

Sampling from the Sherrington-Kirkpatrick Gibbs measure via algorithmic stochastic localization

Ahmed El Alaoui* Andrea Montanari† Mark Sellke‡

Abstract

We consider the Sherrington-Kirkpatrick model of spin glasses at high-temperature and no external field, and study the problem of sampling from the Gibbs distribution μ in polynomial time. We prove that, for any inverse temperature $\beta < 1/2$, there exists an algorithm with complexity $O(n^2)$ that samples from a distribution μ^{alg} which is close in normalized Wasserstein distance to μ . Namely, there exists a coupling of μ and μ^{alg} such that if $(\mathbf{x}, \mathbf{x}^{\text{alg}}) \in \{-1, +1\}^n \times \{-1, +1\}^n$ is a pair drawn from this coupling, then $n^{-1} \mathbb{E}\{\|\mathbf{x} - \mathbf{x}^{\text{alg}}\|_2^2\} = o_n(1)$. The best previous results, by Bauerschmidt and Bodineau [BB19] and by Eldan, Koehler, Zeitouni [EKZ21], implied efficient algorithms to approximately sample (under a stronger metric) for $\beta < 1/4$.

We complement this result with a negative one, by introducing a suitable “stability” property for sampling algorithms, which is verified by many standard techniques. We prove that no stable algorithm can approximately sample for $\beta > 1$, even under the normalized Wasserstein metric.

Our sampling method is based on an algorithmic implementation of stochastic localization, which progressively tilts the measure μ towards a single configuration, together with an approximate message passing algorithm that is used to approximate the mean of the tilted measure.

Contents

1	Introduction	2
2	Main results	4
2.1	Sampling algorithm for $\beta < 1/2$	4
2.1.1	Approximating the mean of the Gibbs measure	4
2.1.2	Sampling via stochastic localization	5
2.1.3	Theoretical guarantee	6
2.2	Hardness for stable algorithms, for $\beta > 1$	7
2.3	Notations	8
3	Properties of stochastic localization	8
4	Analysis of the sampling algorithm and proof of Theorem 2.1	10
4.1	The planted model and contiguity	10
4.2	Approximate Message Passing	12
4.3	Natural Gradient Descent	19
4.4	Continuous limit and proof of Theorem 2.1	22

*Department of Statistics and Data Science, Cornell University

†Department of Electrical Engineering and Department of Statistics, Stanford University

‡Department of Mathematics, Stanford University

5	Algorithmic stability and disorder chaos	26
5.1	Algorithmic stability: Proof of Theorem 2.3	26
5.2	Hardness for stable algorithms: Proof of Theorems 2.5 and 2.6	29
A	Convergence analysis of Natural Gradient Descent	34
A.1	Technical preliminaries	34
A.2	Proof of Lemma 4.10	39

1 Introduction

This Sherrington-Kirkpatrick (SK) Gibbs measure is the probability distribution over $\Sigma_n = \{-1, +1\}^n$ given by

$$\mu_{\mathbf{A}}(\mathbf{x}) = \frac{1}{Z(\beta, \mathbf{A})} \exp \left\{ \frac{\beta}{2} \langle \mathbf{x}, \mathbf{A} \mathbf{x} \rangle \right\}, \tag{1.1}$$

where $\beta \geq 0$ is an inverse temperature parameter and $\mathbf{A} \sim \text{GOE}(n)$; i.e., \mathbf{A} is symmetric, $A_{ij} \sim \text{N}(0, 1/n)$ i.i.d. for $i \leq j \leq n$ and $A_{ii} \sim \text{N}(0, 2/n)$, $i \leq n$. The parameter β is fixed and we will leave implicit the dependence of μ upon β , unless mentioned otherwise.

In this paper, we consider the problem of efficiently sampling from the Sherrington-Kirkpatrick spin glass measure. Namely, we seek a randomized algorithm that accepts as input \mathbf{A} and generates $\mathbf{x}^{\text{alg}} \sim \mu_{\mathbf{A}}^{\text{alg}}$, such that: (i) The algorithm runs in polynomial time (for any \mathbf{A}); (ii) The distribution $\mu_{\mathbf{A}}^{\text{alg}}$ is close to $\mu_{\mathbf{A}}$ for typical realizations of \mathbf{A} . Given a bounded distance $\text{dist}(\mu, \nu)$ between probability distributions μ, ν , the second condition can be formalized by requiring $\mathbb{E}\{\text{dist}(\mu_{\mathbf{A}}, \mu_{\mathbf{A}}^{\text{alg}})\} = o_n(1)$.

Gibbs sampling (also known in this context as Glauber dynamics) provides an algorithm to approximately sample from $\mu_{\mathbf{A}}$. However, standard techniques to bound its mixing time (e.g., Dobrushin condition [AH87]) only imply polynomial mixing for a vanishing interval of temperatures $\beta = O(n^{-1/2})$. By contrast, physicists [SZ81, MPV87] predict fast convergence to equilibrium (at least for certain observables) for all $\beta < 1$.

Significant progress on this question was achieved only recently. In [BB19], Bauerschmidt and Bodineau showed that, for $\beta < 1/4$, the measure $\mu_{\mathbf{A}}$ can be decomposed into a log-concave mixture of product measures. They use this decomposition to prove that $\mu_{\mathbf{A}}$ satisfies a log-Sobolev inequality, although not for the Dirichlet form of Glauber dynamics¹. Eldan, Koehler, Zeitouni [EKZ21] prove that, in the same region $\beta < 1/4$, $\mu_{\mathbf{A}}$ satisfies a Poincaré inequality for the Dirichlet form of Glauber dynamics. Hence Glauber dynamics mixes in $O(n^2)$ spin flips in total variation distance. This mixing time estimate was improved to $O(n \log n)$ by [AJK⁺21] using a modified log Sobolev inequality, see also [CE22, Corollary 51]. The aforementioned results apply deterministically to any matrix \mathbf{A} satisfying $\beta(\lambda_{\max}(\mathbf{A}) - \lambda_{\min}(\mathbf{A})) \leq 1 - \varepsilon$ (for some constant $\varepsilon > 0$).

For *spherical* spin glasses, it is shown in [GJ19] that Langevin dynamics have a polynomial spectral gap at high temperature. Meanwhile [BAJ18] proves that at sufficiently low temperature, the mixing times of Glauber and Langevin dynamics are exponentially large in Ising and spherical spin glasses, respectively.

In this paper we develop a different approach which is not based on a Monte Carlo Markov Chain strategy. We build on the well known remark that approximate sampling can be reduced to approximate computation of expectations of the measure $\mu_{\mathbf{A}}$, and of a family of measures obtained from $\mu_{\mathbf{A}}$. One well known method to achieve this reduction is via sequential sampling [JVV86, CDHL05, BD11]. A sequential

¹We note in passing that their result immediately suggests a sampling algorithm: sample from the log-concave mixture using Langevin dynamics, and then sample from the corresponding component using the product form.

sampling approach to $\mu_{\mathbf{A}}$ would proceed as follows. Order the variables $x_1, \dots, x_n \in \{-1, +1\}$ arbitrarily. At step i compute the marginal distribution of x_i , conditional to x_1, \dots, x_{i-1} taking the previously chosen values: $p_s^{(i)} := \mu_{\mathbf{A}}(x_i = s | x_1, \dots, x_{i-1})$, $s \in \{-1, +1\}$. Fix $x_i = +1$ with probability $p_{+1}^{(i)}$ and $x_i = -1$ with probability $p_{-1}^{(i)}$.

