

Information-theoretically Optimal Sparse PCA

Yash Deshpande

Department of Electrical Engineering
Stanford, CA.

Andrea Montanari

Departments of Electrical Engineering and Statistics
Stanford, CA.

Abstract—Sparse Principal Component Analysis (PCA) is a dimensionality reduction technique wherein one seeks a low-rank representation of a data matrix with additional sparsity constraints on the obtained representation. We consider two probabilistic formulations of sparse PCA: a spiked Wigner and spiked Wishart (or spiked covariance) model. We analyze an Approximate Message Passing (AMP) algorithm to estimate the underlying signal and show, in the high dimensional limit, that the AMP estimates are information-theoretically optimal. As an immediate corollary, our results demonstrate that the posterior expectation of the underlying signal, which is often intractable to compute, can be obtained using a polynomial-time scheme. Our results also effectively provide a *single-letter characterization* of the sparse PCA problem.

I. INTRODUCTION

Suppose we are given data $Y_\lambda \in \mathbb{R}^{n \times n}$ distributed according to the following *spiked Wigner* model:

$$Y_\lambda = \sqrt{\frac{\lambda}{n}} x x^\top + Z. \quad (1)$$

Here $x \in \mathbb{R}^n$, and each coordinate x_i is an independent Bernoulli random variable with probability ε , denoted by $x_i \sim \text{Ber}(\varepsilon)$. $Z \in \mathbb{R}^{n \times n}$ is a symmetric matrix where $(Z_{ij})_{i \leq j}$ are i.i.d $\text{N}(0, 1)$ variables, independent of x . Analogously, consider the following *spiked Wishart* model:

$$Y_\lambda = \sqrt{\frac{\lambda}{n}} u v^\top + Z. \quad (2)$$

Here $u \in \mathbb{R}^m$, with i.i.d coordinates $u_i \sim \text{N}(0, 1)$ and $v \in \mathbb{R}^n$ with i.i.d Bernoulli coordinates $v_j \sim \text{Ber}(\varepsilon)$. Further, $Z \in \mathbb{R}^{m \times n}$ is a matrix with $Z_{ij} \sim \text{N}(0, 1)$ i.i.d. random variables.

In either case, our data consists of a *sparse, rank-one* matrix observed through Gaussian noise. We let X denote the clean, underlying signal ($x x^\top$ or $u v^\top$ for the spiked Wigner or Wishart model respectively). Our task is to estimate the signal X from the data Y_λ in the high dimension asymptotic where $n \rightarrow \infty, m \rightarrow \infty$ with $m/n \rightarrow \alpha \in (0, \infty)$. This paper focuses on estimation in the sense of the mean squared error, defined for an estimator $\hat{X}(Y_\lambda)$ as:

$$\text{mse}(\hat{X}, \lambda) \equiv \frac{1}{n^2} \mathbb{E} \left\{ \|\hat{X} - X\|_F^2 \right\}. \quad (3)$$

It is well-known [1] that the mean squared error is minimized by the estimator $\hat{X} = \mathbb{E}\{X|Y_\lambda\}$, i.e. the conditional expectation of the signal given the observations. Consequently,

the minimum mean squared error (MMSE) is given by:

$$\text{M-mmse}(\lambda, n) \equiv \frac{1}{n^2} \mathbb{E} \left\{ \|X - \mathbb{E}\{X|Y_\lambda\}\|_F^2 \right\}. \quad (4)$$

In this paper, we analyze an iterative scheme called *approximate message passing* (AMP) to estimate the clean signal X . The machinery of approximate message passing reduces the high-dimensional *matrix* problem in models (1), (2) to the following simpler *scalar* denoising problem:

$$Y_\lambda = \sqrt{\lambda} X_0 + N, \quad (5)$$

where $X_0 \sim \text{Ber}(\varepsilon)$ and $N \sim \text{N}(0, 1)$ are independent. The scalar MMSE [2] in estimating X_0 from Y_λ is given by:

$$\text{S-mmse}(X_0, \lambda) = \mathbb{E} \left\{ (X_0 - \mathbb{E}\{X_0|Y_\lambda\})^2 \right\}.$$

Our main results, characterize the optimal mean squared error $\text{M-mmse}(\lambda, n)$ in the large n asymptotic, when $\varepsilon > \varepsilon_c \approx 0.05$, and establish that AMP achieves this fundamental limit. For the spiked Wigner model we prove the following.

