
All-or-nothing statistical and computational phase transitions in sparse spiked matrix estimation

Jean Barbier

International Center for Theoretical Physics
Strada Costiera 11, 34151 Trieste, Italy
jbarbier@ictp.it

Nicolas Macris

Ecole Polytechnique Fédérale de Lausanne
CH 1015 Lausanne, Switzerland
nicolas.macris@epfl.ch

Cynthia Rush

Department of Statistics, Columbia University
New York, NY 10025
cynthia.rush@columbia.edu

Abstract

We determine statistical and computational limits for estimation of a rank-one matrix (the spike) corrupted by an additive gaussian noise matrix, in a sparse limit, where the underlying hidden vector (that constructs the rank-one matrix) has a number of non-zero components that scales sub-linearly with the total dimension of the vector, and the signal-to-noise ratio tends to infinity at an appropriate speed. We prove explicit low-dimensional variational formulas for the asymptotic mutual information between the spike and the observed noisy matrix and analyze the approximate message passing algorithm in the sparse regime. For Bernoulli and Bernoulli-Rademacher distributed vectors, and when the sparsity and signal strength satisfy an appropriate scaling relation, we find all-or-nothing phase transitions for the asymptotic minimum and algorithmic mean-square errors. These jump from their maximum possible value to zero, at well defined signal-to-noise thresholds whose asymptotic values we determine exactly. In the asymptotic regime the statistical-to-algorithmic gap diverges indicating that sparse recovery is hard for approximate message passing.

1 Introduction and setting

In modern machine learning and high dimensional statistics one often faces regression, classification, or estimation tasks, where the dimension of the feature vectors is much larger than the effective underlying dimensionality of the structure at hand. For example, hand-written MNIST digits are presented as vectors consisting of 28×28 pixels, in other words, they are binary vectors with 784 dimensions, whereas [1, 2] estimate their effective dimension to be in the orders of 10 's. Similarly the ISOMAP face database consists of images (256 levels of gray) of size 64×64 , i.e., vectors in \mathbb{R}^{4096} , whereas the correct intrinsic dimension is only 3 (for the vertical, horizontal pause and lighting direction). Natural images, which are generally sparse in a wavelet basis [3], are another popular example of low effective dimensionality. For natural images, a very simple model of low-dimensional structure, namely vectors with a sparse number of non-zero components, has proven immensely useful for studying these types of data structures and has led to the development of the whole area of compressed sensing [4, 5]. Similarly, matrix completion can be performed successfully when the number of sampled matrix elements is much smaller than the total number of elements, as long as one assumes the matrix is low-rank [6].

These and other developments have amply justified the “bet on sparsity principle”, which, in a nutshell, says that intrinsic low-dimensionality is often a crucial ingredient for the interpretability of high dimensional statistical models [7, 8]. In this context, it is of great importance to determine computational limits of estimation and to establish fundamental information theoretical (i.e., statistical) limits as benchmarks. Broadly speaking, exact results in the direction of computational or information theoretic limits usually fall in two categories. The first direction, traditional in statistics and computer science, derives *finite size* bounds on thresholds marking the onset of feasible signal recovery or learning [9, 10]. Such results usually leave out exact constants or do not always give the exact asymptotics. The second approach, is an *average case* approach (in the spirit of the statistical mechanics treatment of high dimensional systems), that models feature vectors by a *random ensemble*, taken as a set of random vectors with independently identically distributed (i.i.d.) components, and a small but fixed fraction of non-zero components. For example, the distribution might be a Bernoulli distribution, denoted $\text{Ber}(\rho_n)$ with $0 < \rho_n < 1$ and $\rho_n \rightarrow \rho > 0$ fixed, as the dimension of the vectors $n \rightarrow +\infty$. In Bayesian settings with known priors and hyper-parameters this approach has been highly successful, yielding exact formulas for the mutual information and minimum mean-square error (MMSE), as well as exact expressions (with constants) for statistical and computational message passing phase transition thresholds in the limit of *infinite dimensions* [11]. While the mathematical analysis of this approach is well developed in compressed sensing, generalized linear estimation, or rank-one noisy matrix and tensor estimation [12–23], the cited works all fall short of addressing the “true” sparse limit where $\rho_n \rightarrow 0$ instead of the limit being fixed (i.e., $\rho_n \rightarrow \rho > 0$) as $n \rightarrow +\infty$; to be more precise we manage to tackle the regime $\rho_n = \Omega(n^{-\beta})$ for $\beta \in [0, 1/6)$ for the information-theoretic analysis, and $\rho_n = \Omega((\ln n)^{-\alpha})$ for any positive fixed α for the algorithmic results. The terminology “true sparsity” is employed in order to emphasize this contrast. To the best of our knowledge the only works addressing this “true” sparse limit, in the average case approach for statistical phase transitions, are [24, 25] which consider linear regression.

In this work, we address the issue of “true” sparsity in the average case approach for the problem of rank-one matrix estimation from noisy observations of the entries. Low-rank matrix estimation (or factorization) is an important problem with numerous applications in image processing, principal component analysis (PCA), machine learning, DNA microarray data, and tensor decompositions. We determine information theoretic limits of the problem as well as computational limits of an approximate message passing algorithm [26–31] for signal estimation in the case of a noisy symmetric rank-one matrix model. Let us now introduce the model.

Setting: In the *sparse spiked Wigner matrix model* we consider a sparse signal-vector $\mathbf{X} = (X_1, \dots, X_n) \in \mathbb{R}^n$ with i.i.d. components distributed according to $P_{X,n} = \rho_n p_X + (1 - \rho_n)\delta_0$, where δ_0 is the Dirac mass at zero and $\rho_n \in (0, 1]^{\mathbb{N}}$ is a sequence of weights that will eventually tend to 0; the signal has in expectation a *sub-linear* number $n\rho_n$ of non-zero components. For the distribution p_X we assume that *i*) it is independent of n , *ii*) it has finite support in an interval $[-S, S]$, *iii*) it has second moment equal to 1 (without loss of generality). One has access to the symmetric data matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ with noisy entries

$$\mathbf{W} = \sqrt{\frac{\lambda_n}{n}} \mathbf{X} \otimes \mathbf{X} + \mathbf{Z}, \text{ or componentwise } W_{ij} = \sqrt{\frac{\lambda_n}{n}} X_i X_j + Z_{ij}, \quad 1 \leq i < j \leq n \quad (1)$$

where $\lambda_n > 0$ controls the strength of the signal and the noise is i.i.d. gaussian $Z_{ij} \sim \mathcal{N}(0, 1)$ for $i < j$ and symmetric $Z_{ij} = Z_{ji}$. Notice that the matrix \mathbf{W} can be viewed as a sum of a gaussian matrix from the Wigner ensemble perturbed by a rank-one matrix, $\mathbf{X} \mathbf{X}^\top$ (the “spike”). We focus, in particular, on binary \mathbf{X} generated with i.i.d. Bernoulli entries $X_i \sim P_{X,n} = \text{Ber}(\rho_n)$, or Bernoulli-Rademacher entries, $X_i \sim P_{X,n} = (1 - \rho_n)\delta_0 + \rho_n \frac{1}{2}(\delta_{-1} + \delta_1)$. In the Bayesian setting, we suppose that the prior $P_{X,n}$ and hyper-parameters are known. As we will see, when $\rho_n \rightarrow 0$, non-trivial estimation is possible only when $\lambda_n \rightarrow +\infty$.

The goal is to estimate the sparse spike $\mathbf{X} \otimes \mathbf{X}$ from the data \mathbf{W} . In the spiked Wigner model with linear sparsity, a class of polynomial-time algorithms, referred to as approximate message passing or AMP, have been shown to provide Bayes-optimal signal estimation for some problem settings asymptotically as $n \rightarrow +\infty$ [32–34]. Moreover, AMP algorithms have been applied successfully for signal recovery to a number of other low-rank matrix estimation problems [35–38] and, based on bold conjectures from the statistical physics literature, it is suggested that the estimation performance of AMP is the best among polynomial-time algorithms. Again, AMP is also provably optimal in some

parameters regimes. In this work, we study the properties of an AMP algorithm designed for signal estimation for the spiked Wigner matrix model in the sub-linear sparsity regime and compare its performance to benchmarks established by the information theoretic limits. This analysis provides a better understanding of the computational vs. theoretical gaps posed by the problem.

Some background and related work: In recent years, there has been much progress in understanding such spiked matrix models, which have played a crucial role in the analysis of threshold phenomena in high-dimensional statistical models for almost two decades, but most of this work has focused on standard settings, by which we mean problem settings where the distribution P_X is fixed independent of the problem dimension n . This means that the expected number of non-zero components of X , even if “small”, will scale *linearly* with n . Early rigorous results found in [39] determined the location of the information theoretic phase transition point in a spiked covariance model using spectral methods, and [40, 41] did the same for the Wigner case. More recently, the information theoretic limits and those of hypothesis testing have been derived, with the additional structure of sparse vectors, for large but finite sizes [42–44]. A lot of efforts have also been devoted to computational aspects of sparse PCA with many remarkable results [33, 43–50]. The picture that has emerged is that the information theoretic and computational phase transition regimes are not on the same scale and that the computational-to-statistical gap diverges in the limit of vanishing sparsity. However, the exact thresholds with constants as well as the behaviour of the mean-square errors remained unknown.

Using heuristic methods from the statistical physics of spin glass theory (the so-called replica method [51]), the authors of [52] observed an interesting phenomenology of the information theoretical and computational limits with sharp phase transitions as $n \rightarrow +\infty$. The rigorous mathematical theory of these phase transitions is now largely under control. On one hand, an approximate message passing algorithm for signal recovery can be rigorously analyzed via its state evolution [27, 28, 53], and on the other hand, the asymptotic mutual information per variable between the hidden spike and data matrices has been rigorously computed in a series of works using various methods (cavity method, spatial coupling, interpolation methods, PDE techniques) [16–23, 54–57]. The information theoretic phase transitions are then signaled by singularities, as a function of the signal strength, in the limit of the mutual information per variable when $n \rightarrow +\infty$. The phase transition also manifests itself as a jump discontinuity in the minimum mean-square error (MMSE)¹. Once the mutual information is known, it is usually possible to deduce the MMSE using so-called I-MMSE relations [58, 59]. Essentially, the MMSE can be accessed by differentiating the mutual information with respect to the signal-to-noise strength. Closed form expressions for the asymptotic mutual information therefore allow to benchmark the fundamental information theoretical limits of estimation. We also point the reader towards the works [60–62] which derive limits of detecting the presence of a spike in a noisy matrix, rather than estimating it.

Finally, similar phase transitions in sub-linear sparsity regimes for binary signals have been studied in the context of high-dimensional linear regression or compressed sensing for support recovery [24, 25]. These works focus on the MMSE and prove the occurrence of the 0 – 1 phase transition, which they called an “all-or-nothing” phenomenon. We note that our approach is technically very different in that it determines the variational expressions for mutual informations and finds the transitions as a consequence. Moreover these works do not deal with algorithmic phase transitions, while we consider here the one of AMP.

Our contributions: We provide new results in sparse limits along two main lines:

- The exact statistical threshold for the sharp all-or-nothing statistical transition at the level of the MMSE. This follows from a rigorous derivation of the mutual information in the form of a variational problem.
- The AMP algorithmic threshold and all-or-nothing transition at the level of the AMP mean-square error. This follows from a “finite sample” analysis of the approximate message passing algorithm, allowing to rigorously track its performance in sparse regimes.

Let us explain these contributions in detail.

¹This is the generic singularity and one speaks of a first order transition. In special cases the MMSE may be continuous with a higher discontinuous derivative of the mutual information.

In this work, we identify the correct *scaling regimes* of vanishing sparsity and diverging signal strength in which non-trivial information theoretic and algorithmic AMP phase transitions occur. Moreover, we determine the statistical-to-algorithmic gap in the scaling regime. These scalings, thresholds, as well as formulas for the mutual information, were first heuristically and numerically derived in [52] using the non-rigorous replica method of spin-glass theory and the state evolution equations for AMP. However, it must be stressed that, not only were these calculations far from rigorous, but more importantly the limit $n \rightarrow +\infty$ is taken first for a fixed parameter $\rho_n = \rho$, and the sparse limit $\rho \rightarrow 0_+$ is taken only after. Although the thresholds found in this way agree with our derivations, this is far from evident a priori. In contrast, our results are entirely rigorous and valid in the truly sparse limit. Therefore the picture found in [52] is fully vindicated. In addition, we also establish that the MMSE and AMP phase transitions are of the all-or-nothing type, a novelty of the present work.

The information theoretic analysis is done via the adaptive interpolation method [19, 20, 22], first introduced in the non-sparse matrix estimation problems, to provide for the sparse limit, closed form expressions of the mutual information in terms of low-dimensional variational expressions (theorem 1 in section 2). That the adaptive interpolation method can be extended to the sparse limit is interesting and not a priori obvious. Using the I-MMSE relation and the solution of the variational problems for Bernoulli and Bernoulli-Rademacher distributions of the sparse signal, we then find that the MMSE displays an all-or-nothing phase transition (corollary 1) and we determine the exact threshold (with constants).

A useful property of AMP is that in the large system limit $n \rightarrow +\infty$, its performance can be exactly characterized and rigorously analyzed through its so-called *state evolution*. When $\rho_n \rightarrow \rho > 0$, the validity of the state evolution analysis for AMP for low-rank matrix estimation follows from the standard AMP theory [27, 28] (with some additional work needed to deal with technicalities relating to the algorithm's initialization [34]), however, in the sub-linear sparsity regime considered here, proving the validity of the state evolution characterization requires a new and non-trivial analysis using “finite sample” techniques, first developed in [63]. We find that the algorithmic MSE, denoted MSE_{AMP} displays an all-or-nothing transition as well and we determine the scaling of the threshold (the constant being obtained numerically). Interestingly, the transition is on a very different signal-to-noise scale as compared to the MMSE (theorem 2 found in section 3).

Let us describe in a bit more detail the sparse regimes we study and the corresponding thresholds. To gain some intuition, we first note that for sub-linear sparsity, phase transitions can appear only if the signal strength tends to infinity. This can be seen from the following heuristic argument: notice that the total signal-to-noise ratio per non-zero component² scales as $(\lambda_n/n)\rho_n^2 n^2 / (\rho_n n) = \lambda_n \rho_n$, meaning that $\lambda_n \rightarrow +\infty$ is necessary in order to have enough energy to estimate the non-zero components. Our analysis shows that non-trivial information theoretic and AMP phase transitions occur at different scales:

- **Statistical phase transition regime:** While our results are more general (see appendix A and theorem 3) our main interest is in a regime of the form

$$\lambda_n = 4\gamma |\ln \rho_n| \rho_n^{-1}, \quad \rho_n = \Omega(n^{-\beta}), \quad (2)$$

for $\beta, \gamma \in \mathbb{R}_{>0}$ and β small enough. We prove that in this regime a phase transition occurs as function of γ .

- **Algorithmic AMP phase transition regime:** We control the performance of AMP for a number of time-iterations $t = o(\frac{\ln n}{\ln \ln n})$ and rigorously prove that the all-or-nothing transition occurs for

$$\lambda_n = w \rho_n^{-2}, \quad \rho_n = \Omega((\ln n)^{-\alpha}), \quad (3)$$

where $w, \alpha \in \mathbb{R}_{\geq 0}$ are fixed constants (note that we can take any $\alpha > 1$). Controlling the AMP iterations in this regime is already highly non-trivial, however, we conjecture that the result still holds when $\rho_n = \Omega(n^{-\beta})$ for $\beta > 0$ small enough, but refining the analysis in appendix K to find the stronger result is left for future work.

²In more detail, this is equal to the signal-to-noise ratio per observation $(\lambda_n/n)\rho_n^2$ times the number of observations $\Theta(n^2)$ divided by the expected number of non-zero components $\rho_n n$.

The relation $\lambda_n \sim \rho_n^{-2}$ for the AMP threshold was obtained in [52] based on a stability analysis of the linearized state evolution. However, we recall that in their setting $\rho_n = \rho$, $n \rightarrow +\infty$, and not only is the sparse limit $\rho \rightarrow 0_+$ taken after the high-dimensional limit, but also the AMP iterations are not controlled. In appendix G in the supplementary material we provide a simpler alternative argument that does not require linearizing the recursion.

We focus in particular on binary signals with $P_{X,n}$ equal to $\text{Ber}(\rho_n)$ or Bernoulli-Rademacher $(1 - \rho_n)\delta_0 + \rho_n\frac{1}{2}(\delta_{-1} + \delta_1)$. For these distributions we prove the existence of all-or-nothing transitions for the MMSE and MSE_{AMP} for the specific sparsity regimes stated above. This is illustrated in figures 1 and 2, found in sections 2 and 3, which display, for the Bernoulli prior, the explicit asymptotic values to which the finite n mutual information and MMSE converge. The results are similar for the Bernoulli-Rademacher distribution. In figure 1, we see that as $\rho_n \rightarrow 0_+$ the (suitably normalized) mutual information approaches the broken line with an angular point at $\lambda/\lambda_c(\rho_n) = 1$ where $\lambda_c(\rho_n) = 4|\ln \rho_n|/\rho_n$. Moreover the (suitably normalized) MMSE tends to its maximum possible value 1 for $\lambda/\lambda_c(\rho_n) < 1$, develops a jump discontinuity at $\lambda/\lambda_c(\rho_n) = 1$, and takes the value 0 when $\lambda/\lambda_c(\rho_n) > 1$ as $\rho_n \rightarrow 0$. In figure 2, we observe the same behavior for MSE_{AMP} as a function of $\lambda/\lambda_{\text{AMP}}(\rho_n)$, but now the algorithmic threshold is $\lambda_{\text{AMP}}(\rho_n) = 1/(e\rho_n^2)$, where the constant $1/e$ is approximated numerically. Note that the same asymptotic behavior is observed in the related problem of finding a small hidden community in a graph, see figure 5 in [64].

2 Statistical phase transition

The phase transition manifests itself as a singularity (more precisely a discontinuous first order derivative) in the mutual information $I(\mathbf{X} \otimes \mathbf{X}; \mathbf{W}) = H(\mathbf{W}) - H(\mathbf{W}|\mathbf{X} \otimes \mathbf{X})$. Note that because the data \mathbf{W} depends on \mathbf{X} only through $\mathbf{X} \otimes \mathbf{X}$ we have $H(\mathbf{W}|\mathbf{X} \otimes \mathbf{X}) = H(\mathbf{W}|\mathbf{X})$ and therefore $I(\mathbf{X} \otimes \mathbf{X}; \mathbf{W}) = I(\mathbf{X}; \mathbf{W})$. From now on we use the form $I(\mathbf{X}; \mathbf{W})$.

To state the result, we define the *potential function*:

$$i_n^{\text{pot}}(q, \lambda, \rho) \equiv \frac{\lambda}{4}(q - \rho)^2 + I_n(X; \sqrt{\lambda q}X + Z), \quad (4)$$

where $I_n(X; \sqrt{\lambda q}X + Z)$ is the mutual information for a scalar gaussian channel, with $X \sim P_{X,n}$ and $Z \sim \mathcal{N}(0, 1)$. The mutual information I_n is indexed by n because of its dependence on $P_{X,n}$.

Theorem 1 (Mutual information for the sparse spiked Wigner model). *Let the sequences λ_n and ρ_n verify (2) with $\beta \in [0, 1/6)$ and $\gamma > 0$. There exists $C > 0$ independent of n such that*

$$\frac{1}{\rho_n |\ln \rho_n|} \left| \frac{1}{n} I(\mathbf{X}; \mathbf{W}) - \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda_n, \rho_n) \right| \leq C \frac{(\ln n)^{1/3}}{n^{(1-6\beta)/7}}. \quad (5)$$

The mutual information is thus given, to leading order, by a one-dimensional variational problem

$$I(\mathbf{X}; \mathbf{W}) = n\rho_n |\ln \rho_n| \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda_n, \rho_n) + \text{correction terms}.$$

The factor $\rho_n |\ln \rho_n|$ is related to the entropy (in nats) of the support of the signal given by $-\rho_n \ln \rho_n + (1 - \rho_n) \ln(1 - \rho_n)$, which behaves like $n\rho_n |\ln \rho_n|$ for $\rho_n \rightarrow 0_+$. In particular, for both the *Bernoulli* and *Bernoulli-Rademacher* distributions an analytical solution of the variational problem, given in appendix F, shows that $(\rho_n |\ln \rho_n|)^{-1} \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda_n, \rho_n)$ tends to the singular function $\gamma \mathbb{I}(\gamma \leq 1) + \mathbb{I}(\gamma \geq 1)$ as $n \rightarrow +\infty$ and $\rho_n \rightarrow 0$, where we recall that $\lambda_n = 4\gamma |\ln \rho_n| \rho_n^{-1}$. See figure 1. Let us mention that $\beta < 1/6$ is probably not a fundamental limit to the validity of the result but rather is an artefact of the sub-optimality of our proof technique.

We now turn to the consequences for the MMSE. It is convenient to work with the ‘‘matrix’’ MMSE defined as $\text{MMSE}((X_i X_j)_{i < j} | \mathbf{W}) \equiv \mathbb{E} \| (X_i X_j)_{i < j} - \mathbb{E}[(X_i X_j)_{i < j} | \mathbf{W}] \|_{\mathbb{F}}^2$. This quantity satisfies the I-MMSE relation [58, 59] (see also appendix I for a self-contained derivation),

$$\frac{d}{d\lambda_n} \frac{1}{n} I(\mathbf{X}; \mathbf{W}) = \frac{1}{2n^2} \text{MMSE}((X_i X_j)_{i < j} | \mathbf{W}).$$

In appendix J we prove:

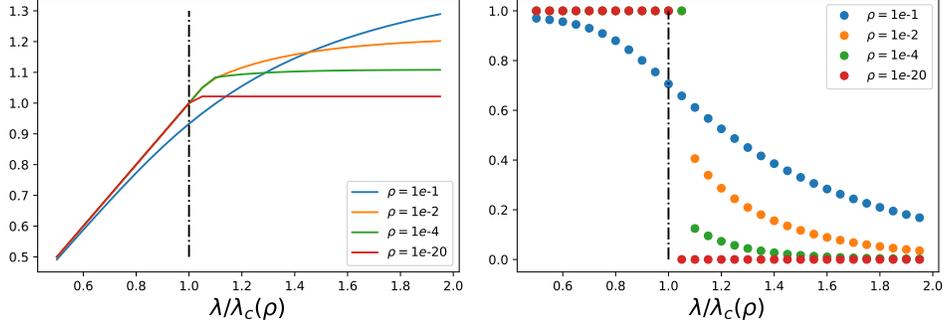


Figure 1: A sequence of suitably normalized asymptotic mutual information $(\rho |\ln \rho|)^{-1} \inf_{q \in [0, \rho]} i_n^{\text{pot}}(q, \lambda, \rho)$ (left) and associated minimum mean-square error (MMSE) $\rho^{-2} \frac{d}{d\lambda} \inf_{q \in [0, \rho]} i_n^{\text{pot}}(q, \lambda, \rho)$ (right) curves as a function of $\lambda/\lambda_c(\rho)$ with $\lambda_c(\rho) = 4|\ln \rho|/\rho$ for the model $X_i \sim \text{Ber}(\rho)$ and various $\rho = \rho_n$ values (that can be converted to signal sizes through $\rho_n = \Omega(n^{-\beta})$ given a sparsity scaling β) using the potential function defined in (4). These are the curves towards which, respectively, the finite size mutual information $(n\rho |\ln \rho|)^{-1} I(\mathbf{X}; \mathbf{W})$ and minimum mean-square error $(n\rho)^{-2} \text{MMSE}((X_i X_j)_{i < j} | \mathbf{W})$, converge: see theorem 1 and corollary 1. In the sparse limit $\rho \rightarrow 0$, the MMSE curves approach a 0–1 phase transition with the discontinuity at $\lambda = \lambda_c(\rho)$. This corresponds to an angular point for the mutual information (by the I-MMSE relation).

Corollary 1 (Minimum mean-square error for the sparse spiked Wigner model). *Let $\frac{1}{2}m_n(\lambda, \rho_n) \equiv \rho_n^{-2} \frac{d}{d\lambda} \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda, \rho_n)$. Let $\epsilon > 0$ and sequences λ_n and ρ_n verifying (2) with $\beta \in [0, 1/13]$. There exists $C' > 0$ independent of n such that*

$$m_n(\lambda_n + \epsilon, \rho_n) - \frac{C' (\ln n)^{4/3}}{\epsilon n^{(1-13\beta)/7}} \leq \frac{\text{MMSE}((X_i X_j)_{i < j} | \mathbf{W})}{(n\rho_n)^2} \leq m_n(\lambda_n - \epsilon, \rho_n) + \frac{C' (\ln n)^{4/3}}{\epsilon n^{(1-13\beta)/7}}.$$

Concretely the derivative $(d/d\lambda_n) \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda_n, \rho_n)$ is computed, using the envelope theorem [65], as $(\partial/\partial \lambda_n) i_n^{\text{pot}}(q_n^*, \lambda_n, \rho)$ where $q_n^* = q_n^*(\lambda_n, \rho_n)$ is the solution of the variational problem, which is unique almost everywhere (except at the phase transition point, see e.g. [15] for such proofs). For Bernoulli and Bernoulli-Rademacher distributions, we easily compute the limiting behavior $m_n(\lambda_n, \rho_n)$ from the solution of the variational problem stated above, and find that $(n\rho_n)^{-2} \text{MMSE}((X_i X_j)_{i < j} | \mathbf{W})$ tends to $\mathbb{I}(\gamma \leq 1)$ as $n \rightarrow +\infty$.

Figure 1 shows the mutual information and MMSE computed from the numerical solution of the variational problem for a sequence of $\text{Ber}(\rho_n)$ distributions. We check that the limiting curves are indeed approached as $\rho_n \rightarrow 0$ and, in particular, the suitably rescaled MMSE displays the all-or-nothing transition at $\lambda/\lambda_c(\rho_n) = 1$ as $n \rightarrow +\infty$ with $\lambda_c(\rho_n) = 4|\ln \rho_n|/\rho_n$. For the Bernoulli-Rademacher distribution the transition location is the same, suggesting that the hardness of the inference is only related, for discrete priors, to the recovery of the support. For more generic distributions than these two cases the situation is richer. Although one generically observes phase transitions *in the same scaling regime*, the limiting curves appear to be more complicated than the simple staircase shape and the jumps are not necessarily located at $\gamma = 1$. A classification of these transitions is an interesting problem that is out of the scope of this paper.

3 AMP algorithmic phase transition

Approximate message passing (AMP) is a low complexity algorithm that iteratively updates estimates of the unknown signal, which, in the case of the spiked Wigner model is \mathbf{X} , from the noisy data \mathbf{W} . The iterative estimates are denoted $\{\mathbf{x}^t\}_{t \geq 1}$. Let $\mathbf{A} \equiv \mathbf{W}/\sqrt{n}$ and initialize with $f_0(\mathbf{x}^0)$ independent of \mathbf{W} , such that $\langle f_0(\mathbf{x}^0), \mathbf{X} \rangle > 0$. Then let $\mathbf{x}^1 = \mathbf{A}f_0(\mathbf{x}^0)$, and for $t \geq 1$, compute

$$\mathbf{x}^{t+1} = \mathbf{A}f_t(\mathbf{x}^t) - \mathbf{b}_t f_{t-1}(\mathbf{x}^{t-1}), \quad \mathbf{b}_t = \frac{1}{n} \sum_{i=1}^n f'_t(x_i^t), \quad (6)$$

where the scalar function $f_t : \mathbb{R} \rightarrow \mathbb{R}$ is applied elementwise to vector input, i.e., $f_t(\mathbf{x}) = (f_t(x_1), \dots, f_t(x_n))$ for a vector $\mathbf{x} \in \mathbb{R}^n$, and its exact value is given in what follows (in (9)). We

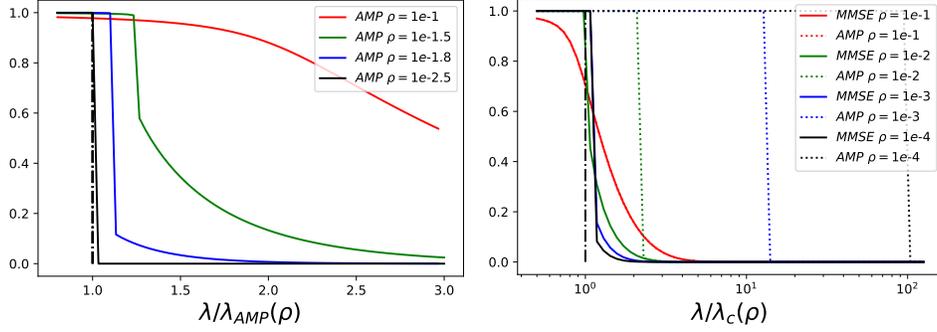


Figure 2: Left: The mean-square error towards which the suitably normalized matrix-MSE of the AMP algorithm, MSE_{AMP} , converges for various sparsity levels, see theorem 2. An all-or-nothing transition appears as $\rho = \rho_n \rightarrow 0$ at $\lambda_{\text{AMP}}(\rho) = 1/(e\rho)^2$. Comparing to figure 1 the transition becomes sharper much faster as ρ decreases. Right: Horizontal axis is on a log scale. The statistical-to-algorithmic gap diverges as $\rho \rightarrow 0$.

refer to the functions $\{f_t\}_{t \geq 0}$ as “denoisers”, for reasons that will become clear momentarily. Notice that (6) gives both matrix estimates $\hat{\mathbf{X}} \hat{\mathbf{X}}^\top = f_t(\mathbf{x}^t)[f_t(\mathbf{x}^t)]^\top$ and signal estimates $\hat{\mathbf{X}} = f_t(\mathbf{x}^t)$.