We follow a different route, which is similar in spirit, but that we find more convenient technically, and of potential practical interest. Our approach is motivated by the stochastic localization process [Eld20]. Given any probability measure μ on \mathbb{R}^n with finite second moment, positive time $t > 0$, and vector $\mathbf{y} \in \mathbb{R}^n$, define the tilted measure

$$\mu_{\mathbf{y},t}(\mathrm{d}\mathbf{x}) := \frac{1}{Z(\mathbf{y})} e^{\langle \mathbf{y}, \mathbf{x} \rangle - \frac{t}{2} \|\mathbf{x}\|_2^2} \mu(\mathrm{d}\mathbf{x}), \quad (1.2)$$

and let its mean vector be

$$\mathbf{m}(\mathbf{y}, t) := \int_{\mathbb{R}^n} \mathbf{x} \mu_{\mathbf{y},t}(\mathrm{d}\mathbf{x}). \quad (1.3)$$

Consider the stochastic differential equation² (SDE)

$$\mathrm{d}\mathbf{y}(t) = \mathbf{m}(\mathbf{y}(t), t) \mathrm{d}t + \mathrm{d}\mathbf{B}(t), \quad \mathbf{y}(0) = 0, \quad (1.4)$$

where $(\mathbf{B}(t))_{t \geq 0}$ is a standard Brownian motion in \mathbb{R}^n . Then, the measure-valued process $(\mu_{\mathbf{y}(t),t})_{t \geq 0}$ is a martingale and (almost surely) $\mu_{\mathbf{y}(t),t} \Rightarrow \delta_{\mathbf{x}^*}$ as $t \rightarrow \infty$, for some random \mathbf{x}^* (i.e. the measure localizes). As a consequence of the martingale property, $\mathbb{E}[\int \varphi(\mathbf{x}) \mu_{\mathbf{y}(t),t}(\mathrm{d}\mathbf{x})]$ is a constant for any bounded continuous function φ , whence $\mathbb{E}[\varphi(\mathbf{x}^*)] = \int \varphi(\mathbf{x}) \mu(\mathrm{d}\mathbf{x})$. In other words, \mathbf{x}^* is a sample from μ . For further information on this process, we refer to Section 3.

In order to use this process as an algorithm to sample from the SK measure $\mu = \mu_{\mathbf{A}}$, we need to overcome two problems:

- *Discretization.* We need to discretize the SDE (1.4) in time, and still guarantee that the discretization closely tracks the original process. This is of course possible only if the map $\mathbf{y} \mapsto \mathbf{m}(\mathbf{y}, t)$ is sufficiently regular.
- *Mean computation.* We need to be able to compute the mean vector $\mathbf{m}(\mathbf{y}, t)$ efficiently. To this end, we use an approximate message passing (AMP) algorithm for which we can leverage earlier work [DAM17] to establish that $\|\mathbf{m}(\mathbf{y}) - \widehat{\mathbf{m}}_{\text{AMP}}(\mathbf{y})\|_2^2/n = o_n(1)$ along the algorithm trajectory. (Note that the SK measure is supported on vectors with $\|\mathbf{x}\|_2^2 = n$, and hence the quadratic component of the tilt in Eq. (1.2) drops out. We will therefore write $\mathbf{m}(\mathbf{y})$ or $\mathbf{m}(\mathbf{A}, \mathbf{y})$ instead $\mathbf{m}(\mathbf{y}, t)$ for the mean of the Gibbs measure.)

To our knowledge, ours is the first algorithmic implementation of the stochastic localization process, although a recent paper by Nam, Sly and Zhang [NSZ22] uses this process (without naming it as such) to show that the Ising measure on the infinite regular tree is a factor of IID process up to a constant factor away from the Kesten–Stigum, or “reconstruction”, threshold. Their construction can easily be transformed into a sampling algorithm.

In order to state our results, we define the normalized 2-Wasserstein distance between two probability measures μ, ν on \mathbb{R}^n with finite second moments as

$$W_{2,n}(\mu, \nu)^2 = \inf_{\pi \in \mathcal{C}(\mu, \nu)} \frac{1}{n} \mathbb{E}_{\pi} \left[\|\mathbf{X} - \mathbf{Y}\|_2^2 \right], \quad (1.5)$$

where the infimum is over all couplings $(\mathbf{X}, \mathbf{Y}) \sim \pi$ with marginals $\mathbf{X} \sim \mu$ and $\mathbf{Y} \sim \nu$.

In this paper, we establish two main results.

²If μ has finite variance, then $\mathbf{y} \rightarrow \mathbf{m}(\mathbf{y}, t)$ is Lipschitz and so this SDE is well posed with unique strong solution.

Sampling algorithm for $\beta < 1/2$. We prove that the strategy outlined above yields an algorithm with complexity $O(n^2)$, which samples from a distribution $\mu_{\mathbf{A}}^{\text{alg}}$ such that $W_{2,n}(\mu_{\mathbf{A}}^{\text{alg}}, \mu_{\mathbf{A}}) = o_{\mathbb{P},n}(1)$.

Hardness for stable algorithms, for $\beta > 1$. We prove that no algorithm satisfying a certain *stability* property can sample from the SK measure (under the same criterion $W_{2,n}(\mu_{\mathbf{A}}^{\text{alg}}, \mu_{\mathbf{A}}) = o_{\mathbb{P},n}(1)$) for $\beta > 1$, i.e., when replica symmetry is broken. Roughly speaking, stability formalizes the notion that the algorithm output behaves continuously with respect to the matrix \mathbf{A} .

It is worth pointing out that we expect our algorithm to be successful (in the sense described above) for all $\beta < 1$ and that closing the gap between $\beta = 1/2$ and $\beta = 1$ should be within reach of existing techniques, at the price of a longer technical argument. We expound on this point in Remark 2.1 further below, and in Section 4.3.

The hardness results for $\beta > 1$ are proven using the notion of disorder chaos, in a similar spirit to the use of the *overlap gap property* for random optimization, estimation, and constraint satisfaction problems [GS14, RV17, GS17, CGPR19, GJ21, GJW20, Wei22, GK21, BH21, GJW21, HS21]. While the overlap gap property has been used to rule out stable algorithms for this class of problems, and variants have been used to rule out efficient sampling by specific Markov chain algorithms, to the best of our knowledge we are the first to rule out stable sampling algorithms using these ideas. In sampling there is no hidden solution or set of solutions to be found, and therefore no notion of an overlap gap in the most natural sense. Instead, we argue directly that the distribution to be sampled from is unstable in a $W_{2,n}$ sense at low temperature, and hence cannot be approximated by any stable algorithm.

The rest of the paper is organized as follows. In Section 2 we formally state our results. In Section 3 we collect some useful properties of the stochastic localization process, and we present the analysis of our algorithm in Section 4. Finally, the proof of hardness under stability is given in Section 5.

2 Main results

2.1 Sampling algorithm for $\beta < 1/2$

In this section we describe the sampling algorithm, and formally state the result of our analysis. As pointed out in the introduction, a main component is the computation of the mean of the tilted SK measure:

$$\mu_{\mathbf{A},\mathbf{y}}(\mathbf{x}) := \frac{1}{Z(\mathbf{A},\mathbf{y})} \exp \left\{ \frac{\beta}{2} \langle \mathbf{x}, \mathbf{A}\mathbf{x} \rangle + \langle \mathbf{y}, \mathbf{x} \rangle \right\}, \quad \mathbf{x} \in \{-1, +1\}^n. \quad (2.1)$$

We describe the algorithm to approximate this mean in Section 2.1.1, the overall sampling procedure (which uses this estimator as a subroutine) in Section 2.1.2, and our Wasserstein-distance guarantee in Section 2.1.3.

