The Annals of Statistics
2021, Vol. 49, No. 1, 321–345
https://doi.org/10.1214/20-AOS1958

ESTIMATION OF LOW-RANK MATRICES VIA APPROXIMATE MESSAGE PASSING

BY ANDREA MONTANARI[1] AND RAMJI VENKATARAMANAN[2]

[1]*Department of Electrical Engineering and Department of Statistics, Stanford University, montanari@stanford.edu*
[2]*Department of Engineering, University of Cambridge, ramji.v@eng.cam.ac.uk*

Consider the problem of estimating a low-rank matrix when its entries are perturbed by Gaussian noise, a setting that is also known as "spiked model" or "deformed random matrix." If the empirical distribution of the entries of the spikes is known, optimal estimators that exploit this knowledge can substantially outperform simple spectral approaches. Recent work characterizes the asymptotic accuracy of Bayes-optimal estimators in the high-dimensional limit. In this paper, we present a practical algorithm that can achieve Bayes-optimal accuracy above the spectral threshold. A bold conjecture from statistical physics posits that no polynomial-time algorithm achieves optimal error below the same threshold (unless the best estimator is trivial).

Our approach uses Approximate Message Passing (AMP) in conjunction with a spectral initialization. AMP algorithms have proved successful in a variety of statistical estimation tasks, and are amenable to exact asymptotic analysis via state evolution. Unfortunately, state evolution is uninformative when the algorithm is initialized near an unstable fixed point, as often happens in low-rank matrix estimation problems. We develop a new analysis of AMP that allows for spectral initializations, and builds on a decoupling between the outlier eigenvectors and the bulk in the spiked random matrix model.

Our main theorem is general and applies beyond matrix estimation. However, we use it to derive detailed predictions for the problem of estimating a rank-one matrix in noise. Special cases of this problem are closely related—via universality arguments—to the network community detection problem for two asymmetric communities. For general rank-one models, we show that AMP can be used to construct confidence intervals and control false discovery rate.

We provide illustrations of the general methodology by considering the cases of sparse low-rank matrices and of block-constant low-rank matrices with symmetric blocks (we refer to the latter as to the "Gaussian block model").

1. Introduction. The "spiked model" is the simplest probabilistic model of a data matrix with a latent low-dimensional structure. Consider, to begin with, the case of a symmetric matrix. The data are written as the sum of a low-rank matrix (the signal) and Gaussian component (the noise):

$$\boldsymbol{A} = \sum_{i=1}^{k} \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^{\mathsf{T}} + \boldsymbol{W}. \tag{1.1}$$

Here, $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k$ are nonrandom numbers, $\boldsymbol{v}_i \in \mathbb{R}^n$ are nonrandom vectors and $\boldsymbol{W} \sim \mathsf{GOE}(n)$ is a matrix from the Gaussian orthogonal ensemble.[1] The asymmetric (rectangular)

Received May 2018; revised August 2019.
MSC2020 subject classifications. Primary 62H99; secondary 62F15, 62E20.
Key words and phrases. Low-rank matrix estimation, approximate message passing, spectral initialization.

[1]Recall that this means that $\boldsymbol{W} = \boldsymbol{W}^{\mathsf{T}}$, and the entries $(W_{ij})_{i\leq j\leq n}$ are independent with $(W_{ii})_{i\leq n} \sim_{\text{iid}} \mathsf{N}(0, 2/n)$ and $(W_{ij})_{i<j\leq n} \sim_{\text{iid}} \mathsf{N}(0, 1/n)$.

version of the same model is also of interest. In this case, we observe $\boldsymbol{A} \in \mathbb{R}^{n\times d}$ given by

$$\boldsymbol{A} = \sum_{i=1}^{k} \sqrt{\lambda_i}\boldsymbol{u}_i\boldsymbol{v}_i^{\mathsf{T}} + \boldsymbol{W}, \tag{1.2}$$

where $\boldsymbol{W}$ is a noise matrix with entries $(W_{ij})_{i\le n, j\le d} \sim_{\text{iid}} \mathsf{N}(0, 1/n)$. An important special case assumes $\boldsymbol{u}_i \sim \mathsf{N}(0, \boldsymbol{I}_n/n)$. In this case,[2] the rows of $\boldsymbol{A}$ are i.i.d. samples from a high-dimensional Gaussian $\boldsymbol{a}_i \sim \mathsf{N}(\boldsymbol{0}, \boldsymbol{\Sigma})$ where $\boldsymbol{\Sigma} = (\sum_{i=1}^{k} \lambda_i \boldsymbol{v}_i\boldsymbol{v}_i^{\mathsf{T}} + \boldsymbol{I}_d)/n$. Theoretical analysis of this spiked covariance model has led to a number of important statistical insights [32, 33].

Within probability theory, the spiked model (1.1) is also known as "deformed GOE" or "deformed Wigner random matrix," and the behavior of its eigenvalues and eigenvectors has been studied in exquisite detail [3, 4, 8, 9, 14, 24, 36]. The most basic phenomenon unveiled by this line of work is the so-called BBAP phase transition, first discovered in the physics literature [28], and named after the authors of [3]. Let k_* be the number of rank-one terms with $|\lambda_i| > 1$. Then the spectrum of $\boldsymbol{A}$ is formed by a bulk of eigenvalues in the interval $[-2, 2]$ (whose distribution follows Wigner's semicircle), plus k_* outliers that are in one-to-one correspondence with the large rank-one terms in (1.1). The eigenvectors associated to the outliers exhibit a significant correlation with the corresponding vectors $\boldsymbol{v}_i$. To simplify the discussion, in the rest of this Introduction we will assume that $\lambda_i \ge 0$ for all i.

The spiked model (1.1), (1.2) and their generalizations have also been studied from a statistical perspective [31, 48]. A fundamental question in this context is to estimate the vectors $\boldsymbol{v}_i$ from a single realization of the matrix $\boldsymbol{A}$. It is fair to say that this question is relatively well understood when the vectors $\boldsymbol{v}_i$ are unstructured, for example, they are a uniformly random orthonormal set (distributed according to the Haar measure). In this case, and in the high-dimensional limit $n, d \to \infty$, the best estimator of vector $\boldsymbol{v}_i$ is the ith eigenvector of $\boldsymbol{A}$. Random matrix theory provides detailed information about its asymptotic properties.

This paper is concerned with the case in which the vectors $\boldsymbol{v}_i$ are structured, for example, they are sparse, or have bounded entries. This structure is not captured by spectral methods, and other approaches lead to significantly better estimators. This scenario is relevant for a broad range of applications, including sparse principal component analysis [17, 33, 53], nonnegative principal component analysis [38, 43], community detection under the stochastic block model [1, 16, 45], and so on. Understanding what are optimal ways of exploiting the structure of signals is—to a large extent—an open problem.

Significant progress has been achieved recently under the assumption that the vectors $(v_{1,j}, \dots, v_{k,j}) \in \mathbb{R}^k$ (i.e., the k-dimensional vectors obtained by taking the jth component of the vectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$) are approximately i.i.d. (across $j \in \{1, \dots, n\}$) with some common distribution $\mu_{\boldsymbol{U}}$ on $\mathbb{R}^k$. This is, for instance, the case if each $\boldsymbol{v}_\ell$ has i.i.d. components, and distinct vectors are independent (but mutual independence between $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$ is not required). Following heuristic derivations using statistical physics methods (see, e.g., [40]), closed form expressions have been rigorously established for the Bayes-optimal estimation error in the limit $n \to \infty$ (with λ_i's fixed). We refer to [16, 17] for special cases and to [5, 37, 39, 42] for an increasingly general theory.

Unfortunately, there is no general algorithm that computes the Bayes-optimal estimator and is guaranteed to run in polynomial time. Markov Chain Monte Carlo can have exponentially large mixing time and is difficult to analyze [26]. Variational methods are nonconvex and do not come with consistency guarantees [12]. Classical convex relaxations do not generally achieve the Bayes optimal error, since they incorporate limited prior information [30].

[2]For the formal analysis of this model, it will be convenient to consider the case of deterministic vectors $\boldsymbol{u}_i$, $\boldsymbol{v}_i$ satisfying suitable asymptotic conditions. However, these conditions hold almost surely, for example, $\boldsymbol{u}_i \sim \mathsf{N}(0, \boldsymbol{I}_n/n)$.

In the positive direction, approximate message passing (AMP) algorithms have been successfully applied to a number of low-rank matrix estimation problems [25, 35, 43, 47, 50]. In particular, AMP was proved to achieve the Bayes optimal estimation error in special cases of the model (1.1), in the high-dimensional limit $n \to \infty$ [17, 18]. In fact, a bold conjecture from statistical physics suggests that the estimation error achieved by AMP is the same that can be achieved by the optimal polynomial-time algorithm.

An important feature of AMP is that it admits an exact characterization in the limit $n \to \infty$ that goes under the name of *state evolution* [7, 13, 22]. There is however one notable case in which the state evolution analysis of AMP falls short of its goal: when AMP is initialized near an unstable fixed point. This is typically the case for the problem of estimating the vectors $\boldsymbol{v}_i$'s in the spiked model (1.1). (We refer to the next section for a discussion of this point.)

In order to overcome this problem, we propose a two-step algorithm:

1. We compute the principal eigenvectors $\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_{k_*}$ of $\boldsymbol{A}$, which correspond to the outlier eigenvalues.
2. We run AMP with an initialization that is correlated with these eigenvectors.

Our main result (Theorem 5) is a general asymptotically exact analysis of this type of procedure. The analysis applies to a broad class of AMP algorithms, with initializations that are obtained by applying separable functions to the eigenvectors $\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_{k_*}$ (under some technical conditions). Let us emphasize that *our core technical result (state-evolution analysis) is completely general and applies beyond low-rank matrix estimation.*

The rest of the paper is organized as follows.

- Section 2 applies our main results to the problem of estimating a rank-one matrix in Gaussian noise (the case $k = 1$ of the model (1.1)). We compute the asymptotic empirical distribution of our estimator. In particular, this characterizes the asymptotics of all sufficiently regular separable losses.

  We then illustrate how this state evolution analysis can be used to design specific AMP algorithms, depending on what prior knowledge we have about the entries of $\boldsymbol{v}_1$. In a first case study, we only know that $\boldsymbol{v}_1$ is sparse, and analyze an algorithm based on iterative soft thresholding. In the second, we assume that the empirical distribution of the entries of $\boldsymbol{v}_1$ is known, and develop a Bayes-AMP algorithm. The asymptotic estimation error achieved by Bayes-AMP coincides (in certain regimes) with the Bayes-optimal error (see Corollary 2.3). When this is not the case, no polynomial-time algorithm is known that outperforms our method.
- Section 3 shows how AMP estimates can be used to construct confidence intervals and p-values. In particular, we prove that the resulting p-values are asymptotically valid on the nulls, which in turn can be used to establish asymptotic false discovery rate control using a Benjamini–Hochberg procedure.
- Section 4 generalizes the analysis of Section 2 to the case of rectangular matrices. This allows, in particular, to derive optimal AMP algorithms for the spiked covariance model. The theory for rectangular matrices is completely analogous to the one for symmetric ones, and indeed can be established via a reduction to symmetric matrices.
- Section 5 discusses a new phenomenon arising in case of degeneracies between the values $\lambda_1, \dots, \lambda_k$. For the sake of concreteness, we consider the case $\boldsymbol{A} = \lambda \boldsymbol{A}_0 + \boldsymbol{W}$, where $\boldsymbol{A}_0$ is a rank-k matrix obtained as follows. We partition $\{1, \dots, n\}$ in $q = k + 1$ groups and set $A_{0,ij} = k/n$ if i, j belong to the same group and $A_{0,ij} = -1/n$ otherwise. Due to its close connections with the stochastic block model of random graphs, we refer to this as to the "Gaussian block model."

  It turns out that in such degenerate cases, the evolution of AMP estimates does not concentrate around a deterministic trajectory. Nevertheless, state evolution captures the

asymptotic behavior of the algorithm in terms of a random initialization (whose distribution is entirely characterized) plus a deterministic evolution.

- Section 6 presents our general result in the case of a symmetric matrix $\boldsymbol{A}$ distributed according to the model (1.1). Our theorems provide an asymptotic characterization of a general AMP algorithm in terms of a suitable state evolution recursion. A completely analogous result holds for rectangular matrices. The corresponding statement is presented in the Supplementary Material.
- Section 7 provides an outline of the proofs of our main results. Earlier state evolution results do not allow to rigorously analyze AMP unless its initialization is independent from the data matrix $\boldsymbol{A}$. In particular, they do not allow to analyze the spectral initialization used in our algorithm. In order to overcome this challenge, we prove a technical lemma (Lemma B.3 in the Supplementary Material [44]) that specifies an approximate representation for the conditional distribution of $\boldsymbol{A}$ given its leading outlier eigenvectors and the corresponding eigenvalues. Namely, $\boldsymbol{A}$ can be approximated by a sum of rank-one matrices, corresponding to the outlier eigenvectors, plus a projection of a new random matrix $\boldsymbol{A}^{\text{new}}$ independent of $\boldsymbol{A}$. We leverage this explicit independence to establish state evolution for our algorithm.

Complete proofs of the main results are deferred to the Appendices A and B in the Supplementary Material [44]. For the reader's convenience, we present separate proofs for the case of rank $k = 1$, and then for the general case, which is technically more involved. The proofs concerning the examples in Section 2 and 4 are also presented in the Appendices.

As mentioned above, while several of our examples concern low-rank matrix estimation, the main result in Section 6 is significantly more general, and is potentially relevant to a broad range of applications in which AMP is run in conjunction with a spectral initialization.

2. Estimation of symmetric rank-one matrices. In order to illustrate our main result (to be presented in Section 6), we apply it to the problem of estimating a rank-one symmetric matrix in Gaussian noise. We will begin with a brief heuristic discussion of AMP and its application to rank-one matrix estimation. The reader is welcome to consult the substantial literature on AMP for further background [6, 7, 11, 29].

2.1. *Main ideas and heuristic justification.* Let $\boldsymbol{x}_0 = \boldsymbol{x}_0(n) \in \mathbb{R}^n$ be a sequence of signals indexed by the dimension n, satisfying the following conditions:

(i) Their rescaled ℓ_2-norms converge $\lim_{n\to\infty} \|\boldsymbol{x}_0(n)\|_2/\sqrt{n} = 1$;
(ii) The empirical distributions of the entries of $\boldsymbol{x}_0(n)$ converges weakly to a probability distribution ν_{X_0} on $\mathbb{R}$, with unit second moment.

We then consider the following spiked model, for $\boldsymbol{W} \sim \mathsf{GOE}(n)$:

$$\boldsymbol{A} = \frac{\lambda}{n}\boldsymbol{x}_0\boldsymbol{x}_0^{\mathsf{T}} + \boldsymbol{W}. \tag{2.1}$$

Given one realization of the matrix $\boldsymbol{A}$, we would like to estimate the signal $\boldsymbol{x}_0$. Note that this matrix is of the form (1.1) with $k = 1$, $\lambda_1 = \lambda\|\boldsymbol{x}_0(n)\|_2^2/n \to \lambda$ and $\boldsymbol{v}_1 = \boldsymbol{x}_0(n)/\|\boldsymbol{x}_0(n)\|_2$.

In order to discuss informally the main ideas in AMP, assume for a moment to be given an additional noisy observation of $\boldsymbol{x}_0$, call it $\boldsymbol{y} \in \mathbb{R}^n$, which is independent of $\boldsymbol{A}$ (i.e., independent of $\boldsymbol{W}$, since $\boldsymbol{x}_0$ is deterministic). More specifically, assume $\boldsymbol{y} \sim \mathsf{N}(\mu_0\boldsymbol{x}_0, \sigma_0^2\boldsymbol{I}_n)$. How can we denoise this observation, and incorporate the quadratic observation $\boldsymbol{A}$ in (2.1)?

A first idea would be to denoise $\boldsymbol{y}$, using an entrywise scalar denoiser $f_0 : \mathbb{R} \to \mathbb{R}$. We denote the vector obtained by applying f_0 componentwise by $f_0(\boldsymbol{y})$. Of course, the choice

of f_0 depends on our knowledge of $\boldsymbol{x}_0$. For instance, if we know that $\boldsymbol{x}_0$ is sparse, then we could apply componentwise soft thresholding:

$$f_0(y_i) = \eta(y_i; \tau), \tag{2.2}$$

where $\eta(x;\tau) = \mathrm{sign}(x)(|x| - \tau)_+$, and τ is a suitable threshold level. Classical theory guarantees the accuracy of such a denoiser [20, 21].

