithm, both of which incorrectly assume a Gaussian setup. The first result exploits the theory of spherical integrals and of low-rank matrix perturbations; the idea behind the second one is to design and analyze an artificial AMP which, by taking advantage of the flexibility in the denoisers, is able to “correct” the mismatch. Armed with these sharp asymptotic characterizations, we unveil a rich and often unexpected phenomenology. For example, despite AMP is in principle designed to efficiently compute the Bayes estimator, the former is *outperformed* by the latter in terms of mean-square error. We show that this performance gap is due to an incorrect estimation of the signal norm. In fact, when the SNR is large enough, the overlaps of the AMP and the Bayes estimator coincide, and they even match those of optimal estimators taking into account the structure of the noise.

1 Introduction

The estimation of a low-rank matrix from noisy data is a central problem in machine learning, and it appears, e.g., in sparse principal component analysis (PCA) [50, 88], community detection [1, 62], and group synchronization [69]. In this paper, we consider the prototypical task of recovering a symmetric rank-1 matrix $\mathbf{X}\mathbf{X}^\top$ from noisy observations of the form

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

Here, $\lambda_* > 0$ is the signal-to-noise ratio (SNR) which quantifies the signal strength, and $\mathbf{Z} \in \mathbb{R}^{N \times N}$ is a random matrix that captures the noise. A natural estimator of \mathbf{X} is given by the principal eigenvector of \mathbf{Y} . Its performance and, more generally, the behavior of the eigenvalues and eigenvectors of (1) has been studied in exquisite detail in statistics [49, 68] and random matrix theory [7, 8, 21, 22, 26, 40, 52, 6]. Going beyond the spectral estimator given by the principal eigenvector, approximate message

passing (AMP) constitutes a popular family of iterative algorithms. The reason for this popularity lies in two especially attractive features: (i) AMP algorithms can be tailored to exploit knowledge on the structure of the signal; and (ii) under suitable assumptions, their performance in the high-dimensional limit is accurately tracked by a deterministic recursion known as *state evolution* [18, 24, 48]. Using the state evolution machinery, it has been showed that AMP achieves Bayes-optimal performance in some Gaussian models [29, 30, 61, 13], and a bold conjecture from statistical physics is that AMP is optimal in the class of polynomial-time algorithms for a large class of inference problems with random design. AMP algorithms have been recently considered for several practical scenarios, including genetics [86], inpainting [66] and MRI image recovery [76].

A general class of noise models that has received attention in the literature is given by the family of *rotationally invariant matrices*. This is much milder than requiring \mathbf{Z} to be Gaussian: it only imposes that the matrix of eigenvectors is uniformly random, and allows for an arbitrary spectrum. Hence, \mathbf{Z} can capture the complex correlation structure which often occurs in applications (e.g., recommender systems [20] and bioinformatics [72]). However, estimating noise statistics from data is costly or, for problems involving large-scale datasets (computational genomics is a paradigmatic example [36, 2]), may be impossible. Thus, one natural idea is to simply assume Gaussian statistics for \mathbf{Z} . The case in which \mathbf{Z} is actually Gaussian has been thoroughly studied [53, 28, 55, 54, 12, 61]. Beyond Gaussianity, a rapidly growing literature is focusing on rotationally invariant models assuming *perfect knowledge of the statistics of the structured matrix appearing in the problem* (such as noise in inference, a sensing, data, or coding matrix in regression tasks, weight matrices in neural networks, or a matrix of interactions in spin glass models) [67, 23, 37, 38, 42, 65, 56, 57, 74, 78, 79, 16, 33, 66, 71, 47]. However, despite this impressive progress when the noise statistics is known, low-rank estimation in a *mismatched setting with partial to no knowledge of the statistics of the rotationally invariant noise matrix remains poorly understood*. The goal of this work is thus to shed light into the following fundamental question:

Suppose that the noise statistics is unknown or unreachible, and hence Gaussian noise is naively assumed. What is the impact of this mismatch on the overall performance of inference methods?

1.1 Summary of contributions

In this paper, we provide rigorous performance guarantees for two inference methods which incorrectly assume Gaussian noise statistics: a Bayes estimator, and an AMP algorithm. Then, by exploiting these sharp analytical characterizations, we describe a number of surprising effects coming from numerical simulations. Our main findings are detailed below.

Theoretical results. (i) We give a closed-form expression for the mean-square error (MSE) of the Bayes estimator which samples from the *mismatched* posterior (cf. Theorem 1). Under a certain concentration assumption, we also present an asymptotic result on the overlap of such estimator. (ii) We provide a state evolution analysis (cf. Theorem 2) for the *Gaussian* AMP algorithm which is designed for Gaussian noise.

Numerical results. The mismatched Bayes and AMP estimators display surprising behaviors – already for the case of a spherically distributed signal – and the two performance metrics (MSE and overlap) exhibit a remarkably different phenomenology: (i) As for the MSE, the Gaussian AMP is *outperformed* by the Bayes estimator. Here, the surprise comes from the fact that AMP algorithms are, in principle, designed to sample from the posterior distribution (and this is often what happens in the matched case). (ii) In contrast, when the SNR is large enough, the overlaps of the two mismatched estimators *coincide*, and they even match the overlap of estimators which exploit the noise statistics, namely the optimal spectral method which minimizes the MSE and the correct AMP designed in [37]. (iii) Under certain conditions, the MSE of the mismatched Bayes estimator *matches* that of a Gaussian spectral method with no information on the noise structure. (iv) The mismatched estimators are *outperformed* – in terms of MSE – by the optimal spectral method and the correct AMP, whose performance coincide. (v) Finally, the MSE curves of the Bayes and Gaussian spectral methods exhibit a striking non-monotone behavior.

1.2 Related work

The impact of mismatch. Given the practical relevance of understanding the effect of mismatch in statistical inference, a line of work has approached the issue from various angles. Regression being

probably the most paradigmatic model of inference task, mismatch has been thoroughly studied in this context. In information theory, the trend was initiated in [83, 85]. Studies for M-estimation and robust statistics followed based on statistical mechanics approaches [3], on the analysis of approximate message passing algorithms [25, 77, 32, 43], on Gordon's convex min-max theorem [81, 80] and the leave-one-out method [35, 34]. Concerning the analysis of Bayesian approaches to mismatched linear regression, we refer to [27, 63, 75, 11] or [9] for a review.

In contrast, the problem studied in the present paper, namely low-rank matrix inference with mismatch, has received attention only recently. It is clear that in some way, related questions were analyzed in the aforementioned random matrix theory literature, although from a rather different perspective. Concerning the precise issue of the performance degradation due to mismatch in Bayesian inference, an exception (and inspiration for our work) is the recent paper [71]. However, the authors of [71] focus on mismatch in the signal-to-noise ratio for the Gaussian noise setting; their results can thus be recovered as a special case of the present ones. To the best of our knowledge, this is the first work that considers mismatch in the noise statistics for low-rank matrix estimation.

Approximate message passing. AMP algorithms have been applied to a wide range of inference problems. Examples include estimation in linear models [19, 18, 31], generalized linear models [73, 13, 58, 59], and low-rank matrix recovery with Gaussian noise [17, 29, 41, 51, 55, 61], see also the survey [39]. A general AMP iteration for rotationally invariant matrices has been recently analyzed in [37, 87], and by providing suitable instances of this abstract iteration, AMP algorithms have been developed for low-rank [37, 87, 60] and generalized linear models [82]. Furthermore, an AMP-based method which uses the classical idea of empirical Bayes to reduce the high-dimensional noise in PCA is proposed in [86], which also provides applications to genetics. However, the existing results cannot be applied to the mismatched setting considered in this work, since the Gaussian AMP does not contain the right Onsager corrections. In fact, the algorithm designer assumes Gaussian statistics for the noise and, therefore, constructs an AMP algorithm with the Onsager correction suitable for Gaussian noise. Finally, we remark that the performance of the Gaussian AMP is numerically compared with that of the correct AMP (exploiting the knowledge of noise statistics) in [87]. Our Theorem 2 provides rigorous foundations for such a comparison.