Theorem 1. *There exists an $\varepsilon_c \in (0, 1)$ such that for all $\varepsilon > \varepsilon_c$, and every $\lambda \geq 0$ the squared error of AMP iterates \hat{X}^t satisfies the following:*

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \text{mse}(\hat{X}^t, \lambda) = \lim_{n \rightarrow \infty} \text{M-mmse}(\lambda, n).$$

Further, the limit on the RHS above satisfies, for every $\lambda > 0$:

$$\lim_{n \rightarrow \infty} \text{M-mmse}(\lambda, n) = \varepsilon^2 - y_*^2 / \lambda^2,$$

where $y_* = y_*(\lambda)$ solves $y_* = \lambda(\varepsilon - \text{S-mmse}(X_0, y_*))$.

Some remarks are in order:

Remark I.1. The combination of Theorem 1 and Eq.(10) effectively yield a *single-letter characterization* of Model (1), connecting the limiting matrix MMSE with the MMSE of a calibrated scalar denoising problem $\text{S-mmse}(X_0, y_*(\lambda))$.

Remark I.2. It is straightforward to establish that $\varepsilon_c < 1$. However, numerically we obtain that $\varepsilon_c \approx 0.05$. Thus, for most values of ε , our results completely characterize the spiked Wigner model.

Background and Motivation

Probabilistic models similar to Eqs.(1), (2) have been the focus of much recent work in random matrix theory [3]–[7]. The focus in this literature is to analyze the limiting distribution of the eigenvalues of the matrix Y_λ / \sqrt{n} and, in particular, identifying regimes in which this distribution differs

from that of the pure noise matrix Z . The typical picture that emerges from this line of work is that a *phase transition* occurs at a well-defined critical signal-to-noise ratio $\lambda_c = \lambda_c(\alpha, \varepsilon)$:

Above the threshold $\lambda > \lambda_c$: there exists an outlier eigenvalue and the principal eigenvector corresponding to this outlier has a positive correlation with the signal. For instance, in the spiked Wigner case, letting $\hat{x}_1(Y_\lambda)$ denote the normalized principal eigenvector of Y_λ we obtain $\langle \hat{x}_1(Y_\lambda), x \rangle / \sqrt{n\varepsilon} \geq \delta(\varepsilon) > 0$ asymptotically.

Below the threshold $\lambda < \lambda_c$: the spectral distribution of the observation Y_λ is indistinguishable from that of the pure noise Z . Furthermore, the principal eigenvector is asymptotically *orthogonal* to the signal factors. For the spiked Wigner case, this implies that $\langle \hat{x}_1(Y_\lambda), x \rangle / \sqrt{n\varepsilon} \rightarrow 0$ asymptotically.

This phase transition phenomenon has been demonstrated under considerably fewer assumptions than we make in Eqs. (1), (2). We refer the interested reader to [4], [8] and the references therein for further details.

It is clear from these results that vanilla PCA, which involves using the principle eigenvector is ineffective in estimating the underlying clean signal X when $\lambda < \lambda_c$. Indeed PCA only makes use of the fact that the underlying signal is *low-rank*, or in fact rank-one in our case. Since we make additional sparsity assumptions in our models (1), (2) it is natural to ask if this can be leveraged when we have a small signal-to-noise ratio λ . In the last decade, a considerable amount of work in the statistics community has studied this problem. Our spiked Wishart model Eq. (2) is a special case of the *spiked covariance model* in statistics, first introduced by Johnstone and Lu [9], [10]. Johnstone and Lu proposed a simple diagonal thresholding scheme that estimates the support of v using the largest diagonal entries of the Gram matrix $Y_\lambda^\top Y_\lambda$. An M-estimator for the underlying factors was proposed by [11]. A number of other practical algorithms [12]–[14] have also been proposed to outperform diagonal thresholding.

Some recent work [15], [16] has focused on the *support recovery* guarantees for such algorithms, or estimating consistently the positions of non-zeros in v . Let $k = n\varepsilon$ denote the expected size of the support of v . Amini and Wainwright [17] proved that unless $k \leq cm/\log n$, *no algorithm* would be able to consistently estimate the support of v due to information-theoretic obstructions. They further demonstrate that a (computationally intractable) algorithm that searches through all possible k -sized subsets of rows of the data matrix can recover the support provided $k \leq c'm/\log n$.