A key property of AMP is that, asymptotically as $n \rightarrow \infty$, a deterministic, scalar recursion referred to as *state evolution* exactly characterizes its performance, in the sense that the estimates x_i^t converge to random variables with mean and variance governed by the state evolution. For the sub-linear sparsity regime, we introduce an n -dependent state evolution, reflecting that our sparsity level ρ_n and signal strength λ_n both now change as n grows. We will show, based on measure concentration arguments, that the usual asymptotic characterization also gives a finite sample approximation,

meaning that for any n fixed but large, x_i^t is *approximately* distributed as $x_i^t \stackrel{d}{\approx} \mu_t^n X_0^n + \sqrt{\tau_t^n} Z$ where μ_t^n and τ_t^n are characterized by the state evolution below with $X_0^n \sim P_{X,n}$ independent of standard gaussian Z . The n -dependent state evolution is defined as follows: for $t \geq 1$,

$$\mu_1^n = \sqrt{\lambda_n} \langle f_0(\mathbf{x}^0), \mathbf{X} \rangle / n, \quad \tau_1^n = \|f_0(\mathbf{x}^0)\|^2 / n, \quad (7)$$

$$\mu_{t+1}^n = \sqrt{\lambda_n} \mathbb{E} \left\{ X_0^n f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z) \right\}, \quad \tau_{t+1}^n = \mathbb{E} \left\{ [f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z)]^2 \right\}, \quad (8)$$

where we include the n superscript to emphasize the dependence.

A well-motivated choice of denoiser functions $\{f_t\}_{t \geq 0}$ are the conditional expectation denoisers. Namely, given that we have knowledge of the prior distribution of the signal elements, and considering the approximate characterization of the estimate x_i^t via the state evolution, the Bayes-optimal way to update our signal estimate at any iteration is the following: for $t \geq 1$,

$$f_t(x) = \mathbb{E} \{ X_0^n \mid \mu_t^n X_0^n + \sqrt{\tau_t^n} Z = x \}, \quad (9)$$

with $X_0^n \sim P_{X,n}$ independent of standard gaussian Z . Strictly speaking, $f_t(\cdot)$ also has an n -dependency, so to be consistent we should label $f_t(\cdot) \equiv f_t^n(\cdot)$, however we drop this for simplicity. With this choice of denoiser function, the state evolution (8) simplifies: by the Law of Total Expectation, $\mathbb{E} \{ X_0^n f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z) \} = \mathbb{E} \{ [f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z)]^2 \}$, thus $\mu_t^n = \sqrt{\lambda_n} \tau_t^n$, and so

$$\tau_{t+1}^n = \mathbb{E} \left\{ \left[\mathbb{E} \{ X_0^n \mid \sqrt{\lambda_n} \tau_t^n X_0^n + \sqrt{\tau_t^n} Z \} \right]^2 \right\}. \quad (10)$$

The performance guarantees given by the state evolution are stated informally in what follows, with a more formal result given in appendix K. The proof extends and refines³ the finite sample analysis of AMP given in [63, Theorem 1]. These guarantees concern the convergence of the empirical

³The result in [63] is a general AMP algorithm with a “rectangular” structure that does not cover the “symmetric” AMP in (6). However, extensions of this result to the symmetric case are straightforward, as discussed in [63, Section 1], but technical. Moreover, the dependence on n for the state evolution requires that these values are tracked carefully through the proof, whereas this was not done in [63], as these values were assumed to be universal constants. For simplicity of exposition in this document, we do not elaborate further on these technicalities at this point and put these details in appendix K.

distribution of x_i^t to its approximating distribution determined by the state evolution and specifically apply to the AMP algorithm using the denoiser in (9). For all order 2 pseudo-Lipschitz functions⁴, denoted $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}$ with Lipschitz constant $L_\psi > 0$, we have that for $\epsilon \in (0, 1)$ and $t \geq 1$,

$$\mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n \psi(X_i, f_t(x_i^t)) - \mathbb{E}\left\{\psi(X_0^n, f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z))\right\}\right| \geq \epsilon\right) \leq CC_t \exp\left\{\frac{-cc_t n \epsilon^2}{L_\psi^2 \gamma_n^t}\right\} \quad (11)$$

where $\mathbf{X} = (X_1, \dots, X_n)$ is the true signal and C, C_t, c, c_t are universal constants not depending on n or ϵ , but with C_t, c_t depending on the iteration t and whose exact value is given in theorem 2. Finally, γ_n^t characterizes the way the bound depends on the state evolution parameters and its exact value is given in (14). We want to consider, specifically, the vector-MSE and matrix-MSE of AMP, namely $\frac{1}{n} \|\mathbf{X} - f_t(\mathbf{x}^t)\|^2$ and $\frac{1}{n^2} \|\mathbf{X} \mathbf{X}^\top - f_t(\mathbf{x}^t)[f_t(\mathbf{x}^t)]^\top\|_F^2$, for any $t \geq 1$.

Theorem 2 (Finite sample state evolution). *Consider AMP in (6) using the conditional expectation denoiser in (9). Then for $\epsilon \in (0, 1)$ and $t \geq 1$, let $\text{bound}_t \equiv CC_t \exp\{-cc_t n \epsilon^2 / \gamma_n^t\}$, then*

$$\mathbb{P}\left(\left|\frac{1}{n} \|\mathbf{X} - f_t(\mathbf{x}^t)\|^2 - (\rho_n - \tau_{t+1}^n)\right| \geq \epsilon\right) \leq \text{bound}_t, \quad (12)$$

$$\mathbb{P}\left(\left|\frac{1}{n^2} \|\mathbf{X} \mathbf{X}^\top - f_t(\mathbf{x}^t)[f_t(\mathbf{x}^t)]^\top\|_F^2 - (\rho_n^2 - (\tau_{t+1}^n)^2)\right| \geq \epsilon\right) \leq \text{bound}_t, \quad (13)$$

where $X_0^n \sim P_{X,n}$ and τ_t^n is defined in (10). The values C, c are universal constants not depending on n or ϵ with C_t, c_t given by $C_t = C_1^t (t!)^{C_2}, c_t = [c_1^t (t!)^{c_2}]^{-1}$. Finally,

$$\begin{aligned} \gamma_n^t &\equiv \lambda_n^{2t-1} (\nu^n + \tau_1^n) (\nu^n + \tau_1^n + \tau_2^n) \cdots (\nu^n + \sum_{i=1}^t \tau_i^n) \\ &\quad \times \max\{1, \hat{b}_1\} \max\{1, \hat{b}_2\} \cdots \max\{1, \hat{b}_{t-1}\}, \end{aligned} \quad (14)$$

where ν^n is the variance factor of sub-Gaussian \mathbf{X}^n (for $P_{X,n} = \text{Ber}(\rho_n)$ we have $\nu^n \leq 1/4$; see lemma 14 in the supplementary material and $\hat{b}_t = \mathbb{E}\{f_t'(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z)\}$).

Theorem 2 follows from the finite sample guarantees given in (11), and, in appendix K, we discuss in more detail the proof of theorem 2 and result 11. We make a few remarks on the result here.

Remark 1: ρ_n normalization and all-or-nothing transition. To be consistent with the previously stated results, we could renormalize the MSEs as follows and the result still holds as

$$\begin{aligned} \mathbb{P}\left(\left|\frac{1}{\rho_n n} \|\mathbf{X} - f_t(\mathbf{x}^t)\|^2 - \left(1 - \frac{\tau_{t+1}^n}{\rho_n}\right)\right| \geq \epsilon\right) &\leq CC_t \exp\{-cc_t n \rho_n^2 \epsilon^2 / \gamma_n^t\}, \\ \mathbb{P}\left(\left|\frac{1}{(\rho_n n)^2} \|\mathbf{X} \mathbf{X}^\top - f_t(\mathbf{x}^t)[f_t(\mathbf{x}^t)]^\top\|_F^2 - \left(1 - \left(\frac{\tau_{t+1}^n}{\rho_n}\right)^2\right)\right| \geq \epsilon\right) &\leq CC_t \exp\{-cc_t n \rho_n^4 \epsilon^2 / \gamma_n^t\}. \end{aligned}$$

In appendix G we show that $\tau_{t+1}^n / \rho_n \rightarrow 0$ for $\lambda_n \rho_n^2 \rightarrow 0$ and $\tau_{t+1}^n / \rho_n \rightarrow 1$ for $\lambda_n \rho_n^2 \rightarrow +\infty$. This is consistent with the numerics on figure 2 where we see a transition for $\lambda_n \rho_n^2 = 1/e^2$.

Remark 2: AMP regime and statistical-to-algorithmic gap. We apply theorem 2 in the regime where $t = o\left(\frac{\ln n}{\ln \ln n}\right)$, which, as discussed in [63], is the regime where the state evolution predictions are meaningful with respect to the values of C_t, c_t and the constraints they specify on how large t can be compared to the dimension n . In our work, we also have constraints related to the γ_n^t value in (14) that appears in the denominator of the rate of concentration. Considering these constraints, we apply theorem 2 for signal strength and sparsity scaling like $\lambda_n \rho_n^2 = w$ and $\rho_n = \Omega((\ln n)^{-\alpha})$ with $w, \alpha \in \mathbb{R}_+$, and show that the above probabilities indeed tend to zero as $n \rightarrow +\infty$. Appendix L provides the details of this calculation.

Note that since theorem 1 and corollary 1 hold for $\rho_n = \Omega(n^{-\beta})$ and thus for $\rho_n = \Omega((\ln n)^{-\alpha})$ as well, then both the statistical and algorithmic transitions (and therefore the statistical-to-computational gap) are proven for $\rho_n = \Omega((\ln n)^{-\alpha})$.

Remark 3: λ_n, τ^n dependence. The λ_n dependence in γ_n^t defined in (14) comes from the (pseudo-) Lipschitz constants L_f in (11). The dependence on the Lipschitz constants, and on the state evolution

⁴For any $n, m \in \mathbb{N}_{>0}$, a function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is *pseudo-Lipschitz of order 2* if there exists a constant $L > 0$ such that $\|\phi(\mathbf{x}) - \phi(\mathbf{y})\| \leq L(1 + \|\mathbf{x}\| + \|\mathbf{y}\|)\|\mathbf{x} - \mathbf{y}\|$ for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.

parameters τ_t^n , was not stated explicitly in the original concentration bound in [63, Theorem 1] as the authors assume these values do not change with n and, thus, can be absorbed into the universal constants. By examining the proof of [63, Theorem 1], one gets that the dependence takes the form in (14). More details on how we arrive at the rates in theorem 2 can be found in appendix K.

Remark 4: Algorithm initialization. We assume that the AMP algorithm in (6) was initialized with $f_0(\mathbf{x}^0)$ independent of \mathbf{W} such that $\langle f_0(\mathbf{x}^0), \mathbf{X} \rangle > 0$. The second condition ensures that $\mu_0^n \neq 0$ (which would mean $\mu_t^n = 0$ for all $t \geq 0$). If $P_{X,n}$ is $\text{Ber}(\rho_n)$, one could use, for example, $f_0(\mathbf{x}^0) = \mathbf{1}$, since the mean of the signal elements is positive. However, if $P_{X,n}$ is Bernoulli-Rademacher, a more complicated initialization procedure is needed since initializing in this way would cause the algorithm to get stuck in an unstable fixed point. We refer the reader to [34] for a discussion of an appropriate spectral initialization for this setting. However, such an initialization violates the assumption of independence with \mathbf{W} . The theoretical idea in [34] that allows one to get around this dependence is to analyze AMP in (6) with a matrix $\tilde{\mathbf{A}}$ that is an *approximate* representation of the conditional distribution of \mathbf{A} given the initialization, and then to show that with high probability the two algorithms will be close each other. We believe that incorporating these ideas with the finite sample guarantee in (11) would be straightforward, and theorem 2 could be extended to the setting of AMP with a spectral initialization.

Broader impact

One cannot underestimate the relevance of sparse estimation in modern technology, and although this work is valid within the limits of a theoretical model, it participates towards better fundamental understanding of necessary resources in terms of energy and quantity of data when this data is sparse. Besides radical transitions in behaviour under small changes of control parameters, we also show that an estimation task can become computationally hard or impossible, even with (practically) unbounded signal strengths. Broadly speaking, such results provide guidelines for better design and less wasteful engineering systems.

Acknowledgments

J.B. acknowledges discussions with Galen Reeves during his visit of Duke University. C.R. acknowledges support from NSF CCF #1849883 and N.M. from Swiss National Foundation for Science grant number 200021E 17554.

References

- [1] J. Costa and A. Hero. Learning intrinsic dimension and intrinsic entropy of high-dimensional datasets. In *European Signal Processing Conference (EUSIPCO), Vienna, Austria, 2004*. Zenodo, sep 2004.
- [2] M. Hein and J.-Y. Audibert. Intrinsic dimensionality estimation of submanifolds in rd.
- [3] S. Mallat. *A Wavelet Tour of Signal Processing (Third Edition): The Sparse Way*. Academic Press, Boston, third edition, 2009.
- [4] E. J. Candès, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [5] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [6] E. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- [7] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference and prediction*. Springer, 2 edition, 2009.
- [8] I. Rish and G. Grabarnik. *Sparse Modeling: Theory, Algorithms, and Applications*. CRC Press, Inc., USA, 1st edition, 2014.

- [9] T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall/CRC, 2015.
- [10] M. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
- [11] L. Zdeborová and F. Krzakala. Statistical physics of inference: thresholds and algorithms. *Advances in Physics*, 65(5), Aug 2016.
- [12] J. Barbier, M. Dia, N. Macris, and F. Krzakala. The Mutual Information in Random Linear Estimation. In *54th Annual Allerton Conference on Communication, Control, and Computing*, September 2016.
- [13] J. Barbier, N. Macris, M. Dia, and F. Krzakala. Mutual information and optimality of approximate message-passing in random linear estimation. *IEEE Transactions on Information Theory*, 2020.
- [14] G. Reeves and H. D. Pfister. The replica-symmetric prediction for compressed sensing with gaussian matrices is exact. In *2016 IEEE International Symposium on Information Theory (ISIT)*, July 2016.
- [15] J. Barbier, F. Krzakala, N. Macris, L. Miolane, and L. Zdeborová. Optimal errors and phase transitions in high-dimensional generalized linear models. *Proceedings of the National Academy of Sciences*, 116(12):5451–5460, 2019.
- [16] M. Lelarge and L. Miolane. Fundamental limits of symmetric low-rank matrix estimation. *Probability Theory and Related Fields*, 173(3-4):859–929, 2018.
- [17] L. Miolane. Fundamental limits of low-rank matrix estimation: The non-symmetric case. *ArXiv e-prints*, February 2017.
- [18] J. Barbier, M. Dia, N. Macris, F. Krzakala, T. Lesieur, and L. Zdeborová. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. In *Advances in Neural Information Processing Systems (NIPS) 29*, pages 424–432. 2016.
- [19] J. Barbier and N. Macris. The adaptive interpolation method: a simple scheme to prove replica formulas in bayesian inference. *Probability Theory and Related Fields*, Oct 2018.
- [20] J. Barbier and N. Macris. The adaptive interpolation method for proving replica formulas. applications to the curie–weiss and wigner spike models. *Journal of Physics A: Mathematical and Theoretical*, 52(29):294002, jun 2019.
- [21] T. Lesieur, L. Miolane, M. Lelarge, F. Krzakala, and L. Zdeborová. Statistical and computational phase transitions in spiked tensor estimation. In *IEEE International Symposium on Information Theory (ISIT), 2017*.
- [22] J. Barbier, N. Macris, and L. Miolane. The Layered Structure of Tensor Estimation and its Mutual Information. In *55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, September 2017.
- [23] J.-C. Mourrat. Hamilton-jacobi equations for finite-rank matrix inference. *arXiv preprint arXiv:1904.05294*, 2019.
- [24] D. Gamarnik and I. Zadik. High dimensional regression with binary coefficients. estimating squared error and a phase transtition. In *Conference on Learning Theory*, pages 948–953, 2017.
- [25] G. Reeves, J. Xu, and I. Zadik. The all-or-nothing phenomenon in sparse linear regression. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2652–2663. PMLR, 25–28 Jun 2019.
- [26] Y. Kabashima. A CDMA multiuser detection algorithm on the basis of belief propagation. *Journal of Physics A: Mathematical and General*, 36(43):11111–11121, oct 2003.
- [27] M. Bayati and A. Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Trans. on Information Theory*, 2011.
- [28] D. L. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- [29] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová. Probabilistic reconstruction in compressed sensing: algorithms, phase diagrams, and threshold achieving matrices. *J. Stat. Mech. Theory Exp.*, (8), 2012.

- [30] A. Montanari. Graphical models concepts in compressed sensing. In Y. C. Eldar and G. Kutyniok, editors, *Compressed Sensing*, pages 394–438. Cambridge University Press, 2012.
- [31] S. Rangan. Generalized approximate message passing for estimation with random linear mixing. In *Proc. IEEE Int. Symp. Inf. Theory*, pages 2168–2172, 2011.
- [32] Y. Deshpande and A. Montanari. Finding hidden cliques of size $\sqrt{N/e}$ in nearly linear time. *Foundations of Computational Mathematics*, 15(4):1069–1128, 2015.
- [33] Y. Deshpande and A. Montanari. Information-theoretically optimal sparse pca. In *2014 IEEE International Symposium on Information Theory*, pages 2197–2201. IEEE, 2014.
- [34] A. Montanari and R. Venkataramanan. Estimation of low-rank matrices via approximate message passing. *arXiv preprint arXiv:1711.01682*, 2017.
- [35] J. Vila, P. Schniter, and J. Meola. Hyperspectral unmixing via turbo bilinear approximate message passing. *IEEE Transactions on Computational Imaging*, 1(3):143–158, 2015.
- [36] A. K. Fletcher and S. Rangan. Iterative reconstruction of rank-one matrices in noise. *Information and Inference: A Journal of the IMA*, 7(3):531–562, 2018.
- [37] J. T. Parker, P. Schniter, and V. Cevher. Bilinear generalized approximate message passing. part i: Derivation. *IEEE Transactions on Signal Processing*, 62(22):5839–5853, 2014.
- [38] A. Montanari and E. Richard. Non-negative principal component analysis: Message passing algorithms and sharp asymptotics. *IEEE Transactions on Information Theory*, 62(3):1458–1484, 2015.
- [39] J. Baik, G. B. Arous, and S. Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Annals of Probability*, page 1643, 2005.
- [40] S. Péché. The largest eigenvalue of small rank perturbations of hermitian random matrices. *Probability Theory and Related Fields*, 134(1):127–173, 2006.
- [41] D. Féral and S. Péché. The largest eigenvalue of rank one deformation of large wigner matrices. *Communications in mathematical physics*, 272(1):185–228, 2007.
- [42] A. A. Amini and M. J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. *Ann. Statist.*, 37(5B):2877–2921, 10 2009.
- [43] M. Brennan, G. Bresler, and W. Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 48–166. PMLR, 06–09 Jul 2018.
- [44] D. Gamarnik, A. Jagannath, and S. Sen. The overlap gap property in principal submatrix recovery. *arXiv preprint arXiv:1908.09959*, 2019.
- [45] T. Cai, Z. Ma, and Y. Wu. Optimal estimation and rank detection for sparse spiked covariance matrices. *Probability Theory and Related Fields*, 161(3):781–815, Apr 2015.
- [46] R. Krauthgamer, B. Nadler, and D. Vilenchik. Do semidefinite relaxations solve sparse pca up to the information limit? *Ann. Statist.*, 43(3):1300–1322, 06 2015.
- [47] Y. Deshpande and A. Montanari. Sparse pca via covariance thresholding. *Journal of Machine Learning Research*, 17(141):1–41, 2016.
- [48] T. Wang, Q. Berthet, and R. J. Samworth. Statistical and computational trade-offs in estimation of sparse principal components. *Ann. Statist.*, 44(5):1896–1930, 10 2016.
- [49] Q. Berthet and P. Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pages 1046–1066, Princeton, NJ, USA, 12–14 Jun 2013. PMLR.
- [50] T. Ma and A. Wigderson. Sum-of-squares lower bounds for sparse pca. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS’15*, pages 1612–1620, Cambridge, MA, USA, 2015. MIT Press.
- [51] M. Mezard and A. Montanari. *Information, physics and computation*. Oxford University Press, 2009.

- [52] T. Lesieur, F. Krzakala, and L. Zdeborová. Constrained low-rank matrix estimation: phase transitions, approximate message passing and applications. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(7):073403, jul 2017.
- [53] E. Bolthausen. An iterative construction of solutions of the tap equations for the sherrington–kirkpatrick model. *Communications in Mathematical Physics*, 325(1):333–366, 2014.
- [54] S. B. Korada and N. Macris. Exact solution of the gauge symmetric p-spin glass model on a complete graph. *Journal of Statistical Physics*, 136(2):205–230, 2009.
- [55] F. Krzakala, J. Xu, and L. Zdeborová. Mutual information in rank-one matrix estimation. In *2016 IEEE Information Theory Workshop (ITW)*, pages 71–75. IEEE, 2016.
- [56] A. El Alaoui and F. Krzakala. Estimation in the spiked wigner model: a short proof of the replica formula. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 1874–1878. IEEE, 2018.
- [57] J. Barbier, C. Luneau, and N. Macris. Mutual information for low-rank even-order symmetric tensor factorization. *arXiv preprint arXiv:1904.04565*, 2019.
- [58] D. Guo, S. Shamai, and S. Verdú. Mutual information and minimum mean-square error in gaussian channels. *IEEE Trans. on Information Theory*, 51(4):1261–1282, April 2005.
- [59] D. Guo, Y. Wu, S. S. Shitz, and S. Verdú. Estimation in gaussian noise: Properties of the minimum mean-square error. *IEEE Transactions on Information Theory*, 57(4):2371–2385, 2011.
- [60] A. Perry, A. S. Wein, A. S. Bandeira, A. Moitra, et al. Optimality and sub-optimality of pca i: Spiked random matrix models. *The Annals of Statistics*, 46(5):2416–2451, 2018.
- [61] A. E. Alaoui, F. Krzakala, and M. I. Jordan. Finite size corrections and likelihood ratio fluctuations in the spiked wigner model. *arXiv preprint arXiv:1710.02903*, 2017.
- [62] A. E. Alaoui and M. I. Jordan. Detection limits in the high-dimensional spiked rectangular model. In *Conference On Learning Theory, COLT 2018, Stockholm, Sweden, 6-9 July 2018.*, pages 410–438, 2018.
- [63] C. Rush and R. Venkataramanan. Finite sample analysis of approximate message passing algorithms. *IEEE Trans. Information Theory*, 64(11):7264–7286, 2018.
- [64] A. Montanari. Finding one community in a sparse graph. *Journal of Statistical Physics*, 161(2):273–299, 2015.
- [65] P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- [66] A. Javanmard and A. Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *J. Infor. & Inference*, 2:115, 2013.
- [67] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford University Press, 2013.

A General results on the mutual information

In this appendix we give a more general form of theorem 1 in section 2. Our analysis by the adaptive interpolation method works for any regime where the sequences λ_n and ρ_n verify:

$$C \leq \lambda_n \rho_n = O(n^\gamma) \quad \text{for some constants } \gamma \in [0, 1/2) \quad \text{and } C > 0. \quad (15)$$

Of course this contains the regime (2) as a special case. Our general result is a statement on the smallness of

$$\Delta I_n \equiv \frac{1}{\rho_n |\ln \rho_n|} \left| \frac{1}{n} I(\mathbf{X}; \mathbf{W}) - \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda_n, \rho_n) \right|.$$

The analysis of section B leads to the following general theorem.

Theorem 3 (Sparse spiked Wigner model). *Let the sequences λ_n and ρ_n verify (15) and let $\alpha > 0$. There exists a constant $C > 0$ independent of n , such that the mutual information for the Wigner spike model verifies*

$$\Delta I_n \leq \frac{C}{|\ln \rho_n|} \max \left\{ \frac{1}{n^\alpha}, \frac{\lambda_n}{n \rho_n}, \left(\frac{\lambda_n^4}{n^{1-4\alpha} \rho_n^2} (1 + \lambda_n \rho_n^2) \right)^{1/3} \right\}.$$

In particular, choosing $\lambda_n = \Theta(|\ln \rho_n|/\rho_n)$ (which is the appropriate scaling to observe a phase transition),

$$\Delta I_n \leq C \max \left\{ \frac{1}{n^\alpha |\ln \rho_n|}, \frac{1}{n \rho_n^2}, \left(\frac{|\ln \rho_n|}{n^{1-4\alpha} \rho_n^6} \right)^{1/3} \right\}.$$

If, in addition, we set $\rho_n = \Omega(n^{-\beta})$ for $\beta \geq 0$ (which is the regime in (2)), then we have

$$\Delta I_n \leq C \max \left\{ \frac{1}{n^\alpha \ln n}, \frac{1}{n^{1-2\beta}}, \left(\frac{\ln n}{n^{1-4\alpha-6\beta}} \right)^{1/3} \right\}.$$

This bound vanishes as n grows if $\beta \in [0, 1/6)$ and $\alpha \in (0, (1-6\beta)/4]$. The final bound is optimized (up to polylog factors) by setting $\alpha = (1-6\beta)/7$. In this case (again, when $\lambda_n = \Theta(|\ln \rho_n|/\rho_n)$ and $\rho_n = \Omega(n^{-\beta})$),

$$\Delta I_n \leq C \frac{(\ln n)^{1/3}}{n^{(1-6\beta)/7}}.$$

B Information theoretic analysis by the adaptive interpolation method

In this section we provide the essential architecture for the proof of theorem 1 which relies on the adaptive interpolation method [19, 20]. The proof requires concentration properties for “free energies” and “overlaps” which are deferred to appendices C and D. When no confusion is possible we use the notation $\mathbb{E}\|\mathbf{A}\|^2 = \mathbb{E}[\|\mathbf{A}\|^2]$.

B.1 The interpolating model.