2 Setup of the problem

2.1 Random matrix theory preliminaries

We start with some useful notions of random matrix theory. Given a probability measure ρ of compact support $K \subseteq \mathbb{R}$, we let $H_\rho : \mathbb{R} \setminus K \mapsto \mathbb{R}$ be the *Hilbert transform* of ρ : $H_\rho(z) := \int_K \rho(d\gamma)(z - \gamma)^{-1}$. We define $\bar{\gamma} := \max K$, $\underline{\gamma} := \min K$, $\bar{h} := \lim_{z \downarrow \bar{\gamma}} H_\rho(z)$ and $\underline{h} := \lim_{z \uparrow \underline{\gamma}} H_\rho(z)$, where the limits exist due to the monotonicity of H_ρ but may be infinite. As H_ρ is a bijection between $\mathbb{R} \setminus K$ and its image $(\underline{h}, \bar{h}) \setminus \{0\}$, its inverse exists and it is denoted by $K_\rho : (\underline{h}, \bar{h}) \setminus \{0\} \mapsto \mathbb{R} \setminus K$. The *R-transform* of ρ is $R_\rho : (\underline{h}, \bar{h}) \setminus \{0\} \mapsto \mathbb{R} \setminus K$ given by $R_\rho(x) := K_\rho(x) - 1/x$. The coefficients $\{\bar{\kappa}_k\}_{k \geq 1}$ of the Taylor series of R_ρ are the *free cumulants* associated to ρ , i.e., $R_\rho(x) = \sum_{k=0}^{\infty} \bar{\kappa}_{k+1} x^k$, and they can be computed from the moments of ρ , see e.g. Section 2.5 of [64]. Furthermore, $R'_\rho(x)$ and $H'_\rho(z)$ denote the derivatives of the *R-transform* and the Hilbert transform of ρ , respectively.

2.2 Model of mismatched low-rank matrix estimation

We consider the problem of estimating a rank-one informative *spike* $\mathbf{X}\mathbf{X}^\top$ corrupted by an additive symmetric noise matrix $\mathbf{Z} \in \mathbb{R}^{N \times N}$ from data \mathbf{Y} generated as in (1). The scalar $\lambda_* \in \mathbb{R}_{\geq 0}$ is the signal-to-noise ratio (SNR), and the noise \mathbf{Z} can be decomposed as $\mathbf{Z} = \mathbf{O}\mathbf{\Sigma}\mathbf{O}^\top$, where $\mathbf{\Sigma} := \text{diag}(\gamma_1, \dots, \gamma_N) \in \mathbb{R}^{N \times N}$ is a diagonal matrix containing the eigenvalues of \mathbf{Z} and \mathbf{O} is some orthogonal matrix. We also denote by $\bar{\gamma}_N$ and $\underline{\gamma}_N$ respectively $\max\{\gamma_1, \dots, \gamma_N\}$ and $\min\{\gamma_1, \dots, \gamma_N\}$. Similarly, $\bar{\nu}_N$ and $\underline{\nu}_N$ denote the largest and smallest eigenvalues of the data matrix \mathbf{Y} . Throughout the paper, the main technical assumption will be the following.

Assumption 1. *The signal \mathbf{X} has norm \sqrt{N} . The random noise matrix \mathbf{Z} is rotationally invariant and is independent of \mathbf{X} . Moreover, the empirical measure $N^{-1} \sum_{i \leq N} \delta_{\gamma_i}$ of eigenvalues of \mathbf{Z} converges weakly towards a limiting measure ρ of compact support $K \subseteq \mathbb{R}$. Finally, $\bar{\gamma}_N$ and $\underline{\gamma}_N$ converge a.s. to $\bar{\gamma} = \max K$ and $\underline{\gamma} = \min K$, respectively.*

The noise matrix \mathbf{Z} being rotationally invariant means that \mathbf{O} is a random orthogonal matrix (i.e., sampled from the Haar measure) independent of Σ . Under Assumption 1, we have that $\bar{\nu}_N$ and $\underline{\nu}_N$ converge a.s. to finite limits, which we denote by $\bar{\nu}$ and $\underline{\nu}$, respectively. These have an explicit form in terms of ρ and λ_* given in [21, Theorem 2.1] and which we reproduce for convenience in Appendix A. We remark that the assumption on the ground-truth signal is rather mild: \mathbf{X} can have any distribution over the sphere of radius \sqrt{N} , and it might even be deterministic. For example, $\mathbf{X} \in \{-1, 1\}^N$ could be a uniform vector over the centred N -dimensional hyper-cube of sides $2\sqrt{N}$.

For the analysis of the Bayes estimator, we require a second assumption on the asymptotic eigenvalue distribution ρ . This ensures that the limit for the overlap of the spectral estimators with the signal is continuous in the SNR, whenever a small additive Wigner noise of variance $\epsilon \geq 0$ is added to \mathbf{Y} .

Assumption 2. Let ρ_ϵ be the spectral density of the free convolution of ρ and a semicircle law of radius $2\epsilon \geq 0$. Let H_{ρ_ϵ} be the associated Hilbert transform and $\bar{\gamma}_\epsilon$ the rightmost point of the support of ρ_ϵ . Then, we assume that $\lim_{z \downarrow \bar{\gamma}_\epsilon} H_{\rho_\epsilon}'(z) = -\infty$ for all $\epsilon \geq 0$.

For the definition of the free convolution and its link to random matrix theory, we refer the reader to [4, 70]. We remark that Assumption 2 is satisfied by a wide class of random matrices. In particular, by combining Theorem 2.2 of [10] with Proposition 2.4 of [21], it suffices that the support of the limiting spectral measure ρ is (i) compact, (ii) connected, and (iii) it has a proper decay rate at its edges (i.e., ρ is of Jacobi type), see Assumption 2.1 of [10]. These conditions can be easily checked for many random matrix models, including the ones discussed in the examples of Section 4.

Note that, if ρ is the semicircle law of radius 2, the noise is asymptotically equal in distribution to a standard Wigner matrix \mathbf{W} with density $\sim \exp(-\frac{N}{4}\text{Tr}\mathbf{W}^2)$. As the elements of this type of symmetric matrices are i.i.d. Gaussian, in this case we say that there is *Gaussian noise*. For other limiting ρ , the elements of $\mathbf{Z} \neq \mathbf{W}$ remain correlated and we say that there is *structured noise*. We denote by \mathbf{W} a sequence of standard Wigner matrices while we reserve \mathbf{Z} for a generic sequence of rotationally invariant matrices.

Sources of mismatch. In this paper, we consider what happens when there is a *mismatch* between the true noise statistics and the assumptions on the noise statistics made in the inference algorithm. In particular, we study the case in which the noise is assumed to be Gaussian. Gaussianity is in fact the most standard assumption made when precise knowledge of the noise structure is lacking. We also consider the case in which the SNR is estimated incorrectly, i.e., the statistician assumes that the data is generated according to $(\sqrt{\lambda}/N)\mathbf{X}\mathbf{X}^\top + \mathbf{W}$, where $\lambda \neq \lambda_*$ and λ_* is the correct SNR present in (1). Our goal is to quantify how these sources of ignorance affect the algorithmic performance.