Since we consider $\varepsilon = \Theta(1)$ and $m = \Theta(n)$, in our case $k = \Theta(m)$ and consequently, estimating the support correctly is impossible. It is for this reason that we instead focus on another natural figure-of-merit: the mean squared error, defined in Eq. (3) above. Somewhat surprisingly, we are able to prove (for a regime $\varepsilon_c < \varepsilon \leq 1$) that a *computationally efficient* algorithm asymptotically achieves the information-theoretically optimal mean squared error *for any signal-to-noise ratio* λ .

Other related work

Rangan et al. [18] considered a model similar to Eq. (2) with general structural assumptions on the factors u and v . They proposed an approximate message passing algorithm analogous to the one we analyze and characterize its high-dimensional behavior. Based on non-rigorous but powerful tools from statistical physics, they conjecture that AMP asymptotically achieves the (optimal) performance of the joint MMSE estimator of u and v . In the restricted setting of sparse PCA, we rigorously confirm this conjecture, and validate the statistical physics arguments.

A model similar to Eq. (1) was considered by [19], motivated by the “planted clique” problem in theoretical computer science. The sparsity regime of interest in this work was $k = O(\sqrt{n})$, with a focus on recovering the “clique”, analogous to support recovery in the spiked covariance model.

Organization

The paper is organized as follows. In Section II we give details of the AMP algorithm and formally state our results. For brevity, we only provide the proof of one of our main results in Section III.

II. ALGORITHM AND MAIN RESULTS

In the interest of exposition, we restrict ourselves to the spiked Wigner model (1) and defer the discussion of the Wishart model (2) to Section II-D.

A. Approximate Message Passing

Approximate message passing (AMP) is a low complexity iterative algorithm that produces iterates $x^t, \hat{x}^t \in \mathbb{R}^n$. For a data matrix A we define for $t \geq 0$:

$$x^{t+1} = Ax^t - b_t \hat{x}^{t-1} \quad (6)$$

$$\hat{x}^t = f_t(x^t). \quad (7)$$

Here $f_t : \mathbb{R} \rightarrow \mathbb{R}$ are scalar functions and $\{b_t\}_{t \geq 0}$ is a sequence of scalars. Here and below, for a scalar function f , we define its extension to \mathbb{R}^n by applying it component-wise, i.e. $f : \mathbb{R}^n \rightarrow \mathbb{R}^n, v \mapsto f(v) = (f(v_1), f(v_2), \dots, f(v_n))^\top$. We further define the matrix estimate $\hat{X}^t \equiv \hat{x}^t (\hat{x}^t)^\top$. For the complete description of the algorithm, we refer the reader to Algorithm 1 below, which provides prescriptions for the functions f_t and the scalars b_t .

B. State evolution

The key property of approximate message passing is that it admits an *asymptotically exact* characterization in the high-dimensional limit where $n \rightarrow \infty$. The iterates x_t^t converge as $n \rightarrow \infty$ to Gaussian random variables with a prescribed mean and variance. These prescribed mean and variance parameters evolve according to *deterministic* recursions, jointly termed “state evolution”. We define for $t \geq 0$:

$$\mu_{t+1} = \sqrt{\lambda} \mathbb{E}\{X_0 f_t(\mu_t X_0 + \sqrt{\tau_t} Z)\} \quad (8)$$

$$\tau_{t+1} = \mathbb{E}\{f_t(\mu_t X_0 + \sqrt{\tau_t} Z)^2\}, \quad (9)$$

Algorithm 1 Symmetric Bayes-optimal AMP

Input: Data Y_λ as in Eq. (1)