Let $\epsilon \in [s_n, 2s_n]$, for a sequence s_n tending to zero as $s_n = n^{-\alpha}/2 \in (0, 1/2)$, for $\alpha > 0$ chosen later. Let $q_n : [0, 1] \times [s_n, 2s_n] \mapsto [0, \rho_n]$ and set

$$R_n(t, \epsilon) \equiv \epsilon + \lambda_n \int_0^t ds q_n(s, \epsilon).$$

Consider the following interpolating estimation model, where $t \in [0, 1]$, with accessible data $(W_{ij}(t))_{i,j}$ and $(\tilde{W}_i(t, \epsilon))_i$ obtained through

$$\begin{cases} W_{ij}(t) = W_{ji}(t) &= \sqrt{(1-t) \frac{\lambda_n}{n}} X_i X_j + Z_{ij}, & 1 \leq i < j \leq n, \\ \tilde{\mathbf{W}}(t, \epsilon) &= \sqrt{R_n(t, \epsilon)} \mathbf{X} + \tilde{\mathbf{Z}}, \end{cases}$$

with standard gaussian noise $\tilde{\mathbf{Z}} \sim \mathcal{N}(0, \mathbf{I}_n)$, and $Z_{ij} = Z_{ji} \sim \mathcal{N}(0, 1)$. The posterior associated with this model reads (here $\| - \|$ is the ℓ_2 norm)

$$\begin{aligned} dP_{n,t,\epsilon}(\mathbf{x} | \mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)) &= \frac{1}{\mathcal{Z}_{n,t,\epsilon}(\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon))} \left(\prod_{i=1}^n dP_{X,n}(x_i) \right) \\ &\times \exp \left\{ \sum_{i < j}^n \left((1-t) \frac{\lambda_n}{n} \frac{x_i^2 x_j^2}{2} - \sqrt{(1-t) \frac{\lambda_n}{n}} x_i x_j W_{ij}(t) \right) + R_n(t, \epsilon) \frac{\|\mathbf{x}\|^2}{2} \right. \\ &\quad \left. - \sqrt{R_n(t, \epsilon)} \mathbf{x} \cdot \tilde{\mathbf{W}}(t, \epsilon) \right\}. \end{aligned}$$

The normalization factor $\mathcal{Z}_{n,t,\epsilon}(\dots)$ is also called partition function. We also define the mutual information density for the interpolating model

$$i_n(t, \epsilon) \equiv \frac{1}{n} I(\mathbf{X}; (\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon))). \quad (16)$$

The (n, t, ϵ, R_n) -dependent Gibbs-bracket (that we simply denote $\langle - \rangle_t$ for the sake of readability) is defined for functions $A(\mathbf{x}) = A$

$$\langle A(\mathbf{x}) \rangle_t = \int dP_{n,t,\epsilon}(\mathbf{x} | \mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)) A(\mathbf{x}). \quad (17)$$

Lemma 1 (Boundary values). *The mutual information for the interpolating model verifies*

$$\begin{cases} i_n(0, \epsilon) = \frac{1}{n} I(\mathbf{X}; \mathbf{W}) + O(\rho_n s_n), \\ i_n(1, \epsilon) = I_n(X; \{\lambda_n \int_0^1 dt q_n(t, \epsilon)\}^{1/2} X + Z) + O(\rho_n s_n). \end{cases} \quad (18)$$

where $I_n(X; \{\lambda_n \int_0^1 dt q_n(t, \epsilon)\}^{1/2} X + Z)$ is the mutual information for a scalar gaussian channel with input $X \sim P_{X,n}$ and noise $Z \sim \mathcal{N}(0, 1)$.

Proof. We start with the chain rule for mutual information:

$$i_n(0, \epsilon) = \frac{1}{n} I(\mathbf{X}; \mathbf{W}(0)) + \frac{1}{n} I(\mathbf{X}; \tilde{\mathbf{W}}(0, \epsilon) | \mathbf{W}(0)).$$

Note that, by the definition of $\mathbf{W}(t)$,

$$I(\mathbf{X}; \mathbf{W}(0)) = I(\mathbf{X}; \mathbf{W}).$$

Moreover we claim $\frac{1}{n} I(\mathbf{X}; \tilde{\mathbf{W}}(0, \epsilon) | \mathbf{W}(0)) = O(\rho_n s_n)$, which yields the first identity in (18). This claim simply follows from the I-MMSE relation (appendix I) and $R_n(0, \epsilon) = \epsilon$:

$$\frac{d}{d\epsilon} \frac{1}{n} I(\mathbf{X}; \tilde{\mathbf{W}}(0, \epsilon) | \mathbf{W}(0)) = \frac{1}{2n} \text{MMSE}(\mathbf{X} | \tilde{\mathbf{W}}(0, \epsilon), \mathbf{W}(0)) \leq \frac{\rho_n}{2}. \quad (19)$$

The last inequality above is true because $\text{MMSE}(\mathbf{X} | \tilde{\mathbf{W}}(0, \epsilon), \mathbf{W}(0)) \leq \mathbb{E} \|\mathbf{X} - \mathbb{E} \mathbf{X}\|^2 = n \text{Var}(X_1) \leq n \rho_n$, as the components of \mathbf{X} are i.i.d. from $P_{X,n}$. Therefore $\frac{1}{n} I(\mathbf{X}; \tilde{\mathbf{W}}(0, \epsilon) | \mathbf{W}(0))$ is $\frac{\rho_n}{2}$ -Lipschitz in $\epsilon \in [s_n, 2s_n]$. Moreover, we have that $I(\mathbf{X}; \tilde{\mathbf{W}}(0, 0) | \mathbf{W}(0)) = 0$. This implies the claim.

The proof of the second identity in (18) again starts from the chain rule for mutual information

$$i_n(1, \epsilon) = \frac{1}{n} I(\mathbf{X}; \tilde{\mathbf{W}}(1, \epsilon)) + \frac{1}{n} I(\mathbf{X}; \mathbf{W}(1) | \tilde{\mathbf{W}}(1, \epsilon)).$$

Note that $I(\mathbf{X}; \mathbf{W}(1) | \tilde{\mathbf{W}}(1, \epsilon)) = 0$ as $\mathbf{W}(1)$ does not depend on \mathbf{X} . Moreover,

$$\begin{aligned} \frac{1}{n} I(\mathbf{X}; \tilde{\mathbf{W}}(1, \epsilon)) &= I_n(X; \sqrt{R_n(1, \epsilon)} X + Z) \\ &= I_n(X; \{\lambda_n \int_0^1 dt q_n(t, \epsilon)\}^{1/2} X + Z) + O(\rho_n s_n). \end{aligned}$$

because $I_n(X; \sqrt{\gamma} X + Z)$ is a $\frac{\rho_n}{2}$ -Lipschitz function of γ , by an application of the I-MMSE relation (appendix I) $\frac{d}{d\gamma} I_n(X; \sqrt{\gamma} X + Z) = \text{MMSE}(X | \sqrt{\gamma} X + Z) / 2 \leq \text{Var}(X) / 2 \leq \rho_n / 2$. \square

B.2 Fundamental sum rule.

Proposition 1 (Sum rule). *The mutual information verifies the following sum rule:*

$$\frac{1}{n} I(\mathbf{X}; \mathbf{W}) = i_n^{\text{pot}}(\int_0^1 dt q_n(t, \epsilon); \lambda_n, \rho_n) + \frac{\lambda_n}{4} (\mathcal{R}_1 - \mathcal{R}_2 - \mathcal{R}_3) + O(\rho_n s_n) + O\left(\frac{\lambda_n}{n}\right) \quad (20)$$

with non-negative “remainders” that depend on (n, ϵ, R_n) ,

$$\begin{cases} \mathcal{R}_1 \equiv \int_0^1 dt (q_n(t, \epsilon) - \int_0^1 ds q_n(s, \epsilon))^2, \\ \mathcal{R}_2 \equiv \int_0^1 dt \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle_t)^2 \rangle_t, \\ \mathcal{R}_3 \equiv \int_0^1 dt (q_n(t, \epsilon) - \mathbb{E} \langle Q \rangle_t)^2, \end{cases} \quad (21)$$

where $Q = \frac{1}{n} \mathbf{x} \cdot \mathbf{X}$ is called the overlap. The constants in the $O(\dots)$ terms are independent of n, t, ϵ .

Proof. By the fundamental theorem of calculus $i_n(0, \epsilon) = i_n(1, \epsilon) - \int_0^1 dt \frac{d}{dt} i_n(t, \epsilon)$. Note that $i_n(0, \epsilon)$ and $i_n(1, \epsilon)$ are given by (18). The t -derivative of the interpolating mutual information is simply computed combining the I-MMSE relation with the chain rule for derivatives

$$\frac{d}{dt} i_n(t, \epsilon) = -\frac{\lambda_n}{2} \frac{1}{n^2} \sum_{i < j} \mathbb{E} [(X_i X_j - \langle x_i x_j \rangle_t)^2] + \frac{\lambda_n q_n(t, \epsilon)}{2} \frac{1}{n} \mathbb{E} \|\mathbf{X} - \langle \mathbf{x} \rangle_t\|^2 \quad (22)$$

$$= -\frac{\lambda_n}{4} \frac{1}{n^2} \mathbb{E} \|\mathbf{X} \otimes \mathbf{X} - \langle \mathbf{x} \otimes \mathbf{x} \rangle_t\|_{\mathbb{F}}^2 + \frac{\lambda_n q_n(t, \epsilon)}{2} \frac{1}{n} \mathbb{E} \|\mathbf{X} - \langle \mathbf{x} \rangle_t\|^2 + O\left(\frac{\lambda_n}{n}\right). \quad (23)$$

The correction term in (23) comes from completing the diagonal terms in the sum $\sum_{i<j}$ in order to construct the matrix-MMSE for $\mathbf{X} \otimes \mathbf{X}$, namely the first term on the r.h.s. of (23). This expression can be simplified by application of the Nishimori identities (appendix H contains a proof of these general identities). Starting with the second term (a vector-MMSE)

$$\begin{aligned} \frac{1}{n} \mathbb{E} \|\mathbf{X} - \langle \mathbf{x} \rangle_t\|^2 &= \mathbb{E} [\|\mathbf{X}\|^2 + \|\langle \mathbf{x} \rangle_t\|^2 - 2\mathbf{X} \cdot \langle \mathbf{x} \rangle_t] \\ &= \frac{1}{n} \mathbb{E} [\|\mathbf{X}\|^2 - \mathbf{X} \cdot \langle \mathbf{x} \rangle_t] = \rho_n - \mathbb{E} \langle Q \rangle_t, \end{aligned} \quad (24)$$

where we used $\mathbb{E} \|\mathbf{X}\|^2 = n\rho_n$ and the Nishimori identity $\mathbb{E} \|\langle \mathbf{x} \rangle_t\|^2 = \mathbb{E} [\mathbf{X} \cdot \langle \mathbf{x} \rangle_t]$. By similar manipulations we obtain for the matrix-MMSE

$$\frac{1}{n^2} \text{MMSE}(\mathbf{X} \otimes \mathbf{X} | \tilde{\mathbf{W}}(t, \epsilon), \mathbf{W}(t)) = \frac{1}{n^2} \mathbb{E} \|\mathbf{X} \otimes \mathbf{X} - \langle \mathbf{x} \otimes \mathbf{x} \rangle_t\|_{\mathbb{F}}^2 = \rho_n^2 - \mathbb{E} \langle Q^2 \rangle_t. \quad (25)$$

From (18), (23), (24), (25) and the fundamental theorem of calculus we deduce

$$\begin{aligned} \frac{1}{n} I(\mathbf{X}; \mathbf{W}) &= I_n(X; \{\lambda_n \int_0^1 dt q_n(t, \epsilon)\}^{1/2} X + Z) \\ &\quad + \frac{\lambda_n}{4} \int_0^1 dt \left\{ \rho_n^2 - \mathbb{E} \langle Q^2 \rangle_t - 2q_n(t, \epsilon)(\rho_n - \mathbb{E} \langle Q \rangle_t) \right\} + O(\rho_n s_n) + O\left(\frac{\lambda_n}{n}\right). \end{aligned}$$

The terms on the r.h.s can be re-arranged so that the potential (4) appears, and this gives immediately the sum rule (20). \square

Theorem 1 follows from the upper and lower bounds proven below, and applied for $s_n = \frac{1}{2}n^{-\alpha}$.

B.3 Upper bound: linear interpolation path.

Proposition 2 (Upper bound). *We have*

$$\frac{1}{n} I(\mathbf{X}; \mathbf{W}) \leq \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda_n, \rho_n) + O(\rho_n s_n) + O\left(\frac{\lambda_n}{n}\right).$$

Proof. Fix $q_n(t, \epsilon) = q_n \in [0, \rho_n]$ a constant independent of ϵ, t . The interpolation path $R_n(t, \epsilon)$ is therefore a simple linear function of time. From (21) \mathcal{R}_1 cancels and since \mathcal{R}_2 and \mathcal{R}_3 are non-negative we get from Proposition (1)

$$\frac{1}{n} I(\mathbf{X}; \mathbf{W}) \leq i_n^{\text{pot}}(q, \lambda_n, \rho_n) + O(\rho_n s_n) + O\left(\frac{\lambda_n}{n}\right).$$

Note that the error terms $O(\dots)$ are bounded independently of q_n . Therefore optimizing the r.h.s over the free parameter $q_n \in [0, \rho_n]$ yields the upper bound. \square

B.4 Lower bound: adaptive interpolation path.

We start with a definition: the map $\epsilon \mapsto R_n(t, \epsilon)$ is called *regular* if it is a \mathcal{C}^1 -diffeomorphism whose jacobian is greater or equal to one for all $t \in [0, 1]$.

Proposition 3 (Lower bound). *Consider sequences λ_n and ρ_n satisfying $c_1 \leq \lambda_n \rho_n \leq c_2 n^\gamma$ for some constants positive constant c_1, c_2 and $\gamma \in [0, 1/2[$. Then*

$$\frac{1}{n} I(\mathbf{X}; \mathbf{W}) \geq \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda_n, \rho_n) + O(\rho_n s_n) + O\left(\frac{\lambda_n}{n}\right) + O\left(\left(\frac{\lambda_n \rho_n}{n s_n^4}\right)^{1/3}\right). \quad (26)$$

Proof. First note that the regime (2) for the sequences λ_n, ρ_n satisfies the more general condition assumed in this lemma (this is the condition in theorem 3 of appendix A). Assume for the moment that the map $\epsilon \mapsto R_n(t, \epsilon)$ is regular. Then, based on Proposition 7 and identity (38) (appendix

D), we have a bound on the overlap fluctuation. Namely, for some numerical constant $C \geq 0$ independent of n

$$\begin{aligned} \frac{\lambda_n}{s_n} \int_{s_n}^{2s_n} d\epsilon \mathcal{R}_2 &= \frac{\lambda_n}{s_n} \int_{s_n}^{2s_n} d\epsilon \int_0^1 dt \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle_{n,t,R_n(t,\epsilon)})^2 \rangle_{n,t,R_n(t,\epsilon)} \\ &\leq C \left(\frac{\lambda_n^4 \rho_n}{n s_n^4} \right)^{1/3}. \end{aligned} \quad (27)$$

Using this concentration result, and $\mathcal{R}_1 \geq 0$, and averaging the sum rule (20) over $\epsilon \in [s_n, 2s_n]$ (recall the error terms are independent of ϵ) we find

$$\begin{aligned} \frac{I(\mathbf{X}; \mathbf{W})}{n} &\geq \frac{1}{s_n} \int_{s_n}^{2s_n} d\epsilon i_n^{\text{pot}} \left(\int_0^1 dt q_n(t, \epsilon), \lambda_n, \rho_n \right) - \frac{\lambda_n}{4} \frac{1}{s_n} \int_{s_n}^{2s_n} d\epsilon \int_0^1 dt (q_n(t, \epsilon) - \mathbb{E} \langle Q \rangle_t)^2 \\ &\quad + O(\rho_n s_n) + O\left(\frac{\lambda_n}{n}\right) + O\left(\left(\frac{\lambda_n^4 \rho_n}{n s_n^4}\right)^{1/3}\right). \end{aligned} \quad (28)$$

At this stage it is natural to see if we can choose $q_n(t, \epsilon)$ to be the solution of $q_n(t, \epsilon) = \mathbb{E} \langle Q \rangle_t$. Setting $F_n(t, R_n(t, \epsilon)) \equiv \mathbb{E} \langle Q \rangle_{n,t,R_n(t,\epsilon)}$, we recognize a *first order ordinary differential equation*

$$\frac{d}{dt} R_n(t, \epsilon) = F_n(t, R_n(t, \epsilon)) \quad \text{with initial condition} \quad R_n(0, \epsilon) = \epsilon. \quad (29)$$

As $F_n(t, R_n(t, \epsilon))$ is C^1 with bounded derivative w.r.t. its second argument the Cauchy-Lipschitz theorem implies that (29) admits a unique global solution $R_n^*(t, \epsilon) = \epsilon + \int_0^t ds q_n^*(s, \epsilon)$, where $q_n^* : [0, 1] \times [s_n, 2s_n] \mapsto [0, \rho_n]$. Note that any solution must satisfy $q_n^*(t, \epsilon) \in [0, \rho_n]$ because $\mathbb{E} \langle Q \rangle_{n,t,\epsilon} \in [0, \rho_n]$ as can be seen from a Nishimori identity (appendix H) and (24).

We check that R_n^* is regular. By Liouville's formula the jacobian of the flow $\epsilon \mapsto R_n^*(t, \epsilon)$ satisfies

$$\frac{d}{d\epsilon} R_n^*(t, \epsilon) = \exp \left\{ \int_0^t ds \frac{d}{dR} F_n(s, R) \Big|_{R=R_n^*(s,\epsilon)} \right\}.$$

Applying repeatedly the Nishimori identity of Lemma 7 (appendix H) one obtains (this computation does not present any difficulty and can be found in section 6 of [19])

$$\frac{d}{dR} F_n(s, R) = \frac{1}{n} \sum_{i,j=1}^n \mathbb{E} \left[(\langle x_i x_j \rangle_{n,s,R} - \langle x_i \rangle_{n,s,R} \langle x_j \rangle_{n,s,R})^2 \right] \geq 0 \quad (30)$$

so that the flow has a jacobian greater or equal to one. In particular it is locally invertible (surjective). Moreover it is injective because of the unicity of the solution of the differential equation, and therefore it is a C^1 -diffeomorphism. Thus $\epsilon \mapsto R_n^*(t, \epsilon)$ is regular. With the choice R_n^* , i.e., by suitably *adapting* the interpolation path, we cancel \mathcal{R}_3 . This yields

$$\begin{aligned} \frac{1}{n} I(\mathbf{X}; \mathbf{W}) &\geq \frac{1}{s_n} \int_{s_n}^{2s_n} d\epsilon i_n^{\text{pot}} \left(\int_0^1 dt q_n^*(t, \epsilon), \lambda_n, \rho_n \right) + O(\dots) \\ &\geq \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, \lambda_n, \rho_n) + O(\dots) \end{aligned}$$

where the $O(\dots)$ is a shorthand notation for the three error terms in (28). This the desired result. \square

C Concentration of free energy

For this appendix it is convenient to use the language of statistical mechanics.

C.1 Statistical mechanics notations

We express the posterior of the interpolating model

$$\begin{aligned} dP_{n,t,\epsilon}(\mathbf{x} | \mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)) &= \frac{1}{\mathcal{Z}_{n,t,\epsilon}(\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon))} \\ &\quad \times \left(\prod_{i=1}^n dP_{X,n}(x_i) \right) \exp \left\{ -\mathcal{H}_{n,t,\epsilon}(\mathbf{x}, \mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)) \right\} \end{aligned} \quad (31)$$

with normalization constant (partition function) $\mathcal{Z}_{n,t,\epsilon}$ and “hamiltonian”

$$\mathcal{H}_{n,t,\epsilon}(\mathbf{x}, \mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)) = \mathcal{H}_{n,t,\epsilon}(\mathbf{x}, \mathbf{X}, \mathbf{Z}, \tilde{\mathbf{Z}}) \quad (32)$$

$$\begin{aligned} &\equiv \sum_{i < j}^n \left((1-t) \frac{\lambda_n}{n} \frac{x_i^2 x_j^2}{2} - \sqrt{(1-t) \frac{\lambda_n}{n}} x_i x_j W_{ij}(t) \right) + R_n(t, \epsilon) \frac{\|\mathbf{x}\|^2}{2} - \sqrt{R_n(t, \epsilon)} \mathbf{x} \cdot \tilde{\mathbf{W}}(t, \epsilon) \\ &= (1-t) \lambda_n \sum_{i < j}^n \left(\frac{x_i^2 x_j^2}{2n} - \frac{x_i x_j X_i X_j}{n} - \frac{x_i x_j Z_{ij}}{\sqrt{n(1-t)\lambda_n}} \right) \\ &\quad + R_n(t, \epsilon) \left(\frac{\|\mathbf{x}\|^2}{2} - \mathbf{x} \cdot \mathbf{X} - \frac{\mathbf{x} \cdot \tilde{\mathbf{Z}}}{\sqrt{R_n(t, \epsilon)}} \right). \end{aligned} \quad (33)$$

It will also be convenient to work with “free energies” rather than mutual informations. The free energy $F_n(t, \epsilon)$ and (its expectation $f_n(t, \epsilon)$) for the interpolating model is simply minus the (expected) log-partition function:

$$F_{n,t,\epsilon}(\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)) \equiv -\frac{1}{n} \ln \mathcal{Z}_{n,t,\epsilon}(\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)), \quad (34)$$

$$f_n(t, \epsilon) \equiv \mathbb{E} F_{n,t,\epsilon}(\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)). \quad (35)$$

The expectation \mathbb{E} carries over the data. The averaged free energy is related to the mutual information $i_n(t, \epsilon)$ given by (16) through

$$i_n(t, \epsilon) = f_n(t, \epsilon) + \frac{n-1}{n} \frac{\rho^2 \lambda (1-t)}{4} + \frac{\rho R_n(t, \epsilon)}{2}. \quad (36)$$

C.2 Free energy concentration

In this section we prove a concentration identity for the free energy (34) onto its average (35).

Proposition 4 (Free energy concentration for the spiked Wigner model). *We have*

$$\mathbb{E} \left[\left(F_{n,t,\epsilon}(\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon)) - f_n(t, \epsilon) \right)^2 \right] \leq \frac{2\rho_n S^2}{n} \left((2s_n + \lambda_n \rho_n)^2 + S^4 \right) + \frac{3}{2} \frac{\lambda_n \rho_n^2}{n} + 2 \frac{s_n \rho_n}{n}.$$

Considering sequences λ_n and ρ_n verifying (15) and with $s_n = (1/2)n^{-\alpha} \rightarrow 0_+$ the bound simplifies to $C(S)\lambda_n^2 \rho_n^3 / n$ with positive constant $C(S) \leq \frac{5}{2} + 8S^2 + 2S^6$.

The proof is based on two classical concentration inequalities,

Proposition 5 (Gaussian Poincaré inequality). *Let $\mathbf{U} = (U_1, \dots, U_N)$ be a vector of N independent standard normal random variables. Let $g : \mathbb{R}^N \rightarrow \mathbb{R}$ be a continuously differentiable function. Then*

$$\text{Var}(g(\mathbf{U})) \leq \mathbb{E} \|\nabla g(\mathbf{U})\|^2.$$

Proposition 6 (Efron-Stein inequality). *Let $\mathcal{U} \subset \mathbb{R}$, and a function $g : \mathcal{U}^N \rightarrow \mathbb{R}$. Let $\mathbf{U} = (U_1, \dots, U_N)$ be a vector of N independent random variables with law P_U that take values in \mathcal{U} . Let $\mathbf{U}^{(i)}$ a vector which differs from \mathbf{U} only by its i -th component, which is replaced by U'_i drawn from P_U independently of \mathbf{U} . Then*

$$\text{Var}(g(\mathbf{U})) \leq \frac{1}{2} \sum_{i=1}^N \mathbb{E}_{\mathbf{U}} \mathbb{E}_{U'_i} [(g(\mathbf{U}) - g(\mathbf{U}^{(i)}))^2].$$

We start by proving the concentration w.r.t. the gaussian variables. It is convenient to make explicit the dependence of the partition function of the interpolating model in the independent quenched variables instead of the data: $\mathcal{Z}_{n,t,\epsilon}(\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{Z}}) = \mathcal{Z}_{n,t,\epsilon}(\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon))$.

Lemma 2 (Concentration w.r.t. the gaussian variables). *We have*

$$\mathbb{E} \left[\left(\frac{1}{n} \ln \mathcal{Z}_{n,t,\epsilon}(\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{Z}}) - \frac{1}{n} \mathbb{E}_{\mathbf{Z}, \tilde{\mathbf{Z}}} \ln \mathcal{Z}_{n,t,\epsilon}(\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{Z}}) \right)^2 \right] \leq \frac{3}{2} \frac{\lambda_n \rho_n^2}{n} + 2 \frac{s_n \rho_n}{n}.$$

Proof. Fix all variables except $\mathbf{Z}, \tilde{\mathbf{Z}}$. Let $g(\mathbf{Z}, \tilde{\mathbf{Z}}) \equiv -\frac{1}{n} \ln \mathcal{Z}_{n,t,\epsilon}(\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{Z}})$ be the free energy seen as a function of the gaussian variables only. The free energy gradient reads $\mathbb{E}\|\nabla g\|^2 = \mathbb{E}\|\nabla_{\mathbf{Z}} g\|^2 + \mathbb{E}\|\nabla_{\tilde{\mathbf{Z}}} g\|^2$. Let us denote $\mathcal{H}(t) \equiv \mathcal{H}_{n,t,\epsilon}$ the interpolating Hamiltonian (32).

$$\begin{aligned} \mathbb{E}\|\nabla_{\mathbf{Z}} g\|^2 &= \frac{1}{n^2} \mathbb{E}\|\langle \nabla_{\mathbf{Z}} \mathcal{H}(t) \rangle_t\|^2 = \frac{(1-t)\lambda_n}{n^3} \sum_{i<j} \mathbb{E}[\langle x_i x_j \rangle_t^2] \leq \frac{(1-t)\lambda_n}{n^3} \sum_{i<j} \mathbb{E}[\langle x_i x_j \rangle_t^2] \\ &\stackrel{\text{N}}{=} \frac{(1-t)\lambda_n}{n^3} \sum_{i<j} \mathbb{E}[(X_i X_j)^2] \leq \frac{\lambda_n \rho_n^2}{2n} \end{aligned}$$

where we used a Nishimori identity for the last equality. Similarly, and using $\lambda_n \rho_n \geq 1$ and $s_n < 1/2$,

$$\mathbb{E}\|\nabla_{\tilde{\mathbf{Z}}} g\|^2 = \frac{R(\epsilon)}{n^2} \mathbb{E}\|\langle \mathbf{x} \rangle_t\|^2 \leq \frac{R(\epsilon)}{n^2} \mathbb{E}\|\langle \|\mathbf{x}\|^2 \rangle_t \stackrel{\text{N}}{=} \frac{R(\epsilon)}{n^2} \mathbb{E}\|\mathbf{X}\|^2 \leq \frac{(2s_n + \rho_n \lambda_n) \rho_n}{n}.$$

Therefore Proposition 5 directly implies the stated result. \square

We now consider the fluctuations due to the signal realization:

Lemma 3 (Concentration w.r.t. the spike). *We have*

$$\mathbb{E}\left[\left(-\frac{1}{n} \mathbb{E}_{\mathbf{Z}, \tilde{\mathbf{Z}}} \ln \mathcal{Z}_{n,t,\epsilon}(\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{Z}}) - f_n(t, \epsilon)\right)^2\right] \leq \frac{2\rho_n S^2}{n} \left((2s_n + \lambda_n \rho_n)^2 + S^4\right).$$

Proof. Let $g(\mathbf{X}) \equiv -\frac{1}{n} \mathbb{E}_{\mathbf{Z}, \tilde{\mathbf{Z}}} \ln \mathcal{Z}_{n,t,\epsilon}(\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{Z}})$. Define $\mathbf{X}^{(i)}$ as a vector with same entries as \mathbf{X} except the i -th one that is replaced by X'_i drawn independently from $P_{X,n}$. Let us estimate $(g(\mathbf{X}) - g(\mathbf{X}^{(i)}))^2$ by interpolation. Let $\mathcal{H}(t, s\mathbf{X} + (1-s)\mathbf{X}^{(i)})$ be the interpolating Hamiltonian (32) with \mathbf{X} replaced by $s\mathbf{X} + (1-s)\mathbf{X}^{(i)}$. Then

$$\begin{aligned} \mathbb{E}[(g(\mathbf{X}) - g(\mathbf{X}^{(i)}))^2] &= \mathbb{E}\left[\left(\int_0^1 ds \frac{dg}{ds}(s\mathbf{X} + (1-s)\mathbf{X}^{(i)})\right)^2\right] \\ &= \frac{1}{n^2} \mathbb{E}\left[\left(\int_0^1 ds \left\langle \frac{d\mathcal{H}}{ds}(t, s\mathbf{X} + (1-s)\mathbf{X}^{(i)}) \right\rangle_t\right)^2\right] \\ &= \frac{1}{n^2} \mathbb{E}\left[\left(\left(X_i - X'_i\right) \left\langle R_\epsilon(t) x_i + \frac{1-t}{n} x_i \sum_{j(\neq i)} X_j x_j \right\rangle_t\right)^2\right] \\ &\leq \frac{2}{n^2} \mathbb{E}\left[\left(X_i - X'_i\right)^2 \left(\langle x_i \rangle_t^2 (2s_n + \lambda_n \rho_n)^2 + \frac{1}{n^2} \sum_{j,k(\neq i)} X_j X_k \langle x_i x_j \rangle_t \langle x_i x_k \rangle_t\right)\right] \\ &\leq \frac{2}{n^2} \mathbb{E}[(X_i - X'_i)^2] \left(S^2 (2s_n + \lambda_n \rho_n)^2 + S^6\right) \\ &\leq \frac{4\rho_n S^2}{n^2} \left((2s_n + \lambda_n \rho_n)^2 + S^4\right). \end{aligned}$$

We used $(a+b)^2 \leq 2(a^2 + b^2)$ for the second inequality and $\mathbb{E}[(X_i - X'_i)^2] = 2\text{Var}(X_i) \leq 2\rho_n$. Therefore Proposition 6 implies the claim. \square

D Overlap concentration: proof of inequality (27)

The derivations below will apply for any $t \in [0, 1]$ so we drop all un-necessary notations and indices. Only the dependence of the free energies in $R(\epsilon) \equiv R_n(t, \epsilon)$ matters, so we denote $F(R(\epsilon)) \equiv F_{n,t,\epsilon}(\mathbf{W}(t), \tilde{\mathbf{W}}(t, \epsilon))$ and $f(R(\epsilon)) \equiv f_n(t, \epsilon)$.