2.3 Three classes of inference procedures for PCA

Mismatched Bayes estimator. The statistician assumes Gaussian noise and that the signal \mathbf{X} has no specific structure, i.e., it is uniformly distributed on the N -dimensional sphere $\mathbb{S}^{N-1}(\sqrt{N})$. The mismatched posterior distribution reads, using the Gaussian log-likelihood $-\frac{N}{4}\text{Tr}(\mathbf{Y} - \frac{\sqrt{\lambda}}{N}\mathbf{x}\mathbf{x}^\top)^2$,

$$P_{\text{mis}}(d\mathbf{x} | \mathbf{Y}) = \frac{1}{Z_N(\mathbf{Y})} \exp\left(-\frac{\sqrt{\lambda}}{2}\langle \mathbf{x}, \mathbf{Y}\mathbf{x} \rangle\right) \mu_N(d\mathbf{x}), \quad (2)$$

where $Z_N(\mathbf{Y})$ is the normalization constant and μ_N is the uniform measure over $\mathbb{S}^{N-1}(\sqrt{N})$. In case the signal \mathbf{X} lies on the sphere (by Assumption 1) but is not uniformly distributed on it, then this leads to a third source of mismatch. However, the MSE formula of Theorem 1 does *not* depend on it; e.g. if $\mathbf{X} \in \{-1, 1\}^N$ the error remains the same. The origin of this is the uniformly spherical prior used in the posterior (2) which is the most uninformative one.

The Bayes estimators we analyze are obtained as the posterior means

$$M_{\text{mis}}(\mathbf{Y}) := \int \mathbf{x}\mathbf{x}^\top P_{\text{mis}}(d\mathbf{x} | \mathbf{Y}), \quad (3)$$

which may not be practical to compute. Notice that, if the noise is Wigner, then the likelihood $\exp(\dots)$ in (2) is correct. This Bayes-optimal case has already been rigorously characterized (see, for example, [12, 14, 54]), and it will serve us as a base case for comparison with our mismatched setting.

The performance of a sequence of estimators $M_N(\mathbf{Y}) \in \mathbb{R}^{N \times N}$ of the spike $\mathbf{X}\mathbf{X}^\top$ is quantified via the asymptotic mean-square error (MSE):

$$\text{MSE}(M) := \lim_{N \rightarrow +\infty} \frac{1}{2N^2} \mathbb{E} \|\mathbf{X}\mathbf{X}^\top - M_N(\mathbf{Y})\|_F^2, \quad (4)$$

where $\|\cdot\|_F^2$ denotes the Frobenius norm and \mathbb{E} the expectation over the spike $\mathbf{X}\mathbf{X}^\top$ and the noise \mathbf{Z} . We also consider another performance measure which is insensitive to the norm of an estimator $\hat{\mathbf{x}}$ of \mathbf{X} , namely, the rescaled overlap:

$$\text{Overlap}(\hat{\mathbf{x}}) := \lim_{N \rightarrow +\infty} \frac{|\langle \mathbf{X}, \hat{\mathbf{x}}(\mathbf{Y}) \rangle|}{\|\hat{\mathbf{x}}(\mathbf{Y})\| \cdot \|\mathbf{X}\|}, \quad (5)$$

where $\|\mathbf{X}\| = \sqrt{N}$ due to the spherical constraint. This alignment measure is sometimes also referred to as *cosine distance*. For both AMP and spectral estimators, this overlap will almost surely be deterministic. If the norm of the estimator vanishes, by convention we fix the overlap to zero. Note that the definition (5) is not meaningful for the Bayes case. In fact, if one defines $\hat{\mathbf{x}}_{\text{mis}} := \int \mathbf{x} P_{\text{mis}}(d\mathbf{x} | \mathbf{Y})$, then this quantity is zero by sign symmetry ($\pm \mathbf{x}$ have the same posterior weight). Thus, the correct definition of the overlap for the mismatched Bayes estimator is

$$\text{Overlap}_{\text{mis}} := \lim_{N \rightarrow +\infty} \left(\frac{1}{N} \frac{1}{\|M_{\text{mis}}(\mathbf{Y})\|_F} \int P_{\text{mis}}(d\mathbf{x} | \mathbf{Y}) \langle \mathbf{X}, \mathbf{x} \rangle^2 \right)^{1/2}, \quad (6)$$

while for the ‘‘one-shot’’ AMP and spectral estimators we use the former definition (5).

Approximate message passing. We remark that, as opposed to our analysis of the Bayes estimator, the state evolution characterization of AMP (Theorem 2) does not require \mathbf{X} to be uniformly distributed on $\mathbb{S}^{N-1}(\sqrt{N})$. The AMP analysis can accommodate more generic (potentially mismatched) prior distributions and, in fact, AMP algorithms are well equipped to exploit structure in the signal. For simplicity, we assume to have access to an initialization $\hat{\mathbf{x}}^1 \in \mathbb{R}^N$, which is independent of the noise \mathbf{Z} and has a strictly positive correlation with \mathbf{X} , i.e.,

$$(\mathbf{X}, \hat{\mathbf{x}}^1) \xrightarrow{W_2} (X, \hat{x}_1), \quad \mathbb{E}[X \hat{x}_1] := \epsilon > 0, \quad \mathbb{E}[\hat{x}_1^2] \leq 1, \quad (7)$$

where $(\mathbf{X}, \hat{\mathbf{x}}^1) \xrightarrow{W_2} (X, \hat{x}_1)$ denotes convergence of the joint empirical distribution of $(\mathbf{X}, \hat{\mathbf{x}}^1)$ to the random variable (X, \hat{x}_1) in Wasserstein-2 (W_2) distance. Then, for $t \geq 1$ the AMP iteration reads

$$\mathbf{x}^t = \mathbf{Y} \hat{\mathbf{x}}^t - \beta_t \hat{\mathbf{x}}^{t-1}, \quad \hat{\mathbf{x}}^{t+1} = h_{t+1}(\mathbf{x}^t), \quad (8)$$

where we assume that $\hat{\mathbf{x}}^0 = \mathbf{0}$. Here, the function $h_{t+1}: \mathbb{R} \rightarrow \mathbb{R}$ is applied component-wise, and it can be chosen in order to exploit prior information about the signal; $\beta_1 = 0$ and, for $t \geq 2$, $\beta_t = \langle h'_t(\mathbf{x}^{t-1}) \rangle$, where h'_t denotes the derivative of h_t . The AMP estimator of \mathbf{X} is $\hat{\mathbf{x}}^t$, and the one of the spike $\mathbf{X}\mathbf{X}^\top$ is $M_{\text{AMP}}^t = \hat{\mathbf{x}}^t (\hat{\mathbf{x}}^t)^\top$. We refer to this algorithm as *Gaussian AMP*, since this is the AMP that is implemented for Gaussian noise (and known SNR). For a discussion on how (8) can be derived, we refer the interested reader to the review [39]. We note that the initialization (7) is impractical and one can design AMP iterations which are initialized with the eigenvector $\bar{\mathbf{v}}_N$ of the data matrix \mathbf{Y} associated to the largest eigenvalue $\bar{\nu}_N$, see [61, 60, 87]. We are also able to provide a state evolution result for the Gaussian AMP with spectral initialization, and this is discussed in Appendix C.4.

Spectral estimators. Finally, we consider spectral estimators of the form $C \bar{\mathbf{v}}_N \bar{\mathbf{v}}_N^\top$, where $\bar{\mathbf{v}}_N$ is the unit-norm eigenvector associated to the largest eigenvalue $\bar{\nu}_N(\mathbf{Y})$ and $C > 0$ is a scaling constant taking into account the spectrum of the data matrix. The asymptotic description of $\bar{\mathbf{v}}_N$ and $\bar{\nu}_N$ for additive and multiplicative low-rank perturbations of rotationally invariant matrices has been obtained in [21, 22, 6]. In particular, by Theorem 2.2 of [22], the squared overlap $\langle \mathbf{X}, \bar{\mathbf{v}}_N \rangle^2 / N$ converges a.s. to a constant $C(\rho, \lambda_*)$, whose explicit expression is recalled in Appendix A. We study two variants:

- The *optimal spectral estimator (OptSpec)* is given by $M_{\text{OS}} := C_{\text{OS}} \bar{\mathbf{v}}_N \bar{\mathbf{v}}_N^\top$, where $C_{\text{OS}}(\rho, \lambda_*) := C(\rho, \lambda_*)$ is the optimal scaling that minimizes the MSE.
- The *Gaussian mismatched spectral estimator (GauSpec)* is given by $M_{\text{GS}} := C_{\text{GS}} \bar{\mathbf{v}}_N \bar{\mathbf{v}}_N^\top$, where C_{GS} is the optimal scaling if the noise was Gaussian and the SNR was equal to $\lambda > 0$. By letting ρ_{SC} be the standard semi-circle law, we have $C_{\text{GS}}(\lambda) := C(\rho_{\text{SC}}, \lambda) = 1 - 1/\lambda$.