2.1.1 Approximating the mean of the Gibbs measure

We will denote our approximation of the mean of the Gibbs measure $\mu_{\mathbf{A},\mathbf{y}}$ by $\widehat{\mathbf{m}}(\mathbf{A}, \mathbf{y})$, while the actual mean will be $\mathbf{m}(\mathbf{A}, \mathbf{y})$.

The algorithm to compute $\widehat{\mathbf{m}}(\mathbf{A}, \mathbf{y})$ is given in Algorithm 1, and is composed of two phases:

1. An Approximate Message Passing (AMP) algorithm is run for K_{AMP} iterations and constructs a first estimate of the mean. We denote by $\text{AMP}(\mathbf{A}, \mathbf{y}; k)$ the estimate produced after k AMP iterations

$$\text{AMP}(\mathbf{A}, \mathbf{y}; k) := \widehat{\mathbf{m}}^k. \quad (2.2)$$

Algorithm 1: MEAN OF THE TILTED GIBBS MEASURE

Input: Data $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{y} \in \mathbb{R}^n$, parameters $\beta, \eta > 0$, $q \in (0, 1)$, iteration numbers K_{AMP} , K_{NGD} .

- 1 $\widehat{\mathbf{m}}^{-1} = \mathbf{z}^0 = 0$,
- 2 **for** $k = 0, \dots, K_{\text{AMP}} - 1$ **do**
- 3 $\widehat{\mathbf{m}}^k = \tanh(\mathbf{z}^k)$, $\mathbf{b}_k = \frac{\beta^2}{n} \sum_{i=1}^n (1 - \tanh^2(z_i^k))$,
- 4 $\mathbf{z}^{k+1} = \beta \mathbf{A} \widehat{\mathbf{m}}^k + \mathbf{y} - \mathbf{b}_k \widehat{\mathbf{m}}^{k-1}$,
- 5 **end**
- 6 $\mathbf{u}^0 = \mathbf{z}^{K_{\text{AMP}}}$,
- 7 **for** $k = 0, \dots, K_{\text{NGD}} - 1$ **do**
- 8 $\mathbf{u}^{k+1} = \mathbf{u}^k - \eta \cdot \nabla \mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^{+,k}; \mathbf{y}, q)$,
- 9 $\widehat{\mathbf{m}}^{+,k+1} = \tanh(\mathbf{u}^{k+1})$,
- 10 **end**
- 11 **return** $\widehat{\mathbf{m}}^{+,K_{\text{NGD}}}$

2. Natural gradient descent (NGD) is run for K_{NGD} iterations with initialization given by vector computed at the end of the first phase. This phase attempts to minimize the following version of the TAP free energy (for a specific value of q):

$$\mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q) := -\frac{\beta}{2} \langle \mathbf{m}, \mathbf{A} \mathbf{m} \rangle - \langle \mathbf{y}, \mathbf{m} \rangle - \sum_{i=1}^n h(m_i) - \frac{n\beta^2(1-q)(1+q-2Q(\mathbf{m}))}{4}, \quad (2.3)$$

$$Q(\mathbf{m}) = \frac{1}{n} \|\mathbf{m}\|^2, \quad h(m) = -\frac{1+m}{2} \log\left(\frac{1+m}{2}\right) - \frac{1-m}{2} \log\left(\frac{1-m}{2}\right). \quad (2.4)$$

The second stage is motivated by the TAP (Thouless-Anderson-Palmer) equations for the Gibbs mean of a high-temperature spin glass [MPV87, Tal10]. Essentially by construction, stationary points for the function $\mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q)$ satisfy the TAP equations, and we show in Lemma 4.11 that the first stage above constructs an approximate stationary point for $\mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q)$. The effect of the second stage is therefore numerically small, but it turns out to reduce the error incurred by discretizing time in line 6 of Algorithm 2.

Let us emphasize that this two-stage construction is considered for technical reasons. Indeed a simpler algorithm, that runs AMP for a larger number of iteration, and does not run NGD at all, is expected to work but our arguments do not go through. The hybrid algorithm above allows us to exploit known properties of AMP (precise analysis via state evolution) and of $\mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q)$ (Lipschitz continuity of the minimizer in \mathbf{y}).

2.1.2 Sampling via stochastic localization

Our sampling algorithm is presented as Algorithm 2. The algorithm makes uses of constants $q_k := q_k(\beta, t)$. With $W \sim \mathcal{N}(0, 1)$ a standard Gaussian, these constants are defined for $k, \beta, t \geq 0$ by the recursion

$$q_{k+1} = \mathbb{E} \left\{ \tanh(\beta^2 q_k + t + \sqrt{\beta^2 q_k + t} W)^2 \right\}, \quad q_0 = 0, \quad q_* = \lim_{k \rightarrow \infty} q_k. \quad (2.5)$$

This iteration can be implemented via a one-dimensional integral, and the limit q_* is approached exponentially fast in k (see Lemma 4.5 below). The values $q_*(\beta, t = \ell\delta)$ for $\ell \in \{0, \dots, L\}$ can be precomputed and are independent of the input \mathbf{A} . For the sake of simplicity, we will neglect errors in this calculation.

Algorithm 2: APPROXIMATE SAMPLING FROM THE SK GIBBS MEASURE

Input: Data $\mathbf{A} \in \mathbb{R}^{n \times n}$, parameters $(\beta, \eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta)$

- 1 $\hat{\mathbf{y}}_0 = 0$,
- 2 **for** $\ell = 0, \dots, L - 1$ **do**
- 3 Draw $\mathbf{w}_{\ell+1} \sim \mathcal{N}(0, \mathbf{I}_n)$ independent of everything so far;
- 4 Set $q = q_*(\beta, t = \ell\delta)$;
- 5 Set $\widehat{\mathbf{m}}(\mathbf{A}, \hat{\mathbf{y}}_\ell)$ the output of Algorithm 1, with parameters $(\beta, \eta, q, K_{\text{AMP}}, K_{\text{NGD}})$;
- 6 Update $\hat{\mathbf{y}}_{\ell+1} = \hat{\mathbf{y}}_\ell + \widehat{\mathbf{m}}(\mathbf{A}, \hat{\mathbf{y}}_\ell) \delta + \sqrt{\delta} \mathbf{w}_{\ell+1}$
- 7 **end**
- 8 Set $\widehat{\mathbf{m}}(\mathbf{A}, \hat{\mathbf{y}}_L)$ the output of Algorithm 1, with parameters $(\eta, q, K_{\text{AMP}}, K_{\text{NGD}})$;
- 9 Draw $\{x_i^{\text{alg}}\}_{i \leq n}$ conditionally independent with $\mathbb{E}[x_i^{\text{alg}} | \mathbf{y}, \{\mathbf{w}_\ell\}] = \widehat{m}_i(\mathbf{A}, \hat{\mathbf{y}}_L)$
- 10 **return** \mathbf{x}^{alg}