However, $f_0(\boldsymbol{y})$ does not exploit the observation $\boldsymbol{A}$ in any way. We could try to improve this estimate by multiplying $f_0(\boldsymbol{y})$ by $\boldsymbol{A}$:

$$\boldsymbol{x}^1 = \boldsymbol{A} f_0(\boldsymbol{y}) = \frac{\lambda}{n}\langle \boldsymbol{x}_0, f_0(\boldsymbol{y})\rangle \boldsymbol{x}_0 + \boldsymbol{W} f_0(\boldsymbol{y}). \tag{2.3}$$

It is not hard to see that the second term is a centered Gaussian vector whose entries have variance close to $\|f_0(\boldsymbol{y})\|^2/n \to \sigma_1^2 \equiv \mathbb{E}\{f_0(\mu_0 X_0 + \sigma_0 G)^2\}$, while the first term is essentially deterministic by the law of large numbers. We thus obtain that $\boldsymbol{x}^1$ is approximately $\mathsf{N}(\mu_1 \boldsymbol{x}_0, \sigma_1^2 \boldsymbol{I}_n)$, where

$$\mu_1 = \lambda \mathbb{E}\{X_0 f_0(\mu_0 X_0 + \sigma_0 G)\}, \qquad \sigma_1^2 = \mathbb{E}\{f_0(\mu_0 X_0 + \sigma_0 G)^2\}. \tag{2.4}$$

Here, expectation is taken with respect to $X_0 \sim \nu_{X_0}$ independent of $G \sim \mathsf{N}(0,1)$. This analysis also suggests how to design the function f_0: ideally, it should maximize the signal-to-noise ratio (SNR) μ_1^2/σ_1^2. Of course, the precise choice of f_0 depends on our prior knowledge of $\boldsymbol{x}_0$. For instance, if we know the law ν_{X_0}, we can maximize this ratio by taking $f_0(y) = \mathbb{E}\{X_0 \mid \mu_0 X_0 + \sigma_0 G = y\}$.

At this point, it would be tempting to iterate the above procedure, and consider the nonlinear power iteration

$$\boldsymbol{x}_{\mathrm{PI}}^{t+1} = \boldsymbol{A} f_t(\boldsymbol{x}_{\mathrm{PI}}^t), \tag{2.5}$$

for a certain sequence of functions $f_t : \mathbb{R} \to \mathbb{R}$. (As above, $f_t(\boldsymbol{x}_{\mathrm{PI}}^t)$ is the vector obtained by applying f_t componentwise to $\boldsymbol{x}_{\mathrm{PI}}^t$, and we will use superscripts to indicate the iteration number.) While this approach has been studied in the literature [15, 34, 52], sharp results could only be established in a high SNR regime where $\lambda = \lambda(n) \to \infty$ at a sufficiently fast rate. Indeed, analyzing the recursion (2.5) is difficult because $\boldsymbol{A}$ is correlated with $\boldsymbol{x}_{\mathrm{PI}}^t$ (unlike in equation (2.3)), and hence the simple calculation that yields equation (2.4) is no longer permitted. This problem is compounded by the fact that we do not have an additional observation $\boldsymbol{y}$ independent of $\boldsymbol{A}$, and instead we plan to use a spectral initialization $\boldsymbol{x}^0 \propto \boldsymbol{\varphi}_1$ that depends on the top eigenvector $\boldsymbol{\varphi}_1$ of $\boldsymbol{A}$. As a consequence, even the first step of the analysis (given in equation (2.4)) is no longer obvious.

Let us emphasize that these difficulties are not a limitation of the proof technique. For $t > 1$, the iterates (2.5) are no longer Gaussian or centered around $\mu_t \boldsymbol{x}_0$, for some scaling factor μ_t. This can be easily verified by considering, for instance, the function $f_t(x) = x^2$ (we refer to [6] which carries out the calculation for such an example).

AMP solves the correlation problem in nonlinear power iteration by modifying equation (2.5): namely, we subtract from from $\boldsymbol{A} f_t(\boldsymbol{x}_{\mathrm{L}}^t)$ the part that is correlated to the past iterates. Let $\mathfrak{S}_t \equiv \sigma(\{\boldsymbol{x}_{\mathrm{L}}^0, \boldsymbol{x}_{\mathrm{L}}^2, \ldots, \boldsymbol{x}_{\mathrm{L}}^t\})$ be the σ-algebra generated by iterates up to time t. The correction that compensates for correlations is most conveniently explained by using the following Long AMP recursion, introduced in [11]:

$$\boldsymbol{x}_{\mathrm{L}}^{t+1} = \boldsymbol{A} f_t(\boldsymbol{x}_{\mathrm{L}}^t) - \mathbb{E}\{\boldsymbol{W} f_t(\boldsymbol{x}_{\mathrm{L}}^t) | \mathfrak{S}_t\} + \overline{\alpha}_t \boldsymbol{x}_0 + \sum_{s=0}^{t} \alpha_{t,s} \boldsymbol{x}_{\mathrm{L}}^s \tag{2.6}$$

$$= \sum_{s=0}^{t} \alpha_{t,s} \boldsymbol{x}_{\mathrm{L}}^s + \left(\overline{\alpha}_t + \frac{\lambda}{n}\langle \boldsymbol{x}_0, f_t(\boldsymbol{x}_{\mathrm{L}}^t)\rangle\right)\boldsymbol{x}_0 + \boldsymbol{W} f_t(\boldsymbol{x}_{\mathrm{L}}^t) - \mathbb{E}\{\boldsymbol{W} f_t(\boldsymbol{x}_{\mathrm{L}}^t) | \mathfrak{S}_t\}, \tag{2.7}$$

where $(\overline{\alpha}_t)_{0\le t}$, $(\alpha_{t,s})_{0\le s\le t}$ are suitable sequences of deterministic numbers. In words, the new vector $\boldsymbol{x}_{\mathrm{L}}^{t+1}$ is a linear combination of iterates up to time t, plus a term $\boldsymbol{x}_0(\overline{\alpha}_t + \lambda\langle \boldsymbol{x}_0, f_t(\boldsymbol{x}_{\mathrm{L}}^t)\rangle/n)$ that is essentially deterministic, plus a random term $(\boldsymbol{W} f_t(\boldsymbol{x}_{\mathrm{L}}^t) - \mathbb{E}\{\boldsymbol{W} f_t(\boldsymbol{x}_{\mathrm{L}}^t)|\mathfrak{S}_t\})$ that is uncorrelated with the past. If the past iterates $(\boldsymbol{x}_{\mathrm{L}}^s)_{0\le s\le t}$ are jointly Gaussian, then the first two components (linear and deterministic) are also jointly Gaussian with $(\boldsymbol{x}_{\mathrm{L}}^s)_{0\le s\le t}$. Since the third (random) term is uncorrelated with the past iterates, it can be shown by induction that the sequence $(\boldsymbol{x}_{\mathrm{L}}^t)_{0\le t\le T}$ is approximately Gaussian as $n\to\infty$, for any fixed t (in the sense of finite dimensional marginals), and centered around $\boldsymbol{x}_0$; see [11].

At first sight, this might appear as a mathematical trick, with no practical implications. Indeed equation (2.6) does not provide an algorithm. We are explicitly using the true signal $\boldsymbol{x}_0$ which we are supposed to estimate, and the expectation $\mathbb{E}\{\boldsymbol{W} f_t(\boldsymbol{x}_{\mathrm{L}}^t)|\mathfrak{S}_t\}$ is, at best, hard to compute. However, it turns out that (for a certain choice of the numbers $(\overline{\alpha}_t)_{0\le t}$, $(\alpha_{t,s})_{0\le s\le t}$), the term subtracted from $\boldsymbol{A} f_t(\boldsymbol{x}_{\mathrm{L}}^t)$ in equation (2.6) can be approximated by $\mathsf{b}_t f_{t-1}(\boldsymbol{x}_{\mathrm{L}}^{t-1})$ with a coefficient b_t that can be computed easily. We will not try to justify this approximation here (see, for instance, [11]). We will instead use the resulting algorithm (given below in equation (2.8)) as the starting point of our analysis.

2.2. *General analysis.* Motivated by the discussion in the previous section, we consider the following general algorithm for rank-one matrix estimation in the model (2.1). In order to estimate $\boldsymbol{x}_0$, we compute the principal eigenvector of $\boldsymbol{A}$, to be denoted by $\boldsymbol{\varphi}_1$, and apply the following iteration, with initialization $\boldsymbol{x}^0 = \sqrt{n}\boldsymbol{\varphi}_1$:

$$\boldsymbol{x}^{t+1} = \boldsymbol{A} f_t(\boldsymbol{x}^t) - \mathsf{b}_t f_{t-1}(\boldsymbol{x}^{t-1}), \quad \mathsf{b}_t = \frac{1}{n}\sum_{i=1}^n f_t'(x_i^t). \tag{2.8}$$

Here, $f_t(\boldsymbol{x}) = (f_t(x_1), \dots, f_t(x_n))^{\mathsf{T}}$ is a separable function for each t. As mentioned above, we can think of this iteration as an approximation of equation (2.6) where all the terms except the first one have been estimated by $-\mathsf{b}_t f_{t-1}(\boldsymbol{x}^{t-1})$. The fact that this is an accurate estimate for large n is far from obvious, but can be established by induction over t [11].

Note that $\boldsymbol{x}_0$ can be estimated from the data $\boldsymbol{A}$ only up to an overall sign (since $\boldsymbol{x}_0$ and $-\boldsymbol{x}_0$ give rise to the same matrix $\boldsymbol{A}$ as per equation (2.1)). In order to resolve this ambiguity, we will assume, without loss of generality, that $\langle \boldsymbol{x}_0, \boldsymbol{\varphi}_1\rangle \ge 0$.

THEOREM 1. *Consider the $k = 1$ spiked matrix model of equation (2.1), with $\boldsymbol{x}_0(n) \in \mathbb{R}^n$ a sequence of vectors satisfying assumptions (i), (ii) above, and $\lambda > 1$. Consider the AMP iteration in equation (2.8) with initialization $\boldsymbol{x}^0 = \sqrt{n}\boldsymbol{\varphi}_1$ (where, without loss of generality $\langle \boldsymbol{x}_0, \boldsymbol{\varphi}_1\rangle \ge 0$). Assume $f_t : \mathbb{R} \to \mathbb{R}$ to be Lipschitz continuous for each $t \in \mathbb{N}$.*

Let $(\mu_t, \sigma_t)_{t\ge 0}$ be defined via the recursion

$$\mu_{t+1} = \lambda\mathbb{E}\big[X_0 f_t(\mu_t X_0 + \sigma_t G)\big], \tag{2.9}$$

$$\sigma_{t+1}^2 = \mathbb{E}\big[f_t(\mu_t X_0 + \sigma_t G)^2\big], \tag{2.10}$$

where $X_0 \sim \nu_{X_0}$ and $G \sim \mathsf{N}(0,1)$ are independent, and the initial condition is $\mu_0 = \sqrt{1-\lambda^{-2}}$, $\sigma_0 = 1/\lambda$.

Then, for any function $\psi : \mathbb{R}\times\mathbb{R} \to \mathbb{R}$ with $|\psi(\boldsymbol{x}) - \psi(\boldsymbol{y})| \le C(1 + \|\boldsymbol{x}\|_2 + \|\boldsymbol{y}\|_2)\|\boldsymbol{x} - \boldsymbol{y}\|_2$ for a universal constant $C > 0$, the following holds almost surely for $t \ge 0$:

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^n \psi\big(x_{0,i}, x_i^t\big) = \mathbb{E}\big\{\psi\big(X_0, \mu_t X_0 + \sigma_t G\big)\big\}. \tag{2.11}$$

The proof of this theorem is presented in Appendix A in the Supplementary Material [44].

One peculiarity of our approach is that we do not commit to a specific choice of the nonlinearities f_t, and instead develop a sharp asymptotic characterization for any—sufficiently regular—nonlinearity. A poor choice of the functions f_t might result in large estimation error, and yet Theorem 1 will continue to hold.

On the other hand, the state evolution characterization can be used to design optimal nonlinearities in a principled way. Given equations (2.9) and (2.10), the general principle is quite transparent. The optimal nonlinearity is defined in terms of a scalar denoising problem. For $X_0 \sim \nu_{X_0}$ and $G \sim \mathsf{N}(0,1)$ independent, consider the problem of estimating X_0 from the noisy observation $Y = \mu_t X_0 + \sigma_t G$. At step t, f_t should be constructed as to maximize the ratio $\mathbb{E}[X_0 f_t(\mu_t X_0 + \sigma_t G)]/\mathbb{E}[f_t(\mu_t X_0 + \sigma_t G)^2]^{1/2}$. Two specific instantiations of this principle are given in Sections 2.3 and 2.4.

REMARK 2.1. The state evolution recursion of equations (2.9), (2.10) in Theorem 1 was already derived by Fletcher and Rangan in [25]. However, as explained in [25], Section 5.3, their results only apply to cases in which AMP can be initialized in a way that: (i) has positive correlation with the spike $\boldsymbol{x}_0$ (and this correlation does not vanish as $n \to \infty$); (ii) is independent of $\boldsymbol{A}$.

Theorem 1 analyzes an algorithm which does not require such an initialization, and hence applies more broadly.

2.3. *The case of a sparse spike.* In some applications, we might know that the spike $\boldsymbol{x}_0$ is sparse. We consider a simple model in which $\boldsymbol{x}_0$ is known to have at most $n\varepsilon$ nonzero entries for some $\varepsilon \in (0,1)$.

Because of its importance, the use of nonlinear power iteration methods for this problem has been studied by several authors in the past [34, 41, 52]. However, none of these works obtains precise asymptotics in the moderate SNR regime (i.e., for λ, ε of order one). In contrast, sharp results can be obtained by applying Theorem 1. Here, we will limit ourselves to taking the first steps, deferring a more complete analysis to future work. We focus on the case of symmetric matrices for simplicity (cf. equation (2.1)), but a generalization to rectangular matrices is straightforward along the lines of Section 4.

The sparsity assumption implies that the random variable X_0 entering the state evolution recursion in equation (2.9) should satisfy $\nu_{X_0}(\{0\}) \geq 1 - \varepsilon$. Classical theory for the sparse sequence model [20, 21] suggests taking f_t to be the soft thresholding denoiser $f_t(x) = \eta(x;\tau_t)$, for $(\tau_t)_{t\geq 0}$ a well-chosen sequence of thresholds. The resulting algorithm reads

$$\begin{aligned} \boldsymbol{x}^{t+1} &= \boldsymbol{A}\hat{\boldsymbol{x}}^t - \mathsf{b}_t \hat{\boldsymbol{x}}^{t-1}, \qquad \hat{\boldsymbol{x}}^t = \eta\big(\boldsymbol{x}^t;\tau_t\big), \\ \mathsf{b}_t &= \frac{1}{n}\big\|\hat{\boldsymbol{x}}^t\big\|_0, \end{aligned} \tag{2.12}$$

where $\|\boldsymbol{v}\|_0$ is the number of nonzero entries of vector $\boldsymbol{v}$. The initialization is, as before $\boldsymbol{x}^0 = \sqrt{n}\boldsymbol{\varphi}_1$. The algorithm alternates soft thresholding, to produce sparse estimates and power iteration, with the crucial correction term $-\mathsf{b}_t\hat{\boldsymbol{x}}^{t-1}$.

Theorem 1 can be directly applied to characterize the performance of this algorithm for any fixed distribution ν_{X_0} of the entries of $\boldsymbol{x}_0$. For instance, we obtain the following exact prediction for the asymptotic correlation between estimates $\hat{\boldsymbol{x}}^t$ and the signal $\boldsymbol{x}_0$:

$$\lim_{n\to\infty} \frac{|\langle \hat{\boldsymbol{x}}^t(\boldsymbol{A}), \boldsymbol{x}_0\rangle|}{\|\hat{\boldsymbol{x}}^t(\boldsymbol{A})\|_2 \|\boldsymbol{x}_0\|_2} = \frac{\mu_{t+1}}{\lambda\sigma_{t+1}}. \tag{2.13}$$

For a given distribution ν_{X_0}, it is easy to compute μ_t, σ_t using equation (2.9) with $f_t(\boldsymbol{x}^t) = \eta(\boldsymbol{x}^t;\tau_t)$.