For both OptSpec and GauSpec, the estimator of \mathbf{X} is given by $\hat{\mathbf{x}}_{\text{GS}} = \hat{\mathbf{x}}_{\text{OS}} = \hat{\mathbf{x}}_{\text{Spec}} := \sqrt{N} \bar{\mathbf{v}}_N$. Notice that, while OptSpec requires knowing exactly the statistics of the noise and the SNR, GauSpec

represents the spectral estimator that would be used by a statistician who assumes the noise to be Gaussian and the SNR λ (instead of λ_*). An application of Theorem 2.2 in [22] gives that

$$\text{MSE}(M_{\text{OS}}) = \frac{1}{2}(1 - C^2(\rho, \lambda_*)), \quad \text{and} \quad \text{MSE}(M_{\text{GS}}) = \frac{1}{2}(1 + (1 - 1/\lambda)^2 - 2(1 - 1/\lambda)C(\rho, \lambda_*)).$$

Furthermore, the respective overlaps coincide and are given by $\sqrt{C_{\text{OS}}}$. Note that $\text{MSE}(M_{\text{GS}}) = \text{MSE}(M_{\text{OS}}) + \frac{1}{2}(C(\rho, \lambda_*) - (1 - 1/\lambda))^2$. Thus, as expected, we have $\text{MSE}(M_{\text{OS}}) \leq \text{MSE}(M_{\text{GS}})$. We also remark that, as the noise matrix is rotationally invariant, both the MSE and the overlap of the spectral estimators do not depend on the ground-truth signal distribution.

3 Main results: performance of inference algorithms

3.1 Mismatched Bayes estimator

Consider the following functions of $\lambda, \lambda_* > 0$ and of the asymptotic spectral noise density ρ :

$$\begin{aligned} M(\lambda, \lambda_*) &:= \left(1 - \frac{1}{\sqrt{\lambda\lambda_*}}\right) \left(1 - \frac{1}{\lambda_*} R'_\rho\left(\frac{1}{\sqrt{\lambda_*}}\right)\right) \mathbf{1}(\bar{h}\sqrt{\lambda_*} \geq 1 \cap \lambda\lambda_* > 1), \\ Q(\lambda, \lambda_*) &:= \left(1 - \frac{1}{\sqrt{\lambda\lambda_*}}\right)^2 \mathbf{1}(\bar{h}\sqrt{\lambda_*} \geq 1 \cap \lambda\lambda_* > 1) + \left(1 - \frac{\bar{h}}{\sqrt{\lambda}}\right)^2 \mathbf{1}(\bar{h}\sqrt{\lambda_*} < 1 \cap \sqrt{\lambda} > \bar{h}), \end{aligned} \tag{9}$$

where $\mathbf{1}(E)$ stands for the indicator function of the event E . The idea is that $Q(\lambda, \lambda_*)$ and $M(\lambda, \lambda_*)$ represent the squared norm of the mismatched Bayes estimator and its inner product with the signal, respectively. This leads to our main result on the mismatched Bayes MSE, which is stated below.

Theorem 1 (Performance of mismatched Bayes estimator). *Consider a spiked model (1). Let Assumptions 1 and 2 hold. Then, the MSE of the mismatched Bayes estimator (3) is given by*

$$\text{MSE}(M_{\text{mis}}) = \frac{1}{2}(1 + Q(\lambda, \lambda_*) - 2M(\lambda, \lambda_*)). \tag{10}$$

As a consequence of the auxiliary lemmas used in the proof of Theorem 1 and presented in Appendix B, we have that $N^{-2} \int P_{\text{mis}}(d\mathbf{x} | \mathbf{Y})(\mathbf{X}, \mathbf{x})^2$ converges almost surely to $M(\lambda, \lambda_*)$. Furthermore, for noise matrices regularized by arbitrarily small Wigner matrices, by using similar techniques as in [15] we can prove that $\|M_{\text{mis}}\|_{\text{F}}/N$ converges almost surely to $Q^{1/2}(\lambda, \lambda_*)$. We conjecture that this convergence holds even without the Wigner regularization. By assuming this conjecture, we have the following almost sure convergence result for the overlap:

$$\text{Overlap}_{\text{mis}} = \frac{M^{1/2}(\lambda, \lambda_*)}{Q^{1/4}(\lambda, \lambda_*)}, \tag{11}$$

which holds whenever the denominator is non-zero. When comparing the overlaps of different methods in Section 4, we will use (11) for the mismatched Bayes estimator.

Let us highlight the following property of the MSE (10) when there is no mismatch in the SNR:

$$\text{If } \lambda = \lambda_* \text{ and } \bar{h} \geq 1, \text{ then } \text{MSE}(M_{\text{mis}}) = \text{MSE}(M_{\text{GS}}). \tag{12}$$

However, in general $\text{MSE}(M_{\text{mis}})$ and $\text{MSE}(M_{\text{GS}})$ differ. E.g., when the assumed SNR λ is different from the real one λ_* , the norm of M_{mis} incorporates information about λ_* while the one for M_{GS} only depends on λ . Finally, we note that, from a statistical view-point, Assumption 1 can be replaced by the following: \mathbf{X} is uniformly distributed on $\mathbb{S}^{N-1}(\sqrt{N})$ and is independent of \mathbf{Z} , which may be a generic symmetric matrix with converging empirical spectral density.

Sketch of the proof of Theorem 1. The goal is to evaluate the asymptotic values of $\mathbb{E}\|M_{\text{mis}}\|_{\text{F}}^2$ and $\text{tr}\mathbb{E}[M_{\text{mis}}\mathbf{X}\mathbf{X}^\top]$, from which both the MSE and the mean overlap of the Bayes estimator can be obtained. Our strategy is to first compute the log-moment generating function f_{mis} of the mismatched Bayes model. Then, the above quantities of interest can be accessed by taking derivatives of f_{mis} with respect to appropriate parameters. However, as the noise in the inference model is not Gaussian, $\mathbb{E}\|M_{\text{mis}}\|_{\text{F}}^2$ cannot be computed in the standard way using an ‘‘I-MMSE’’ type of formula [46] (like it is done, e.g., in [71]). To address this issue, we introduce a generalized model for which the noise matrix is given by the original one plus an independent Wigner matrix multiplied by a small parameter ϵ . We then compute the log-moment generating function of this generalized model using the theory of low-rank perturbations of rotationally invariant random matrices [21] and of the low-rank spherical integral [45]. Finally, the desired results are obtained by a limiting argument in $\epsilon \rightarrow 0$.