Let \mathcal{L} be the $R(\epsilon)$ -derivative of the Hamiltonian (32) divided by n :

$$\mathcal{L}(\mathbf{x}, \mathbf{X}, \tilde{\mathbf{Z}}) = \mathcal{L} \equiv \frac{1}{n} \frac{d\mathcal{H}_{n,t,\epsilon}}{dR(\epsilon)} = \frac{1}{n} \left(\frac{\|\mathbf{x}\|^2}{2} - \mathbf{x} \cdot \mathbf{X} - \frac{\mathbf{x} \cdot \tilde{\mathbf{Z}}}{2\sqrt{R(\epsilon)}} \right). \quad (37)$$

The overlap fluctuations are upper bounded by those of \mathcal{L} , which are easier to control, as

$$\mathbb{E}\langle(Q - \mathbb{E}\langle Q \rangle_t)^2\rangle_t \leq 4 \mathbb{E}\langle(\mathcal{L} - \mathbb{E}\langle \mathcal{L} \rangle_t)^2\rangle_t. \quad (38)$$

The bracket is again the expectation w.r.t. the posterior of the interpolating model (17). A detailed derivation of this inequality can be found in appendix E and involves only elementary algebra using the Nishimori identity and integrations by parts w.r.t. the gaussian noise $\tilde{\mathbf{Z}}$.

We have the following identities: for any given realisation of the quenched disorder

$$\frac{dF}{dR(\epsilon)} = \langle \mathcal{L} \rangle_t, \quad (39)$$

$$\frac{1}{n} \frac{d^2 F}{dR(\epsilon)^2} = -\langle (\mathcal{L} - \langle \mathcal{L} \rangle_t)^2 \rangle_t + \frac{1}{4n^2 R(\epsilon)^{3/2}} \langle \mathbf{x} \rangle_t \cdot \tilde{\mathbf{Z}}. \quad (40)$$

The gaussian integration by part formula (59) with hamiltonian (32) yields

$$\frac{\mathbb{E}\langle \tilde{\mathbf{Z}} \cdot \mathbf{x} \rangle_t}{\sqrt{R(\epsilon)}} = \mathbb{E}\langle \|\mathbf{x}\|^2 \rangle_t - \mathbb{E}\|\langle \mathbf{x} \rangle_t\|^2 \stackrel{\text{N}}{=} \mathbb{E}\langle \|\mathbf{x}\|^2 \rangle_t - \mathbb{E}\langle \mathbf{X} \cdot \mathbf{x} \rangle_t = \mathbb{E}\langle \|\mathbf{x}\|^2 \rangle_t - n \mathbb{E}\langle Q \rangle_t. \quad (41)$$

Therefore averaging (39) and (40) we find

$$\frac{df}{dR(\epsilon)} = \mathbb{E}\langle \mathcal{L} \rangle_t \stackrel{\text{N}}{=} -\frac{1}{2} \mathbb{E}\langle Q \rangle_t, \quad (42)$$

$$\frac{1}{n} \frac{d^2 f}{dR(\epsilon)^2} = -\mathbb{E}\langle (\mathcal{L} - \langle \mathcal{L} \rangle_t)^2 \rangle_t + \frac{1}{4n^2 R(\epsilon)} \mathbb{E}\langle \|\mathbf{x} - \langle \mathbf{x} \rangle_t\|^2 \rangle_t. \quad (43)$$

We always work under the assumption that the map $\epsilon \in [s_n, 2s_n] \mapsto R(\epsilon) \in [R(s_n), R(2s_n)]$ is regular, and do not repeat this assumption in the statements below. The concentration inequality (27) is a direct consequence of the following result (combined with Fubini's theorem):

Proposition 7 (Total fluctuations of \mathcal{L}). *Let the sequences λ_n and ρ_n verify (15). Then*

$$\int_{s_n}^{2s_n} d\epsilon \mathbb{E}\langle (\mathcal{L} - \mathbb{E}\langle \mathcal{L} \rangle_t)^2 \rangle_t \leq C \left(\frac{\lambda_n \rho_n}{n s_n} (1 + \lambda_n \rho_n^2) \right)^{1/3}$$

for a constant $C > 0$ that is independent of n , as long as the r.h.s. is $\omega(1/n)$.

The proof of this proposition is broken in two parts, using the decomposition

$$\mathbb{E}\langle (\mathcal{L} - \mathbb{E}\langle \mathcal{L} \rangle_t)^2 \rangle_t = \mathbb{E}\langle (\mathcal{L} - \langle \mathcal{L} \rangle_t)^2 \rangle_t + \mathbb{E}[\langle (\mathcal{L} \rangle_t - \mathbb{E}\langle \mathcal{L} \rangle_t)^2].$$

Thus it suffices to prove the two following lemmas. The first lemma expresses concentration w.r.t. the posterior distribution (or ‘‘thermal fluctuations’’) and is a direct consequence of concavity properties of the average free energy and the Nishimori identity.

Lemma 4 (Thermal fluctuations of \mathcal{L}). *We have*

$$\int_{s_n}^{2s_n} d\epsilon \mathbb{E}\langle (\mathcal{L} - \langle \mathcal{L} \rangle_t)^2 \rangle_t \leq \frac{\rho_n}{n} \left(1 + \frac{\ln 2}{4} \right).$$

Proof. We emphasize again that the interpolating free energy (16) is here viewed as a function of $R(\epsilon)$. In the argument that follows we consider derivatives of this function w.r.t. $R(\epsilon)$. By (43)

$$\begin{aligned} \mathbb{E}\langle (\mathcal{L} - \langle \mathcal{L} \rangle_t)^2 \rangle_t &= -\frac{1}{n} \frac{d^2 f}{dR(\epsilon)^2} + \frac{1}{4n^2 R(\epsilon)} (\mathbb{E}\langle \|\mathbf{x}\|^2 \rangle_t - \mathbb{E}\|\langle \mathbf{x} \rangle_t\|^2) \\ &\leq -\frac{1}{n} \frac{d^2 f}{dR(\epsilon)^2} + \frac{\rho_n}{4n\epsilon}, \end{aligned} \quad (44)$$

where we used $R(\epsilon) \geq \epsilon$ and $\frac{1}{n} \mathbb{E}\langle \|\mathbf{x}\|^2 \rangle_t \stackrel{\text{N}}{=} \mathbb{E}[X_1^2] = \rho_n$. We integrate this inequality over $\epsilon \in [s_n, 2s_n]$. Recall the map $\epsilon \mapsto R(\epsilon)$ has a Jacobian ≥ 1 , is \mathcal{C}^1 and has a well defined \mathcal{C}^1 inverse

since we have assumed that it is regular. Thus integrating (44) and performing a change of variable (to get the second inequality) we obtain

$$\begin{aligned} \int_{s_n}^{2s_n} d\epsilon \mathbb{E} \langle (\mathcal{L} - \langle \mathcal{L} \rangle_t)^2 \rangle_t &\leq -\frac{1}{n} \int_{s_n}^{2s_n} d\epsilon \frac{d^2 f}{dR(\epsilon)^2} + \frac{\rho_n}{4n} \int_{s_n}^{2s_n} \frac{d\epsilon}{\epsilon} \\ &\leq -\frac{1}{n} \int_{R(s_n)}^{R(2s_n)} dR(\epsilon) \frac{d^2 f}{dR(\epsilon)^2} + \frac{\rho_n}{4n} \int_{s_n}^{2s_n} \frac{d\epsilon}{\epsilon} \\ &= \frac{1}{n} \left(\frac{df}{dR(\epsilon)}(R(s_n)) - \frac{df}{dR(\epsilon)}(R(2s_n)) \right) + \frac{\rho_n}{4n} \ln 2. \end{aligned}$$

We have $|f'(R(\epsilon))| = |\mathbb{E} \langle Q \rangle_t / 2| \leq \rho_n / 2$ so the first term is certainly smaller in absolute value than ρ_n / n . This concludes the proof of Lemma 4. \square

The second lemma expresses the concentration w.r.t. the quenched disorder variables and is a consequence of the concentration of the free energy onto its average (w.r.t. the quenched variables).

Lemma 5 (Quenched fluctuations of \mathcal{L}). *Let the sequences λ_n and ρ_n verify (15). Then*

$$\int_{s_n}^{2s_n} d\epsilon \mathbb{E} [(\langle \mathcal{L} \rangle_t - \mathbb{E} \langle \mathcal{L} \rangle_t)^2] \leq C \left(\frac{\lambda_n \rho_n}{n s_n} (1 + \lambda_n \rho_n^2) \right)^{1/3}$$

for a constant $C > 0$ that is independent of n , as long as the r.h.s. is $\omega(1/n)$.

Proof. Consider the following functions of $R(\epsilon)$:

$$\begin{aligned} \tilde{F}(R(\epsilon)) &\equiv F(R(\epsilon)) + S \frac{\sqrt{R(\epsilon)}}{n} \sum_{i=1}^n |\tilde{Z}_i|, \\ \tilde{f}(R(\epsilon)) &\equiv \mathbb{E} \tilde{F}(R(\epsilon)) = f(R(\epsilon)) + S \sqrt{R(\epsilon)} \mathbb{E} |\tilde{Z}_1|. \end{aligned} \quad (45)$$

Because of (40) we see that the second derivative of $\tilde{F}(R(\epsilon))$ w.r.t. $R(\epsilon)$ is negative so that it is concave. Note $F(R(\epsilon))$ itself is not necessarily concave in $R(\epsilon)$, although $f(R(\epsilon))$ is. Concavity of $f(R(\epsilon))$ is not obvious from (43) (obtained from differentiating $\mathbb{E} \langle \mathcal{L} \rangle_t$ w.r.t. $R(\epsilon)$) but can be seen from (61) (obtained instead by differentiating $-\frac{1}{2} \mathbb{E} \langle Q \rangle_t$ which reads $\frac{d}{dR(\epsilon)} \mathbb{E} \langle Q \rangle_t = -2 \frac{d^2}{dR(\epsilon)^2} f \geq 0$).

Evidently $\tilde{f}(R(\epsilon))$ is concave too. Concavity then allows to use the following standard lemma:

Lemma 6 (A bound for concave functions). *Let $G(x)$ and $g(x)$ be concave functions. Let $\delta > 0$ and define $C_\delta^-(x) \equiv g'(x - \delta) - g'(x) \geq 0$ and $C_\delta^+(x) \equiv g'(x) - g'(x + \delta) \geq 0$. Then*

$$|G'(x) - g'(x)| \leq \delta^{-1} \sum_{u \in \{x-\delta, x, x+\delta\}} |G(u) - g(u)| + C_\delta^+(x) + C_\delta^-(x).$$

First, from (45) we have

$$\tilde{F}(R(\epsilon)) - \tilde{f}(R(\epsilon)) = F(R(\epsilon)) - f(R(\epsilon)) + S \sqrt{R(\epsilon)} A_n \quad (46)$$

with $A_n \equiv \frac{1}{n} \sum_{i=1}^n |\tilde{Z}_i| - \mathbb{E} |\tilde{Z}_1|$. Second, from (39), (42) we obtain for the $R(\epsilon)$ -derivatives

$$\tilde{F}'(R(\epsilon)) - \tilde{f}'(R(\epsilon)) = \langle \mathcal{L} \rangle_t - \mathbb{E} \langle \mathcal{L} \rangle_t + \frac{S A_n}{2 \sqrt{R(\epsilon)}}. \quad (47)$$

From (46) and (47) it is then easy to show that Lemma 6 implies

$$\begin{aligned} |\langle \mathcal{L} \rangle_t - \mathbb{E} \langle \mathcal{L} \rangle_t| &\leq \delta^{-1} \sum_{u \in \{R(\epsilon)-\delta, R(\epsilon), R(\epsilon)+\delta\}} (|F(u) - f(u)| + S |A_n| \sqrt{u}) \\ &\quad + C_\delta^+(R(\epsilon)) + C_\delta^-(R(\epsilon)) + \frac{S |A_n|}{2 \sqrt{\epsilon}} \end{aligned} \quad (48)$$

where $C_\delta^-(R(\epsilon)) \equiv \tilde{f}'(R(\epsilon) - \delta) - \tilde{f}'(R(\epsilon)) \geq 0$ and $C_\delta^+(R(\epsilon)) \equiv \tilde{f}'(R(\epsilon)) - \tilde{f}'(R(\epsilon) + \delta) \geq 0$. We used $R(\epsilon) \geq \epsilon$ for the term $S |A_n| / (2 \sqrt{\epsilon})$. Note that δ will be chosen later on strictly smaller

than s_n so that $R(\epsilon) - \delta \geq \epsilon - \delta \geq s_n - \delta$ remains positive. Remark that by independence of the noise variables $\mathbb{E}[A_n^2] = (1 - 2/\pi)/n \leq 1/n$. We square the identity (48) and take its expectation. Then using $(\sum_{i=1}^p v_i)^2 \leq p \sum_{i=1}^p v_i^2$, and that $R(\epsilon) \leq 2s_n + \lambda_n \rho_n$, as well as the free energy concentration Proposition 4 (under the assumption that λ_n and ρ_n verify (15)),

$$\begin{aligned} \frac{1}{9} \mathbb{E}[(\langle \mathcal{L} \rangle_t - \mathbb{E} \langle \mathcal{L} \rangle_t)^2] &\leq \frac{3}{n\delta^2} \left(C\lambda_n^2 \rho_n^3 + S(2s_n + \lambda_n \rho_n + \delta) \right) \\ &\quad + C_\delta^+(R(\epsilon))^2 + C_\delta^-(R(\epsilon))^2 + \frac{S}{4n\epsilon}. \end{aligned} \quad (49)$$

Recall $|C_\delta^\pm(R(\epsilon))| = |\tilde{f}'(R(\epsilon) \pm \delta) - \tilde{f}'(R(\epsilon))|$. By (42), (45) and $R(\epsilon) \geq \epsilon$ we have

$$|\tilde{f}'(R(\epsilon))| \leq \frac{1}{2} \left(\rho_n + \frac{S}{\sqrt{R(\epsilon)}} \right) \leq \frac{1}{2} \left(\rho_n + \frac{S}{\sqrt{\epsilon}} \right) \quad (50)$$

Thus, as $\epsilon \geq s_n$,

$$|C_\delta^\pm(R(\epsilon))| \leq \rho_n + \frac{S}{\sqrt{\epsilon - \delta}} \leq \rho_n + \frac{S}{\sqrt{s_n - \delta}}.$$

We reach

$$\begin{aligned} &\int_{s_n}^{2s_n} d\epsilon \{C_\delta^+(R(\epsilon))^2 + C_\delta^-(R(\epsilon))^2\} \\ &\leq \left(\rho_n + \frac{S}{\sqrt{s_n - \delta}} \right) \int_{s_n}^{2s_n} d\epsilon \{C_\delta^+(R(\epsilon)) + C_\delta^-(R(\epsilon))\} \\ &\leq \left(\rho_n + \frac{S}{\sqrt{s_n - \delta}} \right) \int_{R(s_n)}^{R(2s_n)} dR(\epsilon) \{C_\delta^+(R(\epsilon)) + C_\delta^-(R(\epsilon))\} \\ &= \left(\rho_n + \frac{S}{\sqrt{s_n - \delta}} \right) \left[(\tilde{f}(R(s_n) + \delta) - \tilde{f}(R(s_n) - \delta)) \right. \\ &\quad \left. + (\tilde{f}(R(2s_n) - \delta) - \tilde{f}(R(2s_n) + \delta)) \right] \end{aligned}$$

where we used that the Jacobian of the \mathcal{C}^1 -diffeomorphism $\epsilon \mapsto R(\epsilon)$ is ≥ 1 (by regularity) for the second inequality. The mean value theorem and (50) imply $|\tilde{f}(R(\epsilon) - \delta) - \tilde{f}(R(\epsilon) + \delta)| \leq \delta(\rho_n + \frac{S}{\sqrt{s_n - \delta}})$. Therefore

$$\int_{s_n}^{2s_n} d\epsilon \{C_\delta^+(R(\epsilon))^2 + C_\delta^-(R(\epsilon))^2\} \leq 2\delta \left(\rho_n + \frac{S}{\sqrt{s_n - \delta}} \right)^2.$$

Set $\delta = \delta_n = o(s_n)$. Thus, integrating (49) over $\epsilon \in [s_n, 2s_n]$ yields

$$\begin{aligned} &\int_{s_n}^{2s_n} d\epsilon \mathbb{E}[(\langle \mathcal{L} \rangle_t - \mathbb{E} \langle \mathcal{L} \rangle_t)^2] \\ &\leq \frac{27s_n}{n\delta_n^2} \left(C\lambda_n^2 \rho_n^3 + S(2s_n + \lambda_n \rho_n + \delta_n) \right) + 18\delta_n \left(\rho_n + \frac{S}{\sqrt{s_n - \delta_n}} \right)^2 + \frac{9S \ln 2}{4n} \\ &\leq \frac{Cs_n \lambda_n \rho_n}{n\delta_n^2} (1 + \lambda_n \rho_n^2) + \frac{C\delta_n}{s_n} + \frac{C}{n} \end{aligned}$$

where the constant C is generic, and may change from place to place. Finally we optimize the bound choosing $\delta_n^3 = s_n^2 \lambda_n \rho_n (1 + \lambda_n \rho_n^2)/n$. We verify the condition $\delta_n = o(s_n)$: we have $(\delta_n/s_n)^3 = O(\lambda_n \rho_n (1 + \lambda_n \rho_n^2)/(ns_n))$ which, by (15), indeed tends to 0_+ for an appropriately chosen sequence s_n . So the dominating term δ_n/s_n gives the result. \square

E Proof of inequality (38)

Let us drop the index in the bracket $\langle - \rangle_t$ and simply denote $R \equiv R_n(t, \epsilon)$. We start by proving the identity

$$-2 \mathbb{E} \langle Q(\mathcal{L} - \mathbb{E} \langle \mathcal{L} \rangle) \rangle = \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle)^2 \rangle + \mathbb{E} \langle (Q - \langle Q \rangle)^2 \rangle. \quad (51)$$

Using the definitions $Q \equiv \frac{1}{n} \mathbf{x} \cdot \mathbf{X}$ and (37) gives

$$\begin{aligned} 2 \mathbb{E} \langle Q(\mathcal{L} - \mathbb{E}(\mathcal{L})) \rangle &= \mathbb{E} \left[\frac{1}{n} \langle Q \|\mathbf{x}\|^2 \rangle - 2 \langle Q^2 \rangle - \frac{1}{n\sqrt{R}} \langle Q(\tilde{\mathbf{Z}} \cdot \mathbf{x}) \rangle \right] \\ &\quad - \mathbb{E} \langle Q \rangle \mathbb{E} \left[\frac{1}{n} \langle \|\mathbf{x}\|^2 \rangle - 2 \langle Q \rangle - \frac{1}{n\sqrt{R}} \langle \tilde{\mathbf{Z}} \cdot \mathbf{x} \rangle \right]. \end{aligned} \quad (52)$$

The gaussian integration by part formula (59) with Hamiltonian (32) yields

$$\frac{1}{n\sqrt{R}} \mathbb{E} \langle Q(\tilde{\mathbf{Z}} \cdot \mathbf{x}) \rangle = \frac{1}{n} \mathbb{E} \langle Q \|\mathbf{x}\|^2 \rangle - \frac{1}{n} \mathbb{E} \langle Q(\mathbf{x} \cdot \langle \mathbf{x} \rangle) \rangle \stackrel{N}{=} \frac{1}{n} \mathbb{E} \langle Q \|\mathbf{x}\|^2 \rangle - \mathbb{E}[\langle Q \rangle^2].$$

For the last equality we used the Nishimori identity as follows

$$\frac{1}{n} \mathbb{E} \langle Q(\mathbf{x} \cdot \langle \mathbf{x} \rangle) \rangle = \frac{1}{n^2} \mathbb{E} \langle (\mathbf{x} \cdot \mathbf{X})(\mathbf{x} \cdot \langle \mathbf{x} \rangle) \rangle \stackrel{N}{=} \frac{1}{n^2} \mathbb{E} \langle (\mathbf{X} \cdot \mathbf{x})(\mathbf{X} \cdot \langle \mathbf{x} \rangle) \rangle = \mathbb{E}[\langle Q \rangle^2].$$

Note that we already proved (41), namely

$$\frac{1}{n\sqrt{R}} \mathbb{E} \langle \tilde{\mathbf{Z}} \cdot \mathbf{x} \rangle = \frac{1}{n} \mathbb{E} \langle \|\mathbf{x}\|^2 \rangle - \mathbb{E} \langle Q \rangle.$$

Therefore (52) finally simplifies to

$$\begin{aligned} 2 \mathbb{E} \langle Q(\mathcal{L} - \mathbb{E}(\mathcal{L})) \rangle &= \mathbb{E}[\langle Q \rangle^2] - 2 \mathbb{E} \langle Q^2 \rangle + \mathbb{E}[\langle Q \rangle]^2 \\ &= -(\mathbb{E} \langle Q^2 \rangle - \mathbb{E}[\langle Q \rangle]^2) - (\mathbb{E} \langle Q^2 \rangle - \mathbb{E}[\langle Q \rangle]^2). \end{aligned}$$

which is identity (51).

This identity implies the inequality

$$2 |\mathbb{E} \langle Q(\mathcal{L} - \mathbb{E}(\mathcal{L})) \rangle| = 2 |\mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle)(\mathcal{L} - \mathbb{E}(\mathcal{L})) \rangle| \geq \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle)^2 \rangle$$

and an application of the Cauchy-Schwarz inequality gives

$$2 \{ \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle)^2 \rangle \mathbb{E} \langle (\mathcal{L} - \mathbb{E}(\mathcal{L}))^2 \rangle \}^{1/2} \geq \mathbb{E} \langle (Q - \mathbb{E} \langle Q \rangle)^2 \rangle.$$

This ends the proof of (38).

F Heuristic derivation of the information theoretic phase transition

In this section we analyze the potential function in order to heuristically locate the information theoretic transition in the special case of the spiked Wigner model with Bernoulli prior $P_X = \text{Ber}(\rho)$. The main hypotheses behind this computation are *i*) that the SNR $\lambda = \lambda(\rho)$ varies with ρ as $\lambda = 4\gamma |\ln \rho| / \rho$ with $\gamma > 0$ and independent of ρ ; that *ii*) in this SNR regime the potential possesses only two minima $\{q^+, q^-\}$ that approach, as $\rho \rightarrow 0_+$, the boundary values $q^- = o(\rho / |\ln \rho|)$ and $q^+ \rightarrow \rho$. For the Bernoulli prior the potential explicitly reads

$$\begin{aligned} i_n^{\text{pot}}(q, \lambda, \rho) &\equiv \frac{\lambda(q^2 + \rho^2)}{4} - (1 - \rho) \mathbb{E} \ln \left\{ 1 - \rho + \rho e^{-\frac{1}{2}\lambda q + \sqrt{\lambda q Z}} \right\} - \rho \mathbb{E} \ln \left\{ 1 - \rho + \rho e^{\frac{1}{2}\lambda q + \sqrt{\lambda q Z}} \right\}. \end{aligned}$$

We used that

$$I(X; \sqrt{\gamma}X + Z) = -\mathbb{E} \ln \int dP_X(x) e^{-\frac{1}{2}\gamma x^2 + \gamma X x + \sqrt{\gamma} Z x} + \frac{1}{2} \mathbb{E}[X^2] \gamma. \quad (53)$$

Let us compute this function around its assumed minima. Starting with $q^- = o(\rho / |\ln \rho|)$ (this means that this quantity goes to 0_+ faster than $\rho / |\ln \rho|$ as ρ vanishes) we obtain at leading order

after a careful Taylor expansion in $\lambda q^- \rightarrow 0_+$ (the symbol \approx means equality up to lower order terms as $\rho \rightarrow 0_+$)

$$i_n^{\text{pot}}(q^-, \lambda, \rho) \approx \frac{\lambda(q^-)^2}{4} + \frac{\lambda\rho^2}{4} - \frac{\rho(\lambda q^-)^2}{8} \approx \frac{\lambda\rho^2}{4} = \gamma\rho|\ln\rho|. \quad (54)$$

For the other minimum $q^+ \rightarrow \rho$, because $\lambda q^+ \rightarrow +\infty$ the Z contribution in the exponentials appearing in the potential can be dropped due to the presence of the square root. We obtain at leading order

$$i_n^{\text{pot}}(q^+, \lambda, \rho) \approx 2\gamma\rho|\ln\rho| - \ln\{1 + \rho^{1+2\gamma}\} - \rho\ln\{1 + \rho^{1-2\gamma}\}.$$

Here there are two cases to consider: $\gamma > 1/2$ and $0 < \gamma \leq 1/2$. We start with $\gamma > 1/2$. In this case the potential simplifies to

$$i_n^{\text{pot}}(q^+, \lambda, \rho) \approx \rho|\ln\rho|.$$

Now for $0 < \gamma \leq 1/2$ we have

$$i_n^{\text{pot}}(q^+, \lambda, \rho) \approx 2\gamma\rho|\ln\rho|.$$

The information theoretic threshold $\lambda_c = \lambda_c(\rho)$ is defined as the first non-analyticity in the mutual information. In the present setting this corresponds to a discontinuity of the first derivative w.r.t. the SNR of the mutual information (and we therefore speak about a “first-order phase transition”). By the I-MMSE formula this threshold manifests itself as a discontinuity in the MMSE. In the high sparsity regime $\rho \rightarrow 0_+$ the transition is actually as sharp as it can be with a 0–1 behavior. This translates, at the level of the potential, as the SNR threshold where its minimum is attained at q^- just below and instead at q^+ just above. So we equate $\lim_{\rho \rightarrow 0_+} i_n^{\text{pot}}(q^-, \lambda_c, \rho) = \lim_{\rho \rightarrow 0_+} i_n^{\text{pot}}(q^+, \lambda_c, \rho)$ and solve for λ_c . This is only possible, under the constraint $\gamma > 0$ independent of ρ , in the case $\gamma > 1/2$ and gives $\gamma = 1$ which is the claimed information theoretic threshold $\lambda_c(\rho) = 4|\ln\rho|/\rho$. Repeating this analysis for the Bernoulli-Rademacher prior $P_X = (1 - \rho)\delta_0 + \frac{1}{2}\rho(\delta_{-1} + \delta_1)$ leads the same threshold, which suggests that the transition is only related (for discrete priors) to the recovery of the support of the signal.

Another piece of information gained from this analysis is that around the transition the mutual information divided by n is $\Theta(\rho|\ln\rho)$. Therefore the proper normalization for the mutual information is $(n\rho|\ln\rho|)^{-1}I(\mathbf{X}; \mathbf{W})$ for it to have a well defined non trivial limit in the regime $\rho \rightarrow 0_+$.

Finally for $\gamma \leq 1$ the minimum of the potential is attained at q^- and the rescaled mutual information $(n\rho|\ln\rho|)^{-1}I(\mathbf{X}; \mathbf{W})$ equals γ as seen from (54). If instead $\gamma \geq 1$ the minimum is attained at q^+ and the mutual information instead saturates to 1, so we get the asymptotic singular function $\gamma\mathbb{I}(\gamma \leq 1) + \mathbb{I}(\gamma \geq 1)$.