We can also use Theorem 1 to characterize the minimax behavior over $n\varepsilon$-sparse vectors. We sketch the argument next: similar arguments were developed in [19, 22] in the context of compressed sensing. The basic idea is to lower bound the signal-to-noise ratio (SNR) $\mu_{t+1}^2/\sigma_{t+1}^2$ iteratively as a function of the SNR at the previous iteration, over the set of probability distributions $\mathcal{F}_\varepsilon = \{\nu_{X_0} : \nu_{X_0}(\{0\}) \geq 1-\varepsilon, \int x^2 \nu_{X_0}(\mathrm{d}x) = 1\}$. As shown in Appendix E.1 in the Supplementary Material [44], it is sufficient to consider the extremal points of the set $\mathcal{F}_\varepsilon$, which are given by the three-points priors:

$$\pi_{p,a_1,a_2} \equiv (1-\varepsilon)\delta_0 + \varepsilon p \delta_{a_1} + \varepsilon(1-p)\delta_{a_2}, \qquad pa_1^2 + (1-p)a_2^2 = 1, \quad p \in [0,1]. \tag{2.14}$$

We then define the following SNR maps:

$$S_*(\gamma, \theta; \nu_{X_0}) \equiv \frac{[\mathbb{E}\{X_0 \eta(\sqrt{\gamma} X_0 + G; \theta)\}]^2}{\mathbb{E}\{\eta(\sqrt{\gamma} X_0 + G; \theta)^2\}}, \tag{2.15}$$

$$S(\gamma; \theta) \equiv \inf\big\{S_*(\gamma, \theta; \pi_{p,a_1,a_2}) : pa_1^2 + (1-p)a_2^2 = 1, p \in [0,1]\big\}. \tag{2.16}$$

The interpretation of these quantities is as follows: $\gamma \mapsto S_*(\gamma, \theta; \nu_{X_0})$ describes the evolution of the signal-to-noise ratio after one step of AMP, when the signal distribution is ν_{X_0}; the map $\gamma \mapsto S(\gamma; \theta)$ is the same evolution, for the least favorable prior, which can be taken of the form π_{p,a_1,a_2}.

Notice that the function $S_*(\gamma, \theta; \pi_{p,a_1,a_2})$ can be evaluated by performing a small number (six, to be precise) of Gaussian integrals. The function S is defined by a two-dimensional optimization problem, which can be computed numerically quite efficiently.

We define the sequences $(\underline{\gamma}_t)_{t\geq 0}$, $(\theta_t)_{t\geq 0}$ by setting $\underline{\gamma}_0 = \lambda^2 - 1$, and then recursively

$$\underline{\gamma}_{t+1} = \lambda^2 S(\underline{\gamma}_t; \theta_t), \qquad \theta_t = \arg\max_{\theta \in [0,\infty]} S(\underline{\gamma}_t; \theta). \tag{2.17}$$

The next proposition provides the desired lower bound for the signal-to-noise ratio over the class of sparse vectors.

PROPOSITION 2.1. *Assume the setting of Theorem 1, and furthermore,* $\|\boldsymbol{x}_0(n)\|_0 \leq n\varepsilon$. *Let* $(\hat{\boldsymbol{x}}^t = \hat{\boldsymbol{x}}^t(\boldsymbol{A}))_{t\geq 0}$ *be the sequence of estimates produced by the AMP iteration equation* (2.12) *with initialization* $\boldsymbol{x}^0 = \sqrt{n}\boldsymbol{\varphi}_1$, *and thresholds* $\tau_t = \theta_t \hat{\sigma}_t$ *where* $\hat{\sigma}_t$ *is a estimator of* σ_t *from data* $\boldsymbol{x}^0, \ldots, \boldsymbol{x}^t$ *such that* $\hat{\sigma}_t \xrightarrow{a.s.} \sigma_t$. (*For instance, take* $\hat{\sigma}_t^2 \equiv \|f_{t-1}(\boldsymbol{x}^{t-1})\|_2^2/n$ *for* $t \geq 1$. *For* $t = 0$, *take* $\hat{\sigma}_0^2 \equiv 1/\hat{\lambda}$, *where* $\hat{\lambda}$ *is given in equation* (3.1).)

Then for any fixed $t \geq 0$ *we have, almost surely,*

$$\lim_{n\to\infty} \frac{|\langle \hat{\boldsymbol{x}}^t(\boldsymbol{A}), \boldsymbol{x}_0\rangle|}{\|\hat{\boldsymbol{x}}^t(\boldsymbol{A})\|_2 \|\boldsymbol{x}_0\|_2} = \frac{\mu_{t+1}}{\lambda\sigma_{t+1}} \geq \frac{\sqrt{\underline{\gamma}_{t+1}}}{\lambda}. \tag{2.18}$$

Here, $(\mu_{t+1}, \sigma_{t+1})$ *are recursively defined as follows, starting from* $\mu_0 = \sqrt{1-\lambda^{-2}}$ *and* $\sigma_0^2 = \lambda^{-2}$:

$$\mu_{t+1} = \lambda \mathbb{E}\big\{X_0 \eta(\mu_t X_0 + \sigma_t G; \theta_t \sigma_t)\big\}, \qquad \sigma_{t+1}^2 = \mathbb{E}\big\{\eta(\mu_t X_0 + \sigma_t G; \theta_t \sigma_t)^2\big\}. \tag{2.19}$$

The proof of Proposition 2.1 is given in Appendix E in the Supplementary Material [44]. The proposition reduces the analysis of algorithm (2.12) to the study of a one-dimensional recursion $\underline{\gamma}_{t+1} = \lambda^2 S(\underline{\gamma}_t; \theta_t)$, which is much simpler. We defer this analysis to future work. We emphasize that the AMP algorithm in equation (2.12) with thresholds $\tau_t = \theta_t \hat{\sigma}_t$ does not require knowledge of either the sparsity level ε or the SNR parameter λ—these quantities are only required to compute the sequence of lower bounds $(\underline{\gamma}_t)_{t\geq 0}$.

2.4. *Bayes-optimal estimation.* As a second application of Theorem 1, we consider the case in which the asymptotic empirical distribution ν_{X_0} of the entries of $\boldsymbol{x}_0$ is known. This case is of special interest because it provides a lower bound on the error achieved by any AMP algorithm.

To simplify some of the formulas below, we assume here a slightly different normalization for the initialization, but otherwise we use the same algorithm as in the general case, namely

$$\boldsymbol{x}^0 = \sqrt{n\lambda^2(\lambda^2-1)}\boldsymbol{\varphi}_1, \tag{2.20}$$

$$\boldsymbol{x}^{t+1} = \boldsymbol{A} f_t(\boldsymbol{x}^t) - \mathsf{b}_t f_{t-1}(\boldsymbol{x}^{t-1}), \quad \mathsf{b}_t = \frac{1}{n}\sum_{i=1}^{n} f_t'(x_i^t). \tag{2.21}$$

In order to define the optimal nonlinearity, consider again the scalar denoising problem of estimating X_0 from the noisy observation $Y = \sqrt{\gamma}X_0 + G$ (note that $X_0, G \in \mathbb{R}$ are scalar random variables). The minimum mean square error is

$$\mathsf{mmse}(\gamma) = \mathbb{E}\big\{\big[X_0 - \mathbb{E}(X_0|\sqrt{\gamma}X_0 + G)\big]^2\big\}. \tag{2.22}$$

With these notation, we can introduce the state evolution recursion

$$\gamma_0 = \lambda^2 - 1, \tag{2.23}$$

$$\gamma_{t+1} = \lambda^2\big\{1 - \mathsf{mmse}(\gamma_t)\big\}. \tag{2.24}$$

These describe the evolution of the effective signal-to-noise ratio along the algorithm execution.

The optimal nonlinearity $f_t(\cdot)$ after t iterations is the minimum mean square error denoiser for signal-to-noise ratio γ_t:

$$f_t(y) \equiv \lambda F(y;\gamma_t), \tag{2.25}$$

$$F(y;\gamma) \equiv \mathbb{E}\{X_0 \mid \gamma X_0 + \sqrt{\gamma}G = y\}. \tag{2.26}$$

After t iterations, we produce an estimate of $\boldsymbol{x}_0$ by computing $\hat{\boldsymbol{x}}^t(\boldsymbol{A}) \equiv f_t(\boldsymbol{x}^t)/\lambda = F(\boldsymbol{x}^t;\gamma_t)$. We will refer to this choice as to Bayes AMP.

REMARK 2.2. Implementing the Bayes-AMP algorithm requires to approximate the function $F(y;\gamma)$ of equation (2.26). This amounts to a one-dimensional integral and can be done very accurately by standard quadrature methods: a simple approach that works well in practice is to replace the measure ν_{X_0} by a combination of finitely many point masses. Analogously, the function $\mathsf{mmse}(\gamma)$ (which is needed to compute the sequence γ_t), can be computed by the same method.[3]

We are now in position to state the outcome of our analysis for Bayes AMP, whose proof is deferred to Appendix F in the Supplementary Material [44].

THEOREM 2. *Consider the spiked matrix model* (2.1), *with* $\boldsymbol{x}_0(n) \in \mathbb{R}^n$ *a sequence of vectors satisfying assumptions* (i), (ii) *above, and* $\lambda > 1$. *Let* $(\boldsymbol{x}^t)_{t\geq 0}$ *be the sequence of iterates generated by the Bayes AMP algorithm defined in equation* (2.8), *with initialization* (2.20), *and optimal choice of the nonlinearity defined by equation* (2.25). *Assume* $F(\cdot;\gamma) : \mathbb{R} \to \mathbb{R}$ *to be Lipschitz continuous for any* $\gamma \in (0, \lambda^2]$. *Finally, define state evolution by equations* (2.23), (2.24).

[3]AMP noes not require high accuracy in the approximations of the nonlinear functions f_t. As shown several times in the Appendices (see, e.g., Appendix A) the algorithm is stable with respect to perturbations of f_t.

Then, for any function $\psi : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ *with* $|\psi(\boldsymbol{x}) - \psi(\boldsymbol{y})| \le C(1 + \|\boldsymbol{x}\|_2 + \|\boldsymbol{y}\|_2)\|\boldsymbol{x} - \boldsymbol{y}\|_2$ *for a universal constant* $C > 0$, *the following holds almost surely for* $t \ge 0$:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \psi\big(x_{0,i}, x_i^t\big) = \mathbb{E}\big\{\psi\big(X_0, \gamma_t X_0 + \gamma_t^{1/2} Z\big)\big\}, \tag{2.27}$$

where expectation is taken with respect to $X_0 \sim \nu_{X_0}$ *and* $Z \sim \mathsf{N}(0,1)$ *mutually independent, and we assumed without loss of generality that* $\langle \boldsymbol{\varphi}_1, \boldsymbol{x}_0 \rangle \ge 0$.

In particular, let $\gamma_{\mathrm{ALG}}(\lambda)$ *denote the smallest strictly positive solution of the fixed point equation* $\gamma = \lambda^2[1 - \mathsf{mmse}(\gamma)]$. *Then the AMP estimate* $\hat{\boldsymbol{x}}^t(\boldsymbol{A}) = f_t(\boldsymbol{x}^t)/\lambda$ *achieves*

$$\lim_{t\to\infty} \lim_{n\to\infty} \frac{|\langle \hat{\boldsymbol{x}}^t(\boldsymbol{A}), \boldsymbol{x}_0 \rangle|}{\|\hat{\boldsymbol{x}}^t(\boldsymbol{A})\|_2 \|\boldsymbol{x}_0\|_2} = \frac{\sqrt{\gamma_{\mathrm{ALG}}(\lambda)}}{\lambda}, \tag{2.28}$$

$$\lim_{t\to\infty} \lim_{n\to\infty} \frac{1}{n} \min_{s\in\{+1,-1\}} \big\|s\hat{\boldsymbol{x}}^t(\boldsymbol{A}) - \boldsymbol{x}_0\big\|_2^2 = 1 - \frac{\gamma_{\mathrm{ALG}}(\lambda)}{\lambda^2}. \tag{2.29}$$

Finally, the algorithm has total complexity $O(n^2 \log n)$.

REMARK 2.3. The assumption on $F(\cdot;\gamma) : \mathbb{R} \to \mathbb{R}$ being Lipschitz continuous is required in order to apply our general theory. Note that this is implied by either of the following: (i) $\mathrm{supp}(\nu_{X_0}) \in [-M, M]$ for some constant M; (ii) ν_{X_0} has log-concave density.

It is interesting to compare the above result with the Bayes optimal estimation accuracy. The following statement is a consequence of the results of [39] (see Appendix D in the Supplementary Material [44]).

PROPOSITION 2.2. *Consider the spiked matrix model* (2.1), *with* $\boldsymbol{x}_0(n) \in \mathbb{R}^n$ *a vector with i.i.d. entries with distribution* ν_{X_0} *with bounded support and* $\int x^2 \nu_{X_0}(\mathrm{d}x) = 1$. *Then there exists a countable set* $D \subseteq \mathbb{R}_{\ge 0}$ *such that, for* $\lambda \in \mathbb{R} \setminus D$, *the Bayes-optimal accuracy in the rank-one estimation problem is given by*

$$\lim_{n\to\infty} \sup_{\hat{\boldsymbol{x}}(\cdot)} \mathbb{E}\left\{\frac{\langle \hat{\boldsymbol{x}}(\boldsymbol{A}), \boldsymbol{x}_0\rangle^2}{\|\hat{\boldsymbol{x}}(\boldsymbol{A})\|_2^2 \|\boldsymbol{x}_0\|_2^2}\right\} = \frac{\gamma_{\mathrm{Bayes}}(\lambda)}{\lambda^2}, \tag{2.30}$$

where the supremum is over (*possibly randomized*) *estimators, that is, measurable functions* $\hat{\boldsymbol{x}} : \mathbb{R}^{n\times n} \times [0,1] \to \mathbb{R}^n$, *where* $[0,1]$ *is endowed with the uniform measure. Here,* $\gamma_{\mathrm{Bayes}}(\lambda)$ *is the fixed point of the recursion* (2.24) *that maximizes the following free energy functional*:

$$\Psi(\gamma, \lambda) = \frac{\lambda^2}{4} + \frac{\gamma^2}{4\lambda} - \frac{\gamma}{2} + \mathrm{I}(\gamma), \tag{2.31}$$

where $\mathrm{I}(\gamma) = \mathbb{E} \log \frac{\mathrm{d}p_{Y|X_0}}{\mathrm{d}p_Y}(Y, X_0)$ *is the mutual information for the scalar channel* $Y = \sqrt{\gamma} X_0 + G$, *with* $X_0 \sim \nu_{X_0}$ *and* $G \sim \mathsf{N}(0,1)$ *mutually independent.*

Together with this proposition, Theorem 2 precisely characterizes the gap between Bayes-optimal estimation and message passing algorithms for rank-one matrix estimation. Simple calculus (together with the relation $\mathrm{I}'(\gamma) = \mathsf{mmse}(\gamma)/2$ [27]) implies that the fixed point of the recursion (2.24) coincide with the stationary points of $\gamma \mapsto \Psi(\gamma, \lambda)$. We therefore have the following characterization of the Bayes optimality of Bayes-AMP.

COROLLARY 2.3. *Under the setting of Theorem 2* (*in particular,* $\lambda > 1$), *let the function* $\Psi(\gamma, \lambda)$ *be defined as in equation* (2.31). *Then Bayes-AMP asymptotically achieves the Bayes-optimal error* (*and* $\gamma_{\mathrm{ALG}}(\lambda) = \gamma_{\mathrm{Bayes}}(\lambda)$) *if and only if the global maximum of* $\gamma \mapsto \Psi(\gamma, \lambda)$ *over* $(0, \infty)$ *is also the first stationary point of the same function* (*as* γ *grows*).

As illustrated in Section 2.5, this condition holds for some cases of interest, and hence message passing is asymptotically optimal for these cases.

REMARK 2.4. In some applications, it is possible to construct an initialization $\boldsymbol{x}^0$ that is positively correlated with the signal $\boldsymbol{x}_0$ and independent of $\boldsymbol{A}$. If this is possible, then the spectral initialization is not required and Theorem 2 follows immediately from [7]. For instance, if ν_{X_0} has positive mean, then it is sufficient to initialize $\boldsymbol{x}^0 = \mathbf{1}$. This principle was exploited in [17, 18, 43].