Define $A = Y_\lambda/\sqrt{n}$, and $\hat{x}^0, \hat{x}^{-1} = 0$. For $t \geq 0$ compute

$$\begin{aligned}x^{t+1} &= A\hat{x}^t - \mathbf{b}_t\hat{x}^{t-1} \\ \hat{x}^t &= f_t(x^t) \\ \hat{X}^t &= \hat{x}^t(\hat{x}^t)^\top,\end{aligned}$$

where $f_t(y) : \mathbb{R} \rightarrow \mathbb{R}$ is recursively defined:

$$f_t(y) = \mathbb{E}\{X_0|\mu_t X_0 + \sqrt{\tau_t}Z = y\}.$$

Here $X_0 \sim \text{Ber}(\varepsilon)$ and $Z \sim \text{N}(0, 1)$ are independent, and μ_t, τ_t are defined as in Eqs. (8), (9). Also the scalars \mathbf{b}_t are computed as:

$$\mathbf{b}_t = \frac{1}{n} \sum_{i=1}^n f'_t(x_i^t).$$

where $X_0 \sim \text{Ber}(\varepsilon)$ and $Z \sim \text{N}(0, 1)$ are independent. The recursion is initialized with $\mu_0 = \tau_0 = 0$.

The state evolution recursions succinctly describe the iterates arising in AMP. Formally, we have, for any continuous function $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}$ the following is true wherever the expectation on the right is defined:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i \in n} \psi(x_i, x_i^t) = \mathbb{E}\{\psi(X_0, \mu_t X_0 + \sqrt{\tau_t}Z)\} \text{ a.s.}$$

where μ_t, τ_t are defined by Eqs. (8), (9). This allows us to track the squared error of the AMP estimator accurately, in the high-dimensional limit, and establish its optimality.

Although we define AMP and the corresponding state evolution for general scalar functions f_t , our prescription Algorithm 1 uses specific choices for f_t . In the spiked Wigner case, we choose $f_t(y) = \mathbb{E}\{X_0|\mu_t X_0 + \sqrt{\tau_t}Z = y\}$, the posterior expectation of X_0 , with observation corrupted by Gaussian noise and SNR μ_t^2/τ_t . To stress this fact, we will refer to our algorithms as Bayes-optimal AMP.

C. Main Result

We first define the following regime for ε :

Definition II.1. Let $\varepsilon_* \in (0, 1)$ be the smallest positive real number such that for every $\varepsilon > \varepsilon_*$ the following is true. For every $\lambda > 0$, the equation below has only one solution in $[0, \infty)$:

$$\lambda^{-1}y = \varepsilon - \text{S-mmse}(X_0, y). \quad (10)$$

Here $X_0 \sim \text{Ber}(\varepsilon)$.

With a slight abuse of notation, we denote by $\text{M-mmse}(\lambda)$ the quantity $\lim_{n \rightarrow \infty} \text{M-mmse}(\lambda, n)$, assuming it exists. Also we define the squared error of AMP at iteration t as:

$$\text{MSE}_{\text{AMP}}(\lambda, t) = \frac{1}{n^2} \|\hat{X}^t - X\|_F^2.$$

Notice that $\text{MSE}_{\text{AMP}}(\lambda, t)$ is a random variable that depends on the realization Y_λ . Our first main result strengthens Theorem 1 for the spiked Wigner case:

Theorem 2. Under Model (1) we have $\text{M-mmse}(\lambda) = \lim_{n \rightarrow \infty} \text{M-mmse}(\lambda, n)$ exists for every $\lambda \geq 0$. This limit satisfies, when $\varepsilon > \varepsilon_*$:

$$\text{M-mmse}(\lambda) = \varepsilon^2 - \frac{y_*(\lambda)^2}{\lambda^2}, \quad (11)$$

where $y_*(\lambda)$ is the unique solution to Eq. (10) above. Further, the symmetric Bayes-optimal AMP algorithm 1 satisfies the following almost surely:

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \text{MSE}_{\text{AMP}}(\lambda, t) = \text{M-mmse}(\lambda). \quad (12)$$

Although this result is asymptotic in nature, simulations show that the predictions are accurate on problems of dimension a few thousands (see Figure 1).