G Heuristic derivation of the AMP algorithmic transition

In this section we derive the AMP algorithmic transition for the spiked Wigner model in the Bernoulli case $P_{X,n} = \rho_n\delta_1 + (1 - \rho_n)\delta_0$. The approach can be applied to the Bernoulli-Rademacher case as well (and probably more generically), and leads to the same scaling for the AMP threshold. The derivation starts from the state evolution recursion for the overlap of AMP (10), or equivalently,

$$\tau_{t+1}^n = \mathbb{E}\left\{X_0^n \mathbb{E}\{X_0^n \mid \sqrt{\lambda_n}\tau_t^n X_0^n + \sqrt{\tau_t^n}Z\}\right\}, \quad \tau_0^n = 0,$$

which, in the Bernoulli case, reads as (recall $Z \sim \mathcal{N}(0, 1)$),

$$\tau_{t+1}^n = \mathbb{E}\left\{\frac{\rho_n^2}{\rho_n + (1 - \rho_n) \exp\{-\frac{1}{2}\lambda_n\tau_t^n - \sqrt{\lambda_n\tau_t^n}Z\}}\right\}. \quad (55)$$

Therefore by plugging $\tau_0^n = 0$ in the recursion we get $\tau_1^n = \rho_n^2$, and then

$$\tau_2^n = \mathbb{E}\left\{\frac{\rho_n^2}{\rho_n + (1 - \rho_n) \exp\{-\frac{1}{2}\lambda_n\rho_n^2 - \sqrt{\lambda_n\rho_n^2}Z\}}\right\}.$$

Now depending on $\lambda_n \rho_n^2 \gg 1$ or $\lambda_n \rho_n^2 \ll 1$ the next step of the recursion has two very different behaviors. When $\lambda_n \rho_n^2 \gg 1$, it becomes

$$\lambda_n \rho_n^2 \gg 1 : \quad \tau_2^n \approx \rho_n \quad \text{and thus} \quad \lambda_n \tau_2^n \approx \lambda_n \rho_n \gg 1.$$

Therefore, the recursion will remain stuck in this “reconstruction state” and converges towards $\tau_\infty^n \approx \rho_n$ which yields the minimal value of the MSE:

$$\frac{\text{MSE}_{\text{AMP}}^\infty}{(n\rho_n)^2} = 1 - \left(\frac{\tau_\infty^n}{\rho_n}\right)^2 \approx 0.$$

When $\lambda_n \rho_n^2 \ll 1$,

$$\lambda_n \rho_n^2 \ll 1 : \quad \tau_2^n \approx \rho_n^2 = \tau_1^n.$$

In this case, the recursion converges towards the “no reconstruction state” $\tau_\infty^n \approx \rho_n^2$, which corresponds to the MSE of a random guess (according to the prior) for the spike signal-matrix, i.e., the MSE corresponding to take as estimator $\mathbf{X}' \otimes \mathbf{X}'$ where $\mathbf{X}' \sim P_{X,n}$ is independent from the ground-truth \mathbf{X} :

$$\frac{\text{MSE}_{\text{AMP}}^\infty}{(n\rho_n)^2} = 1 - \left(\frac{\tau_\infty^n}{\rho_n}\right)^2 \approx 1.$$

This reasoning shows that the behavior of the state evolution must change for a scaling $\lambda_n \rho_n^2 = O(1)$. This argument cannot catch the constant $\lambda_n \rho_n^2 \approx 1/e$, which was numerically approximated in [52].

H The Nishimori identity

Lemma 7 (Nishimori identity). *Let (\mathbf{X}, \mathbf{Y}) be a couple of random variables with joint distribution $P(\mathbf{X}, \mathbf{Y})$ and conditional distribution $P(\mathbf{X}|\mathbf{Y})$. Let $k \geq 1$ and let $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(k)}$ be i.i.d. samples from the conditional distribution. We use the bracket $\langle - \rangle$ for the expectation w.r.t. the product measure $P(\mathbf{x}^{(1)}|\mathbf{Y})P(\mathbf{x}^{(2)}|\mathbf{Y}) \dots P(\mathbf{x}^{(k)}|\mathbf{Y})$ and \mathbb{E} for the expectation w.r.t. the joint distribution. Then, for all continuous bounded function g we have*

$$\mathbb{E}\langle g(\mathbf{Y}, \mathbf{x}^{(1)}, \dots, \mathbf{x}^{(k)}) \rangle = \mathbb{E}\langle g(\mathbf{Y}, \mathbf{X}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(k)}) \rangle.$$

Proof. This is a simple consequence of Bayes formula. It is equivalent to sample the couple (\mathbf{X}, \mathbf{Y}) according to its joint distribution or to sample first \mathbf{Y} according to its marginal distribution and then to sample \mathbf{X} conditionally on \mathbf{Y} from the conditional distribution. Thus the two $(k+1)$ -tuples $(\mathbf{Y}, \mathbf{x}^{(1)}, \dots, \mathbf{x}^{(k)})$ and $(\mathbf{Y}, \mathbf{X}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(k)})$ have the same law. \square

I I-MMSE relation

In this appendix we prove the I-MMSE relation of [58, 59] for the convenience of the reader.

Lemma 8 (I-MMSE formula). *Consider a signal $\mathbf{X} \in \mathbb{R}^n$ with $\mathbf{X} \sim P_X$ that has finite support, and gaussian corrupted data $\mathbf{Y} \sim \mathcal{N}(\sqrt{R}\mathbf{X}, \mathbf{I}_n)$ and possibly additional generic data $\mathbf{W} \sim P_{W|X}(\cdot|\mathbf{X})$ with $H(\mathbf{W})$ bounded. The I-MMSE formula linking the mutual information and the MMSE then reads*

$$\frac{d}{dR} I(\mathbf{X}; (\mathbf{Y}, \mathbf{W})) = \frac{d}{dR} I(\mathbf{X}; \mathbf{Y}|\mathbf{W}) = \frac{1}{2} \text{MMSE}(\mathbf{X}|\mathbf{Y}, \mathbf{W}) = \frac{1}{2} \mathbb{E} \|\mathbf{X} - \langle \mathbf{x} \rangle\|^2, \quad (56)$$

where the Gibbs-bracket $\langle - \rangle$ is the expectation acting on $\mathbf{x} \sim P(\cdot|\mathbf{Y}, \mathbf{W})$.

Proof. First note that by the chain rule for mutual information $I(\mathbf{X}; (\mathbf{Y}, \mathbf{W})) = I(\mathbf{X}; \mathbf{Y}|\mathbf{W}) + I(\mathbf{X}; \mathbf{W})$, so the derivatives in (56) are equal. We will now look at $\frac{d}{dR} I(\mathbf{X}; (\mathbf{Y}, \mathbf{W}))$. Since, conditionally on \mathbf{X}, \mathbf{Y} and \mathbf{W} are independent, we have

$$I(\mathbf{X}; (\mathbf{Y}, \mathbf{W})) = H(\mathbf{Y}, \mathbf{W}) - H(\mathbf{Y}, \mathbf{W}|\mathbf{X}) = H(\mathbf{Y}, \mathbf{W}) - H(\mathbf{Y}|\mathbf{X}) - H(\mathbf{W}|\mathbf{X}).$$

With gaussian noise contribution $H(\mathbf{Y}|\mathbf{X}) = \frac{n}{2} \ln(2\pi e)$. Therefore only $H(\mathbf{Y}, \mathbf{W})$ depends on R . Let us then compute, using the change of variable $\mathbf{Y} = \sqrt{R}\mathbf{X} + \mathbf{Z}$,

$$\begin{aligned}
\frac{d}{dR}I(\mathbf{X};(\mathbf{Y}, \mathbf{W})) &= \frac{d}{dR}H(\mathbf{Y}, \mathbf{W}) \\
&= -\frac{d}{dR} \int dP_X(\mathbf{X})d\mathbf{Y}d\mathbf{W}P_{W|\mathbf{X}}(\mathbf{W}|\mathbf{X}) \frac{e^{-\frac{1}{2}\|\mathbf{Y}-\sqrt{R}\mathbf{X}\|^2}}{(2\pi)^{n/2}} \\
&\quad \times \ln \int dP_X(\mathbf{x})P_{W|\mathbf{X}}(\mathbf{W}|\mathbf{x}) \frac{e^{-\frac{1}{2}\|\mathbf{Y}-\sqrt{R}\mathbf{x}\|^2}}{(2\pi)^{n/2}} \\
&= -\int dP_X(\mathbf{X})d\mathbf{Z}d\mathbf{W}P_{W|\mathbf{X}}(\mathbf{W}|\mathbf{X}) \frac{e^{-\frac{1}{2}\|\mathbf{Z}\|^2}}{(2\pi)^{n/2}} \\
&\quad \times \frac{d}{dR} \ln \int dP_X(\mathbf{x})P_{W|\mathbf{X}}(\mathbf{W}|\mathbf{x}) \frac{e^{-\frac{1}{2}\|\mathbf{Z}-\sqrt{R}(\mathbf{x}-\mathbf{X})\|^2}}{(2\pi)^{n/2}} \\
&= \frac{1}{2\sqrt{R}} \mathbb{E}_{\mathbf{X}, \mathbf{Z}, \mathbf{W}|\mathbf{X}} \langle (\mathbf{Z} + \sqrt{R}(\mathbf{X} - \mathbf{x})) \cdot (\mathbf{X} - \mathbf{x}) \rangle \tag{57}
\end{aligned}$$

where $\mathbf{Z} \sim \mathcal{N}(0, \mathbf{I}_n)$ and the bracket notation is the expectation w.r.t. the posterior proportional to

$$dP_X(\mathbf{x})dP_{W|\mathbf{X}}(\mathbf{W}|\mathbf{x})d\mathbf{Z} \exp \left\{ -\frac{1}{2}\|\mathbf{Z} - \sqrt{R}(\mathbf{x} - \mathbf{X})\|^2 \right\}.$$

In (57) the interchange of derivative and integrals is permitted by a standard application of Lebesgue's dominated convergence theorem in the case where the support of P_X is bounded. Now we use the following gaussian integration by part formula: for any bounded function $\mathbf{g} : \mathbb{R}^n \mapsto \mathbb{R}^n$ of a standard gaussian random vector $\mathbf{Z} \sim \mathcal{N}(0, \mathbf{I}_n)$ we obviously have

$$\mathbb{E}[\mathbf{Z} \cdot \mathbf{g}(\mathbf{Z})] = \mathbb{E}[\nabla_{\mathbf{Z}} \cdot \mathbf{g}(\mathbf{Z})]. \tag{58}$$

This formula applied to a Gibbs-bracket associated to a general Gibbs distribution with hamiltonian $\mathcal{H}(\mathbf{x}, \mathbf{Z})$ (depending on the Gaussian noise and possibly other variables) yields

$$\begin{aligned}
\mathbb{E}[\mathbf{Z} \cdot \langle \mathbf{h}(\mathbf{x}) \rangle] &= \mathbb{E} \nabla_{\mathbf{Z}} \cdot \frac{\int dP(\mathbf{x})e^{-\mathcal{H}(\mathbf{x}, \mathbf{Z})} \mathbf{h}(\mathbf{x})}{\int dP(\mathbf{x}')e^{-\mathcal{H}(\mathbf{x}', \mathbf{Z})}} \\
&= -\mathbb{E} \frac{\int dP_X(\mathbf{x})e^{-\mathcal{H}(\mathbf{x}, \mathbf{Z})} \mathbf{h}(\mathbf{x}) \cdot \nabla_{\mathbf{Z}} \mathcal{H}(\mathbf{x}, \mathbf{Z})}{\int dP_X(\mathbf{x}')e^{-\mathcal{H}(\mathbf{x}', \mathbf{Z})}} \\
&\quad + \mathbb{E} \left[\frac{\int dP_X(\mathbf{x})e^{-\mathcal{H}(\mathbf{x}, \mathbf{Z})} \mathbf{h}(\mathbf{x})}{\int dP_X(\mathbf{x}')e^{-\mathcal{H}(\mathbf{x}', \mathbf{Z})}} \cdot \frac{\int dP_X(\mathbf{x})e^{-\mathcal{H}(\mathbf{x}, \mathbf{Z})} \nabla_{\mathbf{Z}} \mathcal{H}(\mathbf{x}, \mathbf{Z})}{\int dP_X(\mathbf{x}')e^{-\mathcal{H}(\mathbf{x}', \mathbf{Z})}} \right] \\
&= -\mathbb{E} \langle \mathbf{h}(\mathbf{x}) \cdot \nabla_{\mathbf{Z}} \mathcal{H}(\mathbf{x}, \mathbf{Z}) \rangle + \mathbb{E} [\langle \mathbf{h}(\mathbf{x}) \rangle \cdot \langle \nabla_{\mathbf{Z}} \mathcal{H}(\mathbf{x}, \mathbf{Z}) \rangle]. \tag{59}
\end{aligned}$$

Applied to (57), where the "hamiltonian" is $\mathcal{H}(\mathbf{x}, \mathbf{Z}) = -\ln P_{W|\mathbf{X}}(\mathbf{W}|\mathbf{x}) + \frac{1}{2}\|\mathbf{Z} - \sqrt{R}(\mathbf{x} - \mathbf{X})\|^2$, this identity gives

$$\begin{aligned}
\frac{d}{dR}I(\mathbf{X};(\mathbf{Y}, \mathbf{W})) &= \frac{1}{2} \mathbb{E} [\langle \|\mathbf{X} - \mathbf{x}\|^2 \rangle + \frac{1}{\sqrt{R}} \nabla_{\mathbf{Z}} \cdot \langle \mathbf{X} - \mathbf{x} \rangle] \\
&= \frac{1}{2} \mathbb{E} [\langle \|\mathbf{X} - \mathbf{x}\|^2 \rangle - \frac{1}{\sqrt{R}} \langle (\mathbf{X} - \mathbf{x}) \cdot (\mathbf{Z} + \sqrt{R}(\mathbf{X} - \mathbf{x})) \rangle] \\
&\quad + \frac{1}{\sqrt{R}} \langle (\mathbf{X} - \mathbf{x}) \rangle \cdot \langle \mathbf{Z} + \sqrt{R}(\mathbf{X} - \mathbf{x}) \rangle] \\
&= \frac{1}{2} \mathbb{E} \|\mathbf{X} - \langle \mathbf{x} \rangle\|^2.
\end{aligned}$$

□

The MMSE cannot increase when the SNR increases. This translates into the concavity of the mutual information of gaussian channels as a function of the SNR.

Lemma 9 (Concavity of the mutual information in the SNR). *Consider the same setting as Lemma 8. Then the mutual informations $I(\mathbf{X}; (\mathbf{Y}, \mathbf{W}))$ and $I(\mathbf{X}; \mathbf{Y} | \mathbf{W})$ are concave in the SNR of the gaussian channel:*

$$\begin{aligned} \frac{d^2}{dR^2} I(\mathbf{X}; (\mathbf{Y}, \mathbf{W})) &= \frac{d^2}{dR^2} I(\mathbf{X}; \mathbf{Y} | \mathbf{W}) \\ &= \frac{1}{2} \frac{d}{dR} \text{MMSE}(\mathbf{X} | \mathbf{Y}, \mathbf{W}) = -\frac{1}{2n} \sum_{i,j=1}^n \mathbb{E} [(\langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle)^2] \leq 0 \end{aligned}$$

where the Gibbs-bracket $\langle - \rangle$ is the expectation acting on $\mathbf{x} \sim P(\cdot | \mathbf{Y}, \mathbf{W})$.

Proof. Set $Q \equiv \mathbf{x} \cdot \mathbf{X} / n$ where $\mathbf{x} \sim P(\cdot | \mathbf{Y}, \mathbf{W})$. From a Nishimori identity $\text{MMSE}(\mathbf{X} | \mathbf{Y}, \mathbf{W}) = \mathbb{E}_{P_X}[X^2] - \mathbb{E}\langle Q \rangle$. Thus by the I-MMSE formula we have, by a calculation similar to (59),

$$-2 \frac{d^2}{dR^2} I(\mathbf{X}; (\mathbf{Y}, \mathbf{W})) = \frac{d \mathbb{E}\langle Q \rangle}{dR} = n \mathbb{E}[\langle Q \rangle \langle \mathcal{L} \rangle - \langle Q \mathcal{L} \rangle] \quad (60)$$

where we have set

$$\mathcal{L} \equiv \frac{1}{n} \left(\frac{1}{2} \|\mathbf{x}\|^2 - \mathbf{x} \cdot \mathbf{X} - \frac{1}{2\sqrt{R}} \mathbf{x} \cdot \mathbf{Z} \right).$$

Now we look at each term on the right hand side of this equality. The calculation of appendix E shows that

$$-\mathbb{E}\langle Q \mathcal{L} \rangle = \mathbb{E}\langle Q^2 \rangle - \frac{1}{2} \mathbb{E}[\langle Q \rangle^2]$$

so it remains to compute

$$\mathbb{E}[\langle Q \rangle \langle \mathcal{L} \rangle] = \mathbb{E} \left[\langle Q \rangle \frac{\langle \|\mathbf{x}\|^2 \rangle}{2n} - \langle Q \rangle^2 - \langle Q \rangle \frac{\mathbf{Z} \cdot \langle \mathbf{x} \rangle}{2n\sqrt{R}} \right].$$

By formulas (58) and (59) in which the Hamiltonian is (32) we have

$$\begin{aligned} -\frac{1}{2n\sqrt{R}} \mathbb{E}[\mathbf{Z} \cdot \langle \mathbf{x} \rangle \langle Q \rangle] &= -\frac{1}{2n\sqrt{R}} \mathbb{E}[\langle Q \rangle \nabla_{\mathbf{Z}} \cdot \langle \mathbf{x} \rangle + \langle \mathbf{x} \rangle \cdot \nabla \langle Q \rangle] \\ &= -\frac{1}{2n} \mathbb{E}[\langle Q \rangle (\langle \|\mathbf{x}\|^2 \rangle - \|\langle \mathbf{x} \rangle\|^2) + \langle \mathbf{x} \rangle \cdot (\langle Q \mathbf{x} \rangle - \langle Q \rangle \langle \mathbf{x} \rangle)] \\ &\stackrel{\text{N}}{=} -\frac{1}{2n} \mathbb{E}[\langle Q \rangle \langle \|\mathbf{x}\|^2 \rangle] + \frac{1}{n} \mathbb{E}[\langle Q \rangle \|\langle \mathbf{x} \rangle\|^2] - \frac{1}{2} \mathbb{E}[\langle Q \rangle^2]. \end{aligned}$$

In the last equality we used the following consequence of the Nishimori identity. Let $\mathbf{x}, \mathbf{x}^{(2)}$ be two replicas, i.e., conditionally (on the data) independent samples from the posterior (31). Then

$$\frac{1}{n} \mathbb{E}[\langle \mathbf{x} \rangle \cdot \langle Q \mathbf{x} \rangle] = \frac{1}{n^2} \mathbb{E}\langle (\mathbf{x}^{(2)} \cdot \mathbf{x})(\mathbf{x} \cdot \mathbf{X}) \rangle \stackrel{\text{N}}{=} \frac{1}{n^2} \mathbb{E}\langle (\mathbf{x}^{(2)} \cdot \mathbf{X})(\mathbf{X} \cdot \mathbf{x}) \rangle = \mathbb{E}[\langle Q \rangle^2].$$

Thus we obtain

$$\begin{aligned} \mathbb{E}[\langle Q \rangle \langle \mathcal{L} \rangle - \langle Q \mathcal{L} \rangle] &= \mathbb{E}\langle Q^2 \rangle - 2\mathbb{E}[\langle Q \rangle^2] + \frac{1}{n} \mathbb{E}[\langle Q \rangle \|\langle \mathbf{x} \rangle\|^2] \\ &= \frac{1}{n^2} \mathbb{E}\langle (\mathbf{x} \cdot \mathbf{X})^2 - 2(\mathbf{x} \cdot \mathbf{X})(\mathbf{x}^{(1)} \cdot \mathbf{X}) + (\mathbf{x} \cdot \mathbf{X})(\mathbf{x}^{(2)} \cdot \mathbf{x}^{(3)}) \rangle \\ &\stackrel{\text{N}}{=} \frac{1}{n^2} \mathbb{E}\langle (\mathbf{x} \cdot \mathbf{x}^{(0)})^2 - 2(\mathbf{x} \cdot \mathbf{x}^{(0)})(\mathbf{x}^{(1)} \cdot \mathbf{x}^{(0)}) + (\mathbf{x} \cdot \mathbf{x}^{(0)})(\mathbf{x}^{(2)} \cdot \mathbf{x}^{(3)}) \rangle \end{aligned}$$

where $\mathbf{x}^{(0)}, \mathbf{x}, \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}$ are replicas and the last equality again follows from a Nishimori identity. Multiplying this identity by n and rewriting the inner products component-wise we get

$$\begin{aligned} \frac{d \mathbb{E}\langle Q \rangle}{dR} &= \frac{1}{n} \sum_{i,j=1}^n \mathbb{E}\langle x_i x_i^{(0)} x_j x_j^{(0)} - 2x_i x_i^{(0)} x_j^{(1)} x_j^{(0)} + x_i x_i^{(0)} x_j^{(2)} x_j^{(3)} \rangle \\ &= \frac{1}{n} \sum_{i,j=1}^n \mathbb{E}[\langle x_i x_j \rangle^2 - 2\langle x_i \rangle \langle x_j \rangle \langle x_i x_j \rangle + \langle x_i \rangle^2 \langle x_j \rangle^2] \end{aligned} \quad (61)$$

Using (60) this ends the proof of the lemma. Note that we have also shown the positivity claimed in (30) of section B. \square

J Proof of corollary 1

The proof of corollary 1 follows from a combination of theorem 1 and the I-MMSE relation (see [58, 59], and also appendix I). Denote

$$M_n(s) \equiv \frac{1}{(n\rho_n)^2} \text{MMSE}((X_i X_j)_{i < j} | \mathbf{W})|_{\lambda_n = s} \quad \text{and} \quad I_n(s) \equiv \frac{1}{n\rho_n^2} I(\mathbf{X}; \mathbf{W})|_{\lambda_n = s}. \quad (62)$$

The I-MMSE relation in its integral formulation implies

$$\frac{I_n(s + \epsilon) - I_n(s)}{\epsilon} = \frac{1}{2\epsilon} \int_s^{s+\epsilon} M_n(\lambda) d\lambda. \quad (63)$$

Because $M_n(s)$ is a non-increasing function (“information can’t hurt”, which is equivalent to the concavity of mutual information in the signal-to-noise ratio, see [58, 59] or lemma 9) the above identity implies

$$\frac{M_n(s + \epsilon)}{2} \leq \frac{I_n(s + \epsilon) - I_n(s)}{\epsilon} \leq \frac{M_n(s)}{2}. \quad (64)$$

Set

$$i_n(s) \equiv \frac{1}{\rho_n^2} \inf_{q \in [0, \rho_n]} i_n^{\text{pot}}(q, s, \rho_n) \quad \text{so that} \quad \frac{1}{2} m_n(s, \rho_n) \equiv \frac{d}{ds} i_n(s).$$

Because $s \mapsto m_n(s, \rho_n)$ is also a non-increasing function (see, e.g., [15]) we obtain similarly

$$\frac{m_n(s + \epsilon, \rho_n)}{2} \leq \frac{i_n(s + \epsilon) - i_n(s)}{\epsilon} \leq \frac{m_n(s, \rho_n)}{2}. \quad (65)$$

Set $c_n \equiv C(\ln n)^{1/3} n^{-(1-6\beta)/7} |\ln \rho_n| / \rho_n$ which is the right-hand side of (5) multiplied by $(\rho_n |\ln \rho_n|) / \rho_n^2$. Theorem 1 then implies

$$\frac{M_n(s + \epsilon)}{2} \leq \frac{i_n(s + \epsilon) - i_n(s) + 2c_n}{\epsilon} \leq \frac{m_n(s, \rho_n)}{2} + \frac{2c_n}{\epsilon}, \quad (66)$$

$$\frac{m_n(s + \epsilon, \rho_n)}{2} - \frac{2c_n}{\epsilon} \leq \frac{i_n(s + \epsilon) - i_n(s) - 2c_n}{\epsilon} \leq \frac{M_n(s)}{2}. \quad (67)$$

Replacing $\rho_n = \Omega(n^{-\beta})$ with $\beta \in [0, 1/13]$ yields the claimed inequality:

$$m_n(s + \epsilon, \rho_n) - \frac{C'}{\epsilon} \frac{(\ln n)^{4/3}}{n^{(1-13\beta)/7}} \leq M_n(s) \leq m_n(s - \epsilon, \rho_n) + \frac{C'}{\epsilon} \frac{(\ln n)^{4/3}}{n^{(1-13\beta)/7}}. \quad (68)$$

K AMP algorithmic phase transition

In this appendix, we prove theorem 2. To do this, we begin by introducing a general ‘symmetric’ AMP algorithm in section K.1 and show it is quite similar to the AMP algorithm in (6). For this symmetric AMP algorithm, we provide finite sample guarantees like those given in [63] for various ‘non-symmetric’ AMP algorithms. However, we have an added challenge in that terms like the Lipschitz constant of the denoiser f_t in (9) and the state evolution values in (10) depend on n and therefore cannot be treated as universal constants in the rate of concentration, as they were in [63]. The main concentration result for the symmetric AMP is given in theorem 4 in section K.1. Then, in section K.2, we use theorem 4 to prove result (11), from which we prove theorem 2.

K.1 Symmetric AMP finite sample guarantees

We begin by analyzing a ‘symmetric’ AMP algorithm, described now, that is similar to the AMP algorithm in (6). Assume the matrix $\mathbf{Z} \sim \text{GOE}(n)$ is an $n \times n$ matrix from the gaussian orthogonal ensemble, i.e. \mathbf{Z} is a symmetric matrix with $\{Z_{ij}\}_{1 \leq i < j \leq n}$ i.i.d. $\mathcal{N}(0, 1/n)$, and $\{Z_{ii}\}_{1 \leq i \leq n}$ i.i.d. $\mathcal{N}(0, 2/n)$. Start with an initial condition $\mathbf{h}^0 \in \mathbb{R}^n$, independent of \mathbf{Z} , and calculate for $t \geq 0$,

$$\mathbf{h}^{t+1} = \mathbf{Z} g_t(\mathbf{h}^t, \mathbf{X}^n) - c_t g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X}^n). \quad (69)$$

In the above, $g_t : \mathbb{R}^2 \rightarrow \mathbb{R}$ is Lipschitz and separable (i.e. it acts component-wise when applied to vectors) and may depend on n through its Lipschitz constant, denoted L_g^n . The function g_t takes as its second argument a random vector $\mathbf{X}^n \in \mathbb{R}^n$ with entries that are i.i.d. $P_{X,n}$, a sub-gaussian distribution, $c_t = \frac{1}{n} \sum_{i=1}^n g'_t(h_i^t, X_i^n)$, the derivative taken with respect to the first argument, and all terms with negative indices take the value 0 (so that, for example, $\mathbf{h}^1 = \mathbf{Z}g_0(\mathbf{h}^0, \mathbf{X}^n)$). The key result, stated in theorem 4 below, is that for each $t \geq 1$, the empirical distribution of the components of \mathbf{h}^t is approximately equal in distribution to a gaussian $\mathcal{N}(0, \sigma_t^n)$ where the variances $\{\sigma_t^n\}_{t \geq 0}$ are defined via the state evolution: initialize with $\sigma_1^n = \|\mathbf{g}_0(\mathbf{h}^0, \mathbf{X}^n)\|^2/n$, calculate for $t \geq 1$,

$$\sigma_{t+1}^n = \mathbb{E} \left[\left(g_t(\sqrt{\sigma_t^n} Z, X_0^n) \right)^2 \right], \quad (70)$$

where the expectation is with respect to standard gaussian Z independent of $X_0^n \sim P_{X,n}$.

Before stating theorem 4 below, we give the assumptions on the model and the functions used to define the AMP. In what follows, $C, c > 0$ are generic positive constants whose values are not exactly specified but do not depend on n .

Random Vectors: The random vector $\mathbf{X}^n \in \mathbb{R}^n$ used in the denoising functions, is assumed to have entries that are i.i.d. according to a sub-gaussian distribution $P_{X,n}$, in particular, $P_{X,n}$ is $\text{Ber}(\rho_n)$ or Bernoulli-Rademacher, where sub-gaussian random variables are defined in lemma 14.

The function g_t : The denoising function $g_t : \mathbb{R}^2 \rightarrow \mathbb{R}$ in (69) is defined as $g_t(\mathbf{h}^t, \mathbf{X}^n) = f_t(\mathbf{h}^t + \sqrt{\lambda_n} \sigma_t^n \mathbf{X}^n)$ where f_t is the conditional expectation denoiser in (9). With this definition, g_t is separable and Lipschitz continuous for each $t \geq 0$, with Lipschitz constant denoted $L_g^n > 0$, that depends on n . By the Lipschitz property, g_t are weakly differentiable and the weak derivatives, denoted by g'_t , are also differentiable.

Theorem 4. Consider the AMP algorithm in (69) for $\mathbf{h}^0 \in \mathbb{R}^m$ independent of \mathbf{A} under the assumptions above. Then, for any (order-2) pseudo-Lipschitz function $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$, $\epsilon \in (0, 1)$, and $t \geq 1$.