3.2 Approximate message passing

Our main result is a characterization of the iterates (8) in the high-dimensional limit. We show that the joint empirical distribution of $(\mathbf{x}^1, \dots, \mathbf{x}^t, \hat{\mathbf{x}}^1, \dots, \hat{\mathbf{x}}^{t+1})$ converges (in W_2) to the random vector $(x_1, \dots, x_t, \hat{x}_1, \dots, \hat{x}_{t+1})$. The law of the random vector $(x_1, \dots, x_t, \hat{x}_1, \dots, \hat{x}_{t+1})$ is captured by a deterministic state evolution (SE) recursion, which can be expressed via a sequence of vectors $\boldsymbol{\mu}_t = (\mu_1, \dots, \mu_t)$ and matrices $\boldsymbol{\Sigma}_t, \boldsymbol{\Delta}_t, \mathbf{B}_t \in \mathbb{R}^{t \times t}$ defined as follows. We start with the initialization

$$\mu_1 = \sqrt{\lambda_*} \epsilon, \quad \boldsymbol{\Sigma}_1 = \bar{\kappa}_2 \mathbb{E}[\hat{x}_1^2], \quad \boldsymbol{\Delta}_1 = \mathbb{E}[\hat{x}_1^2], \quad \mathbf{B}_1 = \bar{\kappa}_1, \quad (13)$$

where λ_* is the true SNR (see (1)), and $\{\bar{\kappa}_k\}_{k \geq 1}$ are the free cumulants associated to the asymptotic spectral measure of the noise ρ . For $t \geq 1$, given $\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t, \boldsymbol{\Delta}_t, \mathbf{B}_t$, let

$$\begin{cases} (z_1, \dots, z_t) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_t) \text{ and independent of } (X, \hat{x}_1), \\ \mathbf{x}_t = z_t + \mu_t X - \bar{\beta}_t \hat{x}_{t-1} + \sum_{i=1}^t (\mathbf{B}_t)_{t,i} \hat{x}_i, \quad \hat{x}_{t+1} = h_{t+1}(x_t), \end{cases} \quad (14)$$

where we have defined $\hat{x}_0 = 0$, $\bar{\beta}_1 = 0$, and for $t \geq 2$, $\bar{\beta}_t = \mathbb{E}[h'_t(x_{t-1})]$. Then, the parameter μ_{t+1} is

$$\mu_{t+1} = \sqrt{\lambda_*} \mathbb{E}[X \hat{x}_{t+1}]. \quad (15)$$

Next, we compute the matrices $\boldsymbol{\Delta}_{t+1}$ and \mathbf{B}_{t+1} as

$$\begin{cases} (\boldsymbol{\Delta}_{t+1})_{i,j} = \mathbb{E}[\hat{x}_i \hat{x}_j], \quad 1 \leq i, j \leq t+1, \\ \mathbf{B}_{t+1} = \sum_{j=0}^t \bar{\kappa}_{j+1} \bar{\boldsymbol{\Phi}}_{t+1}^j \text{ where } (\bar{\boldsymbol{\Phi}}_{t+1})_{i,j} = 0 \text{ if } i \leq j, \text{ and } (\bar{\boldsymbol{\Phi}}_{t+1})_{i,j} = \mathbb{E}[\partial_j \hat{x}_i] \text{ if } i > j, \end{cases} \quad (16)$$

where $\partial_j \hat{x}_i$ denotes the partial derivative $\partial_{z_j} h_i(z_{i-1} + \mu_{i-1} X - \bar{\beta}_{i-1} \hat{x}_{i-2} + \sum_{k=1}^{i-1} (\mathbf{B}_t)_{i-1,k} \hat{x}_k)$. Finally, we define the covariance matrix $\boldsymbol{\Sigma}_{t+1}$ as

$$\boldsymbol{\Sigma}_{t+1} = \sum_{j=0}^{2t} \bar{\kappa}_{j+2} \boldsymbol{\Theta}_{t+1}^{(j)}, \quad \text{with } \boldsymbol{\Theta}_{t+1}^{(j)} = \sum_{i=0}^j (\bar{\boldsymbol{\Phi}}_{t+1})^i \boldsymbol{\Delta}_{t+1} (\bar{\boldsymbol{\Phi}}_{t+1}^\top)^{j-i}. \quad (17)$$

Note that the $t \times t$ top left sub-matrices of $\boldsymbol{\Sigma}_{t+1}, \boldsymbol{\Delta}_{t+1}$ and \mathbf{B}_{t+1} are given by $\boldsymbol{\Sigma}_t, \boldsymbol{\Delta}_t$ and \mathbf{B}_t . At this point, we are ready to state our state evolution characterization of the AMP algorithm (8). The result is presented in terms of pseudo-Lipschitz test functions. A function $\psi : \mathbb{R}^m \rightarrow \mathbb{R}$ is pseudo-Lipschitz of order 2, i.e., $\psi \in \text{PL}(2)$, if there is a constant $C > 0$ such that $\|\psi(\mathbf{x}) - \psi(\mathbf{y})\| \leq C(1 + \|\mathbf{x}\| + \|\mathbf{y}\|)\|\mathbf{x} - \mathbf{y}\|$. We also make the following assumption on the functions $\{h_{t+1}\}_{t \geq 1}$.

Assumption 3. The function $h_{t+1} : \mathbb{R} \rightarrow \mathbb{R}$ is Lipschitz, and the partial derivatives

$$\partial_{z_k} h_{t+1} \left(z_t + \mu_t X - \bar{\beta}_t \hat{x}_{t-1} + \sum_{k=1}^t (\mathbf{B}_t)_{t,k} \hat{x}_k \right)$$

are continuous on a set of probability 1, under the laws of (z_1, \dots, z_t) and $(\hat{x}_1, \dots, \hat{x}_t)$ given in (14).

This requirement covers most practically relevant choices of h_{t+1} (e.g., soft-thresholding or ReLU), and it is rather standard in the AMP literature, see e.g. [39, 60, 82, 87].

Theorem 2 (State evolution of Gaussian AMP). Consider a spiked model (1) and the AMP algorithm (8) initialized as in (7). Let Assumptions 1 and 3 hold, and let $\psi : \mathbb{R}^{2t+2} \rightarrow \mathbb{R}$ be any pseudo-Lipschitz functions of order 2. Then, for each $t \geq 1$, we almost surely have that, as $N \rightarrow +\infty$,

$$\frac{1}{N} \sum_{i=1}^N \psi((\mathbf{x}^1)_i, \dots, (\mathbf{x}^t)_i, (\hat{\mathbf{x}}^1)_i, \dots, (\hat{\mathbf{x}}^{t+1})_i, (\mathbf{X})_i) \rightarrow \mathbb{E} \psi(x_1, \dots, x_t, \hat{x}_1, \dots, \hat{x}_{t+1}, X), \quad (18)$$

where the random variables on the right are defined in (14).

By using Definition 6.7 and Theorem 6.8 of [84], one readily obtains that (18) is equivalent to the convergence of $(\mathbf{x}^1, \dots, \mathbf{x}^t, \hat{\mathbf{x}}^1, \dots, \hat{\mathbf{x}}^{t+1}, \mathbf{X})$ to $(x_1, \dots, x_t, \hat{x}_1, \dots, \hat{x}_{t+1}, X)$ in W_2 distance. Applying (18) to the pseudo-Lipschitz functions $\psi(\hat{x}_{t+1}, X) = (\hat{x}_{t+1} - X)^2$, $\psi(\hat{x}_{t+1}, X) = \hat{x}_{t+1} \cdot X$ and $\psi(\hat{x}_{t+1}, X) = (\hat{x}_{t+1})^2$, we obtain a high-dimensional characterization of the AMP performance.

Corollary 1. Consider the setting of Theorem 2. Then, for each $t \geq 1$, we almost surely have that

$$\text{MSE}(M_{\text{AMP}}^t) := \lim_{N \rightarrow +\infty} \frac{1}{2N^2} \mathbb{E} \|\mathbf{X} \mathbf{X}^\top - \hat{\mathbf{x}}^t (\hat{\mathbf{x}}^t)^\top\|_{\text{F}}^2 = \frac{1}{2} (1 - 2(\mathbb{E}[\hat{x}_t \cdot X])^2 + (\mathbb{E}[(\hat{x}_t)^2])^2),$$

$$\text{Overlap}(\hat{\mathbf{x}}^t) := \lim_{N \rightarrow +\infty} \frac{|\langle \mathbf{X}, \hat{\mathbf{x}}^t \rangle|}{\|\hat{\mathbf{x}}^t\| \cdot \|\mathbf{X}\|} = \frac{|\mathbb{E}[\hat{x}_t \cdot X]|}{\sqrt{\mathbb{E}[(\hat{x}_t)^2]}}.$$

Sketch of the proof of Theorem 2. The key insight is to exploit the flexibility given by the denoisers of the abstract AMP iteration [87] to “correct” the mismatch. These denoisers are denoted by $\{\hat{h}_t\}_{t \geq 2}$ (see (33) and (34) in Appendix C.1) and they serve solely as a proof technique, hence they do not have an impact on the practicality of the algorithm; in contrast, the denoisers $\{h_t\}_{t \geq 2}$ (see (8)) can be chosen by the algorithm designer to take advantage of the information available about the signal. Thus, we can design an *auxiliary AMP* such that (i) it admits a state evolution characterization, and (ii) its iterates are close to (8). This requires the delicate induction argument in Appendix C: in Appendix C.1, we define the auxiliary AMP; in Appendix C.2, we prove a state evolution result for it; in Appendix C.3, we show by induction that (i) the state evolution parameters of the auxiliary AMP coincide with those defined in this section, and (ii) the iterates of the auxiliary AMP are close (in ℓ_2 norm) to the iterates of the AMP (8), thus concluding the proof of Theorem 2.