The core of the sampling procedure is step 6, which is a standard Euler discretization of the SDE (1.4), with step size δ , over the time interval $[0, T]$, $T = L\delta$. The mean of the Gibbs measure $\mathbf{m}(\mathbf{A}, \mathbf{y})$ is replaced by the output of Algorithm 1 which we recall is denoted by $\widehat{\mathbf{m}}(\mathbf{A}, \mathbf{y})$. We reproduce the Euler iteration here for future reference

$$\hat{\mathbf{y}}_{\ell+1} = \hat{\mathbf{y}}_\ell + \widehat{\mathbf{m}}(\mathbf{A}, \hat{\mathbf{y}}_\ell) \delta + \sqrt{\delta} \mathbf{w}_{\ell+1}. \quad (2.6)$$

The output of the iteration is $\widehat{\mathbf{m}}(\mathbf{A}, \hat{\mathbf{y}}_L)$, which should be thought of as an approximation of $\mathbf{m}(\mathbf{A}, \mathbf{y}(T))$, $T = L\delta$, that is the mean of $\mu_{\mathbf{A}, \mathbf{y}(T)}$. According to the discussion in the introduction, for large T , $\mu_{\mathbf{A}, \mathbf{y}(T)}$ concentrates around $\mathbf{x}^* \sim \mu_{\mathbf{A}}$. In other words, $\mathbf{m}(\mathbf{A}, \mathbf{y}(T))$ is close to the corner \mathbf{x}^* of the hypercube. We round its coordinates independently to produce the output \mathbf{x}^{alg} .

2.1.3 Theoretical guarantee

Our main positive result is the following.

Theorem 2.1. *For any $\varepsilon > 0$ and $\beta_0 < 1/2$ there exist $\eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta$ independent of n , so that the following holds for all $\beta \leq \beta_0$. The sampling algorithm 2 takes as input \mathbf{A} and parameters $(\eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta)$ and outputs a random point $\mathbf{x}^{\text{alg}} \in \{-1, +1\}^n$ with law $\mu_{\mathbf{A}}^{\text{alg}}$ such that with probability $1 - o_n(1)$ over $\mathbf{A} \sim \text{GOE}(n)$,*

$$W_{2,n}(\mu_{\mathbf{A}}^{\text{alg}}, \mu_{\mathbf{A}}) \leq \varepsilon. \quad (2.7)$$

The total complexity of this algorithm is $O(n^2)$.

Remark 2.1. The condition $\beta < 1/2$ arises because our proof requires the Hessian of the TAP free energy to be positive definite at its minimizer. A simple calculation yields

$$\nabla^2 \mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q) = -\beta \mathbf{A} + \mathbf{D}(\mathbf{m}) + \beta^2(1 - q) \mathbf{I}_n, \quad \mathbf{D}(\mathbf{m}) := \text{diag}(\{(1 - m_i^2)^{-1}\}_{i \leq n}). \quad (2.8)$$

A crude bound yields $\nabla^2 \mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q) \succeq -\beta \mathbf{A} + \mathbf{I}_n \succeq (1 - \beta \lambda_{\max}(\mathbf{A})) \mathbf{I}_n$. Since $\text{p-lim}_{n \rightarrow \infty} \lambda_{\max}(\mathbf{A}) = 2$ the desired condition holds trivially for $\beta < 1/2$. However, we expect that a more careful treatment will reveal that the Hessian is locally positive in a neighborhood of the minimizer for all $\beta < 1$.

2.2 Hardness for stable algorithms, for $\beta > 1$

The sampling algorithm 2 enjoys stability properties with respect to changes in the inverse temperature β and the matrix \mathbf{A} which are shared by many natural efficient algorithms. We will use the fact that the actual Gibbs measure does not enjoy this stability property for $\beta > 1$ to conclude that sampling is hard for all stable algorithms.

Throughout this section, we denote the Gibbs and algorithmic output distributions by $\mu_{\mathbf{A},\beta}$ and $\mu_{\mathbf{A},\beta}^{\text{alg}}$ respectively to emphasize the dependence on β .

Definition 2.2. Let $\{\text{ALG}_n\}_{n \geq 1}$ be a family of randomized sampling algorithms, i.e., measurable maps

$$\text{ALG}_n : (\mathbf{A}, \beta, \omega) \mapsto \text{ALG}_n(\mathbf{A}, \beta, \omega) \in [-1, 1]^n,$$

where ω is a random seed (a point in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$). Let \mathbf{A}' and $\mathbf{A} \sim \text{GOE}(n)$ be independent copies of the coupling matrix, and consider perturbations $\mathbf{A}_s = \sqrt{1-s^2}\mathbf{A} + s\mathbf{A}'$ for $s \in [0, 1]$. Finally, denote by $\mu_{\mathbf{A}_s, \beta}^{\text{alg}}$ the law of the algorithm output, i.e., the distribution of $\text{ALG}_n(\mathbf{A}_s, \beta, \omega)$ when $\omega \sim \mathbb{P}$ independent of \mathbf{A}_s, β which are fixed.

We say ALG_n is stable with respect to disorder, at inverse temperature β , if

$$\lim_{s \rightarrow 0} \text{p-lim}_{n \rightarrow \infty} W_{2,n}(\mu_{\mathbf{A},\beta}^{\text{alg}}, \mu_{\mathbf{A}_s,\beta}^{\text{alg}}) = 0. \quad (2.9)$$

We say ALG_n is stable with respect to temperature at inverse temperature β , if

$$\lim_{\beta' \rightarrow \beta} \text{p-lim}_{n \rightarrow \infty} W_{2,n}(\mu_{\mathbf{A},\beta}^{\text{alg}}, \mu_{\mathbf{A},\beta'}^{\text{alg}}) = 0. \quad (2.10)$$

We begin by establishing the stability of the proposed sampling algorithm.

Theorem 2.3 (Stability of the sampling Algorithm 2). *For any $\beta \in (0, \infty)$ and fixed parameters $(\eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta)$, Algorithm 2 is stable with respect to disorder and with respect to temperature.*

This theorem is proved in Section 5.1. As a consequence, the Gibbs measures $\mu_{\mathbf{A},\beta}$ enjoy similar stability properties for $\beta < 1/2$, which amount (as discussed below) to the absence of chaos in both temperature and disorder:

Corollary 2.4. *For any $\beta < 1/2$, the following properties hold for the Gibbs measure $\mu_{\mathbf{A},\beta}$ of the Sherrington-Kirkpatrick model, cf. Eq. (1.1):*

1. $\lim_{s \rightarrow 0} \text{p-lim}_{n \rightarrow \infty} W_{2,n}(\mu_{\mathbf{A},\beta}, \mu_{\mathbf{A}_s,\beta}) = 0.$
2. $\lim_{\beta' \rightarrow \beta} \text{p-lim}_{n \rightarrow \infty} W_{2,n}(\mu_{\mathbf{A},\beta}, \mu_{\mathbf{A},\beta'}) = 0.$

Proof. Take $\varepsilon > 0$ arbitrarily small and choose parameters $(\eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta)$ of Algorithm 2 with the desired tolerance ε so that Theorem 2.1 holds. Combining with Theorem 2.3 using the same parameters $(\eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta)$ implies the result since ε is arbitrarily small. (Recall that $(\eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta)$ can be chosen independent of β for $\beta \leq \beta_0 < 1/2$.) \square

Remark 2.2. We emphasize that Corollary 2.4 makes no reference to the sampling algorithm, and is instead a purely structural property of the Gibbs measure. The sampling algorithm, however, is the key tool of our proof.