However, such a positively correlated initialization is not available in general: the spectral initialization analyzed here aims at overcoming this problem.

REMARK 2.5. No polynomial-time algorithm is known that achieves estimation accuracy superior to the one guaranteed by Theorem 2. In particular, it follows from the optimality of posterior mean with respect to square loss and the monotonicity of the function $\gamma \mapsto \lambda^2\{1 - \mathsf{mmse}(\gamma)\}$ that Bayes AMP is optimal among AMP algorithms. That is, for any other sequence of nonlinearities $f_t(\cdot)$, we have

$$\lim_{n\to\infty} \frac{|\langle f_t(\boldsymbol{x}^t), \boldsymbol{x}_0\rangle|}{\|f_t(\boldsymbol{x}^t)\|_2 \|\boldsymbol{x}_0\|_2} = \frac{\mu_{t+1}}{\lambda\sigma_{t+1}} \le \frac{\sqrt{\gamma_{t+1}}}{\lambda}. \tag{2.32}$$

As further examples, [30] analyzes a semidefinite programming (SDP) algorithm for the special case of a two-points symmetric mixture $\nu_{X_0} = (1/2)\delta_{+1} + (1/2)\delta_{-1}$. Theorem 2 implies that, in this case, message passing is Bayes optimal (since $\gamma_{\mathrm{ALG}} = \gamma_{\mathrm{Bayes}}$ follows from [16]). In contrast, numerical simulations and nonrigorous calculations using the cavity method from statistical physics (see [30]) suggest that SDP is suboptimal.

REMARK 2.6. A result analogous to Theorem 2 for the symmetric two-points distribution $\nu_{X_0} = (1/2)\delta_{+1} + (1/2)\delta_{-1}$ is proved in [46], Theorem 3, in the context of the stochastic block model of random graphs. Note, however, that the approach of [46] requires the graph to have average degree $d \to \infty$, $d = O(\log n)$.

2.5. *An example*: *Two-points distributions*. Theorem 2 is already interesting in very simple cases. Consider the two-points mixture:

$$\nu_{X_0} = \varepsilon\delta_{a_+} + (1-\varepsilon)\delta_{-a_-}, \tag{2.33}$$

$$a_+ = \sqrt{\frac{1-\varepsilon}{\varepsilon}}, \qquad a_- = \sqrt{\frac{\varepsilon}{1-\varepsilon}}. \tag{2.34}$$

Here, the coefficients a_+, a_- are chosen to ensure that $\int x\nu_{X_0}(\mathrm{d}x) = 0$, $\int x^2\nu_{X_0}(\mathrm{d}x) = 1$. The conditional expectation $F(y;\gamma)$ of equation (2.26) can be computed explicitly, yielding

$$F(y;\gamma) = \frac{\varepsilon a_+ e^{a_+ y - \gamma a_+^2/2} - (1-\varepsilon)a_- e^{-a_- y - \gamma a_-^2/2}}{\varepsilon e^{a_+ y - \gamma a_+^2/2} + (1-\varepsilon)e^{-a_- y - \gamma a_-^2/2}}. \tag{2.35}$$

Figure 1 reports the results of numerical simulations with the AMP algorithm described in the previous section. We also plot $\gamma_*(\lambda)/\lambda^2$ as a function of λ, where $\gamma_*(\lambda)$ is the fixed point of the state-evolution equation (2.24). The figure shows plots for four values of $\varepsilon \in (0, 1/2]$. The qualitative behavior depends on the value of ε. For ε close enough to $1/2$, equation (2.24) only has one stable fixed point[4] that is also the minimizer of the free energy functional (2.31). Hence $\gamma_{\mathrm{ALG}}(\gamma) = \gamma_{\mathrm{Bayes}}(\lambda)$ for all values of λ: message passing is always Bayes optimal.

[4]This is proved formally in [16] for $\varepsilon = 1/2$ and holds by a continuity argument for ε close enough to $1/2$. However, here we will limit ourselves to a heuristic discussion based on the numerical solution of equation (2.24).

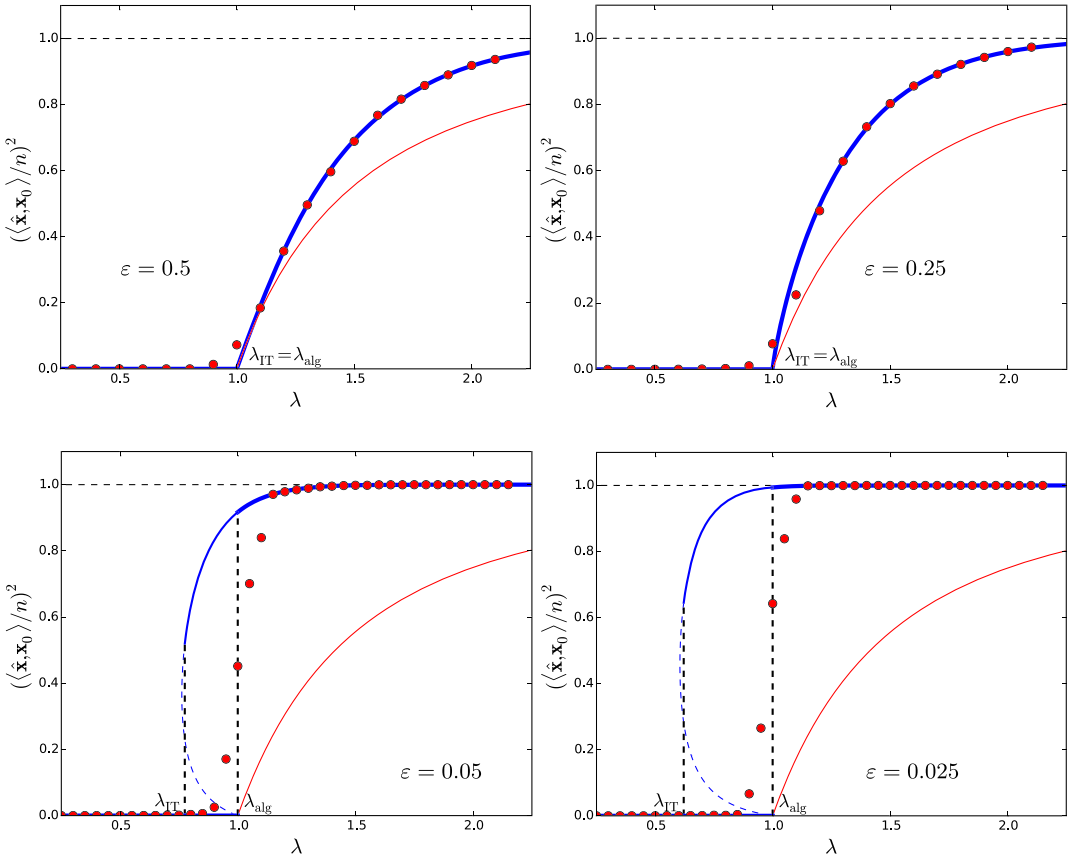

FIG. 1. *Estimation in the single spiked model* (2.1) *with entries of* $\boldsymbol{x}_0$ *following the two-points distribution of equation* (2.33)*, and four different values of the sparsity* $\varepsilon \in \{0.025, 0.05, 0.25, 0.5\}$. *Continuous thick blue line*: *asymptotic accuracy achieved by AMP* (*with spectral initialization*). *Red circles*: *numerical simulations with the AMP algorithm* (*form matrices of dimension* $n = 2000$ *and* $t = 200$ *iterations*). *Continuous thin blue line*: *Bayes optimal estimation accuracy. Dashed blue line*: *other fixed points of state evolution. Red line*: *Accuracy achieved by principal component analysis. Vertical dashed black lines*: *the thresholds* $\lambda_{\rm IT}$ *and* $\lambda_{\rm ALG}$.

For ε small enough, there exists $\lambda_0(\varepsilon) < 1$ such that equation (2.24) has three fixed points for $\lambda \in (\lambda_0(\varepsilon), 1)$: $\gamma_0(\lambda) < \gamma_1(\lambda) < \gamma_2(\lambda)$ whereby $\gamma_0 = 0$ and γ_2 are stable and γ_1 is unstable. AMP is controlled by the smallest stable fixed point, and hence $\gamma_{\rm ALG}(\lambda) = 0$ for all $\lambda < 1$. On the other hand, by minimizing the free energy (2.31) over these fixed points, we obtain that there exists $\lambda_{\rm IT}(\varepsilon) \in (\lambda_0(\varepsilon), 1)$ such that $\gamma_{\rm Bayes}(\lambda) = 0$ for $\lambda < \lambda_{\rm IT}(\varepsilon)$ while $\gamma_{\rm Bayes}(\lambda) = \gamma_2(\lambda)$ for $\lambda > \lambda_{\rm IT}(\varepsilon)$. We conclude that AMP is asymptotically suboptimal for $\lambda \in (\lambda_{\rm IT}(\varepsilon), 1)$, while it is asymptotically optimal for $\lambda \in [0, \lambda_{\rm IT}(\varepsilon))$ and $\lambda \in (1, \infty)$.

3. Confidence intervals, *p*-values, asymptotic FDR control. As an application of Theorem 2, we can construct confidence intervals that achieve a preassigned coverage level $(1 - \alpha)$, where $\alpha \in (0, 1)$. Indeed, Theorem 2 informally states that the AMP iterates $\boldsymbol{x}^t$ are approximately Gaussian with mean (proportional to) the signal $\boldsymbol{x}_0$. This relation can be inverted to construct confidence intervals.

We begin by noting that we do not need to know the signal strength λ. Indeed, for $\lambda > 1$, the latter can be estimated from the maximum eigenvalue of $\boldsymbol{A}$, $\lambda_{\max}(\boldsymbol{A})$, via

$$\hat{\lambda}(\boldsymbol{A}) \equiv \frac{1}{2}\{\lambda_{\max}(\boldsymbol{A}) + \sqrt{\lambda_{\max}(\boldsymbol{A})^2 - 4}\}. \tag{3.1}$$

This is a consistent estimator for $\lambda > 1$, and can replace λ in the iteration of equation (2.8) and initialization (2.20) as well as in the state evolution iteration of equations (2.23) and (2.24). We discuss two constructions of confidence intervals: the first one uses the Bayes AMP algorithm of Section 2.4, and the second instead uses the general algorithm of Section 2.2. The optimality of Bayes AMP translates into shorter confidence intervals but also requires knowledge of the empirical distribution ν_{X_0}.

Bayes-optimal construction. In order to emphasize the fact that we use the estimated λ both in the AMP iteration and in the state evolution recursion, we write $\overline{\boldsymbol{x}}^t$ for the Bayes AMP iterates and $\hat{\gamma}_t$ for the state evolution parameter, instead of $\boldsymbol{x}^t$ and γ_t. We then form the intervals:

$$\hat{J}_i(\alpha; t) = \left[\frac{1}{\hat{\gamma}_t}\overline{x}_i^t - \frac{1}{\sqrt{\hat{\gamma}_t}}\Phi^{-1}\left(1-\frac{\alpha}{2}\right), \frac{1}{\hat{\gamma}_t}\overline{x}_i^t + \frac{1}{\sqrt{\hat{\gamma}_t}}\Phi^{-1}\left(1-\frac{\alpha}{2}\right)\right]. \tag{3.2}$$

We can also define corresponding p-values by

$$p_i = 2\left(1 - \Phi\left(\frac{1}{\sqrt{\hat{\gamma}_t}}|\overline{x}_i^t|\right)\right). \tag{3.3}$$

General construction (no prior knowledge). Given a sequence of Lipschitz functions $f_t : \mathbb{R} \to \mathbb{R}$, we let $\boldsymbol{x}^t$ be the general AMP iterates as per Section 2.2; cf. equation (2.8). In order to form confidence intervals, we need to estimate the parameters μ_t, σ_t. In view of Theorem 1, a possible choice is given by

$$\hat{\sigma}_t^2 \equiv \frac{1}{n}\|f_{t-1}(\boldsymbol{x}^{t-1})\|_2^2, \tag{3.4}$$

$$\hat{\mu}_t^2 \equiv \frac{1}{n}\|\boldsymbol{x}^t\|_2^2 - \frac{1}{n}\|f_{t-1}(\boldsymbol{x}^{t-1})\|_2^2. \tag{3.5}$$

We then construct confidence intervals and p-values

$$\hat{J}_i(\alpha; t) = \left[\frac{1}{\hat{\mu}_t}x_i^t - \frac{\hat{\sigma}_t}{\hat{\mu}_t}\Phi^{-1}\left(1-\frac{\alpha}{2}\right), \frac{1}{\hat{\mu}_t}x_i^t + \frac{\hat{\sigma}_t}{\hat{\mu}_t}\Phi^{-1}\left(1-\frac{\alpha}{2}\right)\right], \tag{3.6}$$

$$p_i(t) = 2\left(1 - \Phi\left(\frac{1}{\hat{\sigma}_t}|x_i^t|\right)\right). \tag{3.7}$$

COROLLARY 3.1. *Consider the spiked matrix model* (2.1), *under the assumptions of Theorem* 1 (*in case of no prior knowledge*) *or Theorem* 2 (*for the Bayes optimal construction*). *Defining the confidence intervals* $\hat{J}_i(\alpha; t)$ *as per equations* (3.2) (3.6), *we have almost surely*

$$\lim_{n\to\infty}\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}\big(x_{0,i} \in \hat{J}_i(\alpha; t)\big) = 1 - \alpha. \tag{3.8}$$

Further assume that the fraction of nonzero entries in the spike is $\|\boldsymbol{x}_0(n)\|_0/n \to \varepsilon \in [0, 1)$, *and* $\nu_{X_0}(\{0\}) = 1 - \varepsilon$. *Then the p-values constructed above are asymptotically valid for the nulls. Namely, let* $i_0 = i_0(n)$ *any index such that* $x_{0,i_0}(n) = 0$. *Then, for any* $\alpha \in [0, 1]$, *and any fixed* $t \geq 0$,

$$\lim_{n\to\infty}\mathbb{P}\big(p_{i_0(n)}(t) \leq \alpha\big) = \alpha. \tag{3.9}$$

The proof of this result is presented in Appendix G in the Supplementary Material [44]. Notice that, by dominated convergence, this corollary also implies validity of the confidence intervals on average, namely $\lim_{n\to\infty}\frac{1}{n}\sum_{i=1}^{n}\mathbb{P}(x_{0,i} \in \hat{J}_i(\alpha; t)) = 1-\alpha$. As mentioned

above, cf. Remark 2.5, the Bayes-optimal construction maximizes the ratio μ_t/σ_t and, therefore, minimizes the length of confidence intervals. This requires however additional knowledge of the empirical distribution ν_{X_0}.

Corollary 3.1 allows to control the probability of false positives when using the p-values p_i; see equation (3.9). We might want to use these p-values to select a subset of variables $\hat{S} \subseteq [p]$ to be considered for further exploration. For such applications, it is common to aim for false discovery rate (FDR) control. The p-values p_i guarantee asymptotic FDR control through a simple Benjamini–Hochberg procedure [10]. For a threshold $s \in [0, 1]$, we define the following estimator of false discovery proportion [23]:

$$\widehat{\mathrm{FDP}}(s; t) \equiv \frac{ns}{1 \vee (\sum_{i=1}^{n} \mathbb{I}_{\{p_i(t) \le s\}})}. \tag{3.10}$$

Using this notion, we define a threshold and a rejection set as follows. Fix $\alpha \in (0, 1)$, let

$$s_*(\alpha; t) \equiv \inf\{s \in [0, 1] : \widehat{\mathrm{FDP}}(s; t) \ge \alpha\}, \qquad \hat{S}(\alpha; t) \equiv \{i \in [n] : p_i(t) < s_*(\alpha; t)\}. \tag{3.11}$$

The false discovery rate for this procedure is defined as usual

$$\mathrm{FDR}(\alpha, t; n) \equiv \mathbb{E}\left\{ \frac{|\hat{S}(\alpha; t) \cap \{i : x_{0,i} = 0\}|}{1 \vee |\hat{S}(\alpha; t)|} \right\}. \tag{3.12}$$

Our next corollary shows that the above procedure is guaranteed to control FDR in an asymptotic sense. Its proof can be found in Appendix H in the Supplementary Material [44].