4 Numerical results and discussion

In all experiments, the density ρ of the eigenvalues of the noise has unit variance and the distribution of the ground-truth signal is uniform on the sphere of radius \sqrt{N} . For the sake of simplicity, we decouple the effect of the mismatch in (i) the noise statistics and (ii) the SNR. More specifically, in the first two examples, we set $\lambda = \lambda_*$, so that only mismatch in the noise statistics is present; and in the last example, the noise is Gaussian, so that there is only SNR mismatch. Additional results when \mathbf{Z} is the free convolution of Rademacher and semicircle spectra are reported in Appendix D.3.

Models of mismatch. (a) *Marcenko-Pastur spectrum.* An example of non-symmetric density is the Marcenko-Pastur law $\rho(x) = \sqrt{x(4-x)}/(2\pi x) \mathbf{1}(x \in [0, 4])$. The R -transform is $R_\rho(x) = 1/(1-x)$, and the results are displayed in Fig. 1a (the same formulas apply if the law is centered).

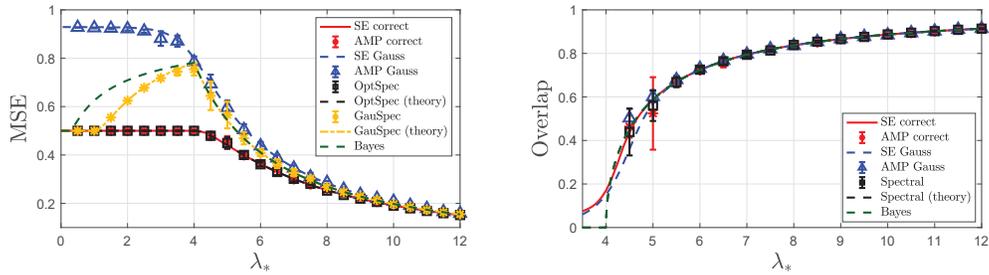
(b) *Uniform spectrum.* We let ρ be the uniform distribution $\mathcal{U}[-\sqrt{3}, \sqrt{3}]$. In this case, we have $R_\rho(x) = \sqrt{3}/\tanh(\sqrt{3}x) - 1/x$, and the results are in Fig. 1b.

(c) *Wigner matrix/semicircle spectrum with mismatched SNR.* We consider a Gaussian noise matrix \mathbf{W} (as assumed by the statistician), but with mismatched SNR by setting $\lambda = 4\lambda_*$, see Fig. 1c.

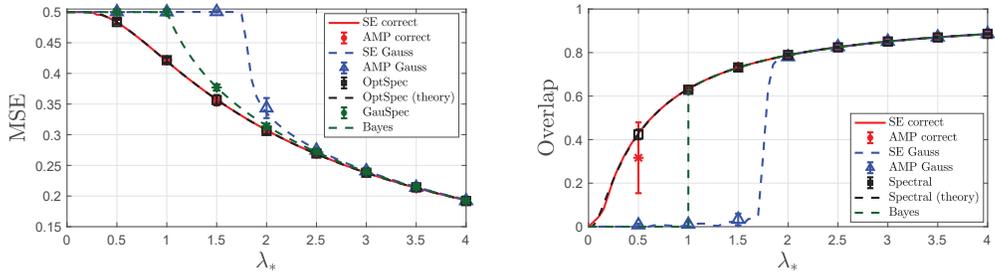
We remark that Assumption 1 is verified for all models. Furthermore, in all the three cases, the support of ρ is a single non-empty interval. Hence, Theorem 2.2 of [10] and Proposition 2.4 of [21] can be used to show that Assumption 2 holds.

Inference algorithms and set-up. The estimators of the spike $\mathbf{X}\mathbf{X}^\top$ are compared in terms of the MSE (4) (plots on the left). The AMP and spectral estimators of \mathbf{X} are compared in terms of the overlap (5), and the Bayes estimator of \mathbf{X} in terms of the overlap (6) (plots on the right). We evaluate those formulas at $N = 8000$ for the various algorithms, and in the $N \rightarrow +\infty$ limit for the theoretical predictions. The *correct AMP* (together with its own SE) is in red. This AMP is *correct* in the sense that the statistician is aware of the noise statistics and, thus, incorporates the right Onsager corrections. The form of the correct AMP and the corresponding state evolution are readily obtained from the results in [37, 87] and, for the reader’s convenience, they are recalled in Appendix D.1. As non-linearities, we use the posterior-mean denoising functions of Section 5.1 in [87], and we estimate the SE parameters consistently from data (see Appendix D.2). The *Gaussian AMP* (together with its own SE) is in blue. This is the AMP algorithm (8), which is chosen when the noise statistics or the SNR are unknown (as it would be optimal for Gaussian noise and known SNR). Its state evolution is given by Theorem 2. As non-linearities $\{h_{t+1}\}_{t \geq 1}$, we use again the posterior-mean denoising functions, and estimate consistently the state evolution parameters from data (see Appendix D.2). We note that the denoisers of the Gaussian AMP depend only on a single iterate, while the denoisers of the correct AMP incorporate all the past iterates. In contrast, the state evolution parameters of the Gaussian AMP at time t depend on *all* the past, in order to correct for the mismatch. The *mismatched Bayes estimator* is plotted in green. Its MSE and overlap are given by (10) and (11), respectively. The *optimal spectral estimator (OptSpec)* is plotted in black, and the *Gaussian mismatched spectral estimator (GauSpec)* is plotted in yellow in Fig. 1a and 1c (where it differs from Bayes) or in green in Fig. 1b (where it coincides with Bayes). The performance measures of both variants of spectral estimators are given at the end of Section 2.3. Each experiment is repeated for 10 independent runs. We report the average and error bars at 1 standard deviation.

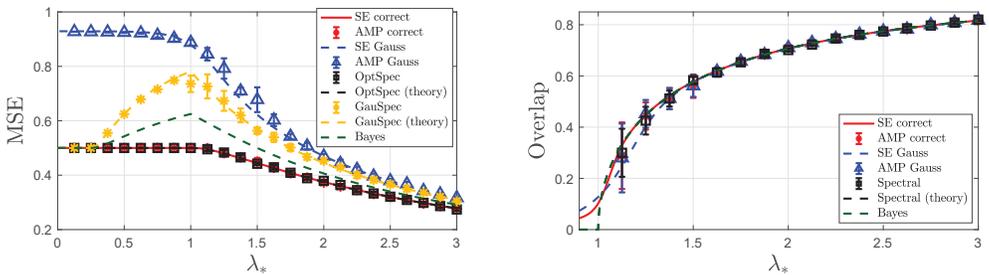
A rich and surprising phenomenology. (i) An intriguing phenomenon is that, in terms of MSE, the *Gaussian AMP* does not perform as well as the *mismatched Bayes estimator*. This effect occurs



(a) Marcenko-Pastur noise.



(b) Uniform eigenvalues distribution in $[-\sqrt{3}, \sqrt{3}]$.



(c) Wigner matrix/semicircle spectrum with mismatched SNR $\lambda = 4\lambda_*$.