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \phi(h_i^t, X_i^n) - \mathbb{E} \left\{ \phi(\sqrt{\sigma_t^n} Z, X_0^n) \right\} \right| \geq \epsilon \right) \leq CC_t \exp \left\{ \frac{-cc_t n \epsilon^2}{L_\phi^2 \tilde{\gamma}_t^n} \right\}, \quad (71)$$

where the expectation is with respect to standard gaussian Z independent of $X_0^n \sim P_{X,n}$, the state evolution values σ_t^n are defined in (70), the constants C_t, c_t are defined in theorem 2, and

$$\tilde{\gamma}_n^{t+1} := \lambda_n^{2t} (\nu^n + \sigma_1^n) (\nu^n + \sigma_1^n + \sigma_2^n) \cdots \left(\nu^n + \sum_{i=1}^{t+1} \sigma_i^n \right) \max\{1, \hat{c}_1\} \max\{1, \hat{c}_2\} \cdots \max\{1, \hat{c}_t\}, \quad (72)$$

where $\hat{c}_t = \mathbb{E}[g'_t(\sqrt{\sigma_t^n} Z, X)]$ and ν is the variance factor of sub-gaussian $X^n \sim P_{X,n}$ (for $P_{X,n} = \text{Ber}(\rho_n)$ we have $\nu^n \leq 1/4$; see lemma 14).

The proof of theorem 4 is given in section K.3. The proof relies heavily on the proof of the finite sample guarantees for various ‘non-symmetric’ AMP algorithms given in [63, theorem 1] and we reference this result throughout. We will use theorem 4 to prove theorem 2, but before doing so, we make a few remarks about extensions of the result and the major differences between theorem 4 and the finite sample guarantees in [63].

Remark 1: Spectral initialization. We assume that the AMP iteration in (69) was initialized with $\mathbf{h}^0 \in \mathbb{R}^n$ independent of \mathbf{Z} . As mentioned previously in section 3, for estimation with a Bernoulli-Rademacher signal prior, one needs to instead use a spectral initialization that will not be independent of the matrix \mathbf{Z} . Theoretically, as introduced in [34], one deals with this dependency by analyzing the AMP iteration using a matrix $\tilde{\mathbf{Z}}$ that is an *approximate* representation of the conditional distribution of \mathbf{Z} given the initialization and then showing that the two algorithms are close each other with high probability. We do not give the details of this rather technical argument here, and instead analyze the simpler case using an independent initialization, though the generalization is likely straightforward.

Remark 2: Rate of the concentration. The rate of concentration depends on λ_n, ρ_n , and the state evolution values, σ_t^n , through $\tilde{\gamma}_n^t$ defined in (72). In particular, the term $\lambda_n^{2(t-1)}$ in $\tilde{\gamma}_n^t$,

appears through the dependency of the rate on the Lipschitz constant of g_t , where $g_t(\mathbf{h}^t, \mathbf{X}^n) = f_t(\mathbf{h}^t + \sqrt{\lambda_n} \sigma_t^n \mathbf{X}^n)$ and f_t is the conditional expectation denoiser in (9). With this definition, $L_g^n = \sqrt{\lambda_n}$. The dependence on these values was not stated explicitly in the concentration bound of [63, theorem 1] as the authors assume that the Lipschitz constant, sparsity, and state evolution terms do not change with n and, thus, can be absorbed into the universal constants.

The presence of these terms in our rate comes from the inductive portion of the proof where one must show that the values $\|g_t(\mathbf{h}^t, \mathbf{X}^n)\|^2/n$ concentrate to known constants. Essentially, this step will add a term $(L_g^n)^2(\nu^n + \sigma_t^n)$ in the rate at each step of the induction. To see this, we point the reader to three facts. First, notice that the approximate distribution of h_i^t is gaussian with variance σ_t^n . Second, it is easy to see that a function $[f(x)]^2$ has the same pseudo-Lipschitz constant as $f(\cdot)$ if $|f(\cdot)|$ is bounded (as n grows), which is the case for g_t in our setting. (More generally, the Lipschitz constant of $[f(x)]^2$ will be no more than L_f^2 .) Finally, we highlight that pseudo-Lipschitz functions taking gaussian and sub-gaussian input concentrate as in [63, Lemma B.4] with $L^2(\nu^n + \sigma_t^n)$ in the denominator of the rate where L is the associated pseudo-Lipschitz constant, ν^n is the sub-gaussian variance factor, and σ_t^n is the gaussian variance. Indeed, we restate [63, Lemma B.4] here for clarity.

Lemma 10. [63, lemma B.4] *Let $Z \in \mathbb{R}^n$ be an i.i.d. standard gaussian vector and $G \in \mathbb{R}^n$ a random vector with entries G_1, \dots, G_n i.i.d. $\sim p_G$, where p_G is sub-gaussian with variance factor ν^n . Then, for any pseudo-Lipschitz function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ with constant L_f^n , non-negative values σ^n , and $0 < \epsilon \leq 1$,*

$$\mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n f(\sqrt{\sigma^n} Z_i, G_i) - \mathbb{E}[f(\sqrt{\sigma^n} Z, G)]\right| \geq \epsilon\right) \leq 2 \exp\left\{\frac{-\kappa n \epsilon^2}{(L_f^n)^2[\nu^n + 4(\nu^n)^2 + \sigma^n + 4(\sigma^n)^2]}\right\}.$$

Since $0 \leq \nu^n, \sigma^n \leq 1$ we drop the squared terms $(\nu^n)^2, (\sigma^n)^2$ from the rate since ν^n, σ^n dominate.

Remark 3: Denoisers The proof of [63, theorem 1] assumes that the weak derivative of the denoiser, g'_t , has bounded derivative everywhere it exists. Here, $g_t(\mathbf{h}^t, \mathbf{X}^n) = f_t(\mathbf{h}^t + \sqrt{\lambda_n} \sigma_t^n \mathbf{X}^n)$ where f_t is the conditional expectation denoiser in (9) and f'_t is given in lemma 16. In particular, $f'_t(x) = \sqrt{\lambda_n} f_t(x)(1 - f_t(x))$, which is not bounded (in n) since λ_n grows with n . However, we can show that $f'_t(x)$ is also Lipschitz, with constant $L_f^2 = \lambda_n$, and we use this fact directly in the proof to get around the boundedness assumption originally used in [63, theorem 1].

K.2 Proving theorem 2

Before we get to the proof of theorem 2, we discuss how we apply the result of theorem 4 to our problem. This will lead to the concentration result in (11), which concerns convergence within pseudo-Lipschitz loss functions of the empirical distribution of x_i^t , the iterate of the AMP algorithm in (6), to its approximating distribution with mean and variance determined by the state evolution. Recall the following definition of a pseudo-Lipschitz function.

Definition 1. *For any $n, m \in \mathbb{N}_{>0}$, a function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is pseudo-Lipschitz of order 2 if there exists a constant $L > 0$ such that $\|\phi(\mathbf{x}) - \phi(\mathbf{y})\| \leq L(1 + \|\mathbf{x}\| + \|\mathbf{y}\|)\|\mathbf{x} - \mathbf{y}\|$ for $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$.*

Now we prove (11). Recall that in our model (1),

$$\frac{1}{\sqrt{n}} \mathbf{W} = \frac{\sqrt{\lambda_n}}{n} \mathbf{X} \otimes \mathbf{X} + \mathbf{Z}, \quad (73)$$

where $\lambda_n > 0$ controls the strength of the signal and the noise is i.i.d. gaussian $Z_{ij} \sim \mathcal{N}(0, 1/n)$ for $i < j$ and symmetric, $Z_{ij} = Z_{ji}$. The AMP algorithm for recovering \mathbf{X} from the data \mathbf{W} is given in (6).

Notice that the AMP algorithm in (6) is similar to (69), the only difference being that the matrix \mathbf{A} in (6) is our data matrix, as opposed to it being $\text{GOE}(n)$ as in (69). If we plug the value of \mathbf{W} from (73) into (6), we find the following iteration: $\mathbf{x}^1 = \frac{\sqrt{\lambda_n}}{n} \mathbf{X} \langle \mathbf{X}, f_0(\mathbf{x}^0) \rangle + \mathbf{Z} f_0(\mathbf{x}^0)$, and for $t \geq 1$,

$$\mathbf{x}^{t+1} = \frac{\sqrt{\lambda_n}}{n} \mathbf{X} \langle \mathbf{X}, f_t(\mathbf{x}^t) \rangle + \mathbf{Z} f_t(\mathbf{x}^t) - \mathbf{b}_t f_{t-1}(\mathbf{x}^{t-1}). \quad (74)$$

Now we define a related iteration to (74) as follows. Initialize with $\mathbf{h}^0 = \mathbf{x}^0$ with denoiser $g_0(\mathbf{h}^0, \mathbf{X}) := f_0(\mathbf{x}^0)$ and $\mathbf{h}^1 = \mathbf{Z}g_0(\mathbf{h}^0, \mathbf{X})$. Then calculate for $t \geq 1$,

$$\mathbf{h}^{t+1} = \mathbf{Z}g_t(\mathbf{h}^t, \mathbf{X}^n) - \mathbf{c}_t g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X}^n), \quad \mathbf{c}_t = \frac{1}{n} \sum_{i=1}^n g'_t(h_i^t, X_i^n), \quad (75)$$

where $g_t(h, X) = f_t(h + \mu_t^n X)$ for $f_t(\cdot)$ the conditional expectation denoiser used in (74) and μ_{t+1}^n calculated from the state evolution in (8) for our original iteration (i.e., the algorithm in (6)). In the above, \mathbf{X}^n is the signal in (73) and we drop the n superscript in what follows. Then, the iteration in (75) takes the exact form of the symmetric AMP in (69) with state evolution given by (70). In particular, the state evolution associated with (75) is $\tau_1^n = \|g_0(\mathbf{h}^0, \mathbf{X})\|^2/n$ and for $t \geq 1$,

$$\tau_{t+1}^n = \mathbb{E} \left[\left(g_t(\sqrt{\sigma_t^n} Z, X) \right)^2 \right] = \mathbb{E} \left[\left(f_t(\sqrt{\sigma_t^n} Z + \mu_t^n X) \right)^2 \right]. \quad (76)$$

The above state evolution is exactly the state evolution for the AMP algorithm in (74) defined in (10). For this reason, we used the τ notation.

As the AMP algorithm in (75) takes the exact form of the symmetric AMP in (69), we can apply theorem 4. The proof idea is to use theorem 4 to give performance guarantees to the algorithm in (75) and then to argue that the algorithm in (74) is asymptotically equivalent to the algorithm in (75) so the performance guarantees hold for (74) as well.

We apply theorem 4 to (75) using the pseudo-Lipschitz function $\phi(h_i^t, X_i^n) = \psi(h_i^t + \mu_t^n X_i^n, X_i^n)$, where ψ is the order 2 pseudo-Lipschitz function in (11), to find that for $t \geq 1$,

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \psi(h_i^t + \mu_t^n X_i, X_i) - \mathbb{E} \left\{ \psi(\sqrt{\tau_t^n} Z + \mu_t^n X, X) \right\} \right| \geq \epsilon \right) \leq CC_t \exp \left\{ \frac{-cc_t n \epsilon^2}{L_\psi^2 \tilde{\gamma}_n^t} \right\}. \quad (77)$$

We have used that $L_\phi = 2L_\psi(1 + \mu_t^n)^2$, which is shown in lemma 17, and that $L_\phi = 2L_\psi(1 + \mu_t^n)^2 \leq \kappa L_\psi$, which follows from the fact that $\mu_t^n \leq \kappa'$ in the regime of interest, as discussed, for example, in (102) in section K.4.

To show how (11) follows from (77), we use the following lemma.

Lemma 11. Define $\text{bound}_t := CC_t \exp \left\{ \frac{-cc_t n \epsilon^2}{L_\psi^2 \tilde{\gamma}_n^t} \right\}$, for $\tilde{\gamma}_n^t$ in (72). Let \mathbf{h}^t be defined by the algorithm in (75) and \mathbf{x}^t by (74). Then for $t \geq 1$, the following are true

$$\mathbb{P} \left(\frac{1}{\sqrt{n}} \|\mathbf{h}^t + \mu_t^n \mathbf{X}\| \geq \kappa_h \right) \leq CC_t e^{-\frac{cc_t n}{\tilde{\gamma}_n^t}}, \quad (78)$$

$$\mathbb{P} \left(\frac{1}{n} \|\mathbf{x}^t - \mathbf{h}^t - \mu_t^n \mathbf{X}\|^2 \geq \frac{\kappa \epsilon^2}{L_\psi^2} \right) \leq \text{bound}_t, \quad (79)$$

$$\mathbb{P} \left(\frac{1}{\sqrt{n}} \|\mathbf{x}^t\| \geq \kappa_x \right) \leq CC_t e^{-\frac{cc_t n}{\tilde{\gamma}_n^t}}, \quad (80)$$

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \psi(X_i, x_i^t) - \psi(X_i, h_i^t + \mu_t^n X_i) \right| \geq \epsilon \right) \leq \text{bound}_t, \quad (81)$$

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \psi(X_i, x_i^t) - \mathbb{E} \left\{ \psi(X_0, \mu_t^n X_0 + \sqrt{\tau_t^n} Z) \right\} \right| \geq \epsilon \right) \leq \text{bound}_t. \quad (82)$$

In (78) and (80) both κ_h and κ_x are universal constants.

The proof of lemma 11 is rather long and technical, so we include it in full detail at the end of the appendix in section K.4 and give a high level sketch here.

The basic idea behind the proof of lemma 11 is that the results in (78) follow from the fact that $h_i^t + \mu_t^n X_i \approx \sqrt{\tau_t^n} Z + \mu_t^n X$ for $X \sim P_{X,n}$ independent of Z standard gaussian by theorem 4. Thus, $\frac{1}{n} \|\mathbf{h}^t + \mu_t^n \mathbf{X}\|^2$ will concentrate to $\tau_t^n + (\mu_t^n)^2 \rho_n$. Then we use concentration to imply boundedness with high probability. The result (80) follows from the same ideas since it can be shown that $x_i^t \approx \sqrt{\tau_t^n} Z + \mu_t^n X$.

Next, results (81) and (82) follow immediately from (78)–(80). To see this, first notice that (82) follows directly from the bounds in (77) and (81) using lemma 19. Next, (81) follows from results (78) – (80). This can be seen by using the following upper bound due to Cauchy-Schwarz,

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=1}^n \psi(X_i, x_i^t) - \psi(X_i, h_i^t + \mu_t^n X_i) \right| &\leq \frac{1}{n} \sum_{i=1}^n \left| \psi(X_i, x_i^t) - \psi(X_i, h_i^t + \mu_t^n X_i) \right| \\ &\leq \frac{L_\psi}{n} \sum_{i=1}^n \left(1 + \|(X_i, x_i^t)\| + \|(X_i, h_i^t + \mu_t^n X_i)\| \right) |x_i^t - h_i^t - \mu_t^n X_i| \\ &\leq \frac{\kappa L_\psi}{\sqrt{n}} \|\mathbf{x}^t - \mathbf{h}^t - \mu_t^n \mathbf{X}\| \sqrt{\left(1 + \frac{2}{n} \|\mathbf{X}\|^2 + \frac{1}{n} \|\mathbf{x}^t\|^2 + \frac{1}{n} \|\mathbf{h}^t + \mu_t^n \mathbf{X}\|^2 \right)}, \end{aligned} \quad (83)$$

and the boundedness of the term $\|\mathbf{X}\|^2/n$. Thus, using $\kappa_B = \sqrt{1 + 4 + \kappa_x^2 + \kappa_h^2} > 0$, a universal constant, by the above bound it follows that

$$\begin{aligned} \mathbb{P}\left(\left| \frac{1}{n} \sum_{i=1}^n \psi(X_i, x_i^t) - \psi(X_i, h_i^t + \mu_t^n X_i) \right| \geq \epsilon \right) \\ \leq \mathbb{P}\left(\frac{1}{\sqrt{n}} \|\mathbf{x}^t - \mathbf{h}^t - \mu_t^n \mathbf{X}\| \sqrt{\left(1 + \frac{2}{n} \|\mathbf{X}\|^2 + \frac{1}{n} \|\mathbf{x}^t\|^2 + \frac{1}{n} \|\mathbf{h}^t + \mu_t^n \mathbf{X}\|^2 \right)} \geq \frac{\kappa \epsilon}{L_\psi} \right) \\ \leq \mathbb{P}\left(\frac{1}{\sqrt{n}} \|\mathbf{x}^t - \mathbf{h}^t - \mu_t^n \mathbf{X}\| \geq \frac{\kappa \epsilon}{\kappa_B L_\psi} \right) + \mathbb{P}\left(\frac{2}{n} \|\mathbf{X}\|^2 \geq 2(1 + \rho_n) \right) \\ \quad + \mathbb{P}\left(\frac{1}{n} \|\mathbf{x}^t\|^2 \geq \kappa_x^2 \right) + \mathbb{P}\left(\frac{1}{n} \|\mathbf{h}^t + \mu_t^n \mathbf{X}\|^2 \geq \kappa_h^2 \right). \end{aligned} \quad (84)$$

Note, we have used $\rho_n \leq 1$ so $\frac{1}{\kappa_B} (1 + 2(1 + \rho_n) + \kappa_x^2 + \kappa_h^2) \leq 1$. Considering the result in (84), we notice that result (81) follows directly from (78)–(80), since by Hoeffding’s inequality (lemma 15),

$$\mathbb{P}\left(\frac{2}{n} \|\mathbf{X}\|^2 \geq 2(1 + \rho_n) \right) \leq \mathbb{P}\left(\left| \frac{1}{n} \|\mathbf{X}\|^2 - \rho_n \right| \geq 1 \right) \leq 2 \exp\{-2n\}. \quad (85)$$

Thus, (81) (hence, (82),) follows easily from (78)–(80) and the main technical piece of proving lemma 11 is then proving results (78)–(80) rigorously. This is done in section K.4.

Now we show that (11) follows from lemma 11 result (82), and then we finally prove theorem 2. Notice that (11) is recovered by applying (82) with pseudo-Lipschitz function $\tilde{\psi}(X_i, x_i^t) = \psi(X_i, f_t(x_i^t))$, as the only difference between (11) and (82) is that x_i^t in (82) is replaced with $f_t(x_i^t)$ in (11). With this choice of pseudo-Lipschitz function, an L_f^2 term is added in the denominator of the rate of concentration, since $L_{\tilde{\psi}} = 3L_\psi \max\{1, L_f\}$, which is shown in lemma 17.

We note that from (14) and (72) it is easy to see that $\tilde{\gamma}_n^t L_f^2 = \tilde{\gamma}_n^t \lambda_n = \gamma_n^t$, noting that $\hat{c}_t = \mathbb{E}[g_t'(\sqrt{\tau_t^n} Z, X)] = \mathbb{E}[f_t'(\sqrt{\tau_t^n} Z + \mu_t^n X)] = \hat{b}_t$. Moreover, the bound on the RHS of (11) equals bound_t defined in theorem 2 when L_ψ is a universal constant.

Now we prove theorem 2 using (11). First, notice that theorem 2 result (12) follows directly from (11) using pseudo-Lipschitz function $\psi(X_i, f_t(x_i^t)) = (X_i - f_t(x_i^t))^2$. This function is pseudo-Lipschitz with constant L_ψ by lemma 16. To see how this proves result (12) in more details, notice that

$$\frac{1}{n} \sum_{i=1}^n \psi(X_i, f_t(x_i^t)) = \frac{1}{n} \sum_{i=1}^n (X_i - f_t(x_i^t))^2 = \frac{1}{n} \|\mathbf{X} - f_t(\mathbf{x}^t)\|^2,$$

and

$$\begin{aligned} \mathbb{E}\left\{ \psi(X_0^n, f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z)) \right\} &= \mathbb{E}\left\{ (X_0^n - f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z))^2 \right\} \\ &= \mathbb{E}\left\{ (X_0^n)^2 \right\} + \mathbb{E}\left\{ [f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z)]^2 \right\} - 2\mathbb{E}\left\{ X_0^n f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z) \right\} = \rho_n - \tau_{t+1}^n, \end{aligned}$$

where the final uses that $\mathbb{E}\{(X_0^n)^2\} = \rho_n$ when $P_{X,n}$ is $\text{Ber}(\rho_n)$ or Bernoulli-Rademacher, the Law of Total Expectation to give $\mathbb{E}\{X_0^n f_t(\sqrt{\lambda_n} \tau_t^n X_0^n + \sqrt{\tau_t^n} Z)\} = \mathbb{E}\{[f_t(\sqrt{\lambda_n} \tau_t^n X_0^n + \sqrt{\tau_t^n} Z)]^2\}$ in the case of the conditional expectation denoiser as in (9), and the state evolution definition in (10).

Now we prove theorem 2 result (13). Now considering the concentration result in (13), notice that

$$\frac{1}{n^2} \|\mathbf{X}\mathbf{X}^T - f_t(\mathbf{x}^t)[f_t(\mathbf{x}^t)]^T\|_F^2 = \frac{1}{n^2} \|\mathbf{X}\|^4 + \frac{1}{n^2} \|f_t(\mathbf{x}^t)\|^4 - \frac{2}{n^2} \langle \mathbf{X}, f_t(\mathbf{x}^t) \rangle^2.$$

Then we will prove the following three results: for bound_t defined in the theorem 2 statement,

$$\mathbb{P}\left(\left|\frac{1}{n^2} \|\mathbf{X}\|^4 - \rho_n^2\right| \geq \epsilon\right) \leq 2e^{-2n\epsilon^2}, \quad (86)$$

$$\mathbb{P}\left(\left|\frac{1}{n^2} \|f_t(\mathbf{x}^t)\|^4 - (\tau_t^n)^2\right| \geq \epsilon\right) \leq \text{bound}_t, \quad (87)$$

$$\mathbb{P}\left(\left|\frac{1}{n^2} \langle \mathbf{X}, f_t(\mathbf{x}^t) \rangle^2 - (\tau_t^n)^2\right| \geq \epsilon\right) \leq \text{bound}_t. \quad (88)$$

Then the final concentration result in (13) follows from lemma 19 as follows:

$$\begin{aligned} & \mathbb{P}\left(\left|\frac{1}{n^2} \|\mathbf{X}\mathbf{X}^T - f_t(\mathbf{x}^t)[f_t(\mathbf{x}^t)]^T\|_F^2 - (\rho_n^2 - (\tau_{t+1}^n)^2)\right| \geq \epsilon\right) \\ &= \mathbb{P}\left(\left|\left(\frac{1}{n^2} \|\mathbf{X}\|^4 - \rho_n^2\right) + \left(\frac{1}{n^2} \|f_t(\mathbf{x}^t)\|^4 - (\tau_{t+1}^n)^2\right) - \left(\frac{2}{n^2} \langle \mathbf{X}, f_t(\mathbf{x}^t) \rangle^2 - 2(\tau_{t+1}^n)^2\right)\right| \geq \epsilon\right) \\ &\leq \mathbb{P}\left(\left|\frac{1}{n^2} \|\mathbf{X}\|^4 - \rho_n^2\right| \geq \frac{\epsilon}{3}\right) + \mathbb{P}\left(\left|\frac{1}{n^2} \|f_t(\mathbf{x}^t)\|^4 - (\tau_{t+1}^n)^2\right| \geq \frac{\epsilon}{3}\right) \\ &\quad + \mathbb{P}\left(\left|\frac{2}{n^2} \langle \mathbf{X}, f_t(\mathbf{x}^t) \rangle^2 - 2(\tau_{t+1}^n)^2\right| \geq \frac{\epsilon}{3}\right). \end{aligned}$$

As a final step, notice that the bounds in (86) - (88) applied to the above give the result in (13).

Now we prove (86) - (88). First we prove (86) using Hoeffding's Inequality, lemma 15,

$$\mathbb{P}\left(\left|\frac{1}{n} \|\mathbf{X}\|^2 - \rho_n\right| \geq \epsilon\right) = \mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n (X_i^2 - \mathbb{E}\{X_i^2\})\right| \geq \epsilon\right) \leq 2e^{-2n\epsilon^2}.$$

Then the result in (86) then follows from the above by lemma 20.

Next, for (87) we apply (11) using the function $\psi(X_i, f_t(x_i^t)) = [f_t(x_i^t)]^2$, which is pseudo-Lipschitz with constant $L_\psi = 2$ by lemma 16), to find

$$\mathbb{P}\left(\left|\frac{1}{n} \|f_t(\mathbf{x}^t)\|^2 - \tau_{t+1}^n\right| \geq \epsilon\right) = \mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n [f_t(x_i^t)]^2 - \tau_{t+1}^n\right| \geq \epsilon\right) \leq \text{bound}_t,$$

where we have used the definition of the state evolution in (10) to give

$$\mathbb{E}\left\{\psi(X_0^n, f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z))\right\} = \mathbb{E}\left\{[f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z)]^2\right\} = \tau_{t+1}^n.$$

Then the result in (87) follows from the above by lemma 20 and the fact that $\tau_t^n \leq \rho_n$.

Finally we prove result (88) by applying (11) using the function $\psi(X_i, f_t(x_i^t)) = X_i f_t(x_i^t)$, which is pseudo-Lipschitz with constant $L_\psi = 2$ by lemma 16, to find

$$\mathbb{P}\left(\left|\frac{1}{n} \langle \mathbf{X}, f_t(\mathbf{x}^t) \rangle - \tau_{t+1}^n\right| \geq \epsilon\right) = \mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n X_i f_t(x_i^t) - \tau_{t+1}^n\right| \geq \epsilon\right) \leq \text{bound}_t,$$

where

$$\mathbb{E}\left\{\psi(X_0^n, f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z))\right\} = \mathbb{E}\left\{X_0^n f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z)\right\} = \tau_{t+1}^n,$$

where the final step uses the Law of Total Expectation to give $\mathbb{E}\{X_0^n f_t(\mu_t^n X_0^n + \sqrt{\tau_t^n} Z)\} = \mathbb{E}\{[f_t(\mu_t X_0^n + \sqrt{\tau_t^n} Z)]^2\}$ in the case of the conditional expectation denoiser as in (9) and the state evolution definition in (10). Then the result in (88) follows from the above by lemma 20 $\tau_t^n \leq \rho_n$.

K.3 Proof of theorem 4

Proof. The proof of theorem 4 proceeds in two steps. In the first step, one studies the conditional distribution of \mathbf{Z} given the output of the algorithm up until iteration t , treating \mathbf{Z} as random and the output as deterministic. In the non-symmetric AMP studied in [63, Theorem 1], the relevant measurement matrix has i.i.d. gaussian entries and this conditional distribution was originally studied in [27]. The result for the case of i.i.d. gaussian \mathbf{Z} is concisely stated in [63, Lemma 4.2]. For the symmetric AMP of (69) that we are interested in, the matrix \mathbf{Z} is $\text{GOE}(n)$ and so this conditioning argument needs to take into account the symmetry of the matrix entries (and consequently the added dependencies). This has been studied in other works that give asymptotic characterizations of the performance of symmetric AMP, for example in [66, Lemma 3], and these results apply directly to our case since this distributional characterization is already non-asymptotic and does not change in our setting. This then allows us to characterize the conditional distribution of the iterates \mathbf{h}^{t+1} , conditional on the previous output of the algorithm. We give this result in Lemma 12 below, but before stating the lemma, we introduce some useful notation.

First, denote $\mathbf{m}^0 := g_0(\mathbf{h}^0, \mathbf{X}^n), \dots, \mathbf{m}^t := g_t(\mathbf{h}^t, \mathbf{X}^n)$ where the terms $g_t(\mathbf{h}^t, \mathbf{X}^n)$ are those used in the symmetric AMP in (69). Then we define \mathcal{S}_0 to be the sigma-algebra generated by $\{g_0(\mathbf{h}^0, \mathbf{X}^n), \mathbf{X}^n\}$ and \mathcal{S}_t for $t \geq 1$ to be the sigma-algebra generated by

$$\mathbf{h}^1, \dots, \mathbf{h}^t, \mathbf{m}^0, \dots, \mathbf{m}^t, \text{ and } \mathbf{X}^n.$$

Using [66, Lemma 3] to characterize the distribution of \mathbf{Z} conditioned on the sigma algebra \mathcal{S}_t , we are able to specify the conditional distributions of \mathbf{h}^{t+1} given \mathcal{S}_t , by observing that conditioning on \mathcal{S}_t for $t \geq 1$ is equivalent to conditioning on the linear constraint⁵

$$\mathbf{Z}\mathbf{M}_t = \mathbf{Y}_t,$$

where $\mathbf{M}_t \in \mathbb{R}^{n \times t}$ and $\mathbf{H}_t \in \mathbb{R}^{n \times t}$ are the matrices

$$\mathbf{M}_t = [\mathbf{m}^0 \mid \dots \mid \mathbf{m}^{t-1}] \quad \text{and} \quad \mathbf{H}_t = [\mathbf{h}^1 \mid \dots \mid \mathbf{h}^t],$$

and $\mathbf{Y}_t \in \mathbb{R}^{n \times t}$ is the matrix $\mathbf{Y}_1 = \mathbf{H}_1$ and $\mathbf{Y}_t = \mathbf{H}_t + [\mathbf{0} \mid \mathbf{M}_{t-1}] \mathbf{C}_t$ for $t \geq 2$, where $\mathbf{C}_t = \text{diag}(c_0, \dots, c_{t-1})$. Note that $[c_1 \mid c_2 \mid \dots \mid c_k]$ denotes a matrix with columns c_1, \dots, c_k .