COROLLARY 3.2. *Consider the spiked matrix model* (2.1), *under the assumptions of Theorem* 1 (*in case of no prior knowledge*) *or Theorem* 2 (*for the Bayes optimal construction*). *Further assume that the fraction of nonzero entries in the spike is* $\|\boldsymbol{x}_0(n)\|_0/n \to \varepsilon \in [0, 1)$, *and* $\nu_{X_0}(\{0\}) = 1 - \varepsilon$. *Then, for any fixed* $t \ge 0$,

$$\lim_{n\to\infty} \mathrm{FDR}(\alpha, t; n) = (1 - \varepsilon)\alpha. \tag{3.13}$$

REMARK 3.1. The procedure defined by threshold and rejection set in equation (3.11) does not assume knowledge of the sparsity level ε. If one knew ε, then an asymptotic false discovery rate of exactly α can be obtained by defining [49]

$$\widehat{\mathrm{FDP}}(s; t) \equiv \frac{n(1 - \varepsilon)s}{1 \vee (\sum_{i=1}^{n} \mathbb{I}_{\{p_i(t) \le s\}})}.$$

With the threshold and rejection set defined as in equation (3.11), such a procedure would have an asymptotic FDR equal to α, and higher power than the procedure using the estimator in equation (3.10).

4. Estimation of rectangular rank-one matrices. The algorithms and analysis developed in previous sections can be generalized to rectangular matrices. We illustrate this by generalizing the rank-one result of Theorem 1. We consider a data matrix $\boldsymbol{A} \in \mathbb{R}^{n\times d}$ given by

$$\boldsymbol{A} = \frac{\lambda}{n} \boldsymbol{u}_0 \boldsymbol{x}_0^{\mathsf{T}} + \boldsymbol{W}, \tag{4.1}$$

where $(\boldsymbol{W}_{ij})_{i\le n, j\le d} \sim_{\mathrm{iid}} \mathsf{N}(0, 1/n)$. To be definite, we will think of sequences of instances indexed by n and assume $n, d \to \infty$ with aspect ratio $d(n)/n \to \alpha \in (0, \infty)$.

We will make the following assumptions on the sequences of vectors $\boldsymbol{u}_0 = \boldsymbol{u}_0(n)$, $\boldsymbol{x}_0 = \boldsymbol{x}_0(n)$:

(i) Their rescaled ℓ_2-norms converge:

$$\lim_{n\to\infty} \|\boldsymbol{u}_0(n)\|_2/\sqrt{n} = 1, \qquad \lim_{n\to\infty} \|\boldsymbol{x}_0(n)\|_2/\sqrt{d(n)} = 1;$$

(ii) The empirical distributions of the entries of $\boldsymbol{x}_0(n)$ and $\boldsymbol{u}_0(n)$ converges weakly to probability distributions ν_{X_0}, ν_{U_0}, on $\mathbb{R}$, with unit second moment.

In analogy with the symmetric case, we initialize the AMP iteration by using the principal right singular vector of $\boldsymbol{A}$, denoted by $\boldsymbol{\varphi}_1$ (which we assume to have unit norm). In the present case, the phase transition for the principal singular vector takes place at $\lambda^2\sqrt{\alpha} = 1$ [2, 48]. Namely, if $\lambda^2\sqrt{\alpha} > 1$ then the correlation between $|\langle \boldsymbol{x}_0, \boldsymbol{\varphi}_1\rangle|/\|\boldsymbol{x}_0\|$ stays bounded away from zero as $n, d \to \infty$.

Setting $\boldsymbol{x}^0 = \sqrt{d}\boldsymbol{\varphi}_1$ and $g_{t-1}(\boldsymbol{u}^{t-1}) = \boldsymbol{0}$, we consider the following AMP iteration:

$$\boldsymbol{u}^t = \boldsymbol{A} f_t(\boldsymbol{x}^t) - \mathsf{b}_t g_{t-1}(\boldsymbol{u}^{t-1}), \quad \mathsf{b}_t = \frac{1}{n}\sum_{i=1}^{d} f_t'(x_i^t), \tag{4.2}$$

$$\boldsymbol{x}^{t+1} = \boldsymbol{A}^{\mathsf{T}} g_t(\boldsymbol{u}^t) - \mathsf{c}_t f_t(\boldsymbol{x}^t), \quad \mathsf{c}_t = \frac{1}{n}\sum_{i=1}^{n} g_t'(u_i^t). \tag{4.3}$$

The asymptotic characterization of this iteration is provided by the next theorem, which generalizes Theorem 1 to the rectangular case.

THEOREM 3. *Consider the $k = 1$ spiked matrix model of equation (4.1), with $n, d \to \infty$, $d/n \to \alpha$. Assume $\boldsymbol{x}_0(n) \in \mathbb{R}^d$, $\boldsymbol{u}_0(n) \in \mathbb{R}^d$ to be two sequences of vectors satisfying assumptions (i), (ii) above, and $\lambda^2\sqrt{\alpha} > 1$. Consider the AMP iteration in equation (2.8) with initialization $\boldsymbol{x}^0 = \sqrt{n}\boldsymbol{\varphi}_1$ (where, without loss of generality $\langle \boldsymbol{x}_0, \boldsymbol{\varphi}_1\rangle \geq 0$). Assume $f_t, g_t : \mathbb{R} \to \mathbb{R}$ to be Lipschitz continuous for each $t \in \mathbb{N}$.*

Let $(\mu_t, \sigma_t)_{t\geq 0}$ be defined via the recursion

$$\mu_{t+1} = \lambda\mathbb{E}\big[U_0 g_t(\overline{\mu}_t U_0 + \overline{\sigma}_t G)\big], \qquad \sigma_{t+1}^2 = \mathbb{E}\big[g_t(\overline{\mu}_t U_0 + \overline{\sigma}_t G)^2\big], \tag{4.4}$$

$$\overline{\mu}_t = \lambda\alpha\mathbb{E}\big[X_0 f_t(\mu_t X_0 + \sigma_t G)\big], \qquad \overline{\sigma}_t^2 = \alpha\mathbb{E}\big[f_t(\mu_t X_0 + \sigma_t G)^2\big], \tag{4.5}$$

where $X_0 \sim \nu_{X_0}$, $U_0 \sim \nu_{U_0}$ and $G \sim \mathsf{N}(0, 1)$ are independent, and the initial condition is

$$\mu_0 = \sqrt{\frac{1-\alpha^{-1}\lambda^{-4}}{1+\lambda^{-2}}}, \qquad \sigma_0 = \sqrt{\frac{\lambda^{-2}+\alpha^{-1}\lambda^{-4}}{1+\lambda^{-2}}}. \tag{4.6}$$

(This is to be substituted in equation (4.5) to yield $\overline{\mu}_0, \overline{\sigma}_0$.)

Then, for any function $\psi : \mathbb{R}\times\mathbb{R} \to \mathbb{R}$ with $|\psi(\boldsymbol{x}) - \psi(\boldsymbol{y})| \leq C(1 + \|\boldsymbol{x}\|_2 + \|\boldsymbol{y}\|_2)\|\boldsymbol{x} - \boldsymbol{y}\|_2$ for a universal constant $C > 0$, the following holds almost surely for $t \geq 0$:

$$\lim_{n\to\infty} \frac{1}{d(n)}\sum_{i=1}^{d(n)} \psi\big(x_{0,i}, x_i^t\big) = \mathbb{E}\big\{\psi(X_0, \mu_t X_0 + \sigma_t G)\big\}, \tag{4.7}$$

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} \psi\big(u_{0,i}, u_i^t\big) = \mathbb{E}\big\{\psi(U_0, \overline{\mu}_t U_0 + \overline{\sigma}_t G)\big\}. \tag{4.8}$$

As a special class of examples covered by this setting, we can consider the case in which we are given i.i.d. Gaussian samples $(\boldsymbol{y}_i)_{i\leq n} \sim \mathsf{N}(\boldsymbol{0}, \boldsymbol{\Sigma})$, with covariance matrix $\boldsymbol{\Sigma} = \rho^2\tilde{\boldsymbol{x}}_0\tilde{\boldsymbol{x}}_0^{\mathsf{T}} + \boldsymbol{I}_d$ where $\tilde{\boldsymbol{x}}_0 = \boldsymbol{x}_0/\sqrt{d}$. Letting $\boldsymbol{A}$ be the matrix with ith row equal to $\boldsymbol{y}_i/\sqrt{n}$, this takes the form of equation (4.1), with $\boldsymbol{u}_0 \sim \mathsf{N}(0, \boldsymbol{I}_n)$, and $\lambda = \rho/\sqrt{\alpha}$. Notice that the

sequence of random Gaussian vectors $\boldsymbol{u}_0(n)$, $n \geq 1$ satisfies conditions (i), (ii) above almost surely, with limit distribution ν_{U_0} equal to the standard Gaussian measure.

In this case, the optimal choice of the function g_t in equation (4.4) is of course linear: $g_t(u) = a_t u$ for some $a_t > 0$. The value of the constant a_t is immaterial, because it only amounts to a common rescaling of the μ_t, σ_t, which can be compensated by a redefinition of f_t in equation (4.5). We set $a_t = \lambda\overline{\mu}_t/(\overline{\mu}_t^2 + \overline{\sigma}_t^2)$. Substituting in equation (4.4), we obtain $\mu_{t+1} = \sigma_{t+1}^2 = \gamma_{t+1}$, where

$$\gamma_{t+1} = \frac{\lambda^2\overline{\gamma}_t^2}{1+\overline{\gamma}_t^2}, \tag{4.9}$$

where $\overline{\gamma}_t = \overline{\mu}_t^2/\overline{\sigma}_t^2$. Taking the ratio of the two equations in (4.5), we obtain

$$\overline{\gamma}_t = \lambda^2\alpha\frac{\mathbb{E}\{X_0 f_t(\gamma_t X + \sqrt{\gamma_t}G)\}^2}{\mathbb{E}\{f_t(\gamma_t X + \sqrt{\gamma_t}G)^2\}}. \tag{4.10}$$

We thus reduced the problem of covariance estimation in the spiked model $\boldsymbol{\Sigma} = \rho^2\tilde{\boldsymbol{x}}_0\tilde{\boldsymbol{x}}_0^{\mathsf{T}} + \boldsymbol{I}_d$, to the analysis of a one-dimensional recursion defined by equations (4.9), (4.10).

5. Degenerate cases and nonconcentration. The spectral initialization at unstable fixed points leads to a new phenomenon that is not captured by previous theory [7]: the evolution of empirical averages (e.g., estimation accuracy) does not always concentrate around a deterministic value. Our main result, Theorem 5 below, provides a description of this phenomenon by establishing a state evolution limit that is *dependent on the random initial condition*. The initial condition converges in distribution to a well-defined limit, which—together with state evolution—yields a complete characterization of the asymptotic behavior of the message passing algorithm.

The nonconcentration phenomenon arises when the deterministic low-rank component in equation (1.1) has degenerate eigenvalues. This is unavoidable in cases in which the underlying low-rank model to be estimated has symmetries.

Here, we illustrate this phenomenon on a simple model that we will refer to as the Gaussian Block Model (GBM). For $q \geq 3$ a fixed integer, let $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$ be a vector of vertex labels with $\sigma_i \in \{1, \dots, q\}$ and consider deterministic matrix $\boldsymbol{A}_0 \in \mathbb{R}^{n\times n}$ (with $\operatorname{rank}(\boldsymbol{A}_0) = q-1$) defined by

$$A_{0,ij} = \begin{cases} (q-1)/n & \text{if } \sigma_i = \sigma_j, \\ -1/n & \text{otherwise.} \end{cases} \tag{5.1}$$

We assume the vertex labeling to be perfectly balanced. That is, $\sum_{i=1}^n \mathbf{1}_{\sigma_i=\sigma} = n/q$ for $\sigma \in \{1, \dots, q\}$: While most of our discussion holds under an approximate balance condition, this assumption avoids some minor technical complications. Notice that $\boldsymbol{A}_0$ is an orthogonal projector on a subspace $\mathcal{V}_n \in \mathbb{R}^n$ of dimension $q-1$. We observe the noisy matrix (with noise $\boldsymbol{W} \sim \mathsf{GOE}(n)$)

$$\boldsymbol{A} = \lambda\boldsymbol{A}_0 + \boldsymbol{W}, \tag{5.2}$$

and would like to estimate $\boldsymbol{A}_0$ from these noisy observations. The matrix $\boldsymbol{A}$ takes the form of equation (1.1) with $k = q-1$, $\lambda_1 = \cdots = \lambda_k = \lambda$ and $\boldsymbol{v}_1, \dots, \boldsymbol{v}_k$ an orthonormal basis of the space $\mathcal{V}_n$. We will assume $\lambda > 1$ so that $k_* = k$. In particular, for $q \geq 3$, the low-rank signal has degenerate eigenvalues.

We use the following AMP algorithm to estimate $\boldsymbol{A}_0$. We compute the top k eigenvectors of $\boldsymbol{A}$, denoted by $\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_k \in \mathbb{R}^n$ and generate $\boldsymbol{x}^t \in \mathbb{R}^{n\times q}$ for $t \geq 0$, according to

$$\boldsymbol{x}^0 = [\sqrt{n}\boldsymbol{\varphi}_1 | \cdots | \sqrt{n}\boldsymbol{\varphi}_k | \mathbf{0}], \tag{5.3}$$

$$\boldsymbol{x}^{t+1} = \boldsymbol{A} f(\boldsymbol{x}^t) - f(\boldsymbol{x}^{t-1})\mathsf{B}_t^{\mathsf{T}}, \tag{5.4}$$

where the "Onsager coefficient" $\mathsf{B}_t \in \mathbb{R}^{q\times q}$ is a matrix given by

$$\mathsf{B}_t = \frac{1}{n}\sum_{i=1}^{n} \frac{\partial f}{\partial \boldsymbol{x}}(\boldsymbol{x}_i^t, y_i). \tag{5.5}$$

Here, $\frac{\partial f}{\partial \boldsymbol{x}} \in \mathbb{R}^{q\times q}$ denotes the Jacobian matrix of the function $f : \mathbb{R}^q \to \mathbb{R}^q$. Furthermore, the function $f : \mathbb{R}^q \to \mathbb{R}^q$ is defined by letting, for $\sigma \in \{1, \dots, q\}$:

$$f(\boldsymbol{z})_\sigma = \lambda\left[\frac{q e^{z_\sigma}}{\sum_{\tau=1}^{q} e^{z_\tau}} - 1\right], \tag{5.6}$$

and $f(\boldsymbol{x})$ is defined for $\boldsymbol{x} \in \mathbb{R}^{n\times q}$ by applying the same function row by row. This choice of the function f corresponds to Bayes-optimal estimation as can be deduced from the state evolution analysis below: we will not discuss this point in detail here.

The output $\boldsymbol{x}^t$ after t iterations of (5.4) can be interpreted as an estimate of the labels $\boldsymbol{\sigma}$ in the following sense. Let $\boldsymbol{x}_0 \in \mathbb{R}^{n\times q}$ be the matrix whose ith row is $\boldsymbol{x}_{0,i} = \boldsymbol{P}^{\perp}\boldsymbol{e}_{\boldsymbol{\sigma}_i}$, with $\boldsymbol{P}^{\perp} \in \mathbb{R}^{q\times q}$ the projector orthogonal to the all ones vector, and $\boldsymbol{e}_1, \dots, \boldsymbol{e}_q$ the canonical basis in $\mathbb{R}^q$. Note that $\boldsymbol{A}_0 = (q/n)\boldsymbol{x}_0\boldsymbol{x}_0^{\mathsf{T}}$. Then $\boldsymbol{x}^t$ is an estimator of $\boldsymbol{x}_0$ (up to a permutation of the labels' alphabet $\{1, \dots, q\}$).