Figure 1: MSE (on the left) and overlap (on the right), as a function of the true SNR λ_* in three mismatched settings.

in all our settings, and it is surprising as the Gaussian AMP (8) is precisely designed to efficiently compute the Bayes estimator (3). In the Bayes-optimal case, it does so (outside of its computational gap) [12]. However, if there is a mismatch, AMP does not compute the (mismatched) posterior mean. This finding adds to the already existing evidence [5] that AMP is poorly understood, and studying the fundamental reasons behind this behavior is an exciting avenue for future research.

(ii) By comparing MSE² and overlap curves, we understand that the discrepancy between Gaussian AMP and Bayes comes from an incorrect estimation of the signal norm. In fact, at large enough SNR, the overlaps of Bayes and Gaussian AMP match, and they even coincide with the overlaps of optimal algorithms (correct AMP and OptSpec), which exploit knowledge about the noise structure. Understanding the origin of the MSE discrepancy (namely, a wrong estimate of the signal length) can potentially lead to procedures which correct for this effect and, hence, reduce the MSE.

(iii) When there is no SNR mismatch and $\bar{h} \geq 1$ (cf. (12)), the MSEs of the Bayes and Gaussian spectral estimators match (see Fig. 1b). Both have no information about the noise structure and about the signal distribution. Yet, this equality is remarkable given that the spectral estimator is the solution of an optimization problem (it is a “zero temperature” estimator in physics parlance), while the Bayes

²The MSE is constructed from the norm $\mathbb{E}\|M(\mathbf{Y})\|_{\mathbb{F}}^2$ and the inner product $\text{tr}\mathbb{E}[M(\mathbf{Y})\mathbf{X}\mathbf{X}^T]$.

estimator aims at computing the mean of a certain posterior distribution (it is a “finite temperature” estimator).

(iv) All mismatched estimators are outperformed by the optimal spectral method and the correct AMP. Both these algorithms take advantage of the noise statistics and achieve the same MSE. This suggests that the two estimators are Bayes-optimal. We remark that solving this conjecture would require understanding the information-theoretic limits of low-rank estimation with structured noise.

(v) Finally, a striking phenomenon – first observed in [71] – is that the Bayes and Gaussian spectral MSE curves may be non-decreasing with the true SNR λ_* , see Figs. 1a and 1c. Going beyond the analysis in [71], we remark that, for large enough SNR, all estimators yield the same overlap and therefore all “point in the correct direction” given by the leading eigenvector, see also part (ii) of this discussion. This links the initial MSE increase to a wrong estimation of the signal’s norm, due to an over-confidence in the data quality (recall that the assumed SNR λ is equal to $4\lambda_*$ in Fig. 1c).

Acknowledgements

M. Mondelli was partially supported by the 2019 Lopez-Loreta Prize. The authors acknowledge discussions with A. Krajenbrink, M. Robinson, A. Depope, N. Macris and F. Pourkamali.

References

- [1] Emmanuel Abbe. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research*, 18(1):6446–6531, 2017.
- [2] Gad Abraham, Yixuan Qiu, and Michael Inouye. Flashpca2: principal component analysis of biobank-scale genotype datasets. *Bioinformatics*, 2017.
- [3] Madhu Advani and Surya Ganguli. Statistical mechanics of optimal convex inference in high dimensions. *Physical Review X*, 6(3):031034, 2016.
- [4] Greg W. Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*. Cambridge university press, 2010.
- [5] Fabrizio Antenucci, Silvio Franz, Pierfrancesco Urbani, and Lenka Zdeborová. Glassy nature of the hard phase in inference problems. *Phys. Rev. X*, 9, Jan. 2019.
- [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] Sayantan Banerjee, Ismaël Castillo, and Subhashis Ghosal. Bayesian inference in high-dimensional models, 2021. [arXiv:2101.04491](https://arxiv.org/abs/2101.04491).
- [10] Zhigang Bao, László Erdős, and Kevin Schnelli. On the support of the free additive convolution. *Journal d’Analyse Mathématique*, 142(1):323–348, 2020.
- [11] Jean Barbier, Wei-Kuo Chen, Dmitry Panchenko, and Manuel Sáenz. Performance of bayesian linear regression in a model with mismatch, 2021. [arXiv:2107.06936](https://arxiv.org/abs/2107.06936).
- [12] 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.
- [13] 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.

- [14] 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):1133–1185, 2019.
- [15] 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.
- [16] Jean Barbier, Nicolas Macris, Antoine Maillard, and Florent Krzakala. The mutual information in random linear estimation beyond i.i.d. matrices. In *International Symposium on Information Theory*, pages 1390–1394, 2018.
- [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*, 2020.
- [18] 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.
- [19] Mohsen Bayati and Andrea Montanari. The LASSO risk for gaussian matrices. *IEEE Transactions on Information Theory*, 58(4):1997–2017, 2012.
- [20] Robert M Bell and Yehuda Koren. Lessons from the netflix prize challenge. *Acm Sigkdd Explorations Newsletter*, 9(2):75–79, 2007.
- [21] 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.
- [22] 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.
- [23] Bhaswar B. Bhattacharya and Subhabrata Sen. High temperature asymptotics of orthogonal mean-field spin glasses. *Journal of Statistical Physics*, 162(1):63–80, 2016.
- [24] 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.
- [25] Jelena Bradic. Robustness in sparse high-dimensional linear models: relative efficiency and robust approximate message passing. *Electronic Journal of Statistics*, 10(2):3894–3944, 2016.
- [26] 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.
- [27] Ismaël Castillo, Johannes Schmidt-Hieber, and Aad Van der Vaart. Bayesian linear regression with sparse priors. *The Annals of Statistics*, 43(5):1986–2018, 2015.
- [28] Yash Deshpande, Emmanuel Abbe, and Andrea Montanari. Asymptotic mutual information for the balanced binary stochastic block model. *Information and Inference: A Journal of the IMA*, 6(2):125–170, 2017.
- [29] Yash Deshpande and Andrea Montanari. Information-theoretically optimal sparse PCA. In *IEEE International Symposium on Information Theory*, pages 2197–2201, 2014.
- [30] 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.
- [31] David L. Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.

- [32] David L. Donoho and Andrea Montanari. High dimensional robust M-estimation: asymptotic variance via approximate message passing. *Probability Theory and Related Fields*, 166(3-4):935–969, 2016.
- [33] Rishabh Dudeja, Junjie Ma, and Arian Maleki. Information theoretic limits for phase retrieval with subsampled haar sensing matrices. *IEEE Transactions on Information Theory*, 66(12):8002–8045, 2020.
- [34] Nouredine El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, 170(1-2):95–175, 2018.
- [35] Nouredine El Karoui, Derek Bean, Peter J. Bickel, Chingway Lim, and Bin Yu. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, 110(36):14557–14562, 2013.
- [36] Barbara E. Engelhardt and Matthew Stephens. Analysis of population structure: a unifying framework and novel methods based on sparse factor analysis. *PLoS genetics*, 6(9):e1001117, 2010.
- [37] Zhou Fan. Approximate message passing algorithms for rotationally invariant matrices. *The Annals of Statistics*, 50(1):197–224, 2022.
- [38] Zhou Fan and Yihong Wu. The replica-symmetric free energy for ising spin glasses with orthogonally invariant couplings, 2021. [arXiv:2105.02797](https://arxiv.org/abs/2105.02797).
- [39] Oliver Y. Feng, Ramji Venkataramanan, Cynthia Rush, and Richard J. Samworth. A unifying tutorial on Approximate Message Passing, 2021. [arXiv:2105.02180](https://arxiv.org/abs/2105.02180).
- [40] 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.
- [41] Alyson K. Fletcher and Sundeep Rangan. Iterative reconstruction of rank-one matrices in noise. *Information and Inference: A Journal of the IMA*, 7(3):531–562, 2018.
- [42] Marylou Gabrié, Andre Manoel, Clément Luneau, Jean Barbier, Nicolas Macris, Florent Krzakala, and Lenka Zdeborová. Entropy and mutual information in models of deep neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12), 2019.
- [43] Cedric Gerbelot, Alia Abbata, and Florent Krzakala. Asymptotic errors for teacher-student convex generalized linear models (or: how to prove Kabashima’s replica formula), 2020. [arXiv:2006.06581](https://arxiv.org/abs/2006.06581).
- [44] Alice Guionnet and Jiaoyang Huang. Large deviations asymptotics of rectangular spherical integral, 2021. [arXiv:2106.07146](https://arxiv.org/abs/2106.07146).
- [45] 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.
- [46] Dongning Guo, Shlomo Shamai, and Sergio Verdú. Mutual information and minimum mean-square error in gaussian channels. *IEEE Transactions on Information Theory*, 51(4):1261–1282, 2005.
- [47] TianQi Hou, YuHao Liu, Teng Fu, and Jean Barbier. Sparse superposition codes under VAMP decoding with generic rotational invariant coding matrices, 2022. [arXiv:2202.04541](https://arxiv.org/abs/2202.04541).
- [48] 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.
- [49] Iain M. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of statistics*, 29(2):295–327, 2001.