We use the notation $\mathbf{m}_{\parallel}^{t+1}$ to denote the projection of \mathbf{m}^{t+1} onto the column space of \mathbf{M}_{t+1} . Let

$$\boldsymbol{\alpha}^{t+1} := (\alpha_0^{t+1}, \alpha_1^{t+1}, \dots, \alpha_t^{t+1})^\top \in \mathbb{R}^{t+1}, \quad (89)$$

be the coefficient vectors of these projections, i.e., $\mathbf{m}_{\parallel}^{t+1} := \sum_{i=0}^t \alpha_i^t \mathbf{m}^i$, meaning $\boldsymbol{\alpha}^t = (\mathbf{M}_{t+1}^\top \mathbf{M}_{t+1})^{-1} \mathbf{M}_{t+1}^\top \mathbf{m}^{t+1}$. The projections of \mathbf{m}^{t+1} onto the orthogonal complement of \mathbf{M}_{t+1} , is denoted by $\mathbf{m}_{\perp}^{t+1} := \mathbf{m}^{t+1} - \mathbf{m}_{\parallel}^{t+1}$. Lemma 13 shows that for large n , the entries of $\boldsymbol{\alpha}$ concentrate around constants. In what follows we show that, for $t \geq 0$, the vector $\boldsymbol{\alpha}^{t+1} \in \mathbb{R}^{t+1}$ in (89) concentrates to the vector

$$\hat{\boldsymbol{\alpha}}^{t+1} := \left[0, \dots, 0, \frac{\sigma_{t+2}^n}{\sigma_{t+1}^n} \right]^\top \in \mathbb{R}^{t+1}, \quad (90)$$

for the state evolution values given in (70). Similarly, Lemma 13 will show that for large n , the norm $\|\mathbf{m}_{\perp}^{t-1}\|^2/n$ concentrates to a constant σ_t^\perp , defined as $\sigma_1^\perp = \sigma_1^n$, and for $t \geq 2$,

$$\sigma_t^\perp := \sigma_t^n \left(1 - \frac{\sigma_t^n}{\sigma_{t-1}^n} \right). \quad (91)$$

With the above notation, we find the following result for the symmetric AMP in (69).

Lemma 12 (Conditional Distribution Lemma). *For the vectors \mathbf{h}^{t+1} defined in (69), the following hold for $t \geq 1$, provided $n > t$ and $\mathbf{M}_t^\top \mathbf{M}_t$ has full column rank.*

$$\mathbf{h}^1|_{\mathcal{S}_0} \stackrel{d}{=} \sqrt{\sigma_1^n} \mathbf{U}_0 + \boldsymbol{\Delta}_0, \quad \text{and} \quad \mathbf{h}^{t+1}|_{\mathcal{S}_t} \stackrel{d}{=} \hat{\alpha}_t^{t+1} \mathbf{h}^t + \sqrt{\sigma_{t+1}^\perp} \mathbf{U}_t + \boldsymbol{\Delta}_t, \quad (92)$$

⁵While conditioning on the linear constraints, we emphasize that only \mathbf{A} is treated as random.

where $U_0, U_t \in \mathbb{R}^n$ are random vectors with elements that are marginally standard gaussian random variables that are independent of the corresponding conditioning sigma-algebras. The terms $\hat{\alpha}_i^t$ for $i \in \{0, 1, \dots, t\}$ are defined in (90) and the terms σ_t^\perp in (91). The deviation terms are $\Delta_t = \mathbf{0}$ and for $t > 0$,

$$\begin{aligned} \Delta_t &= \sum_{r=1}^t (\alpha_r^{t+1} - \hat{\alpha}_r^{t+1}) \mathbf{h}^r + \left[\left(\frac{\|\mathbf{m}_\perp^t\|}{\sqrt{n}} - \sqrt{\sigma_{t+1}^\perp} \right) \mathbf{I} - \frac{\|\mathbf{m}_\perp^t\|}{\sqrt{n}} \mathbf{P}_{\mathbf{M}_t} \right] U_t \\ &+ \mathbf{M}_{t-1} (\mathbf{M}_{t-1}^\top \mathbf{M}_{t-1})^{-1} \left[\mathbf{H}_{t-1}^\top \mathbf{m}_\perp^t - \mathbf{M}_{t-1}^* \left(c_t \mathbf{m}^{t-1} - \sum_{i=1}^{t-1} c_i \alpha_i^t \mathbf{m}^{i-1} \right) \right]. \end{aligned} \quad (93)$$

The second step of the proof is inductive on the iteration t , showing that if the result in (71) holds up to iteration $t-1$ then it will hold at iteration t as well. This is done by showing that the standardized ℓ_2 norms of the terms in the AMP algorithm in (69), like $\frac{1}{n} \|\mathbf{h}^t\|^2$ and $\frac{1}{n} \|\mathbf{m}^t\|^2$, concentrate on deterministic values predicted by the state evolution. This is done by relating these iteration t values to the iteration $t-1$ values through the iteration in (71) and appealing to the conditional distributions from Lemma 12 and your inductive hypothesis. While the proof for the symmetric AMP is largely similar to that for the non-symmetric version, there are additional challenges due to the dependencies created by the symmetry of the $\text{GOE}(n)$ matrix. The details of the inductive proof are quite technical and long and are therefore not included here but we state the result for the symmetric AMP in (69).

For $t \geq 0$, let $\kappa_{-1} = K_{-1} = 1$, and

$$K_t = C(t+1)^5 K_{t-1}, \quad \kappa_t = \frac{\kappa_{t-1}}{c(t+1)^{11}}, \quad (94)$$

where $C, c > 0$ are universal constants (not depending on t, n , or ϵ). To keep the notation compact, we use K, κ, κ' to denote generic positive universal constants whose values may change through the lemma statement.

The result of theorem 4 follows from lemma 13 result (97) below.

Lemma 13. *The following statements hold for $1 \leq t < T^*$ and $\epsilon \in (0, 1)$. Define*

$$\gamma_n^{t+1} := (\nu^n + \sigma_1^n)(\nu^n + \sigma_1^n + \sigma_2^n) \cdots (\nu^n + \sum_{i=1}^{t+1} \sigma_i^n) \times \max\{1, \hat{c}_1\} \max\{1, \hat{c}_2\} \cdots \max\{1, \hat{c}_t\}, \quad (95)$$

where ν is the variance factor of sub-gaussian \mathbf{X}^n (for $P_{X,n} = \text{Ber}(\rho_n)$ we have $\nu^n \leq 1/4$; see lemma 14).

1.
$$\mathbb{P} \left(\frac{1}{n} \|\Delta_t\|^2 \geq \epsilon \right) \leq K(t+1)^2 K_{t-1} \exp \left\{ - \frac{\kappa \kappa_{t-1} n \epsilon}{(t+1)^4 L_g^{4t} \gamma_n^t \max\{1, \hat{c}_t\}} \right\}. \quad (96)$$

2. Denote $\mathbb{E}_\phi := \mathbb{E} \phi_h(\sqrt{\sigma_1^n} \tilde{Z}_1, \dots, \sqrt{\sigma_{t+1}^n} \tilde{Z}_{t+1}, X^n)$. Then for pseudo-Lipschitz function $\phi : \mathbb{R}^{t+1} \rightarrow \mathbb{R}$ with constant L_ϕ we have that

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \phi(h_i^1, \dots, h_i^{t+1}, X_i^n) - \mathbb{E}_\phi \right| \geq \epsilon \right) \leq K(t+1)^3 K_{t-1} \exp \left\{ - \frac{\kappa \kappa_{t-1} n \epsilon^2}{(t+1)^7 L_\phi^2 L_g^{4t} \gamma_n^{t+1}} \right\}. \quad (97)$$

The random variables $\tilde{Z}_0, \dots, \tilde{Z}_t$ are jointly gaussian with zero mean and covariance given by $\mathbb{E}[\tilde{Z}_r \tilde{Z}_t] = \sqrt{\sigma_r^n / \sigma_t^n}$ for $r < t$, and are independent of $\mathbf{X}^n \sim p_{X,n}$.

Denote $\mathbb{E}_{\phi_{Lip}} := \mathbb{E} \phi_{Lip}(\sqrt{\sigma_1^n} \tilde{Z}_1, \dots, \sqrt{\sigma_{t+1}^n} \tilde{Z}_{t+1}, X^n)$. Then for Lipschitz function $\phi_{Lip} : \mathbb{R}^{t+1} \rightarrow \mathbb{R}$ with constant $L_{\phi_{Lip}}$ we have that

$$\begin{aligned} &\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n m_i^0 \phi_{Lip}(h_i^1, \dots, h_i^{t+1}, X_i^n) - \tilde{\sigma} \mathbb{E}_{\phi_{Lip}} \right| \geq \epsilon \right) \\ &\leq K(t+1)^3 K_{t-1} \exp \left\{ - \frac{\kappa \kappa_{t-1} n \epsilon^2}{(t+1)^7 L_\phi^2 L_g^{4t} \gamma_n^{t+1}} \right\}. \end{aligned} \quad (98)$$

The random variables $\tilde{Z}_0, \dots, \tilde{Z}_t$ are as above.

3. Let $L_g > 0$ be the pseudo-Lipschitz constant for the denoiser functions $\{g_t\}_{t \geq 0}$ and let $X_n \doteq c$ be shorthand for

$$\mathbb{P}(|X_n - c| \geq \epsilon) \leq K(t+1)^3 K_{t-1} \exp \left\{ -\frac{\kappa \kappa_{t-1} n \epsilon^2}{(t+1)^7 L_g^{4t+2} \gamma_n^{t+1}} \right\}.$$

For all $0 \leq r \leq t$,

$$\mathbb{P} \left(\left| \frac{1}{n} (\mathbf{h}^{r+1})^* \mathbf{h}^{t+1} - \sigma_{t+1}^n \right| \geq \epsilon \right) \leq K(t+1)^3 K_{t-1} \exp \left\{ -\frac{\kappa \kappa_{t-1} n \epsilon^2}{(t+1)^7 L_g^{4t} \gamma_n^{t+1}} \right\}.$$

$$\frac{1}{n} (\mathbf{m}^0)^* \mathbf{m}^{t+1} \doteq \tilde{\sigma}^n \mathbb{E}[g_{t+1}(\sqrt{\sigma_{t+1}^n} \tilde{Z}_{t+1}, X^n)], \quad \frac{1}{n} (\mathbf{m}^{r+1})^* \mathbf{m}^{t+1} \doteq \sigma_{t+2}^n.$$

$$\mathbb{P} \left(|c_{t+1} - \hat{c}_{t+1}| \geq \epsilon \right) \leq K(t+1)^3 K_{t-1} \exp \left\{ -\frac{\kappa \kappa_{t-1} n \epsilon^2}{(t+1)^7 L_g^{4(t+1)} \gamma_n^{t+1}} \right\}.$$

$$\frac{1}{n} (\mathbf{h}^{t+1})^* \mathbf{m}^{r+1} \doteq \hat{c}_{r+1} \sigma_{t+2}^n, \quad \frac{1}{n} (\mathbf{h}^{r+1})^* \mathbf{m}^{t+1} \doteq \hat{c}_{t+1} \sigma_{t+2}^n.$$

4.

$$\mathbb{P} \left(\frac{1}{n} \mathbf{M}_{t+1}^* \mathbf{M}_{t+1} \text{ is singular} \right) \leq (t+1) K_{t-1} \exp \left\{ -\frac{\kappa_{t-1} \kappa n}{(t+1)^7 L_g^{4t+2} \gamma_n^{t+1}} \right\}. \quad (99)$$

For $\hat{\alpha}^{t+1}$ defined in (90), when the inverse of $\frac{1}{n} \mathbf{M}_{t+1}^* \mathbf{M}_{t+1}$ exists, for $1 \leq i, j \leq t+1$,

$$\mathbb{P} \left(\left| \left[\left(\frac{1}{n} \mathbf{M}_{t+1}^* \mathbf{M}_{t+1} \right)^{-1} - (\mathbf{C}^{t+1})^{-1} \right]_{i,j} \right| \geq \epsilon \right) \leq K K_{t-1} \exp \left\{ -\frac{\kappa \kappa_{t-1} n \epsilon^2}{L_g^{4t+2} \gamma_n^{t+1}} \right\}, \quad (100)$$

$$\mathbb{P} \left(|\alpha_{i-1}^{t+1} - \hat{\alpha}_{i-1}^{t+1}| \geq \epsilon \right) \leq K(t+1)^4 K_{t-1} \exp \left\{ -\frac{\kappa \kappa_{t-1} n \epsilon^2}{(t+1)^9 L_g^{4(t+1)} \gamma_n^{t+1}} \right\}.$$

In the above, the matrix $\mathbf{C}^{t+1} \in \mathbb{R}^{(t+1) \times (t+1)}$ has elements $[\mathbf{C}^{t+1}]_{i,j} = \sigma_{\max\{i,j\}}^n$ for $1 \leq i, j \leq t+1$.

5. With σ_t^\perp defined in (91),

$$\mathbb{P} \left(\left| \frac{1}{n} \|\mathbf{m}_\perp^{t+1}\|^2 - \sigma_{t+2}^\perp \right| \geq \epsilon \right) \leq K(t+1)^5 K_{t-1} \exp \left\{ -\frac{\kappa \kappa_{t-1} n \epsilon^2}{(t+1)^{11} L_g^{4(t+1)} \gamma_n^{t+1}} \right\}. \quad (101)$$

□

K.4 Proof of lemma 11

Proof. To begin with, we prove result (78) then we prove the other results, (79)–(82), inductively.

Before we do so we establish upper and lower bounds on τ_n^t defined in (10). Notice that for the Bernoulli case,

$$\tau_{t+1}^n = \mathbb{E} \left[\frac{\rho_n^2}{\rho_n + (1 - \rho_n) \exp \left\{ -\frac{1}{2} \lambda_n \tau_t^n - \sqrt{\lambda_n \tau_t^n} Z \right\}} \right],$$

where $Z \sim \mathcal{N}(0, 1)$, as shown in appendix G result (55). Therefore, trivially $\tau_{t+1}^n \leq \rho_n$. We also wish to establish a lower bound. First, by Jensen's Inequality applied to the convex function $f(x) = 1/x$ on $x \in (0, \infty)$, we have that

$$\tau_{t+1}^n \geq \frac{\rho_n^2}{\rho_n + (1 - \rho_n) \mathbb{E} \left[\exp \left\{ -\frac{1}{2} \lambda_n \tau_t^n - \sqrt{\lambda_n \tau_t^n} Z \right\} \right]} \stackrel{(a)}{=} \rho_n^2,$$

where step (a) uses that $\mathbb{E}[\exp\{-\frac{1}{2} \lambda_n \tau_t^n - \sqrt{\lambda_n \tau_t^n} Z\}] = 1$ since $\mathbb{E}[\exp\{-tZ\}] = \exp\{\frac{1}{2} t^2\}$. Thus, $\rho_n^2 \leq \tau_{t+1}^n \leq \rho_n$ and, in the regime of interest where $\lambda_n = \kappa \rho_n^{-2}$, using that $\mu_t^n = \sqrt{\lambda_n \tau_t^n}$ by (10), we find $(\mu_t^n)^2 = \lambda_n (\tau_t^n)^2 = \kappa \rho_n^{-2} (\tau_t^n)^2$ and therefore

$$\kappa' \rho_n^2 \leq (\mu_t^n)^2 = \kappa \rho_n^{-2} (\tau_t^n)^2 \leq \kappa. \quad (102)$$

Result (78). We first show that (78) follows immediately from theorem 4. Applying theorem 4 with pseudo-Lipschitz function $\phi(h_i^t, X_i^n) = (h_i^t + \mu_t^n X_i^n)^2$, having constant $L_\phi = 2 \max\{1, (\mu_t^n)^2\}$, as is shown in lemma 17,

$$\mathbb{P}\left(\left|\frac{1}{n}\|\mathbf{h}^t + \mu_t^n \mathbf{X}^n\|^2 - (\mu_t^n)^2 \rho_n - \tau_t^n\right| \geq \epsilon\right) \leq CC_t \exp\left\{\frac{-cc_t n \epsilon^2}{\max\{1, (\mu_t^n)^4\} \tilde{\gamma}_t^n}\right\} \leq \text{bound}_t, \quad (103)$$

where the final inequality follows since $(\mu_t^n)^4 \leq \kappa$ as discussed above in (102). We have also used that $\mathbb{E}[\phi(\sqrt{\tau_t^n} Z, X_0^n)] = \mathbb{E}[(\sqrt{\tau_t^n} Z + \mu_t^n X_0^n)^2] = \tau_t^n + (\mu_t^n)^2 \rho_n$. Then, since $(\mu_t^n)^2 \rho_n \leq 1$, choosing $\kappa_h^2 = 3 > 1 + (\mu_t^n)^2 \rho_n + \tau_t^n$, we recover the first result in (78) with

$$\mathbb{P}\left(\frac{1}{\sqrt{n}}\|\mathbf{h}^t + \mu_t^n \mathbf{X}\| \geq \kappa_h\right) \leq \mathbb{P}\left(\left|\frac{1}{n}\|\mathbf{h}^t + \mu_t^n \mathbf{X}\|^2 - (\mu_t^n)^2 \rho_n - \tau_t^n\right| \geq 1\right).$$

Other results (79)–(82). The proof is inductive on the iteration t . We first show the initialization case $t = 1$. Consider (79), then using the definitions of $\mathbf{x}^1 = \frac{\sqrt{\lambda_n}}{n} \mathbf{X} \langle \mathbf{X}, f_0(\mathbf{x}^0) \rangle + \mathbf{Z} f_0(\mathbf{x}^0)$ from (74) and $\mathbf{h}^1 = \mathbf{Z} g_0(\mathbf{h}^0, \mathbf{X})$ from (75) along with the fact that $f_0(\mathbf{x}^0) = g_0(\mathbf{h}^0, \mathbf{X})$,

$$\mathbf{x}^1 - \mathbf{h}^1 - \mu_1^n \mathbf{X} = \mathbf{X} \left(\frac{\sqrt{\lambda_n}}{n} \langle \mathbf{X}, f_0(\mathbf{x}^0) \rangle - \mu_1^n \right) = 0,$$

where the final inequality follows since $\mu_1^n = \sqrt{\lambda_n} \langle f_0(\mathbf{x}^0), \mathbf{X} \rangle / n$ by (7). Next for result (80), first notice that by the Triangle Inequality, $\|\mathbf{x}^1\| \leq \|\mathbf{x}^1 - \mathbf{h}^1 - \mu_1^n \mathbf{X}\| + \|\mathbf{h}^1 + \mu_1^n \mathbf{X}\|$. Then let $\kappa_x = 2\kappa_h + 2\kappa$ and therefore, by lemma 19,

$$\begin{aligned} \mathbb{P}\left(\frac{1}{\sqrt{n}}\|\mathbf{x}^1\| \geq \kappa_x\right) &\leq \mathbb{P}\left(\frac{1}{\sqrt{n}}\|\mathbf{x}^1 - \mathbf{h}^1 - \mu_1^n \mathbf{X}\| + \frac{1}{\sqrt{n}}\|\mathbf{h}^1 + \mu_1^n \mathbf{X}\| \geq 2\kappa_h + 2\kappa\right) \\ &\leq \mathbb{P}\left(\frac{1}{\sqrt{n}}\|\mathbf{x}^1 - \mathbf{h}^1 - \mu_1^n \mathbf{X}\| \geq \kappa\right) + \mathbb{P}\left(\frac{1}{\sqrt{n}}\|\mathbf{h}^1 + \mu_1^n \mathbf{X}\| \geq \kappa_h\right). \end{aligned}$$

Then the upper bound follows by (78) and (79).

Next, notice that with the bound on $\mathbb{P}(\frac{1}{\sqrt{n}}\|\mathbf{x}^1\| \geq \kappa_x)$ established above, one can prove the $t = 1$ case for (81) and (82), as justified in the work in (83) – (84) along with results (78) and (79).

Now assume that all results (80)–(82) hold up until iteration $t - 1$ and we prove the results for iteration t . As justified in the work in (83) – (84), the results (81) and (82) follow immediately from (78) – (80) so we only aim to prove (79) and (80) here (as (78) was demonstrated above in (103)). We begin by proving (79) which we will then use to prove (80).

Result (79). Next we consider result (79). Using the definitions of \mathbf{x}^{t+1} and \mathbf{h}^{t+1} from (74) and (75) along with Cauchy-Schwarz inequality, we have that

$$\begin{aligned} &\frac{1}{n} \sum_{i=1}^n \left| x_i^t - h_i^t - \mu_t^n X_i \right|^2 \\ &\leq \frac{3}{n} \left| \frac{\sqrt{\lambda_n}}{n} \langle \mathbf{X}, f_{t-1}(\mathbf{x}^{t-1}) \rangle - \mu_t^n \right|^2 \sum_{i=1}^n X_i^2 + \frac{3}{n} \sum_{i=1}^n \left| [\mathbf{Z} f_{t-1}(\mathbf{x}^{t-1})]_i - [\mathbf{Z} g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X})]_i \right|^2 \\ &\quad + \frac{3}{n} \sum_{i=1}^n \left| \mathbf{b}_{t-1} f_{t-2}(x_i^{t-2}) - \mathbf{c}_{t-1} g_{t-2}(h_i^{t-2}, X_i) \right|^2 \\ &= 3 \left| \frac{\sqrt{\lambda_n}}{n} \langle \mathbf{X}, f_{t-1}(\mathbf{x}^{t-1}) \rangle - \mu_t^n \right|^2 \frac{1}{n} \sum_{i=1}^n X_i^2 + \frac{3}{n} \left\| \mathbf{Z} (f_{t-1}(\mathbf{x}^{t-1}) - g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X})) \right\|^2 \\ &\quad + \frac{3}{n} \left\| \mathbf{b}_{t-1} f_{t-2}(\mathbf{x}^{t-2}) - \mathbf{c}_{t-1} g_{t-2}(\mathbf{h}^{t-2}, \mathbf{X}) \right\|^2. \end{aligned} \quad (104)$$

Now we use the upper bounds in (104) along with lemma 19 to give the following upper bound on the probability on the LHS of (79):

$$\begin{aligned} \mathbb{P}\left(\frac{1}{n}\left\|\mathbf{x}^t - \mathbf{h}^t - \mu_t^n \mathbf{X}\right\|^2 \geq \frac{\kappa\epsilon^2}{L_\psi^2}\right) &\leq \mathbb{P}\left(\left|\frac{\sqrt{\lambda_n}}{n}\langle \mathbf{X}, f_{t-1}(\mathbf{x}^{t-1}) \rangle - \mu_t^n\right|^2 \frac{1}{n} \sum_{i=1}^n X_i^2 \geq \frac{\kappa\epsilon}{L_\psi^2}\right) \\ &\quad + \mathbb{P}\left(\frac{1}{n}\left\|\mathbf{Z}(f_{t-1}(\mathbf{x}^{t-1}) - g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X}))\right\|^2 \geq \frac{\kappa\epsilon}{L_\psi^2}\right) \\ &\quad + \mathbb{P}\left(\frac{1}{n}\left\|\mathbf{b}_{t-1}f_{t-2}(\mathbf{x}^{t-2}) - \mathbf{c}_{t-1}g_{t-2}(\mathbf{h}^{t-2}, \mathbf{X})\right\|^2 \geq \frac{\kappa\epsilon}{L_\psi^2}\right). \end{aligned} \tag{105}$$

We label the three terms in the above T_1, T_2, T_3 and provide an upper bound for each.

First consider term T_1 of (105), and recall that $\mu_{t-1}^n = \sqrt{\lambda_n}\tau_{t-1}^n$. Thus, we have the upper bound

$$T_1 \leq \mathbb{P}\left(\left|\frac{1}{n}\langle \mathbf{X}, f_{t-1}(\mathbf{x}^{t-1}) \rangle - \tau_{t-1}^n\right| \geq \frac{\kappa\sqrt{\epsilon}}{\sqrt{(1+\rho_n)\lambda_n L_\psi}}\right) + \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n X_i^2 \geq (1+\rho_n)\right). \tag{106}$$

Notice that we can upper bound the second term in (106) with using $2e^{-2n}$ Hoeffding's inequality (lemma 15) as in (85). We can upper bound the first term in (106) using the induction hypothesis for result (79) for the pseudo-Lipschitz function $\tilde{\psi}(a, b) = af_{t-1}(b)$ with constant $L_{\tilde{\psi}} = L_f$. Thus,

$$\mathbb{P}\left(\left|\frac{1}{n}\langle \mathbf{X}, f_{t-1}(\mathbf{x}^{t-1}) \rangle - \tau_{t-1}^n\right| \geq \frac{\kappa\sqrt{\epsilon}}{\sqrt{(1+\rho_n)\lambda_n L_\psi}}\right) \leq CC_{t-1} \exp\left\{\frac{-cc_{t-1}n\epsilon^2}{L_\psi^2 L_f^2 \lambda_n (1+\rho_n) \tilde{\gamma}_n^{t-1}}\right\}.$$

Finally we notice that the desired result follows since $L_f = \sqrt{\lambda_n}$, proved in lemma 18, and $\lambda_n^2(1-\rho_n)\tilde{\gamma}_n^{t-1} \leq \tilde{\gamma}_n^t$ using the definition of $\tilde{\gamma}_n^t$ in (72).

Now consider term T_2 of (105). First notice that, conditioned on event \mathcal{F}_{t-1} ,

$$\begin{aligned} \frac{1}{\sqrt{n}}\left\|\mathbf{Z}(f_{t-1}(\mathbf{x}^{t-1}) - g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X}))\right\| &\leq \frac{1}{\sqrt{n}}\|\mathbf{Z}\|_{op}\|f_{t-1}(\mathbf{x}^{t-1}) - g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X})\| \\ &\stackrel{(a)}{\leq} \frac{1}{\sqrt{n}}\|\mathbf{Z}\|_{op}\|f_{t-1}(\mathbf{x}^{t-1}) - f_{t-1}(\mathbf{h}^{t-1} + \mu_{t-1}^n \mathbf{X})\| \stackrel{(b)}{\leq} \|\mathbf{Z}\|_{op}\sqrt{\frac{\lambda_n}{n}}\left\|\mathbf{x}^{t-1} - \mathbf{h}^{t-1} - \mu_{t-1}^n \mathbf{X}\right\|, \end{aligned}$$

where step (a) uses that $g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X}) = f_{t-1}(\mathbf{h}^{t-1} + \mu_{t-1}^n \mathbf{X})$ and step (b) uses the Lipschitz property of f_{t-1} . Therefore,

$$\begin{aligned} &\mathbb{P}\left(\frac{1}{\sqrt{n}}\left\|\mathbf{Z}(f_{t-1}(\mathbf{x}^{t-1}) - g_{t-1}(\mathbf{h}^{t-1}, \mathbf{X}))\right\| \geq \frac{\kappa\sqrt{\epsilon}}{L_\psi}\right) \\ &\leq \mathbb{P}\left(\|\mathbf{Z}\|_{op}\sqrt{\frac{\lambda_n}{n}}\left\|\mathbf{x}^{t-1} - \mathbf{h}^{t-1} - \mu_{t-1}^n \mathbf{X}\right\| \geq \frac{\kappa\epsilon}{L_\psi}\right) \\ &\leq \mathbb{P}\left(\frac{1}{\sqrt{n}}\left\|\mathbf{x}^{t-1} - \mathbf{h}^{t-1} - \mu_{t-1}^n \mathbf{X}\right\| \geq \frac{\kappa\epsilon}{L_\psi\sqrt{\lambda_n}}\right) + \mathbb{P}\left(\|\mathbf{Z}\|_{op} \geq \kappa\right) \\ &\leq CC_{t-1} \exp\left\{\frac{-cc_{t-1}n\epsilon^2}{L_\psi^2 \tilde{\gamma}_n^t}\right\} + C \exp\{-cn\}, \end{aligned}$$

where the final inequality follows from the inductive hypothesis for (79) and standard results about tail bounds for operator norms of GOE matrices. In particular, we have used the inductive hypothesis to find

$$\begin{aligned} \mathbb{P}\left(\frac{1}{n}\left\|\mathbf{x}^{t-1} - \mathbf{h}^{t-1} + \mu_{t-1}^n \mathbf{X}\right\|^2 \geq \frac{\kappa\epsilon^2}{L_\psi^2 \lambda_n}\right) &\leq CC_{t-1} \exp\left\{\frac{-cc_{t-1}n\epsilon^2}{L_\psi^2 \lambda_n \tilde{\gamma}_n^{t-1}}\right\} \\ &\leq CC_{t-1} \exp\left\{\frac{-cc_{t-1}n\epsilon^2}{L_\psi^2 \tilde{\gamma}_n^t}\right\}, \end{aligned}$$

where the final inequality follows since $\lambda_n \tilde{\gamma}_n^{t-1} \leq \tilde{\gamma}_n^t$ using the definition of $\tilde{\gamma}_n^t$ in (72).