D. The spiked Wishart model

An asymmetric version of Algorithm 1 can also be written. It involves iterates $u^t, \hat{u}^t \in \mathbb{R}^m$, $v^t, \hat{v}^t \in \mathbb{R}^n$. Define $A = Y_\lambda/\sqrt{m}$, and $\hat{u}^0, \hat{u}^{-1} = 0$. For $t \geq 0$ compute

$$\begin{aligned}u^{t+1} &= A\hat{v}^t - \mathbf{b}_t\hat{u}^t \\ \hat{v}^t &= f_t(v^t) \\ v^t &= A^\top\hat{u}^t - \mathbf{d}_t\hat{v}^{t-1} \\ \hat{u}^t &= g_t(u^t) \\ \hat{X}^t &= \hat{u}^t(\hat{v}^t)^\top,\end{aligned}$$

The following is the analogue of Definition II.1 for the asymmetric Wishart model:

Definition II.2. Let $\tilde{\varepsilon}_*$ be the smallest positive real number such that for every $\varepsilon > \tilde{\varepsilon}_*$ the following is true. For every $\lambda > 0$ equation below has only one solution in $[0, \infty)$:

$$\lambda^{-1}y = \varepsilon - \text{S-mmse}(V, y/(1+y)). \quad (13)$$

Here $V \sim \text{Ber}(\varepsilon)$.

Our second result is for the spiked Wishart model (2):

Theorem 3. The limit $\text{M-mmse}(\lambda) \equiv \lim_{n \rightarrow \infty} \text{M-mmse}(\lambda, n)$ exists for every $\lambda \geq 0$ and, for $\varepsilon > \tilde{\varepsilon}_*$, is given by:

$$\text{M-mmse}(\lambda) = \varepsilon - \frac{y_*(\lambda)^2}{\lambda(1+y_*(\lambda))}, \quad (14)$$

where $y_*(\lambda)$ is the unique solution to Eq. (13). Further, asymmetric Bayes-optimal AMP satisfies the following limits almost surely:

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \text{MSE}_{\text{AMP}}(\lambda, t) = \text{M-mmse}(\lambda). \quad (15)$$

III. PROOF OF THEOREM 2

Owing to space constraints, we restrict ourselves to proving Theorem 2 in this paper. The proof of Theorem 3 follows similar ideas and will be provided in the full version of the




Fig. 1. The solid curves $M\text{-mmse}(\lambda)$ above are computed analytically using Theorem 2. The crosses mark median MSE incurred by AMP in 100 Monte Carlo runs with $n = 2000$ for the spiked Wigner model (1).

present paper. Theorem 2 follows almost immediately from the following two propositions.

Proposition III.1. *Consider the model Eq. (1) with $\varepsilon > \varepsilon_*$, and the approximate message passing orbit obtained by using the recursively defined scalar functions for $t \geq 0$:*

$$f_t(y) = \mathbb{E}\{X_0 | \mu_t X_0 + \sqrt{\tau_t} Z = y\}.$$

Here $X_0 \sim \text{Ber}(\varepsilon)$ and $Z \sim \mathcal{N}(0, 1)$ are independent. Further $\mu_0 = \tau_0 = 0$ and $(\mu_t, \tau_t)_{t \geq 1}$ are defined using the state evolution recursions (8), (9). Then defining $\text{mse}_{\text{AMP}}(\lambda) \equiv \varepsilon^2 - y_*(\lambda)^2/\lambda^2$, the RHS of Eq. (11), the following is true:

$$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \text{MSE}_{\text{AMP}}(\lambda, t) = \text{mse}_{\text{AMP}}(\lambda) \quad (16)$$

$$\int_0^\infty \text{mse}_{\text{AMP}}(\lambda) d\lambda = 4h(\varepsilon). \quad (17)$$

The first limit holds almost surely and in \mathcal{L}_1 and $h(\varepsilon)$ is the binary entropy function $h(\varepsilon) = -\varepsilon \log \varepsilon - (1 - \varepsilon) \log(1 - \varepsilon)$.

Proposition III.2. *For every $\lambda \geq 0$, $M\text{-mmse}(\lambda) = \lim_{n \rightarrow \infty} M\text{-mmse}(\lambda, n)$ exists. Further:*

$$\int_0^\infty M\text{-mmse}(\lambda) d\lambda \geq 4h(\varepsilon), \quad (18)$$

where $h(\varepsilon)$ is defined in Proposition III.1.