Let S_q be the group of $q \times q$ permutation matrices. We evaluate the estimator $\boldsymbol{x}^t$ via the overlap

$$\mathrm{Overlap}_n(\lambda; t) \equiv \max_{\boldsymbol{\Pi}\in\mathsf{S}_q} \frac{\langle \boldsymbol{x}^t, \boldsymbol{x}_0\boldsymbol{\Pi}\rangle}{\|\boldsymbol{x}^t\|_F \|\boldsymbol{x}_0\|_F}, \tag{5.7}$$

where $\langle\cdot,\cdot\rangle$ denotes the Frobenius inner product. In Figure 2, we plot the evolution of the overlap in two sets of numerical simulations, for $q = 3$ and $q = 4$. Each curve is obtained by running AMP (with spectral initialization) on a different realization of the random matrix $\boldsymbol{A}$. The nonconcentration phenomenon is quite clear:

- For fixed number of iterations t and large n, the quantity $\mathrm{Overlap}_n(\lambda; t)$ has large fluctuations, that do not seem to vanish as $n \to \infty$.
- Despite this, the algorithm is effective in reconstructing the signal: after $t = 10$ iterations, the accuracy achieved is nearly independent of the initialization.

The empirical data in Figure 2 are well described by the state evolution prediction that is shown as continuous curves in the same figure. In this case, state evolution operates on the pair of matrices $\boldsymbol{M}_t, \boldsymbol{Q}_t \in \mathbb{R}^{q\times q}$, which are updated according to

$$\boldsymbol{M}_{t+1} = \lambda\mathbb{E}\{f(q\boldsymbol{M}_t\boldsymbol{e}_\sigma + \boldsymbol{Q}_t^{1/2}\boldsymbol{G})\boldsymbol{e}_\sigma^{\mathsf{T}}\boldsymbol{P}^{\perp}\}, \tag{5.8}$$

$$\boldsymbol{Q}_{t+1} = \mathbb{E}\{f(q\boldsymbol{M}_t\boldsymbol{e}_\sigma + \boldsymbol{Q}_t^{1/2}\boldsymbol{G})f(q\boldsymbol{M}_t\boldsymbol{e}_\sigma + \boldsymbol{Q}_t^{1/2}\boldsymbol{G})^{\mathsf{T}}\}, \tag{5.9}$$

where $f : \mathbb{R}^q \to \mathbb{R}^q$ is defined as per equation (5.6), and expectation is with respect to σ uniform in $\{1, \dots, q\}$ independent of $\boldsymbol{G} \sim \mathsf{N}(0, \boldsymbol{I}_q)$. Note that $\boldsymbol{Q}_t$ is symmetric and both $\boldsymbol{Q}_t\mathbf{1} = \boldsymbol{M}_t\mathbf{1} = \mathbf{1}^{\mathsf{T}}\boldsymbol{M}_t = 0$ for all $t \geq 1$.

The state evolution prediction for the present model is provided by the next theorem, which is proved in Appendix I in the Supplementary Material [44].

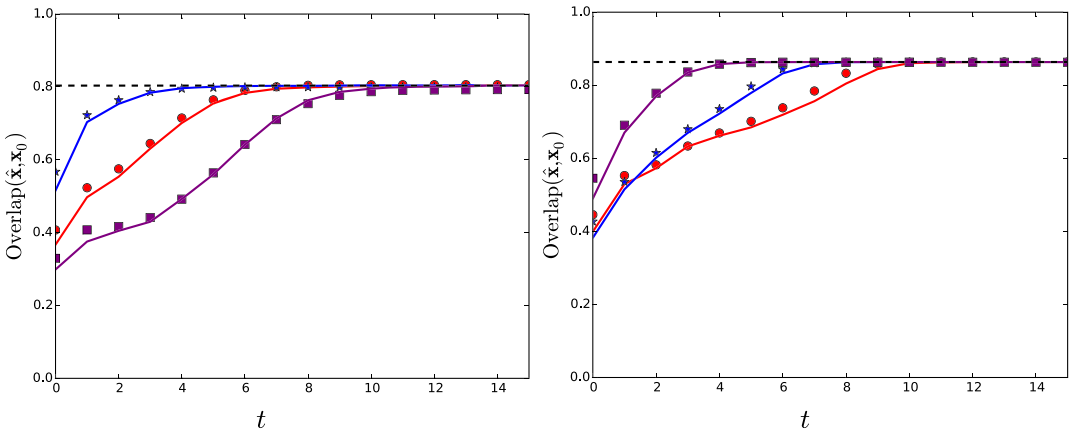

FIG. 2. *Estimation in the Gaussian Block Model of equation (5.2) using the AMP algorithm with spectral initialization of equations (5.3), (5.4). We plot the reconstruction accuracy (overlap) as a function of the number of iterations for $q = 3$, $\lambda = 1.5$, $n = 6000$ (left frame), and $q = 4$, $\lambda = 1.75$, $n = 8000$ (right frame). Each set of symbols corresponds to a different realization of the random matrix $\boldsymbol{A}$, and curves report the corresponding prediction of Theorem 4. Dashed black lines report the Bayes optimal accuracy as per [5, 39].*

THEOREM 4. *Let $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ be the random matrix of equation (5.2) with $\lambda > 1$, and let $\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_k$ be its top k eigenvectors. Denote by $\boldsymbol{x}^t$ the sequence of estimates produced by the AMP algorithm of equation (5.3) with the spectral initialization in equation (5.4).*

Let $\{\boldsymbol{M}_t, \boldsymbol{Q}_t\}_{t\ge 0}$ be the state evolution iterates with initialization $\boldsymbol{M}_0 = (\boldsymbol{x}^0)^{\mathsf{T}}\boldsymbol{x}_0/n$ and $\boldsymbol{Q}_0 = \lambda^{-1}\,\mathrm{diag}(1, 1, \dots, 1, 0)$ Then, for any function $\psi : \mathbb{R}^{2q} \to \mathbb{R}$ with $|\psi(\boldsymbol{x}) - \psi(\boldsymbol{y})| \le C(1 + \|\boldsymbol{x}\|_2 + \|\boldsymbol{y}\|_2)\|\boldsymbol{x} - \boldsymbol{y}\|_2$, we have, almost surely

$$\lim_{n\to\infty}\left|\frac{1}{n}\sum_{i=1}^{n}\psi(\boldsymbol{x}_i^t, \boldsymbol{x}_{0,i}) - \mathbb{E}\{\psi(q\boldsymbol{M}_t\boldsymbol{e}_\sigma + \boldsymbol{Q}_t^{1/2}\boldsymbol{G}, \boldsymbol{P}^{\perp}\boldsymbol{e}_\sigma)\}\right| = 0, \tag{5.10}$$

where expectation is with respect to σ uniform in $\{1, \dots, q\}$ independent of $\boldsymbol{G} \sim \mathsf{N}(0, \boldsymbol{I}_q)$.

Further as $n \to \infty$, $\boldsymbol{M}_0$ converges in distribution as

$$\boldsymbol{M}_0 \overset{\mathrm{d}}{\Rightarrow} \sqrt{q^{-1}(1-\lambda^{-2})}\begin{bmatrix}\boldsymbol{O}^{\mathsf{T}}_{(q-1)\times q}\\ \boldsymbol{0}_{1\times q}\end{bmatrix}, \tag{5.11}$$

where $\boldsymbol{O} \in \mathbb{R}^{q\times(q-1)}$ is Haar distributed orthogonal matrix with column space orthogonal to $\boldsymbol{1}$.

The continuous curves in Figure 2 are obtained as described in the last theorem. For each experiment, we generate a random matrix $\boldsymbol{A}$ according to equation (5.2), compute the spectral initialization of equation (5.3) and set $\boldsymbol{M}_0 = (\boldsymbol{x}^0)^{\mathsf{T}}\boldsymbol{x}_0/n$. We then compute the state evolution sequence $\{(\boldsymbol{M}_t, \boldsymbol{Q}_t)\}_{t\ge 0}$ via equations (5.8), (5.9), and use equation (5.10) to predict the evolution of the overlap. The variability in the initial condition $\boldsymbol{M}_0$ leads to a variability in the predicted trajectory $\{(\boldsymbol{M}_t, \boldsymbol{Q}_t)\}_{t\ge 0}$ that matches well with the empirical data.

Finally, as mentioned above, AMP converges to an accuracy that is roughly independent of the matrix realization for large t, and matches the Bayes optimal prediction of [5, 39]. While a full explanation of this phenomenon goes beyond the scope of the present paper, this behavior can be also explained by Theorem 4: the initialization $\boldsymbol{M}_0$ breaks the symmetry between the q blocks uniformly, as per equation (5.11). Once the symmetry is broken, the state evolution iteration of equations (5.8), (5.9) converges to a fixed point that is unique up to permutations.

6. Main result.

6.1. *Notation and definitions.* We say that a function $\psi : \mathbb{R}^d \to \mathbb{R}$ is pseudo-Lipschitz of order k (and write $\psi \in \mathrm{PL}(k)$) if there exists a constant L such that $|\psi(\boldsymbol{x}) - \psi(\boldsymbol{y})| \le L(1 + (\|\boldsymbol{x}\|/\sqrt{d})^{k-1} + (\|\boldsymbol{y}\|/\sqrt{d})^{k-1})\|\boldsymbol{x} - \boldsymbol{y}\|_2/\sqrt{d}$.

Recall that a sequence of probability distributions ν_n on $\mathbb{R}^m$ *converges weakly* to ν ($\nu_n \overset{w}{\Rightarrow} \nu$) if, for any bounded Lipschitz function $\psi : \mathbb{R}^m \to \mathbb{R}$, $\lim_{n\to\infty} \mathbb{E}\psi(\boldsymbol{X}_n) = \mathbb{E}\psi(\boldsymbol{X})$ where expectation is with respect to $\boldsymbol{X}_n \sim \nu_n$, $\boldsymbol{X} \sim \nu$. Given a (deterministic) sequence of matrices $\boldsymbol{Z}_n \in \mathbb{R}^{n\times d}$ indexed by n (with $d \ge 1$ fixed), we say that the empirical distribution of $\boldsymbol{Z}_n$ converges weakly to a probability distribution ν on $\mathbb{R}^d$ if, letting $\boldsymbol{z}_i = \boldsymbol{Z}_n^{\mathsf{T}}\boldsymbol{e}_i$ denote the ith row of $\boldsymbol{Z}_n$, for each i we have

$$\frac{1}{n}\sum_{i=1}^{n} \delta_{\boldsymbol{z}_{n,i}} \overset{w}{\Rightarrow} \nu. \tag{6.1}$$

Equivalently, $\lim_{n\to\infty} n^{-1}\sum_{i=1}^{n} \psi(\boldsymbol{z}_i) = \mathbb{E}\psi(\boldsymbol{z})$ for $\boldsymbol{z} \sim \nu$ and any bounded Lipschitz function ψ. We apply the same terminology if we are given d vectors $(\boldsymbol{z}_1^{(n)}, \dots, \boldsymbol{z}_d^{(n)})$, where $\boldsymbol{z}_\ell^{(n)} \in \mathbb{R}^n$: in this case $\boldsymbol{Z}_n$ is the matrix with columns $\boldsymbol{z}_1^{(n)}, \dots, \boldsymbol{z}_d^{(n)}$.

Given two probability measures μ (on the space $\mathcal{X}$) and ν (on the space $\mathcal{Y}$), a coupling ρ of μ and ν is a probability distribution on $\mathcal{X} \times \mathcal{Y}$ whose first marginal coincides with μ and second coincides with ν. We denote the set of couplings of μ, ν by $\mathcal{C}(\mu, \nu)$. For $k \ge 1$, the Wasserstein-k (W_k) distance between two probability measures μ, ν on $\mathbb{R}^d$ is defined by

$$W_k(\mu, \nu) \equiv \inf_{\rho\in\mathcal{C}(\mu,\nu)} \mathbb{E}_{(\boldsymbol{X},\boldsymbol{Y})\sim\rho}\{\|\boldsymbol{X} - \boldsymbol{Y}\|_2^k\}^{1/k}, \tag{6.2}$$

where the infimum is over all the couplings of μ and ν. A sequence of probability distributions ν_n on $\mathbb{R}^m$ *converges in* W_k to ν ($\nu_n \overset{W_k}{\Rightarrow} \nu$) if $\lim_{n\to\infty} W_k(\nu_n, \nu) = 0$. An equivalent definition is that, for any $\psi \in \mathrm{PL}(k)$, $\lim_{n\to\infty} \mathbb{E}\psi(\boldsymbol{X}_n) = \mathbb{E}\psi(\boldsymbol{X})$ where expectation is with respect to $\boldsymbol{X}_n \sim \nu_n$, $\boldsymbol{X} \sim \nu$ [51], Theorem 6.9.

Generalizing from the definitions introduced for weak convergence, given sequence of matrices $\boldsymbol{Z}_n \in \mathbb{R}^{n\times d}$ indexed by n (with $d \ge 1$ fixed), we say that the empirical distribution of $\boldsymbol{Z}_n$ converges in W_k to ν (a probability distribution on $\mathbb{R}^d$), if letting $\boldsymbol{z}_i = \boldsymbol{Z}_n^{\mathsf{T}}\boldsymbol{e}_i$ denote the ith row of $\boldsymbol{Z}_n$,

$$\frac{1}{n}\sum_{i=1}^{n} \delta_{\boldsymbol{z}_{n,i}} \overset{W_k}{\Rightarrow} \nu. \tag{6.3}$$

Equivalently, $\lim_{n\to\infty} n^{-1}\sum_{i=1}^{n} \psi(\boldsymbol{z}_i) = \mathbb{E}\psi(\boldsymbol{z})$ for any $\psi \in \mathrm{PL}(k)$ (where $\boldsymbol{z} \sim \nu$). Again the same terminology is used for d-tuples of vectors $(\boldsymbol{z}_1^{(n)}, \dots, \boldsymbol{z}_d^{(n)})$.

We will typically use upper case bold symbols for matrices (e.g., $\boldsymbol{A}$, $\boldsymbol{B}$,...), lower case bold for vectors (e.g., $\boldsymbol{u}$, $\boldsymbol{v}$, ...) and lower case plain font for scalars (e.g., x, y, ...). However, we will often denote random variables and random vectors using upper case.

We often consider vectors (or matrices) whose elements are indexed by arbitrary finite sets. For instance, given finite sets S_1, S_2, $\boldsymbol{Q} \in \mathbb{R}^{S_1\times S_2}$ is a matrix $\boldsymbol{Q} = (Q_{i,j})_{i\in S_1, j\in S_2}$. When there is an obvious ordering of the elements of S_1, S_2, such a matrix is understood to be identified with a matrix in $\mathbb{R}^{n_1\times n_2}$, where $n_i = |S_i|$. For instance, $\mathbb{R}^{[m]\times[n]}$ is identified with $\mathbb{R}^{m\times n}$. Given a vector $\boldsymbol{v} \in \mathbb{R}^m$ an a set $S \subseteq [m]$, we denote by $\boldsymbol{v}_S \in \mathbb{R}^S$ the subvector indexed by elements of S. Analogously, for a matrix $\boldsymbol{M} \in \mathbb{R}^{m\times n}$, we let $\boldsymbol{M}_{R,S} \in \mathbb{R}^{R\times S}$ be the submatrix with row indices in R and column indices in S. If the submatrix includes all the rows, we adopt the shorthand $\boldsymbol{M}_{[m],S}$.

Finally, we adopt the convention that all vectors (including the rows of a matrix) are viewed as column vectors, unless explicitly transposed.

6.2. *Statement of the result*: *Symmetric case*. Recall the spiked model of equation (1.1), which we copy here for the reader's convenience:

$$\boldsymbol{A} = \sum_{i=1}^{k} \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^{\mathsf{T}} + \boldsymbol{W} \equiv \boldsymbol{V} \boldsymbol{\Lambda} \boldsymbol{V}^{\mathsf{T}} + \boldsymbol{W}. \tag{6.4}$$

Here, $\boldsymbol{v}_i \in \mathbb{R}^n$ are nonrandom orthonormal vectors and $\boldsymbol{W} \sim \mathsf{GOE}(n)$. We denote by $\boldsymbol{\varphi}_1, \dots, \boldsymbol{\varphi}_n$ the eigenvectors of $\boldsymbol{A}$, with corresponding eigenvalues $z_1 \geq z_2 \geq \cdots \geq z_n$.