Finally, consider term T_3 of (105). We first give an upper bound using the definition of g_t and the Lipschitz property of f_t with $L_f = \sqrt{\lambda_n}$ as follows:

$$\begin{aligned}
& \left\| \mathbf{b}_{t-1} f_{t-2}(\mathbf{x}^{t-2}) - \mathbf{c}_{t-1} g_{t-2}(\mathbf{h}^{t-2}, \mathbf{X}) \right\| = \left\| \mathbf{b}_{t-1} f_{t-2}(\mathbf{x}^{t-2}) - \mathbf{c}_{t-1} f_{t-2}(\mathbf{h}^{t-2} + \mu_{t-2}^n \mathbf{X}) \right\| \\
& \leq |\mathbf{b}_{t-1}| \left\| f_{t-2}(\mathbf{x}^{t-2}) - f_{t-2}(\mathbf{h}^{t-2} + \mu_{t-2}^n \mathbf{X}) \right\| + |\mathbf{b}_{t-1} - \mathbf{c}_{t-1}| \left\| f_{t-2}(\mathbf{h}^{t-2} + \mu_{t-2}^n \mathbf{X}) \right\| \\
& \leq \lambda_n \left\| \mathbf{x}^{t-2} - \mathbf{h}^{t-2} - \mu_{t-2}^n \mathbf{X} \right\| + |\mathbf{b}_{t-1} - \mathbf{c}_{t-1}| \sqrt{n}.
\end{aligned} \tag{107}$$

In the final step we use the lemma 18 results

$$|\mathbf{b}_{t-1}| \leq \frac{1}{n} \sum_{i=1}^n |f'_{t-1}(x_i^{t-1})| \leq \sqrt{\lambda_n}, \quad \text{and} \quad \left\| f_{t-2}(\mathbf{h}^{t-2} + \mu_{t-2}^n \mathbf{X}) \right\|^2 \leq n.$$

We investigate the term $|\mathbf{b}_{t-1} - \mathbf{c}_{t-1}|$ and recall from their definitions in (6) and (75),

$$\begin{aligned}
|\mathbf{b}_{t-1} - \mathbf{c}_{t-1}| & \leq \frac{1}{n} \sum_{i=1}^n |f'_{t-1}(x_i^{t-1}) - g'_{t-1}(h_i^{t-1}, X_i)| = \frac{1}{n} \sum_{i=1}^n |f'_{t-1}(x_i^{t-1}) - f'_{t-1}(h_i^{t-1} - \mu_{t-1}^n X_i)| \\
& \stackrel{(a)}{=} \frac{\sqrt{\lambda_n}}{n} \sum_{i=1}^n \left| f_{t-1}(x_i^{t-1})(1 - f_{t-1}(x_i^{t-1})) - f_{t-1}(h_i^{t-1} - \mu_{t-1}^n X_i)(1 - f_{t-1}(h_i^{t-1} - \mu_{t-1}^n X_i)) \right| \\
& \stackrel{(b)}{\leq} \frac{\sqrt{\lambda_n}}{n} \sum_{i=1}^n \left| f_{t-1}(x_i^{t-1}) - f_{t-1}(h_i^{t-1} - \mu_{t-1}^n X_i) \right| \stackrel{(c)}{\leq} \frac{\lambda_n}{\sqrt{n}} \left\| \mathbf{x}^{t-1} - \mathbf{h}^{t-1} - \mu_{t-1}^n \mathbf{X} \right\|.
\end{aligned} \tag{108}$$

In the above, step (a) uses lemma 18 for computing the derivative f_t , step (b) uses the bound

$$|f(a)(1 - f(a)) - f(b)(1 - f(b))| \leq |f(a) - f(b)| + |[f(a)]^2 - [f(b)]^2| \leq \kappa |f(a) - f(b)|,$$

for $0 \leq f(a) \leq 1$ for all $a \in \mathbb{R}$, and step (c) uses the Lipschitz property of f_t , namely $L_f = \sqrt{\lambda_n}$, and Cauchy Schwarz to give $\sum_{i=1}^n |a_i| \leq \sqrt{n} \|\mathbf{a}\|$. Plugging the bound in (108) into (107),

$$\begin{aligned}
& \left\| \mathbf{b}_{t-1} f_{t-2}(\mathbf{x}^{t-2}) - \mathbf{c}_{t-1} g_{t-2}(\mathbf{h}^{t-2}, \mathbf{X}) \right\| \\
& \leq \lambda_n \left\| \mathbf{x}^{t-2} - \mathbf{h}^{t-2} - \mu_{t-2}^n \mathbf{X} \right\| + \lambda_n \left\| \mathbf{x}^{t-1} - \mathbf{h}^{t-1} - \mu_{t-1}^n \mathbf{X} \right\|.
\end{aligned}$$

Now we have from lemma 19 that

$$T_3 \leq 2\mathbb{P}\left(\frac{1}{n} \left\| \mathbf{x}^{t-1} - \mathbf{h}^{t-1} - \mu_{t-1}^n \mathbf{X} \right\|^2 \geq \frac{\kappa \epsilon}{L_\psi^2 \lambda_n^2}\right),$$

and the final bound follows from the inductive hypothesis for (81) using that $\lambda_n^2 \tilde{\gamma}_n^{t-1} \leq \tilde{\gamma}_n^t$ using the definition of $\tilde{\gamma}_n^t$ in (72).

Result (80). To complete the proof, we consider result (80). First notice that by the Triangle Inequality, $\|\mathbf{x}^t\| \leq \|\mathbf{x}^t - \mathbf{h}^t - \mu_t^n \mathbf{X}\| + \|\mathbf{h}^t + \mu_t^n \mathbf{X}\|$. Then let $\kappa_x = 2\kappa_h + \frac{2\kappa}{L_\psi}$ and therefore, by lemma 19,

$$\begin{aligned}
\mathbb{P}\left(\frac{1}{\sqrt{n}} \|\mathbf{x}^t\| \geq \kappa_x\right) & = \mathbb{P}\left(\frac{1}{\sqrt{n}} \|\mathbf{x}^t - \mathbf{h}^t - \mu_t^n \mathbf{X}\| + \frac{1}{\sqrt{n}} \|\mathbf{h}^t + \mu_t^n \mathbf{X}\| \geq 2\kappa_h + \frac{2\kappa}{L_\psi}\right) \\
& \leq \mathbb{P}\left(\frac{1}{\sqrt{n}} \|\mathbf{x}^t - \mathbf{h}^t - \mu_t^n \mathbf{X}\| \geq \frac{\kappa}{L_\psi}\right) + \mathbb{P}\left(\frac{1}{\sqrt{n}} \|\mathbf{h}^t + \mu_t^n \mathbf{X}\| \geq \kappa_h\right).
\end{aligned}$$

Then the bound follows by (80) and (79). \square

K.5 Useful lemmas

In this section we introduce a number of technical lemmas that are used to prove our main results. We include proofs only where the proof is non-standard.

Lemma 14. *Recall from [67], that a random variable, X , is sub-gaussian with variance factor ν if $\log \mathbb{E}[e^{t(X - \mathbb{E}[X])}] \leq t^2 \nu / 2$ for all $t \in \mathbb{R}$. When $X \sim p_X$, we have $\nu \leq 1/4$ for $p_X \sim \text{Ber}(\rho)$.*

Proof. By [67, Lemma 2.2 (Hoeffding's lemma)], any bounded random variable $a \leq X \leq b$ is sub-Gaussian with variance factor $(b - a)^2/4$. \square

Lemma 15 (Hoeffding's inequality). *If X_1, \dots, X_n are independent bounded random variables such that $a_i \leq X_i \leq b_i$, then for $\nu = 2[\sum_i (b_i - a_i)^2]^{-1}$, we have $\mathbb{P}(|\frac{1}{n} \sum_{i=1}^n (X_i - \mathbb{E}\{X_i\})| \geq \epsilon) \leq 2e^{-\nu n^2 \epsilon^2}$.*

Lemma 16. *Recall the definition of pseudo-Lipschitz functions of order 2 given in Definition 1. The following functions $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}$ are all pseudo-Lipschitz of order 2 with pseudo-Lipschitz constant 2.*

$$\psi_1(a, b) = (a - b)^2, \quad \psi_2(a, b) = b^2, \quad \psi_3(a, b) = ab. \quad (109)$$

Proof. Verifying the pseudo-Lipschitz property for the functions in (109) is straightforward, so we omit the details. \square

Lemma 17. *Recall the definition of pseudo-Lipschitz functions of order 2 given in Definition 1. Let f_t be the conditional expectation denoiser in (9) with Lipschitz constant L_f^n and let $\psi : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a pseudo-Lipschitz of order 2 function with constant L_ψ . The following functions $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ are all pseudo-Lipschitz of order 2 with the stated pseudo-Lipschitz constants.*

$$\phi_1(a, b) = \psi(a + \mu_t^n b, b), \quad L_{\phi_1} = 2L_\psi(1 + \mu_t^n)^2, \quad (110)$$

$$\phi_2(a, b) = \psi(a, f_t(b)), \quad L_{\phi_2} = 3L_\psi \max\{1, L_f\}, \quad (111)$$

$$\phi_3(a, b) = (\mu_t^n)^{-1} a + b, \quad L_{\phi_3} = \sqrt{2} \max\{1, (\mu_t^n)^{-1}\}, \quad (112)$$

$$\phi_4(a, b) = (a + \mu_t^n b)^2, \quad L_{\phi_4} = 2 \max\{1, (\mu_t^n)^2\}. \quad (113)$$

Proof. For function ϕ_1 in (110), first notice

$$\begin{aligned} |\phi_1(a, b) - \phi_1(\tilde{a}, \tilde{b})| &= |\psi(a + \mu_t^n b, b) - \psi(\tilde{a} + \mu_t^n \tilde{b}, \tilde{b})| \\ &\leq L_\psi(1 + \|(a + \mu_t^n b, b)\| + \|(\tilde{a} + \mu_t^n \tilde{b}, \tilde{b})\|) \times \|(a + \mu_t^n b, b) - (\tilde{a} + \mu_t^n \tilde{b}, \tilde{b})\|. \end{aligned} \quad (114)$$

Next notice

$$\|(a + \mu_t^n b, b) - (\tilde{a} + \mu_t^n \tilde{b}, \tilde{b})\| \leq |a - \tilde{a}| + (1 + \mu_t^n)|b - \tilde{b}| \leq \sqrt{2}(1 + \mu_t^n)\|(a, b) - (\tilde{a}, \tilde{b})\|,$$

and $\|(a + \mu_t^n b, b)\| \leq |a + \mu_t^n b| + |b| \leq |a| + (1 + \mu_t^n)|b| \leq \sqrt{2}(1 + \mu_t^n)\|(a, b)\|$. Thus, from (114), we have result (110):

$$|\phi_1(a, b) - \phi_1(\tilde{a}, \tilde{b})| \leq 2L_\psi(1 + \mu_t^n)^2(1 + \|(a, b)\| + \|(\tilde{a}, \tilde{b})\|) \times \|(a, b) - (\tilde{a}, \tilde{b})\|.$$

For function ϕ_2 in (111), first notice

$$\begin{aligned} |\phi_2(a, b) - \phi_2(\tilde{a}, \tilde{b})| &= |\psi(a, f_t(b)) - \psi(\tilde{a}, f_t(\tilde{b}))| \\ &\leq L_\psi(1 + \|(a, f_t(b))\| + \|(\tilde{a}, f_t(\tilde{b}))\|) \|(a, f_t(b)) - (\tilde{a}, f_t(\tilde{b}))\|. \end{aligned} \quad (115)$$

Next, notice that since $f_t(\cdot)$ is a Lipschitz function with constant L_f ,

$$\begin{aligned} \|(a, f_t(b)) - (\tilde{a}, f_t(\tilde{b}))\|^2 &= |f_t(b) - f_t(\tilde{b})|^2 + |a - \tilde{a}|^2 \\ &\leq L_f^2|b - \tilde{b}|^2 + |a - \tilde{a}|^2 \leq \max\{1, L_f^2\} \|(a, b) - (\tilde{a}, \tilde{b})\|^2, \end{aligned}$$

and since our denoiser of interest f_t in (9) is such that $|f_t(x)| \leq 1$,

$$\|(a, f_t(b))\| \leq |a| + |f_t(b)| \leq |a| + 1 \leq (1 + |a| + |b|) \leq \sqrt{2}(1 + \|(a, b)\|).$$

Thus, from (116),

$$|\tilde{\psi}(a, b) - \tilde{\psi}(\tilde{a}, \tilde{b})| \leq 3L_\psi \max\{1, L_f\} (1 + \|(a, b)\| + \|(\tilde{a}, \tilde{b})\|) \|(a, b) - (\tilde{a}, \tilde{b})\|. \quad (116)$$

Next, the bound for function ϕ_3 in (112) is straightforward:

$$\begin{aligned} |\phi_3(a, b) - \phi_3(\tilde{a}, \tilde{b})| &= \left| (\mu_t^n)^{-1} a + b - (\mu_t^n)^{-1} \tilde{a} - \tilde{b} \right| \\ &\leq (\mu_t^n)^{-1} |a - \tilde{a}| + |b - \tilde{b}| \leq \sqrt{2} \max\{1, (\mu_t^n)^{-1}\} \|(a, b) - (\tilde{a}, \tilde{b})\|. \end{aligned}$$

Finally, for function ϕ_4 in (113), first notice

$$\begin{aligned} |\phi_4(a, b) - \phi_4(\tilde{a}, \tilde{b})| &= \left| (a + \mu_t^n b)^2 - (\tilde{a} + \mu_t^n \tilde{b})^2 \right| \\ &\leq \left| (a + \mu_t^n b) - (\tilde{a} + \mu_t^n \tilde{b}) \right| \left| (a + \mu_t^n b) + (\tilde{a} + \mu_t^n \tilde{b}) \right| \\ &\leq 2 \max\{1, (\mu_t^n)^2\} (\|(a, b)\| + \|(\tilde{a}, \tilde{b})\|) (\|(a, b) - (\tilde{a}, \tilde{b})\|), \end{aligned}$$

where the final inequality uses that $|a + \mu_t^n b| \leq \sqrt{2} \max\{1, \mu_t^n\} \|(a, b)\|$ giving

$$\left| (a + \mu_t^n b) + (\tilde{a} + \mu_t^n \tilde{b}) \right| \leq \sqrt{2} \max\{1, \mu_t^n\} (\|(a, b)\| + \|(\tilde{a}, \tilde{b})\|)$$

and the fact that

$$\left| (a + \mu_t^n b) - (\tilde{a} + \mu_t^n \tilde{b}) \right| \leq |a - \tilde{a}| + \mu_t^n |b - \tilde{b}| \leq \sqrt{2} \max\{1, \mu_t^n\} \|(a, b) - (\tilde{a}, \tilde{b})\|.$$

□

Lemma 18. Recall the definition of pseudo-Lipschitz functions of order 2 given in Definition 1. The conditional expectation denoiser in (9) is Lipschitz with constant $L_f = \sqrt{\lambda_n}$ when $X_0^n \sim P_{X,n}$ and $P_{X,n}$ is either $\text{Ber}(\rho_n)$ or Bernoulli-Rademacher and $\frac{\partial}{\partial x} f_t(x) = \sqrt{\lambda_n} f_t(x)(1 - f_t(x))$. Moreover, the Lipschitz constant can be strengthened to $\sqrt{\lambda_n} \rho_n$ on $x \in (-\infty, \frac{\mu_t^n}{2})$ and $f_t(0) \leq \rho_n$.

Proof. First, recall that $f_t(\cdot)$ is the conditional expectation denoiser given in (9),

$$f_t(x) = \mathbb{E}\{X_0^n \mid \sqrt{\lambda_n} \tau_t^n X_0^n + \sqrt{\tau_t^n} Z = x\}.$$

Notice that for either the Bernoulli or Bernoulli-Rademacher case, we have that $|f_t(x)| \leq 1$ for all $x \in \mathbb{R}$ since $X_0^n \in \{-1, 0, 1\}$.

First consider $P_{X,n} \sim \text{Ber}(\rho_n)$ and we show that $f_t(\cdot)$ is Lipschitz continuous with Lipschitz constant $\sqrt{\lambda_n}$. Let $\phi(x)$ denote the standard gaussian density evaluated at x . First, by Bayes' Rule,

$$\begin{aligned} f_t(x) &= \mathbb{E}\{X_0^n \mid \sqrt{\lambda_n} \tau_t^n X_0^n + \sqrt{\tau_t^n} Z = x\} \\ &= \mathbb{P}(X_0^n = 1 \mid \sqrt{\lambda_n} \tau_t^n X_0^n + \sqrt{\tau_t^n} Z = x) = \frac{\rho_n \phi\left(\frac{x - \sqrt{\lambda_n} \tau_t^n}{\sqrt{\tau_t^n}}\right)}{(1 - \rho_n) \phi\left(\frac{x}{\sqrt{\tau_t^n}}\right) + \rho_n \phi\left(\frac{x - \sqrt{\lambda_n} \tau_t^n}{\sqrt{\tau_t^n}}\right)}. \end{aligned} \quad (117)$$

Now notice that $\frac{\partial}{\partial x} \phi\left(\frac{x-a}{b}\right) = -\frac{(x-a)}{b^2} \phi\left(\frac{x-a}{b}\right)$. Using this and the representation above,

$$\begin{aligned} \frac{\partial}{\partial x} f_t(x) &= \frac{\partial}{\partial x} \left[\frac{\rho_n \phi\left(\frac{x - \sqrt{\lambda_n} \tau_t^n}{\sqrt{\tau_t^n}}\right)}{(1 - \rho_n) \phi\left(\frac{x}{\sqrt{\tau_t^n}}\right) + \rho_n \phi\left(\frac{x - \sqrt{\lambda_n} \tau_t^n}{\sqrt{\tau_t^n}}\right)} \right] \\ &= \frac{-f_t(x)}{\tau_t^n} \left[(x - \sqrt{\lambda_n} \tau_t^n) - \frac{x(1 - \rho_n) \phi\left(\frac{x}{\sqrt{\tau_t^n}}\right) + \rho_n (x - \sqrt{\lambda_n} \tau_t^n) \phi\left(\frac{x - \sqrt{\lambda_n} \tau_t^n}{\sqrt{\tau_t^n}}\right)}{(1 - \rho_n) \phi\left(\frac{x}{\sqrt{\tau_t^n}}\right) + \rho_n \phi\left(\frac{x - \sqrt{\lambda_n} \tau_t^n}{\sqrt{\tau_t^n}}\right)} \right] \\ &= \frac{-f_t(x)(x - \sqrt{\lambda_n} \tau_t^n)}{\tau_t^n} \left[\frac{(1 - \rho_n) \phi\left(\frac{x}{\sqrt{\tau_t^n}}\right) \left[1 - \frac{x}{(x - \sqrt{\lambda_n} \tau_t^n)}\right]}{(1 - \rho_n) \phi\left(\frac{x}{\sqrt{\tau_t^n}}\right) + \rho_n \phi\left(\frac{x - \sqrt{\lambda_n} \tau_t^n}{\sqrt{\tau_t^n}}\right)} \right] \\ &= \sqrt{\lambda_n} f_t(x) (1 - f_t(x)). \end{aligned} \quad (118)$$

Therefore, using (118), we see that $\left| \frac{\partial}{\partial x} f_t(x) \right| \leq \sqrt{\lambda_n}$ and it follows that $f_t(\cdot)$ is Lipschitz continuous with Lipschitz constant $\sqrt{\lambda_n}$.

The fact that $f_t(\cdot)$ is Lipschitz continuous with Lipschitz constant $\sqrt{\lambda_n}$ can be shown similarly for the case where $P_{X,n}$ is Bernoulli-Rademacher.

Finally, notice that from (117) we have

$$f_t(x) = \frac{\rho_n}{(1 - \rho_n) \exp\left\{\frac{1}{2}(\lambda_n \tau_t^n - 2x\sqrt{\lambda_n})\right\} + \rho_n}. \quad (119)$$

Then since $e^x \geq 1 + x$ (which can be seen by showing that $f(x) = e^x - (1 + x)$ has a minimum at $f(0) = 0$),

$$f_t(x) \leq \frac{\rho_n}{(1 - \rho_n)(1 + \frac{1}{2}(\lambda_n \tau_t^n - 2x\sqrt{\lambda_n})) + \rho_n} = \frac{\rho_n}{1 + \frac{1}{2}(1 - \rho_n)(\lambda_n \tau_t^n - 2x\sqrt{\lambda_n})}.$$

The above implies that $f_t(0) \leq \rho_n$, and further, since

$$(1 - \rho_n)(\lambda_n \tau_t^n - 2x\sqrt{\lambda_n}) \geq 0 \quad \text{when} \quad x \leq \frac{\sqrt{\lambda_n} \tau_t^n}{2},$$

we find the bound $0 \leq f_t(x) \leq \rho_n$ when $x \leq \frac{\sqrt{\lambda_n} \tau_t^n}{2}$.

Therefore, by (118), we have $\left| \frac{\partial}{\partial x} f_t(x) \right| \leq \sqrt{\lambda_n} f_t(x) \leq \sqrt{\lambda_n} \rho_n$ and it follows that $f_t(\cdot)$ is Lipschitz continuous with Lipschitz constant $\sqrt{\lambda_n} \rho_n$ on $x \in (-\infty, \frac{\mu_t^n}{2})$. \square

The proof of the following two lemmas can be found in [63, appendix A].

Lemma 19 (Concentration of Sums). *If random variables X_1, \dots, X_M satisfy $P(|X_i| \geq \epsilon) \leq e^{-n\kappa_i \epsilon^2}$ for $1 \leq i \leq M$, then*

$$\mathbb{P}\left(\left|\sum_{i=1}^M X_i\right| \geq \epsilon\right) \leq \sum_{i=1}^M \mathbb{P}\left(|X_i| \geq \frac{\epsilon}{M}\right) \leq M e^{-n(\min_i \kappa_i) \epsilon^2 / M^2}.$$

Lemma 20 (Concentration of Powers). *Assume $c > 0$ and $0 < \epsilon \leq 1$. Then, if $\mathbb{P}(|X_n - c| \geq \epsilon) \leq e^{-\kappa n \epsilon^2}$, it follows that $\mathbb{P}(|X_n^2 - c^2| \geq \epsilon) \leq e^{-\kappa n \epsilon^2 / [1+2c]^2}$.*

L Algorithmic AMP phase transition regime

In this appendix we show that the right-hand side of the bound in theorem 2 for signal strength and sparsity scaling like $\lambda_n \rho_n^2 = w$ and $\rho_n = \Omega((\ln n)^{-\alpha})$ with $w, \alpha \in \mathbb{R}_+$, tends to zero as $n \rightarrow +\infty$. We focus on the Bernoulli prior case but the arguments generalizes to Bernoulli-Rademacher prior.

Let us first upper bound γ_n^t in terms of λ_n and ρ_n in the Bernoulli case. First we use the bound $|f'_t(x)| \leq \sqrt{\lambda_n}$ (see lemma 18) to bound

$$\max\{1, \hat{\mathbf{b}}_1\} \max\{1, \hat{\mathbf{b}}_2\} \cdots \max\{1, \hat{\mathbf{b}}_{t-1}\} \leq \lambda_n^{\frac{t-1}{2}}.$$

From the explicit AMP iteration (see appendix G second formula for example) we have $\tau_t^n \leq \rho_n$. Since $\nu_n \leq 1/4$ (see lemma 14) we get $(\nu^n + \tau_1^n)(\nu^n + \tau_1^n + \tau_2^n) \cdots (\nu^n + \sum_{i=1}^t \tau_i^n) \leq (\frac{1}{4} + \rho_n)(\frac{1}{4} + 2\rho_n) \cdots (\frac{1}{4} + t\rho_n) \leq (\frac{1}{4} + t)!$. Putting everything together we get:

$$\gamma_n^t \leq \left(\frac{1}{4} + t\right)! \lambda_n^{2t-1 + \frac{t-1}{2}}.$$

Now we use the scaling (which is the correct scale for the phase transition to happen) $\lambda_n = w \rho_n^{-2}$ and get:

$$\gamma_n^t \leq \left(\frac{1}{4} + t\right)! \frac{w^{\frac{5t-3}{2}}}{\rho_n^{5t-3}}.$$

Therefore

$$\text{bound}_t \leq CC_t \exp \left\{ -\frac{c}{w^{\frac{5t-3}{2}}} \frac{c_t}{(\frac{1}{4} + t)!} \rho_n^{5t-3} n \epsilon^2 \right\}.$$

Now the t dependence in the constant $c_t = [C^t(t!)^C]^{-1}$ (from now on C is a generic positive constant) and using Stirling's approximation $t! \approx \sqrt{2\pi t} t^t e^{-t}$ this scales at dominant order as $[C^t(t!)^C]^{-1}$. So we have at dominant order

$$\text{bound}_t \approx CC_t \exp\{-C e^{\pm Ct - Ct \ln t} e^{(5t-3) \ln(\rho_n)} e^{\ln n} \epsilon^2\}.$$

Now set the number of iterations to $t = o(\frac{\ln n}{\ln \ln n})$. We get $t \ln t = o(\ln n)$ so $\pm Ct - Ct \ln t = o(\ln n)$ and

$$\text{bound}_t \approx CC_t \exp\{-C e^{-o(\ln n)} e^{o(\frac{\ln n}{\ln \ln n}) \ln(\rho_n)} e^{\ln n} \epsilon^2\}. \quad (120)$$

We set $\rho_n = \Theta(\frac{1}{(\ln n)^\alpha}) = \frac{C}{(\ln n)^\alpha}$. Then $\ln \rho_n = \ln C - \alpha \ln \ln n$ and we get

$$\text{bound}_t \approx CC_t \exp\{-C e^{-o(\ln n)} e^{o(\frac{\ln n}{\ln \ln n})(C - \alpha \ln \ln n)} e^{\ln n} \epsilon^2\}.$$

This leads to

$$\begin{aligned} \text{bound}_t &\approx CC_t \exp\{-C e^{-o(\ln n) + Co(\frac{\ln n}{\ln \ln n}) - \alpha o(\ln n)} e^{\ln n} \epsilon^2\} \\ &\approx CC_t \exp\{-C e^{(\ln n) - (1+\alpha)o(\ln n) + Co(\frac{\ln n}{\ln \ln n})} \epsilon^2\} \\ &\approx CC_t \exp\{-C e^{(\ln n)[1 - (1+\alpha)o(1)] + Co(\frac{1}{\ln \ln n})} \epsilon^2\} \\ &\approx CC_t \exp\{-C n^{1-o_\alpha(1)} \epsilon^2\}. \end{aligned}$$

One can check that the prefactor $C_t = [C^t(t!)^C]$ does not change the dominant order for $t = o(\frac{\ln n}{\ln \ln n})$. This shows that the bound vanishes as $n \rightarrow +\infty$ for $\lambda = w\rho_n^{-2}$ and $\rho_n = \Theta(\frac{1}{(\ln n)^\alpha})$ for any $\alpha \geq 0$. As seen from (120) the bound worsen with decreasing ρ_n . So the result extends to $\rho_n = \Omega(\frac{1}{(\ln n)^\alpha})$.

Note also that in the case of the rescaled bound of remark 1 below theorem 2, the previous derivation is unchanged, up to the constant appearing in the $o_\alpha(1)$ that is changed some other $o_\alpha(1)$ (for n big enough). Indeed, because $\rho_n = \Omega(\frac{1}{(\ln n)^\alpha})$ the ρ_n^2 or ρ_n^4 appearing in the rescaled bound can be absorbed in the $o_\alpha(1)$ of the previous derivation, for n large enough.