The above propositions are proved in Subsections III.1 and III.2 respectively. We first use these to establish Theorem 2. Since the posterior expectation minimizes the mean squared error, we have that $M\text{-mmse}(\lambda, n) \leq \mathbb{E}\{\text{MSE}_{\text{AMP}}(\lambda, t)\}$. Taking the limits $n \rightarrow \infty, t \rightarrow \infty$ in that order, and employing the first claim of Proposition III.1 we have that:

$$M\text{-mmse}(\lambda) \leq \text{mse}_{\text{AMP}}(\lambda).$$

This implies that:

$$\begin{aligned} 4h(\varepsilon) &\leq \int_0^\infty M\text{-mmse}(\lambda) d\lambda \\ &\leq \int_0^\infty \text{mse}_{\text{AMP}}(\lambda) d\lambda = 4h(\varepsilon), \end{aligned}$$

where in the first inequality and the last equality we use Propositions III.1, III.2. This implies that $\text{mse}_{\text{AMP}}(\lambda) = M\text{-mmse}(\lambda)$ for Lebesgue-a.e. λ . Further, as $M\text{-mmse}(\lambda)$ is the pointwise limit of monotone nonincreasing (in λ) functions $M\text{-mmse}(\lambda, n)$ [20], it is monotone nonincreasing, which yields the claim for all $\lambda \in [0, \infty)$. \square

A. Proof of Proposition III.1

Note that:

$$\begin{aligned} \text{MSE}_{\text{AMP}}(\lambda, t) &= \frac{1}{n^2} \|\widehat{X}^t - X\|_F^2 \\ &= \frac{1}{n^2} (\|\widehat{x}^t\|^4 + \|x\|^4 - 2\langle \widehat{x}^t, x \rangle^2). \end{aligned}$$

By the strong law of large numbers, $\|x\|^4/n^2 \rightarrow \varepsilon^2$ almost surely, and in \mathcal{L}_1 . It is not hard to prove that the functions $f_t(y)$ are $\sqrt{\lambda}$ -Lipschitz continuous. Hence, it is a direct consequence of Theorem 1 of [21] that the following limits hold almost surely and in \mathcal{L}_1 :

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n^2} \|\widehat{x}^t\|^4 &= (\mathbb{E}\{f_t(\mu_t X_0 + \sqrt{\tau_t} Z)^2\})^2 \\ \lim_{n \rightarrow \infty} \frac{1}{n^2} \langle \widehat{x}^t, x \rangle^2 &= (\mathbb{E}\{X_0 f_t(\mu_t X_0 + \sqrt{\tau_t} Z)\})^2. \end{aligned}$$

Further, our choice $f_t(y) = \mathbb{E}\{X_0 | \mu_t X_0 + \sqrt{\tau_t} Z = y\}$ yields, by use of the tower property of conditional expectation:

$$\begin{aligned} \mathbb{E}\{X_0 f_t(\mu_t X_0 + \sqrt{\tau_t} Z)\} &= \mathbb{E}\{f_t(\mu_t X_0 + \sqrt{\tau_t} Z)^2\} \\ &= \tau_{t+1}. \end{aligned}$$

It follows that $\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \text{MSE}_{\text{AMP}}(\lambda, t) = \varepsilon^2 - \tau_*^2$ almost surely and in \mathcal{L}_1 where $\tau_* = \tau_*(\lambda)$ denotes the smallest non-negative fixed point of the equation:

$$\tau = \mathbb{E}\left\{\mathbb{E}\{X_0 | \sqrt{\lambda\tau} X_0 + \sqrt{\tau} Z\}^2\right\}. \quad (19)$$

Since the right hand side equals $\varepsilon(1 - \varepsilon)$ at $\tau = 0$ and ε at $\tau = \infty$, at least one fixed point must exist. Hence $\tau_*(\lambda)$ is well defined. Now, note that

$$\begin{aligned} \mathbb{E}\left\{\mathbb{E}\{X_0 | \sqrt{\lambda\tau} X_0 + \sqrt{\tau} Z\}^2\right\} &= \mathbb{E}\left\{\mathbb{E}\{X_0 | \sqrt{\lambda\tau} X_0 + Z\}^2\right\} \\ &= \mathbb{E}\{X_0^2\} - \text{S-mmse}(X_0, \lambda\tau) \\ &= \varepsilon - \text{S-mmse}(X_0, \lambda\tau). \end{aligned}$$

Thus τ_* is a fixed point of Eq. (19) iff $\lambda\tau_* = y_*$ is a fixed point of Eq. (10). It follows from our definition of ε_* that when $\varepsilon > \varepsilon_*$, $\tau_*(\lambda)$ is the unique non-negative fixed point of Eq. (19) and Claim (16) follows. To complete the proof of the proposition, it only remains to show claim (17), for which we have the following