For a sequence of functions $f_t(\cdot) : \mathbb{R}^q \times \mathbb{R} \to \mathbb{R}^q$, we consider the AMP algorithm that produces a sequence of iterates $\boldsymbol{x}^t$ according to the recursion

$$\boldsymbol{x}^{t+1} = \boldsymbol{A} f_t(\boldsymbol{x}^t, \boldsymbol{y}) - f_{t-1}(\boldsymbol{x}^{t-1}, \boldsymbol{y}) \mathsf{B}_t^{\mathsf{T}}. \tag{6.5}$$

Here, $\boldsymbol{y} \in \mathbb{R}^n$ is a fixed vector, and it is understood that $f(\cdot; t)$ is applied row-by-row. Namely, denoting by $\boldsymbol{x}_i^t \in \mathbb{R}^q$ the ith row of $\boldsymbol{x}^t$, the ith row of $f_t(\boldsymbol{x}^t; \boldsymbol{y})$ is given by $f_t(\boldsymbol{x}_i^t, y_i)$. The "Onsager coefficient" $\mathsf{B}_t \in \mathbb{R}^{q \times q}$ is a matrix given by

$$\mathsf{B}_t = \frac{1}{n} \sum_{i=1}^{n} \frac{\partial f_t}{\partial \boldsymbol{x}}(\boldsymbol{x}_i^t, y_i), \tag{6.6}$$

where $\frac{\partial f_t}{\partial \boldsymbol{x}} \in \mathbb{R}^{q \times q}$ denotes the Jacobian matrix of the function $f_t(\cdot, y) : \mathbb{R}^q \to \mathbb{R}^q$. The algorithm is initialized with $\boldsymbol{x}^0 \in \mathbb{R}^{n \times q}$ and $f_{-1}(\boldsymbol{x}^{-1}, \boldsymbol{y}) \in \mathbb{R}^{n \times q}$ is taken to be the all-zeros matrix.

REMARK 6.1. Notice that the present setting generalizes the one of Section 2 in two directions (apart from the more general model for the matrix $\boldsymbol{A}$, cf. equation (6.4)). First, the state of the algorithm is a matrix $\boldsymbol{x}^t \in \mathbb{R}^{n \times q}$ with q an arbitrary fixed integer. While it is natural to take q equal to the number of outliers in the spectrum of $\boldsymbol{A}$ (i.e., $q = k_*$ according to the notation introduced below), we believe that a more general choice of q can be useful for certain applications. Further, the nonlinearity f_t is a function of $\boldsymbol{x}^t$ but also on the independent vector $\boldsymbol{y}$ that can be regarded as side information: again, we believe this additional freedom will be useful for future applications of our main result.

We will make the following assumptions:

(A1) The values $\lambda_i(n)$ have finite limits as $n \to \infty$, that we denote by λ_i. Further, assume there exist k_+, k_- such that $\lambda_1 \geq \dots \lambda_{k_+} > 1 > \lambda_{k_+ + 1}$ and $\lambda_{k - k_-} > -1 > \lambda_{k - k_- + 1} \geq \cdots \geq \lambda_k$. We let $S \equiv (1, \dots, k_+, k - k_- + 1, \dots, k)$, $k_* = k_+ + k_-$ and $\hat{S} \equiv (1, \dots, k_+, n - k_- + 1, \dots, n)$. Further, we let $\boldsymbol{\Lambda}_S$ denote the diagonal matrix with entries $(\boldsymbol{\Lambda}_S)_{ii} = \lambda_i$, $i \in S$.

(A2) Setting $q \geq k_*$, we initialize the iteration (6.5) by setting $\boldsymbol{x}^0 \in \mathbb{R}^{n \times q}$ equal to the matrix with first k_* ordered columns given by $(\sqrt{n} \boldsymbol{\varphi}_i)_{i \in \hat{S}}$, and $\boldsymbol{0}$ for the remaining $q - k_*$ columns.

(A3) The joint empirical distribution of the vectors $(\sqrt{n} \boldsymbol{v}_\ell(n))_{\ell \in S}$, and $\boldsymbol{y}$ has a limit in Wasserstein-2 metric. Namely, if we let $\tilde{\boldsymbol{v}}_i = (\sqrt{n} v_{\ell,i})_{\ell \in S} \in \mathbb{R}^{k_*}$, then there exists a random vector $\boldsymbol{U}$ taking values in $\mathbb{R}^{k_*}$ and a random variable Y, with joint law $\mu_{\boldsymbol{U},Y}$, such that

$$\frac{1}{n} \sum_{i=1}^{n} \delta_{\tilde{\boldsymbol{v}}_i, y_i} \overset{W_2}{\Rightarrow} \mu_{\boldsymbol{U},Y}. \tag{6.7}$$

(A4) The functions $f_t(\cdot, \cdot) : \mathbb{R}^q \times \mathbb{R} \to \mathbb{R}^q$ are Lipschitz continuous.

State evolution operates on the pair of matrices $\boldsymbol{M}_t \in \mathbb{R}^{q\times k_*}$, $\boldsymbol{Q}_t \in \mathbb{R}^{q\times q}$, with $\boldsymbol{Q}_t \succeq \boldsymbol{0}$, evolving according to

$$\boldsymbol{M}_{t+1} = \mathbb{E}\{f_t(\boldsymbol{M}_t\boldsymbol{U} + \boldsymbol{Q}_t^{1/2}\boldsymbol{G}, Y)\boldsymbol{U}^{\mathsf{T}}\}\boldsymbol{\Lambda}_S, \tag{6.8}$$

$$\boldsymbol{Q}_{t+1} = \mathbb{E}\{f_t(\boldsymbol{M}_t\boldsymbol{U} + \boldsymbol{Q}_t^{1/2}\boldsymbol{G}, Y) f_t(\boldsymbol{M}_t\boldsymbol{U} + \boldsymbol{Q}_t^{1/2}\boldsymbol{G}, Y)^{\mathsf{T}}\}, \tag{6.9}$$

where expectation is taken with respect to $(\boldsymbol{U}, Y) \sim \mu_{U,Y}$ independent of $\boldsymbol{G} \sim \mathsf{N}(0, \boldsymbol{I}_q)$. These recursions are initialized with $\boldsymbol{Q}_0$, $\boldsymbol{M}_0$ which will be specified in the statement of Theorem 5 below.

We denote by $\mathcal{R}(\boldsymbol{\Lambda}) \subseteq \mathbb{R}^{S\times[k]}$ the set of orthogonal matrices $\boldsymbol{R}$ (with $\boldsymbol{R}\boldsymbol{R}^{\mathsf{T}} = \boldsymbol{I}_S$) such that $R_{ij} = 0$ if $\lambda_i \neq \lambda_j$ or if $j \notin S$. Notice that the $k_* \times k_*$ submatrix $\boldsymbol{R}_{S,S}$ of $\boldsymbol{R} \in \mathcal{R}(\boldsymbol{\Lambda})$ is a block-diagonal orthogonal matrix, with blocks in correspondence with the degenerate λ_i's. As such, these matrices form a compact group, which we will denote by $\mathcal{R}_*(\boldsymbol{\Lambda}) \subseteq \mathbb{R}^{k_*\times k_*}$. This group can be endowed with the Haar measure, which is just the product of Haar measures over the orthogonal group corresponding to each block. We define the Haar measure on $\mathcal{R}(\boldsymbol{\Lambda})$ by adding $k - k_*$ columns equal to 0 for column indices $j \in [k] \setminus S$.

THEOREM 5. *Let $(\boldsymbol{x}^t)_{t\geq 0}$ be the AMP iterates generated by algorithm (6.5), under assumptions (A1) to (A4), for the spiked matrix model (1.1). For $\eta_n \geq n^{-1/2+\varepsilon}$ such that $\eta_n \to 0$ as $n \to \infty$, define the set of matrices*

$$\mathcal{G}_n(\boldsymbol{\Lambda}) \equiv \Big\{\boldsymbol{Q} \in \mathbb{R}^{S\times[k]} : \min_{\boldsymbol{R}\in\mathcal{R}(\boldsymbol{\Lambda})} \big\|\boldsymbol{Q} - \big(\boldsymbol{I} - \boldsymbol{\Lambda}_S^{-2}\big)^{1/2}\boldsymbol{R}\big\|_F \leq \eta_n\Big\}, \tag{6.10}$$

Let $\boldsymbol{\Omega} \equiv \boldsymbol{\Phi}_{\hat{S}}^{\mathsf{T}}\boldsymbol{V} \in \mathbb{R}^{k_*\times k}$ where $\boldsymbol{\Phi}_{\hat{S}} \in \mathbb{R}^{n\times k_*}$ is the matrix with columns $(\boldsymbol{\varphi}_i)_{i\in\hat{S}}$ and $\boldsymbol{V} \in \mathbb{R}^{n\times k}$ is the matrix with columns $(\boldsymbol{v}_i)_{i\in[k]}$. Denote by $\boldsymbol{\Omega}_0 \in \mathbb{R}^{S\times S}$ the submatrix corresponding to the k_* columns of $\boldsymbol{\Omega}$ with index in S, and let $\tilde{\boldsymbol{\Omega}}_0 = (\boldsymbol{I} - \boldsymbol{\Omega}_0\boldsymbol{\Omega}_0^{\mathsf{T}})^{1/2}$.

Then, for any pseudo-Lipschitz function $\psi : \mathbb{R}^{q+k_*+1} \to \mathbb{R}$, $\psi \in \mathrm{PL}(2)$, the following holds almost surely for $t \geq 0$:

$$\lim_{n\to\infty}\left|\frac{1}{n}\sum_{i=1}^{n}\psi(\boldsymbol{x}_i^t, \tilde{\boldsymbol{v}}_i, y_i) - \mathbb{E}\{\psi(\boldsymbol{M}_t\boldsymbol{U} + \boldsymbol{Q}_t^{1/2}\boldsymbol{G}, \boldsymbol{U}, Y)\}\right| = 0. \tag{6.11}$$

Here, $\tilde{\boldsymbol{v}}_i = (\sqrt{n}v_{\ell,i})_{\ell\in S} \in \mathbb{R}^{k_*}$ and expectation is with respect to $(\boldsymbol{U}, Y) \sim \mu_{U,Y}$ independent of $\boldsymbol{G} \sim \mathsf{N}(0, \boldsymbol{I}_q)$. Finally, $(\boldsymbol{M}_t, \boldsymbol{Q}_t)$ is the state evolution sequence specified by equations (6.8) and (6.9) with initialization $(\boldsymbol{M}_0)_{[k_*],[k_*]} = \boldsymbol{\Omega}_0$, $(\boldsymbol{M}_0)_{[q]\setminus[k_*],[k_*]} = \boldsymbol{0}$, $(\boldsymbol{Q}_0)_{[k_*],[k_*]} = \tilde{\boldsymbol{\Omega}}_0^2$, and $(\boldsymbol{Q}_0)_{i,j} = 0$ if $(i, j) \notin [k_*] \times [k_*]$.

Further, $\mathbb{P}(\boldsymbol{\Omega} \in \mathcal{G}_n(\boldsymbol{\Lambda})) \geq 1 - n^{-A}$ for any $A > 0$ provided $n > n_0(A)$, and $\boldsymbol{\Omega}$ converges in distribution to $(\boldsymbol{I} - \boldsymbol{\Lambda}_S^{-2})^{1/2}\boldsymbol{R}$, with $\boldsymbol{R}$ Haar distributed on $\mathcal{R}(\boldsymbol{\Lambda})$.

The theorem is proved for the case of a rank one spike in Appendix A. The proof for the general case is given in Appendix B. In the following section, we provide a brief overview of the key steps in the proof.

REMARK 6.2. Theorem 5 focuses on the case of symmetric square matrices $\boldsymbol{A}$. However, a standard reduction (see, for instance, [11], Section 6) allows to obtain a completely analogous statement for rectangular matrices, namely $\boldsymbol{A} \in \mathbb{R}^{n\times d}$ with

$$\boldsymbol{A} = \sum_{i=1}^{k}\lambda_i\boldsymbol{u}_i\boldsymbol{v}_i^{\mathsf{T}} + \boldsymbol{W}, \tag{6.12}$$

where $\boldsymbol{W}$ is a noise matrix with independent entries $W_{ij} \sim \mathsf{N}(0, 1/n)$. We already considered the case $k = 1$ of this model in Section 4. Given Theorem 5, the generalization to $k > 1$ rectangular matrices is straightforward: we provide a precise statement in Appendix J in the Supplementary Material [44].

Another generalization of interest would be to non-Gaussian matrices. It might be possible to address this by using the methods of [6].

7. Proof outline. We first consider the rank one spiked model in equation (2.1), and give an outline of the proof of Theorem 1. Letting $\boldsymbol{v} \equiv \frac{\boldsymbol{x}_0}{\sqrt{n}}$, equation (2.1) can be written as

$$\boldsymbol{A} = \lambda \boldsymbol{v}\boldsymbol{v}^{\mathsf{T}} + \boldsymbol{W}. \tag{7.1}$$

Recalling that $(\boldsymbol{\varphi}_1, z_1)$ are the principal eigenvector and eigenvalue of $\boldsymbol{A}$, we write $\boldsymbol{A}$ as the sum of a rank one projection onto the space spanned by $\boldsymbol{\varphi}_1$, plus a matrix that is the restriction of $\boldsymbol{A}$ to the subspace orthogonal to $\boldsymbol{\varphi}_1$. That is,

$$\boldsymbol{A} = z_1 \boldsymbol{\varphi}_1 \boldsymbol{\varphi}_1^{\mathsf{T}} + \boldsymbol{P}^{\perp}(\lambda \boldsymbol{v}\boldsymbol{v}^{\mathsf{T}} + \boldsymbol{W})\boldsymbol{P}^{\perp}, \tag{7.2}$$

where $\boldsymbol{P}^{\perp} = \boldsymbol{I} - \boldsymbol{\varphi}_1 \boldsymbol{\varphi}_1^{\mathsf{T}}$ is the projector onto the space orthogonal to $\boldsymbol{\varphi}_1$. The proof of Theorem 1 is based on an approximate representation of the conditional distribution of $\boldsymbol{A}$ given $(\boldsymbol{\varphi}_1, z_1)$. To this end, we define the matrix

$$\tilde{\boldsymbol{A}} = z_1 \boldsymbol{\varphi}_1 \boldsymbol{\varphi}_1^{\mathsf{T}} + \boldsymbol{P}^{\perp}(\lambda \boldsymbol{v}\boldsymbol{v}^{\mathsf{T}} + \tilde{\boldsymbol{W}})\boldsymbol{P}^{\perp}, \tag{7.3}$$

where $\tilde{\boldsymbol{W}} \sim \mathsf{GOE}(n)$ is independent of $\boldsymbol{W}$.

The proof is based on a key technical lemma (Lemma B.3 in the Supplementary Material [44]) which shows that for large enough n, the conditional distribution of $\boldsymbol{A}$ given $(\boldsymbol{\varphi}_1, z_1)$ is close in (in total variation distance) to that of $\tilde{\boldsymbol{A}}$ with high probability. Given $\tilde{\boldsymbol{A}}$, we consider a sequence of AMP iterates $(\tilde{\boldsymbol{x}}^t)_{t \geq 0}$ obtained by replacing $\boldsymbol{A}$ with $\tilde{\boldsymbol{A}}$ in equation (2.8). That is, we set

$$\tilde{\boldsymbol{x}}^0 = \sqrt{n}\,\mathrm{sign}(\langle \boldsymbol{x}_0, \boldsymbol{\varphi}_1 \rangle)\boldsymbol{\varphi}_1, \qquad \tilde{\boldsymbol{x}}^{t+1} = \tilde{\boldsymbol{A}} f_t(\tilde{\boldsymbol{x}}^t) - \mathsf{b}_t f_{t-1}(\tilde{\boldsymbol{x}}^{t-1}). \tag{7.4}$$

Theorem 1 is proved in three steps:

1. Using the conditional distribution lemma (Lemma B.3 in the Supplementary Material [44]), we show that for any PL(2) test function $\psi : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, almost surely

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \psi(x_i^t, x_{0,i}) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \psi(\tilde{x}_i^t, x_{0,i}), \tag{7.5}$$

whenever the limit on the right exists.