Stability is related to chaos, which is a well studied and important property of spin glasses, see e.g. [Cha09, Che13, Cha14, CHHS15, CP18]. In particular, “disorder chaos” refers to the following phenomenon. Draw $\mathbf{x}^0 \sim \mu_{\mathbf{A},\beta}$ independently of $\mathbf{x}^s \sim \mu_{\mathbf{A}_s,\beta}$, and denote by $\mu_{\mathbf{A},\beta}^{(0,s)} := \mu_{\mathbf{A},\beta} \otimes \mu_{\mathbf{A}_s,\beta}$ their joint distribution. Disorder chaos holds at inverse temperature β if

$$\lim_{s \rightarrow 0} \lim_{n \rightarrow \infty} \mathbb{E} \mu_{\mathbf{A},\beta}^{(0,s)} \left\{ \left(\frac{1}{n} \langle \mathbf{x}^0, \mathbf{x}^s \rangle \right)^2 \right\} = 0. \quad (2.11)$$

Note that disorder chaos is not necessarily a surprising property. For instance when $\beta = 0$, the distribution $\mu_{\mathbf{A}_s,\beta}$ is simply the uniform measure over the hypercube $\{-1, +1\}^n$ for all s , and this example exhibits disorder chaos in the sense of Eq. (2.11). In fact, the SK Gibbs measure exhibits disorder chaos at all $\beta \in [0, \infty)$ [Cha09]. However, for $\beta > 1$, Eq. (2.11) leads to a stronger conclusion.

Theorem 2.5 (Disorder chaos in $W_{2,n}$ distance). *For all $\beta > 1$,*

$$\inf_{s \in (0,1)} \liminf_{n \rightarrow \infty} \mathbb{E} [W_{2,n}(\mu_{\mathbf{A},\beta}, \mu_{\mathbf{A}_s,\beta})] > 0.$$

Finally, we obtain the desired hardness result by reversing the implication in Corollary 2.4: no stable algorithm which can approximately sample from the measure $\mu_{\mathbf{A},\beta}$ in the $W_{2,n}$ sense for $\beta > 1$.

Theorem 2.6. *Fix $\beta > 1$, and let $\{\text{ALG}_n\}_{n \geq 1}$ be a family of randomized algorithms which is stable with respect to disorder as per Definition 2.2 at inverse temperature β . Let $\mu_{\mathbf{A},\beta}^{\text{alg}}$ be the law of the output $\text{ALG}_n(\mathbf{A}, \beta, \omega)$ conditional on \mathbf{A} . Then*

$$\liminf_{n \rightarrow \infty} \mathbb{E} [W_{2,n}(\mu_{\mathbf{A},\beta}^{\text{alg}}, \mu_{\mathbf{A},\beta})] > 0.$$

We refer the reader to Section 5.2 for the proof of this theorem.

2.3 Notations

We use $o_n(1)$ to indicate a quantity tending to 0 as $n \rightarrow \infty$. We use $o_{n,\mathbb{P}}(1)$ for a quantity tending to 0 in probability. If X is a random variable, then $\mathcal{L}(X)$ indicates its law. The quantity $C(\beta)$ refers to a constant depending on β . For $\mathbf{x} \in \mathbb{R}^n$ and $\rho \in \mathbb{R}_{\geq 0}$, we denote the open ball of center \mathbf{x} and radius ρ by $B(\mathbf{x}, \rho) := \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{y} - \mathbf{x}\|_2 < \rho\}$. The uniform distribution on the interval $[a, b]$ is denoted by $\text{Unif}([a, b])$. The set of probability distributions over a measurable space (Ω, \mathcal{F}) is denoted by $\mathcal{P}(\Omega)$.

3 Properties of stochastic localization

We collect in this section the main properties of the stochastic localization process needed for our analysis. To be definite, we will focus on the stochastic localization process for the Gibbs measure (1.1), although most of what we will say generalizes to other probability measures in \mathbb{R}^n , under suitable tail conditions. Throughout this section, the matrix \mathbf{A} is viewed as fixed.

Recalling the tilted measure $\mu_{\mathbf{A},\mathbf{y}}$ of Eq. (1.2), and the SDE of Eq. (1.4), we introduce the shorthand

$$\mu_t = \mu_{\mathbf{A},\mathbf{y}(t)}.$$

The following properties are well known. See for instance [ES22, Propositions 9, 10] or [Eld20]. We provide proofs for the reader’s convenience.

Lemma 3.1. For all $t \geq 0$ and all $\mathbf{x} \in \{-1, +1\}^n$,

$$d\mu_t(\mathbf{x}) = \mu_t(\mathbf{x}) \langle \mathbf{x} - \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}, d\mathbf{B}(t) \rangle. \quad (3.1)$$

As a consequence, for any function $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}^m$, the process $(\mathbb{E}_{\mathbf{x} \sim \mu_t} [\varphi(\mathbf{x})])_{t \geq 0}$ is a martingale.

Proof. Let us evaluate the differential of $\log \mu_t$. By writing Z_t for the normalization constant $Z(\mathbf{y}(t))$ of Eq. (1.2), we get

$$d \log \mu_t(\mathbf{x}) = \langle d\mathbf{y}(t), \mathbf{x} \rangle - d \log Z_t. \quad (3.2)$$

Using Itô's formula for Z_t we have

$$\begin{aligned} dZ_t &= d \sum_{\mathbf{x} \in \{-1, +1\}^n} e^{(\beta/2) \langle \mathbf{x}, \mathbf{A}\mathbf{x} \rangle + \langle \mathbf{y}(t), \mathbf{x} \rangle} \\ &= \sum_{\mathbf{x} \in \{-1, +1\}^n} \left(\langle d\mathbf{y}(t), \mathbf{x} \rangle + \frac{1}{2} \|\mathbf{x}\|^2 dt \right) e^{(\beta/2) \langle \mathbf{x}, \mathbf{A}\mathbf{x} \rangle + \langle \mathbf{y}(t), \mathbf{x} \rangle}. \end{aligned}$$

Therefore, denoting by $[Z]_t$ the quadratic variation process associated to Z_t ,

$$\begin{aligned} d \log Z_t &= \frac{dZ_t}{Z_t} - \frac{1}{2} \frac{d[Z]_t}{Z_t^2} \\ &= \langle d\mathbf{y}(t), \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)} \rangle + \frac{1}{2} \mathbb{E}_{\mu_t} [\|\mathbf{x}\|^2] dt - \frac{1}{2} \|\mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}\|^2 dt \\ &= \langle d\mathbf{y}(t), \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)} \rangle + \frac{n}{2} - \frac{1}{2} \|\mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}\|^2 dt. \end{aligned}$$

Substituting in (3.2) we obtain

$$\begin{aligned} d \log \mu_t(\mathbf{x}) &= \langle d\mathbf{y}(t), \mathbf{x} - \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)} \rangle - \frac{n}{2} dt + \frac{1}{2} \|\mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}\|^2 dt \\ &= \langle d\mathbf{B}_t, \mathbf{x} - \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)} \rangle - \frac{1}{2} \|\mathbf{x} - \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}\|^2 dt. \end{aligned}$$

Applying Itô's formula to $e^{\log \mu_t(\mathbf{x})}$ yields the desired result.