Lemma III.3. *Let $\tau_*(\lambda)$ denote the unique non-negative fixed*

point of Eq. (19). Then

$$\int_0^\infty (\varepsilon^2 - \tau_*(\lambda))^2 d\lambda = 4h(\varepsilon).$$

Proof. Define the function:

$$\begin{aligned} \phi(\lambda, m) &= \frac{m^2}{4} + \frac{\varepsilon^2 \lambda^2}{4} \\ &\quad - \mathbb{E} \log(1 - \varepsilon + \varepsilon \exp\{W(m, \lambda, X_0, Z)\}). \end{aligned}$$

where $W(m, \lambda, x, z) = m\sqrt{\lambda}x - m\sqrt{\lambda}/2 + \lambda^{1/4}m^{1/2}z$. Here $X_0 \sim \text{Ber}(\varepsilon)$, $Z \sim \text{N}(0, 1)$ and are independent. Letting $m_* = \tau_*\sqrt{\lambda}$, it is not hard to show that:

$$\begin{aligned} \left. \frac{\partial \phi}{\partial m} \right|_{m=m_*} &= 0, \\ \left. \frac{\partial \phi}{\partial \lambda} \right|_{m=m_*} &= \frac{1}{4} \left(\varepsilon^2 - \frac{m_*^2}{\lambda} \right) \end{aligned}$$

It follows from the fundamental theorem of calculus that

$$\int_0^\infty (\varepsilon^2 - \tau_*^2) d\lambda = \phi(\lambda, m_*(\lambda))|_0^\infty.$$

It is easy to see that $\phi(0, m_*(0)) = 0$. Further, using the fact that $\tau_*(\lambda) \leq 1$ (as the right hand side of Eq. (19) is bounded by 1), we have that $m_*(\lambda) = O(\sqrt{\lambda})$. Using this, it is not hard to check that $\phi(\lambda, m_*(\lambda)) \rightarrow h(\varepsilon)$ as $\lambda \rightarrow \infty$. This concludes the proof of the lemma. \square

B. Proof of Proposition III.2

We first prove that $\lim_{n \rightarrow \infty} \text{M-mmse}(\lambda, n)$ exists for every $\lambda \geq 0$. Define for $i, j \in [n]$

$$m_{ij}(\lambda, n) \equiv \mathbb{E} \left\{ (X_{ij} - \mathbb{E}\{X_{ij}|Y_\lambda\})^2 \right\}.$$

By the fact that the distribution of (X, Y_λ) is invariant under (identical) row and column permutations, $m_{ij}(\lambda, n) = m_{12}(\lambda, n)$ for every i, j distinct. Consequently:

$$|\text{M-mmse}(\lambda, n) - m_{12}(\lambda, n)| = \frac{1}{n} m_{11}(n, \lambda) \leq \frac{\text{Var}(X_{11})}{n}.$$

Since $\text{Var}(X_{11}) = \varepsilon(1 - \varepsilon) < \infty$ it suffices to prove that $\lim_{n \rightarrow \infty} m_{12}(\lambda, n)$ exists for every $\lambda \geq 0$. To this end, let Y_λ^{n-1} denote the first principal $(n-1) \times (n-1)$ submatrix of Y_λ . Clearly:

$$\begin{aligned} m_{12}(\lambda, n) &\leq \mathbb{E} \left\{ (X_{12} - \mathbb{E}\{X_{12}|Y_\lambda^{n-1}\})^2 \right\} \\ &= m_{12}((n-1)\lambda/n, n-1) \\ &\leq m_{12}(\lambda, n-1), \end{aligned}$$

where the equality follows from the model Eq. (1) and the second inequality from monotonicity of the minimum mean square error in λ [20]. Consequently, for every $\lambda \geq 0$, $m_{12}(\lambda, n)$ is a monotone, bounded sequence and has a limit.