2. Step 1 allows us to establish Theorem 1 by analyzing the modified AMP iteration in equation (7.4). For the modified AMP, the initialization $\tilde{\boldsymbol{x}}^0$ is independent of $\tilde{\boldsymbol{W}}$. Consequently, adapting techniques from standard AMP analysis we show that the following holds almost surely for any PL(2) test function $\psi : \mathbb{R}^3 \to \mathbb{R}$:

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \psi(\tilde{x}_i^t, x_{0,i}, \sqrt{n}\varphi_{1,i}) = \mathbb{E}\{\psi(\alpha_t X_0 + \beta_t L + \tau_t G_0, X_0, L)\}. \tag{7.6}$$

Here, the random variables (X_0, L, G_0) are jointly distributed as follows: $X_0 \sim \nu_{X_0}$ and $G_0 \sim \mathsf{N}(0, 1)$ are independent, $L = \sqrt{1 - \lambda^{-2}} X_0 + \lambda^{-1} G_1$, where $G_1 \sim \mathsf{N}(0, 1)$ is independent of both X_0 and G_0. It is shown in Corollary C.3 that (almost surely) the empirical distribution of $(\boldsymbol{x}_0, \sqrt{n}\boldsymbol{\varphi}_1)$ converges in W_2 to the distribution of (X_0, L). The constants $(\alpha_t, \beta_t, \tau_t)$ in equation (7.6) are iteratively defined using a suitable state evolution recursion given in Appendix A.

3. The proof of Theorem 1 is completed by showing that for $t \geq 0$,

$$\mathbb{E}\{\psi(\alpha_t X_0 + \beta_t L + \tau_t G_0, X_0, L)\} = \mathbb{E}\{\psi(\mu_t X_0 + \sigma_t G, X_0)\}, \tag{7.7}$$

where $(\mu_t, \sigma_t)_{t\geq 0}$ are the state evolution parameters defined in the statement of Theorem 1.

Combining equations (7.5)–(7.7) yields the claim of Theorem 1. The detailed proof of this theorem is given in Appendix A in the Supplementary Material [44].

General case: For the general spiked model equation (1.1), the proof of the state evolution result (equation (6.11) of Theorem 5) is along similar lines. Here, the modified matrix $\tilde{\boldsymbol{A}}$ is defined as

$$\tilde{\boldsymbol{A}} \equiv \sum_{i\in\hat{S}} z_i \boldsymbol{\varphi}_i \boldsymbol{\varphi}_i^{\mathsf{T}} + \boldsymbol{P}^{\perp}\left(\sum_{i=1}^{k} \lambda_i \boldsymbol{v}_i \boldsymbol{v}_i^{\mathsf{T}} + \tilde{\boldsymbol{W}}\right)\boldsymbol{P}^{\perp}, \tag{7.8}$$

where $\boldsymbol{P}^{\perp}$ is the projector onto the orthogonal complement of the space spanned by $(\boldsymbol{\varphi}_i)_{i\in\hat{S}}$, and $\tilde{\boldsymbol{W}} \sim \mathsf{GOE}(n)$ is independent of $\boldsymbol{W}$. (Recall that $\hat{S}$ contains the indices i for which $|\lambda_i| > 1$.) Lemma B.3 shows that with high probability the conditional distributions of $\boldsymbol{A}$ and $\tilde{\boldsymbol{A}}$ are close in total variation distance. We then consider iterates $(\tilde{\boldsymbol{x}})_{t\geq 0}$ generated via the AMP iteration using $\tilde{\boldsymbol{A}}$:

$$\tilde{\boldsymbol{x}}^0 = \sqrt{n}[\boldsymbol{\varphi}_1|\cdots|\boldsymbol{\varphi}_{k_*}|\boldsymbol{0}|\cdots|\boldsymbol{0}], \tag{7.9}$$

$$\tilde{\boldsymbol{x}}^{t+1} = \tilde{\boldsymbol{A}} f_t(\tilde{\boldsymbol{x}}^t, \boldsymbol{y}) - f_{t-1}(\tilde{\boldsymbol{x}}^{t-1}, \boldsymbol{y})\mathsf{B}_t^{\mathsf{T}}. \tag{7.10}$$

Using Lemma B.3, we first show that once the state evolution result equation (6.11) holds for $\tilde{\boldsymbol{x}}^t$, it also holds for $\boldsymbol{x}^t$. The result for $\tilde{\boldsymbol{x}}^t$ is then shown in two steps, which are analogous to equations (7.6) and (7.7) for the rank one case.

Acknowledgements. We thank Leo Miolane for pointing out a gap in an earlier proof of Proposition 2.2. A. M. was partially supported by NSF Grants CCF-1714305 and NSF Grant IIS-1741162. R. V. was partially supported by a Marie Curie Career Integration Grant (Grant Agreement No. 631489) and a Turing Fellowship from the Alan Turing Institute (EPSRC grant reference EP/N510129/1).

SUPPLEMENTARY MATERIAL

Supplement (DOI: 10.1214/20-AOS1958SUPP; .pdf). The supplement [44] contains the proofs of all the results in the main text.

REFERENCES

[1] ABBE, E. (2018). Community detection and stochastic block models: Recent developments. *J. Mach. Learn. Res.* **18** Paper No. 177, 86. MR3827065

[2] BAI, Z. and SILVERSTEIN, J. W. (2010). *Spectral Analysis of Large Dimensional Random Matrices*, 2nd ed. *Springer Series in Statistics*. Springer, New York. MR2567175 https://doi.org/10.1007/978-1-4419-0661-8

[3] BAIK, J., BEN AROUS, G. and PÉCHÉ, S. (2005). Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Ann. Probab.* **33** 1643–1697. MR2165575 https://doi.org/10.1214/009117905000000233

[4] BAIK, J. and SILVERSTEIN, J. W. (2006). Eigenvalues of large sample covariance matrices of spiked population models. *J. Multivariate Anal.* **97** 1382–1408. MR2279680 https://doi.org/10.1016/j.jmva.2005.08.003

[5] BARBIER, J., DIA, M., MACRIS, N., KRZAKALA, F., LESIEUR, T. and ZDEBOROVÁ, L. (2016). Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. In *Advances in Neural Information Processing Systems* 424–432.

[6] BAYATI, M., LELARGE, M. and MONTANARI, A. (2015). Universality in polytope phase transitions and message passing algorithms. *Ann. Appl. Probab.* **25** 753–822. MR3313755 https://doi.org/10.1214/14-AAP1010

[7] BAYATI, M. and MONTANARI, A. (2011). The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Trans. Inf. Theory* **57** 764–785. MR2810285 https://doi.org/10.1109/TIT.2010.2094817

[8] BENAYCH-GEORGES, F. and NADAKUDITI, R. R. (2011). The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Adv. Math.* **227** 494–521. MR2782201 https://doi.org/10.1016/j.aim.2011.02.007

[9] BENAYCH-GEORGES, F. and NADAKUDITI, R. R. (2012). The singular values and vectors of low rank perturbations of large rectangular random matrices. *J. Multivariate Anal.* **111** 120–135. MR2944410 https://doi.org/10.1016/j.jmva.2012.04.019

[10] BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. MR1325392

[11] BERTHIER, R., MONTANARI, A. and NGUYEN, P.-M. (2020). State evolution for approximate message passing with non-separable functions. *Inf. Inference* **9** 33–79. MR4079177 https://doi.org/10.1093/imaiai/iay021

[12] BLEI, D. M., KUCUKELBIR, A. and MCAULIFFE, J. D. (2017). Variational inference: A review for statisticians. *J. Amer. Statist. Assoc.* **112** 859–877. MR3671776 https://doi.org/10.1080/01621459.2017.1285773

[13] BOLTHAUSEN, E. (2014). An iterative construction of solutions of the TAP equations for the Sherrington–Kirkpatrick model. *Comm. Math. Phys.* **325** 333–366. MR3147441 https://doi.org/10.1007/s00220-013-1862-3

[14] CAPITAINE, M., DONATI-MARTIN, C. and FÉRAL, D. (2009). The largest eigenvalues of finite rank deformation of large Wigner matrices: Convergence and nonuniversality of the fluctuations. *Ann. Probab.* **37** 1–47. MR2489158 https://doi.org/10.1214/08-AOP394

[15] CHEN, Y. and CANDÈS, E. J. (2018). The projected power method: An efficient algorithm for joint alignment from pairwise differences. *Comm. Pure Appl. Math.* **71** 1648–1714. MR3847751 https://doi.org/10.1002/cpa.21760

[16] DESHPANDE, Y., ABBE, E. and MONTANARI, A. (2017). Asymptotic mutual information for the balanced binary stochastic block model. *Inf. Inference* **6** 125–170. MR3671474 https://doi.org/10.1093/imaiai/iaw017

[17] DESHPANDE, Y. and MONTANARI, A. (2014). Information-theoretically optimal sparse PCA. In *IEEE International Symposium on Information Theory* (*ISIT*) 2197–2201.

[18] DESHPANDE, Y. and MONTANARI, A. (2015). Finding hidden cliques of size $\sqrt{N/e}$ in nearly linear time. *Found. Comput. Math.* **15** 1069–1128. MR3371378 https://doi.org/10.1007/s10208-014-9215-y

[19] DONOHO, D. L., JOHNSTONE, I. and MONTANARI, A. (2013). Accurate prediction of phase transitions in compressed sensing via a connection to minimax denoising. *IEEE Trans. Inf. Theory* **59** 3396–3433. MR3061255 https://doi.org/10.1109/TIT.2013.2239356

[20] DONOHO, D. L. and JOHNSTONE, I. M. (1994). Minimax risk over l_p-balls for l_q-error. *Probab. Theory Related Fields* **99** 277–303. MR1278886 https://doi.org/10.1007/BF01199026

[21] DONOHO, D. L. and JOHNSTONE, I. M. (1998). Minimax estimation via wavelet shrinkage. *Ann. Statist.* **26** 879–921. MR1635414 https://doi.org/10.1214/aos/1024691081

[22] DONOHO, D. L., MALEKI, A. and MONTANARI, A. (2009). Message passing algorithms for compressed sensing. *Proc. Natl. Acad. Sci. USA* **106** 18914–18919.

[23] EFRON, B. (2010). *Large-Scale Inference*: *Empirical Bayes Methods for Estimation*, *Testing*, *and Prediction*. *Institute of Mathematical Statistics* (*IMS*) *Monographs* **1**. Cambridge Univ. Press, Cambridge. MR2724758 https://doi.org/10.1017/CBO9780511761362

[24] FÉRAL, D. and PÉCHÉ, S. (2007). The largest eigenvalue of rank one deformation of large Wigner matrices. *Comm. Math. Phys.* **272** 185–228. MR2291807 https://doi.org/10.1007/s00220-007-0209-3

[25] FLETCHER, A. K. and RANGAN, S. (2018). Iterative reconstruction of rank-one matrices in noise. *Inf. Inference* **7** 531–562. MR3858334 https://doi.org/10.1093/imaiai/iax014

[26] GAMERMAN, D. and LOPES, H. F. (2006). *Markov Chain Monte Carlo*: *Stochastic Simulation for Bayesian Inference*, 2nd ed. *Texts in Statistical Science Series*. CRC Press/CRC, Boca Raton, FL. MR2260716

[27] GUO, D., SHAMAI, S. and VERDÚ, S. (2005). Mutual information and minimum mean-square error in Gaussian channels. *IEEE Trans. Inf. Theory* **51** 1261–1282. MR2241490 https://doi.org/10.1109/TIT.2005.844072

[28] HOYLE, D. C. and RATTRAY, M. (2004). Principal-component-analysis eigenvalue spectra from data with symmetry-breaking structure. *Phys. Rev. E* **69** 026124.

[29] JAVANMARD, A. and MONTANARI, A. (2013). State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Inf. Inference* **2** 115–144. MR3311445 https://doi.org/10.1093/imaiai/iat004

[30] JAVANMARD, A., MONTANARI, A. and RICCI-TERSENGHI, F. (2016). Phase transitions in semidefinite relaxations. *Proc. Natl. Acad. Sci. USA* **113** E2218–E2223. MR3494080 https://doi.org/10.1073/pnas.1523097113

[31] JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.* **29** 295–327. MR1863961 https://doi.org/10.1214/aos/1009210544

[32] JOHNSTONE, I. M. (2007). High dimensional statistical inference and random matrices. In *International Congress of Mathematicians. Vol. I* 307–333. Eur. Math. Soc., Zürich. MR2334195 https://doi.org/10.4171/022-1/13

[33] JOHNSTONE, I. M. and LU, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *J. Amer. Statist. Assoc.* **104** 682–693. MR2751448 https://doi.org/10.1198/jasa.2009.0121

[34] JOURNÉE, M., NESTEROV, Y., RICHTÁRIK, P. and SEPULCHRE, R. (2010). Generalized power method for sparse principal component analysis. *J. Mach. Learn. Res.* **11** 517–553. MR2600619

[35] KABASHIMA, Y., KRZAKALA, F., MÉZARD, M., SAKATA, A. and ZDEBOROVÁ, L. (2016). Phase transitions and sample complexity in Bayes-optimal matrix factorization. *IEEE Trans. Inf. Theory* **62** 4228–4265. MR3515748 https://doi.org/10.1109/TIT.2016.2556702

[36] KNOWLES, A. and YIN, J. (2013). The isotropic semicircle law and deformation of Wigner matrices. *Comm. Pure Appl. Math.* **66** 1663–1750. MR3103909 https://doi.org/10.1002/cpa.21450

[37] KRZAKALA, F., XU, J. and ZDEBOROVÁ, L. (2016). Mutual information in rank-one matrix estimation. In *IEEE Information Theory Workshop* (*ITW*) 71–75.

[38] LEE, D. D. and SEUNG, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature* **401** 788–791. https://doi.org/10.1038/44565

[39] LELARGE, M. and MIOLANE, L. (2019). Fundamental limits of symmetric low-rank matrix estimation. *Probab. Theory Related Fields* **173** 859–929. MR3936148 https://doi.org/10.1007/s00440-018-0845-x

[40] LESIEUR, T., KRZAKALA, F. and ZDEBOROVÁ, L. (2017). Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications. *J. Stat. Mech. Theory Exp.* **7** 073403, 86. MR3683819 https://doi.org/10.1088/1742-5468/aa7284

[41] MA, Z. (2013). Sparse principal component analysis and iterative thresholding. *Ann. Statist.* **41** 772–801. MR3099121 https://doi.org/10.1214/13-AOS1097

[42] MIOLANE, L. (2017). Fundamental limits of low-rank matrix estimation. arXiv:1702.00473.

[43] MONTANARI, A. and RICHARD, E. (2016). Non-negative principal component analysis: Message passing algorithms and sharp asymptotics. *IEEE Trans. Inf. Theory* **62** 1458–1484. MR3472260 https://doi.org/10.1109/TIT.2015.2457942

[44] MONTANARI, A. and VENKATARAMANAN, R. (2021). Supplement to "Estimation of low-rank matrices via approximate message passing." https://doi.org/10.1214/20-AOS1958SUPP

[45] MOORE, C. (2017). The computer science and physics of community detection: Landscapes, phase transitions, and hardness. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS* **121** 26–61. MR3699594

[46] MOSSEL, E. and XU, J. (2016). Density evolution in the degree-correlated stochastic block model. In *Conference on Learning Theory* 1319–1356.

[47] PARKER, J. T., SCHNITER, P. and CEVHER, V. (2014). Bilinear generalized approximate message passing—Part I: Derivation. *IEEE Trans. Signal Process.* **62** 5839–5853. MR3281527 https://doi.org/10.1109/TSP.2014.2357776

[48] PAUL, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statist. Sinica* **17** 1617–1642. MR2399865

[49] STOREY, J. D. (2002). A direct approach to false discovery rates. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 479–498. MR1924302 https://doi.org/10.1111/1467-9868.00346

[50] VILA, J., SCHNITER, P. and MEOLA, J. (2015). Hyperspectral unmixing via turbo bilinear approximate message passing. *IEEE Trans. Comput. Imaging* **1** 143–158. MR3423296 https://doi.org/10.1109/TCI.2015.2465161

[51] VILLANI, C. (2009). *Optimal Transport*: *Old and New*. *Grundlehren der Mathematischen Wissenschaften* [*Fundamental Principles of Mathematical Sciences*] **338**. Springer, Berlin. MR2459454 https://doi.org/10.1007/978-3-540-71050-9

[52] YUAN, X.-T. and ZHANG, T. (2013). Truncated power method for sparse eigenvalue problems. *J. Mach. Learn. Res.* **14** 899–925. MR3063614

[53] ZOU, H., HASTIE, T. and TIBSHIRANI, R. (2006). Sparse principal component analysis. *J. Comput. Graph. Statist.* **15** 265–286. MR2252527 https://doi.org/10.1198/106186006X113430