Finally, Eq. (3.1) implies that $\mu_t(\mathbf{x})$ is a martingale for every $\mathbf{x} \in \{-1, +1\}^n$. Since $\mathbb{E}_{\mathbf{x} \sim \mu_t} [\varphi(\mathbf{x})]$ is a linear combination of martingales, it is itself a martingale. \square

Lemma 3.2 ([Eld20]). For all $t > 0$,

$$\mathbb{E} \operatorname{cov}(\mu_t) \preceq \frac{1}{t} \mathbf{I}_n. \quad (3.3)$$

Lemma 3.3. For all $t > 0$,

$$W_{2,n}(\mu_{\mathbf{A}}, \mathcal{L}(\mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}))^2 \leq \frac{1}{t}. \quad (3.4)$$

In particular, the mean vector $\mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}$ converges in distribution to a random vector $\mathbf{x}^* \sim \mu_{\mathbf{A}}$ as $t \rightarrow \infty$.

Proof. By Lemma 3.2,

$$\mathbb{E} \left[\mathbb{E}_{\mathbf{x} \sim \mu_t} [\|\mathbf{x} - \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}\|^2] \right] \leq \frac{n}{t},$$

therefore

$$\mathbb{E} \left[W_{2,n}(\mu_t, \delta \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)})^2 \right] \leq \frac{1}{t}.$$

Notice that $(\mu, \nu) \mapsto W_{2,n}^2(\mu, \nu)$ is jointly convex. Since $\mu_{\mathbf{A}} = \mathbb{E}[\mu_t]$, this implies

$$W_{2,n}(\mu_{\mathbf{A}}, \mathcal{L}(\mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}))^2 \leq \mathbb{E} \left[W_{2,n}(\mu_t, \delta \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)})^2 \right] \leq \frac{1}{t}.$$

\square

4 Analysis of the sampling algorithm and proof of Theorem 2.1

This section is devoted to the analysis of Algorithm 2 described in the previous section. An important simplification is obtained by reducing ourselves to working with a corresponding *planted* model. This approach has two advantages: (i) The joint distribution of the matrix \mathbf{A} and the process $(\mathbf{y}(t))_{t \geq 0}$ in (1.4) is significantly simpler in the planted model; (ii) Analysis in the planted model can be cast as a statistical estimation problem. In the latter, Bayes-optimality considerations can be exploited to relate the output of the AMP algorithm $\text{AMP}(\mathbf{A}, \mathbf{y}; k)$ to the true mean vector $\mathbf{m}(\mathbf{A}, \mathbf{y})$.

This section is organized as follows. Section 4.1 introduces the planted model and its relation to the original model. We then analyze the AMP component of our algorithm in Section 4.2, and the NGD component in Section 4.3. Finally, Section 4.4 puts the various elements together and proves Theorem 2.1.

4.1 The planted model and contiguity

Let $\bar{\nu}$ be the uniform distribution over $\{-1, +1\}^n$ and consider the joint distribution of pairs $(\mathbf{x}, \mathbf{A}) \in \{-1, +1\}^n \times \mathbb{R}_{\text{sym}}^{n \times n}$,

$$\mu_{\text{pl}}(d\mathbf{x}, d\mathbf{A}) = \frac{1}{Z_{\text{pl}}} \exp \left\{ -\frac{n}{4} \left\| \mathbf{A} - \frac{\beta \mathbf{x} \mathbf{x}^\top}{n} \right\|_F^2 \right\} \bar{\nu}(d\mathbf{x}) d\mathbf{A}, \quad (4.1)$$

where $d\mathbf{A}$ is the Lebesgue measure over the space of symmetric matrices $\mathbb{R}_{\text{sym}}^{n \times n}$, and the normalizing constant

$$Z_{\text{pl}} := \int \exp \left\{ -\frac{n}{4} \left\| \mathbf{A} - \frac{\beta \mathbf{x} \mathbf{x}^\top}{n} \right\|_F^2 \right\} d\mathbf{A} \quad (4.2)$$

is independent of $\mathbf{x} \in \{-1, +1\}^n$. It is easy to see by construction that the marginal distribution of \mathbf{x} under μ_{pl} is $\bar{\nu}$, and the conditional law $\mu_{\text{pl}}(\cdot | \mathbf{x})$ is a rank-one spiked GOE model with spike $\beta \mathbf{x} \mathbf{x}^\top / n$. Namely, under $\mu_{\text{pl}}(\cdot | \mathbf{x})$, we have

$$\mathbf{A} = \frac{\beta}{n} \mathbf{x} \mathbf{x}^\top + \mathbf{W}, \quad \mathbf{W} \sim \text{GOE}(n). \quad (4.3)$$

On the other hand, $\mu_{\text{pl}}(\cdot | \mathbf{A})$ is the SK measure $\mu_{\mathbf{A}}$.

The marginal of \mathbf{A} under μ_{pl} is not the $\text{GOE}(n)$ distribution μ_{GOE} but takes the form

$$\mu_{\text{pl}}(d\mathbf{A}) = \frac{1}{Z_{\text{pl}}} e^{-\frac{n}{4} \|\mathbf{A}\|_F^2} Z_{\text{SK}}(\mathbf{A}) d\mathbf{A} \quad (4.4)$$

$$= \mu_{\text{GOE}}(d\mathbf{A}) Z_{\text{SK}}(\mathbf{A}), \quad (4.5)$$

where $Z_{\text{SK}}(\mathbf{A})$ is the (rescaled) partition function of the SK measure

$$Z_{\text{SK}}(\mathbf{A}) = 2^{-n} \sum_{\mathbf{x} \in \{-1, +1\}^n} \exp \left\{ \frac{\beta}{2} \langle \mathbf{x}, \mathbf{A} \mathbf{x} \rangle - \frac{\beta^2 n}{4} \right\}. \quad (4.6)$$

By a classical result [ALR87], $Z_{\text{SK}}(\mathbf{A})$ has log-normal fluctuations for all $\beta < 1$:

Theorem 4.1 ([ALR87]). *Let $\beta < 1$, $\mathbf{A} \sim \mu_{\text{GOE}}$ and $\sigma^2 = \frac{1}{4}(-\log(1 - \beta^2) - \beta^2)$. Then*

$$Z_{\text{SK}}(\mathbf{A}) \xrightarrow[n \rightarrow \infty]{d} \exp(W), \quad (4.7)$$

where $W \sim \mathbf{N}(-\sigma^2, 2\sigma^2)$.

Therefore, by Le Cam's first lemma [VdV98, Lemma 6.4], $\mu_{\text{pl}}(d\mathbf{A})$ and $\mu_{\text{GOE}}(d\mathbf{A})$ are mutually contiguous for all $\beta < 1$. For the purpose of our analysis we will need a stronger result about the joint distributions of (\mathbf{A}, \mathbf{y}) under our "random" model and a planted model which we now introduce.