In order to prove the claim 18, we first note that for any finite n the following holds applying the I-MMSE identity of

[2] to the upper triangular portion of X :

$$I(X; Y_\Lambda) = \frac{1}{2n} \int_0^\Lambda \left(\frac{n(n-1)}{2} m_{12}(\lambda, n) + n m_{11}(\lambda, n) \right) d\lambda.$$

For $\Lambda = \infty$ we have that $I(X; Y_\infty) = H(X) - H(X|Y_\infty) = H(X)$ and $H(X) = H(x) = nh(\varepsilon)$. Dividing by n on either side:

$$h(\varepsilon) = \frac{1}{2n^2} \int_0^\infty \left(\frac{n(n-1)}{2} m_{12}(\lambda, n) + n m_{11}(\lambda, n) \right) d\lambda.$$

An application of Fatou's lemma then yields the result. \square

REFERENCES

- [1] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
- [2] D. Guo, S. Shamai, and S. Verdú, "Mutual information and minimum mean-square error in gaussian channels," *Information Theory, IEEE Transactions on*, vol. 51, no. 4, pp. 1261–1282, 2005.
- [3] Z. Füredi and J. Komlós, "The eigenvalues of random symmetric matrices," *Combinatorica*, vol. 1, no. 3, pp. 233–241, 1981.
- [4] A. Knowles and J. Yin, "The isotropic semicircle law and deformation of wigner matrices," *Communications on Pure and Applied Mathematics*, 2013.
- [5] F. Benaych-Georges and R. R. Nadakuditi, "The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices," *Advances in Mathematics*, vol. 227, no. 1, pp. 494–521, 2011.
- [6] J. Baik, G. Ben Arous, and S. Péché, "Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices," *Annals of Probability*, pp. 1643–1697, 2005.
- [7] J. Baik and J. W. Silverstein, "Eigenvalues of large sample covariance matrices of spiked population models," *Journal of Multivariate Analysis*, vol. 97, no. 6, pp. 1382–1408, 2006.
- [8] A. Pizzo, D. Renfrew, and A. Soshnikov, "On finite rank deformations of wigner matrices," in *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, vol. 49, no. 1. Institut Henri Poincaré, 2013, pp. 64–94.
- [9] I. M. Johnstone and A. Y. Lu, "Sparse principal components analysis," *Unpublished manuscript*, 2004.
- [10] —, "On consistency and sparsity for principal components analysis in high dimensions," *Journal of the American Statistical Association*, vol. 104, no. 486, 2009.
- [11] A. d'Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. Lanckriet, "A direct formulation for sparse pca using semidefinite programming," *SIAM review*, vol. 49, no. 3, pp. 434–448, 2007.
- [12] H. Zou, T. Hastie, and R. Tibshirani, "Sparse principal component analysis," *Journal of computational and graphical statistics*, vol. 15, no. 2, pp. 265–286, 2006.
- [13] B. Moghaddam, Y. Weiss, and S. Avidan, "Spectral bounds for sparse pca: Exact and greedy algorithms," in *Advances in neural information processing systems*, 2005, pp. 915–922.
- [14] A. d'Aspremont, F. Bach, and L. E. Ghaoui, "Optimal solutions for sparse principal component analysis," *The Journal of Machine Learning Research*, vol. 9, pp. 1269–1294, 2008.
- [15] R. Krauthgamer, B. Nadler, and D. Vilenchik, "Do semidefinite relaxations really solve sparse pca?" *CoRR*, vol. abs/1306.3690, 2013.
- [16] Y. Deshpande and A. Montanari, "Sparse pca via covariance thresholding," *arXiv preprint arXiv:1311.5179*, 2013.
- [17] A. A. Amini and M. J. Wainwright, "High-dimensional analysis of semidefinite relaxations for sparse principal components," *The Annals of Statistics*, vol. 37, no. 5B, pp. 2877–2921, 2009.
- [18] S. Rangan and A. K. Fletcher, "Iterative estimation of constrained rank-one matrices in noise," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 1246–1250.
- [19] Y. Deshpande and A. Montanari, "Finding hidden cliques of size \sqrt{N}/e in nearly linear time," *arXiv preprint arXiv:1304.7047*, 2013.
- [20] D. Guo, Y. Wu, S. Shamai, and S. Verdú, "Estimation in gaussian noise: Properties of the minimum mean-square error," *Information Theory, IEEE Transactions on*, vol. 57, no. 4, pp. 2371–2385, 2011.
- [21] A. Javanmard and A. Montanari, "State evolution for general approximate message passing algorithms, with applications to spatial coupling," *arXiv preprint arXiv:1211.5164*, 2012.