- [50] Iain M. Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009.
- [51] 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.
- [52] Antti Knowles and Jun Yin. The isotropic semicircle law and deformation of Wigner matrices. *Communications on Pure and Applied Mathematics*, 66(11):1663–1749, 2013.
- [53] Satish Babu Korada and Nicolas Macris. Exact solution of the gauge symmetric p-spin glass model on a complete graph. *Journal of Statistical Physics*, 136(2):205–230, 2009.
- [54] Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. *Probability Theory and Related Fields*, 173(3-4):859–929, 2018.
- [55] 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.
- [56] 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*, 2021.
- [57] 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.
- [58] 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*, 2020.
- [59] Marco Mondelli and Ramji Venkataramanan. Approximate message passing with spectral initialization for generalized linear models. In *International Conference on Artificial Intelligence and Statistics*, 2021.
- [60] Marco Mondelli and Ramji Venkataramanan. PCA initialization for approximate message passing in rotationally invariant models. In *Neural Information Processing Systems*, 2021.
- [61] Andrea Montanari and Ramji Venkataramanan. Estimation of low-rank matrices via approximate message passing. *Annals of Statistics*, 45(1):321–345, 2021.
- [62] Cristopher Moore. The computer science and physics of community detection: landscapes, phase transitions, and hardness, 2017. [arXiv:1702.00467](https://arxiv.org/abs/1702.00467).
- [63] Sumit Mukherjee and Subhabrata Sen. Variational inference in high-dimensional linear regression, 2021. [arXiv:2104.12232](https://arxiv.org/abs/2104.12232).
- [64] Jonathan Novak. Three lectures on free probability. *Random matrix theory, interacting particle systems, and integrable systems*, 65(309-383):13, 2014.
- [65] 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.
- [66] Parthe Pandit, Mojtaba Sahraee-Ardakan, Sundeep Rangan, Philip Schniter, and Alyson K Fletcher. Inference with deep generative priors in high dimensions. *IEEE Journal on Selected Areas in Information Theory*, 1(1):336–347, 2020.
- [67] Giorgio Parisi and Marc Potters. Mean-field equations for spin models with orthogonal interaction matrices. *Journal of Physics A: Mathematical and General*, 28(18):5267, 1995.
- [68] Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, pages 1617–1642, 2007.

- [69] 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.
- [70] Marc Potters and Jean-Philippe Bouchaud. *A first course in random matrix theory*. Cambridge University Press, 2021.
- [71] Farzad Pourkamali and Nicolas Macris. Mismatched estimation of rank-one symmetric matrices under gaussian noise, 2021. [arXiv:2107.08927](https://arxiv.org/abs/2107.08927).
- [72] Alkes L. Price, Nick J. Patterson, Robert M. Plenge, Michael E. Weinblatt, Nancy A. Shadick, and David Reich. Principal components analysis corrects for stratification in genome-wide association studies. *Nature genetics*, 38(8):904–909, 2006.
- [73] Sundeep Rangan. Generalized approximate message passing for estimation with random linear mixing. In *International Symposium on Information Theory*, pages 2168–2172, 2011.
- [74] Sundeep Rangan, Philip Schniter, and Alyson K Fletcher. Vector approximate message passing. *IEEE Transactions on Information Theory*, 65(10):6664–6684, 2019.
- [75] Kolyan Ray and Botond Szabó. Variational Bayes for high-dimensional linear regression with sparse priors. *Journal of the American Statistical Association*, pages 1–12, 2021.
- [76] Subrata Sarkar, Rizwan Ahmad, and Philip Schniter. MRI image recovery using damped denoising vector AMP. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8108–8112, 2021.
- [77] 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.
- [78] 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.
- [79] Koujin Takeda, Shinsuke Uda, and Yoshiyuki Kabashima. Analysis of CDMA systems that are characterized by eigenvalue spectrum. *EPL (Europhysics Letters)*, 76(6):1193, 2006.
- [80] Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Precise error analysis of regularized M-estimators in high dimensions. *IEEE Transactions on Information Theory*, 64(8):5592–5628, 2018.
- [81] Christos Thrampoulidis, Samet Oymak, and Babak Hassibi. Regularized linear regression: A precise analysis of the estimation error. In *Conference on Learning Theory*, 2015.
- [82] Ramji Venkataramanan, Kevin Kögler, and Marco Mondelli. Estimation in rotationally invariant generalized linear models via approximate message passing, 2021. [arXiv:2112.04330](https://arxiv.org/abs/2112.04330).
- [83] Sergio Verdú. Mismatched estimation and relative entropy. *IEEE Transactions on Information Theory*, 56(8):3712–3720, 2010.
- [84] Cédric Villani. *Optimal transport: Old and new*, volume 338. Springer Science & Business Media, 2008.
- [85] Tsachy Weissman. The relationship between causal and noncausal mismatched estimation in continuous-time awgn channels. *IEEE Transactions on Information Theory*, 56(9):4256–4273, 2010.
- [86] Xinyi Zhong, Chang Su, and Zhou Fan. Empirical bayes PCA in high dimensions, 2020. [arXiv:2012.11676](https://arxiv.org/abs/2012.11676).
- [87] Xinyi Zhong, Tianhao Wang, and Zhou Fan. Approximate message passing for orthogonally invariant ensembles: Multivariate non-linearities and spectral initialization, 2021. [arXiv:2110.02318](https://arxiv.org/abs/2110.02318).
- [88] Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *Journal of computational and graphical statistics*, 15(2):265–286, 2006.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] Our model assumptions are clearly explained in Assumptions 1, 2 and 3. To obtain the overlap of the mismatched Bayes estimator, we make a conjecture on the convergence of $\|M_{\text{mis}}\|_{\mathbb{F}}/N$ (see the discussion right after the statement of Theorem 1).
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] We do not see direct paths to applications with negative societal impacts.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] Our model assumptions are clearly explained in Assumptions 1, 2 and 3. To obtain the overlap of the mismatched Bayes estimator, we make a conjecture on the convergence of $\|M_{\text{mis}}\|_{\mathbb{F}}/N$ (see the discussion right after the statement of Theorem 1).
 - (b) Did you include complete proofs of all theoretical results? [Yes] Proof sketches are included right after the corresponding statements. The full proofs of Theorem 1 and 2 are in Appendices B and C, respectively.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] In Section 4, we describe the experimental setup containing the instructions to reproduce the results of Figure 1 (and of Figure 2 contained in Appendix D.3 as well). Additional details are given in Appendices D.1 and D.2. The code that implements the various algorithms is included in the supplementary material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] All details of the experimental setup are described in Section 4 and in Appendices D.1 and D.2.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] Each experiment is repeated for 10 independent runs, and we report the average together with the confidence interval at 1 standard deviation.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A] We did not perform large scale experiments that required to specify this type of details.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A] We are not using existing assets or curating/releasing new assets.
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] We did not use crowdsourcing or conducted research with human subjects.
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]