Recall that $\mathbf{m}(\mathbf{A}, \mathbf{y})$ denotes the mean of the Gibbs measure $\mu_{\mathbf{A}, \mathbf{y}}$ in Eq. (1.2). For a fixed $T \geq 0$, we define two Borel distributions \mathbb{P} and \mathbb{Q} on $(\mathbf{A}, \mathbf{y}) \in \mathbb{R}_{\text{sym}}^{n \times n} \times C([0, T], \mathbb{R}^n)$ as follows:

$$\mathbb{Q} : \begin{cases} \mathbf{A} & \sim \mu_{\text{GOE}}, \\ \mathbf{y}(t) & = \int_0^t \mathbf{m}(\mathbf{A}, \mathbf{y}(s)) ds + \mathbf{B}(t), \quad t \in [0, T], \end{cases} \quad (\text{random}) \quad (4.8)$$

$$\mathbb{P} : \begin{cases} \mathbf{x}_0 & \sim \bar{\nu}, \\ \mathbf{A} & \sim \mu_{\text{pl}}(\cdot | \mathbf{x}_0), \\ \mathbf{y}(t) & = t\mathbf{x}_0 + \mathbf{B}(t), \quad t \in [0, T] \end{cases} \quad (\text{planted}) \quad (4.9)$$

where $(\mathbf{B}(t))_{t \geq 0}$ is a standard Brownian motion in \mathbb{R}^n independent of everything else. Note the SDE defining the process $\mathbf{y} = (\mathbf{y}(t))_{t \in [0, T]}$ in Eq. (4.8) is a restatement of the stochastic localization equation (1.4) applied to the SK measure $\mu_{\mathbf{A}}$.

Proposition 4.2. *For all $T \geq 0$ and $\beta \geq 0$, \mathbb{P} absolutely continuous with respect to \mathbb{Q} and for all $(\mathbf{A}, \mathbf{y}) \in \mathbb{R}_{\text{sym}}^{n \times n} \times C([0, T], \mathbb{R}^n)$,*

$$\frac{d\mathbb{P}}{d\mathbb{Q}}(\mathbf{A}, \mathbf{y}) = Z_{\text{SK}}(\mathbf{A}).$$

Therefore, for all $\beta < 1$, \mathbb{P} and \mathbb{Q} are mutually contiguous. (Namely, for a sequence of events \mathcal{E}_n , $\lim_{n \rightarrow \infty} \mathbb{P}(\mathcal{E}_n) = 0$ if and only if $\lim_{n \rightarrow \infty} \mathbb{Q}(\mathcal{E}_n) = 0$.)

Proof. Fix $\mathbf{x}_0 \in \mathbb{R}^n$. We first calculate the density of the process $\mathbf{y}(t) = t\mathbf{x}_0 + \mathbf{B}(t)$ with respect to Brownian motion. Let \mathbb{W} be the Wiener measure on $C([0, T], \mathbb{R}^n)$. We obtain by Girsanov's theorem that

$$\frac{d\mathbb{P}(\cdot | \mathbf{x}_0)}{d\mathbb{W}}(\mathbf{y}) = e^{\langle \mathbf{x}_0, \mathbf{y}(T) \rangle - T\|\mathbf{x}_0\|^2/2}. \quad (4.10)$$

Notice that the above density only depends on the endpoint $\mathbf{y}(T)$ of the process \mathbf{y} . From this, we obtain an explicit formula for the density of \mathbb{P} with respect to $(d\mathbf{A}) \times \mathbb{W}$:

$$\mathbb{P}(d\mathbf{A}, d\mathbf{y}) = \frac{1}{Z_{\text{pl}}} \left(\int \exp \left\{ -\frac{n}{4} \left\| \mathbf{A} - \frac{\beta \mathbf{x}_0 \mathbf{x}_0^\top}{n} \right\|_F^2 + \langle \mathbf{x}_0, \mathbf{y}(T) \rangle - \frac{T}{2} \|\mathbf{x}_0\|^2 \right\} \bar{\nu}(d\mathbf{x}_0) \right) d\mathbf{A} \mathbb{W}(d\mathbf{y}), \quad (4.11)$$

where $Z_{\text{pl}} = \int e^{-n\|\mathbf{A}\|_F^2/4} d\mathbf{A}$ is given in Eq. (4.2).

Next we derive a similar formula for \mathbb{Q} . Fix a matrix $\mathbf{A} \in \mathbb{R}_{\text{sym}}^{n \times n}$ and let \mathbf{y} be the solution to the SDE in (4.8). Let $(\bar{\mathbf{B}}(t))_{t \geq 0}$ be another standard Brownian motion in \mathbb{R}^n , and consider the process $\bar{\mathbf{y}} = (\bar{\mathbf{y}}(t))_{t \in [0, T]}$ defined by

$$\bar{\mathbf{y}}(t) = t\mathbf{x} + \bar{\mathbf{B}}(t) \quad \text{where } \mathbf{x} \sim \mu_{\mathbf{A}} \text{ independently of } \bar{\mathbf{B}}. \quad (4.12)$$

Then, there exists another Brownian motion $(\mathbf{W}(t))_{t \geq 0}$ adapted to the filtration $(\mathcal{F}_t = \sigma(\bar{\mathbf{y}}(s) : s \leq t))_{t \in [0, T]}$ such that $d\bar{\mathbf{y}}(t) = \mathbf{m}_{\mathbf{A}, \bar{\mathbf{y}}(t)} dt + d\mathbf{W}(t)$ for all $t \in [0, T]$. This is stated as Theorem 7.12 of [LS77], and can be proved directly applying Levy's characterization of Brownian motion to the process $\bar{\mathbf{y}}(t) - \int_0^t \mathbf{m}(\mathbf{A}, \bar{\mathbf{y}}(s)) ds$.

Therefore, the processes $\bar{\mathbf{y}}$ and \mathbf{y} share the same law conditional on \mathbf{A} . Since we computed the law of $\bar{\mathbf{y}}$ in (4.10), we obtain

$$\mathbb{Q}(d\mathbf{A}, d\mathbf{y}) = \frac{1}{Z_{\text{pl}}} \left(\int \exp \left\{ -\frac{n}{4} \|\mathbf{A}\|_F^2 + \langle \mathbf{x}, \mathbf{y}(T) \rangle - \frac{T}{2} \|\mathbf{x}\|^2 \right\} \mu_{\mathbf{A}}(d\mathbf{x}) \right) d\mathbf{A} \mathbb{W}(d\mathbf{y}), \quad (4.13)$$

where Z_{pl} is as above. Since $\mu_{\mathbf{A}}(d\mathbf{x}) = Z_{\text{SK}}(\mathbf{A})^{-1} e^{\beta \langle \mathbf{x}, \mathbf{A}\mathbf{x} \rangle / 2 - \beta^2 n / 4} \bar{\nu}(d\mathbf{x})$, we obtain after simplification

$$\frac{d\mathbb{P}}{d\mathbb{Q}}(\mathbf{A}, \mathbf{y}) = Z_{\text{SK}}(\mathbf{A}). \quad (4.14)$$

Mutual contiguity follows from Theorem 4.1 and Le Cam's first lemma. \square

Therefore, for the remainder of the proof of Theorem 2.1, we work under the ‘‘planted’’ distribution \mathbb{P} . All results proven under \mathbb{P} transfer to \mathbb{Q} by contiguity.

4.2 Approximate Message Passing

In this section we analyze the AMP iteration of 1, which we copy here for the reader's convenience

$$\begin{aligned} \widehat{\mathbf{m}}^{-1} &= \mathbf{z}^0 = 0, \\ \widehat{\mathbf{m}}^k &= \tanh(\mathbf{z}^k), \quad \mathbf{b}_k = \frac{\beta^2}{n} \sum_{i=1}^n (1 - \tanh^2(z_i^k)) \quad \forall k \geq 0, \\ \mathbf{z}^{k+1} &= \beta \mathbf{A} \widehat{\mathbf{m}}^k + \mathbf{y} - \mathbf{b}_k \widehat{\mathbf{m}}^{k-1}. \end{aligned} \quad (4.15)$$

When needed, we will specify the dependence on \mathbf{A}, \mathbf{y} by writing $\widehat{\mathbf{m}}^k = \widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}) = \text{AMP}(\mathbf{A}, \mathbf{y}; k)$ and $\mathbf{z}^k = \mathbf{z}^k(\mathbf{A}, \mathbf{y})$. Throughout this section $(\mathbf{A}, \mathbf{y}) \sim \mathbb{P}$ will be distributed according to the planted model introduced above.

Our analysis will be based on the general state evolution result of [BM11, JM13], which implies the following asymptotic characterization for the iterates. Set $\gamma_0(\beta, t) = 0, \Sigma_{0,i}(\beta, t) = 0$ and recursively define

$$\gamma_{k+1}(\beta, t) = \beta^2 \cdot \mathbb{E} [\tanh(\gamma_k(\beta, t) + t + G_k)], \quad (4.16)$$

$$\Sigma_{k+1,j+1}(\beta, t) = \beta^2 \cdot \mathbb{E} [\tanh(\gamma_k(\beta, t) + t + G_k) \tanh(\gamma_j(\beta, t) + t + G_j)], \quad (4.17)$$

where $(\mathbf{G}_j)_{j \leq k}$ are jointly Gaussian, with zero mean and covariance $\Sigma_{\leq k} + t \mathbf{1}\mathbf{1}^\top$, $\Sigma_{\leq k} := (\Sigma_{ij})_{i,j \leq k}$.

Proposition 4.3 (Theorem 1 of [BM11]). *For $(\mathbf{A}, \mathbf{y}) \sim \mathbb{P}$ and any $k \in \mathbb{Z}_{\geq 0}$, the empirical distribution of the coordinate of the AMP iterates converges almost surely in $W_2(\mathbb{R}^{k+2})$ as follows:*

$$\frac{1}{n} \sum_{i=1}^n \delta_{(z_i^1, \dots, z_i^k, x_i, y_i)} \xrightarrow[n \rightarrow \infty]{W_2} \mathcal{L}(\gamma_{\leq k}(\beta, t)X + \mathbf{G} + Y\mathbf{1}, X, Y), \quad (4.18)$$

$$\gamma_{\leq k}(\beta, t) = (\gamma_1(\beta, t), \dots, \gamma_k(\beta, t)), \quad \mathbf{G} \sim \mathbf{N}(0, \Sigma_{\leq k}). \quad (4.19)$$

On the right-hand side, X is uniformly random in $\{-1, +1\}$, $Y = tX + \sqrt{t}W$ where $W \sim \mathbf{N}(0, 1)$ and X, \mathbf{G}, W are mutually independent.

Remark 4.1. This specific statement follows from [BM11, Theorem 1] by a change of variables, as in [DAM17] or [MV21].

As in [DAM17, Eqs. (69,70)] we argue that the state evolution equations (4.16), (4.17) take a simple form thanks to our specific choice of AMP non-linearity $\tanh(\cdot)$. It will be convenient to use the notations

$$\begin{aligned}\tilde{\gamma}_k(\beta, t) &= \gamma_k(\beta, t) + t, \\ \tilde{\Sigma}_{k,j}(\beta, t) &= \Sigma_{k,j}(\beta, t) + t.\end{aligned}$$

Proposition 4.4. *For any $t \in \mathbb{R}_{\geq 0}$ and $k, j \in \mathbb{Z}_{\geq 0}$,*

$$\Sigma_{k,j}(\beta, t) = \gamma_{k \wedge j}(\beta, t), \quad \text{and} \quad \tilde{\Sigma}_{k,j}(\beta, t) = \tilde{\gamma}_{k \wedge j}(\beta, t).$$

Proof. The two claims are equivalent and we proceed by induction. The base case $k = 0$ holds by definition, so we may assume $\Sigma_{i,j}(\beta, t) = \gamma_{i \wedge j}(\beta, t)$ for $i, j \leq k-1$. Set $Z_j = \gamma_j X + \tilde{G}_j$ where $\tilde{G} \sim \mathbf{N}(0, \tilde{\Sigma}_{\leq k-1})$. Note that, by the induction hypothesis, Z_{k-1} is a sufficient statistic for X given $(Z_j)_{j \leq k-1}$. Using Bayes' rule, and writing $\tilde{\sigma}_{k-1}^2 := \tilde{\Sigma}_{k-1, k-1}$, one easily computes

$$\mathbb{E}[X|Z_{k-1}] = \frac{e^{\tilde{\gamma}_{k-1} Z_{k-1} / \tilde{\sigma}_{k-1}^2} - e^{-\tilde{\gamma}_{k-1} Z_{k-1} / \tilde{\sigma}_{k-1}^2}}{e^{\tilde{\gamma}_{k-1} Z_{k-1} / \tilde{\sigma}_{k-1}^2} + e^{-\tilde{\gamma}_{k-1} Z_{k-1} / \tilde{\sigma}_{k-1}^2}} = \tanh(Z).$$

Therefore using Eq. (4.16), the fact that \tanh is an odd function and $WX \stackrel{d}{=} W$,

$$\begin{aligned}\tilde{\Sigma}_{k,j} &= \mathbb{E} [\mathbb{E}[X|Z_{k-1}] \mathbb{E}[X|Z_{j-1}]] \\ &\stackrel{(a)}{=} \mathbb{E} [X \mathbb{E}[X|Z_{j-1}]] \\ &= \mathbb{E} [X \tanh(\tilde{\gamma}_{j-1} X + \tilde{\sigma}_{j-1}^2 W)] \\ &= \mathbb{E} [\tanh(\tilde{\gamma}_{j-1} + \tilde{\sigma}_{j-1}^2 W)] = \gamma_j,\end{aligned}$$

where in step (a) we crucially used the sufficient statistic property. This completes the inductive step and hence the proof. \square

Define the function $\text{mmse} : \mathbb{R} \rightarrow \mathbb{R}$ given by

$$\text{mmse}(\gamma) \equiv 1 - \mathbb{E} [\tanh(\gamma + \sqrt{\gamma} W)^2] = 1 - \mathbb{E} [\mathbb{E}[X|\gamma X + \sqrt{\gamma} W]^2].$$

It follows from Proposition 4.4 that (4.16) and (4.17) can be expressed just in terms of the sequence $\gamma_k(\beta, t)$ defined by $\gamma_0(t) = 0$ and the recursion

$$\gamma_{k+1}(\beta, t) = \beta^2 (1 - \text{mmse}(\gamma_k(\beta, t) + t)). \quad (4.20)$$

Note that $\gamma_k(\beta, t)$ depends also on β , which is usually treated as constant. The following result details some useful properties of mmse .

Lemma 4.5 (Lemma 6.1 of [DAM17]). *The following properties hold, where $\{\gamma_k(\beta, t)\}_{k \geq 1}$ is as defined by (4.20).*

- (a) mmse is differentiable, strictly decreasing, and convex in $\gamma \in \mathbb{R}_{\geq 0}$.
- (b) $\text{mmse}(0) = 1$, $\text{mmse}'(0) = -1$ and $\lim_{\gamma \rightarrow \infty} \text{mmse}(\gamma) = 0$.

(c) For $t \geq 0$ there exists a non-negative solution $\gamma_* = \gamma_*(\beta, t)$ to the fixed point equation

$$\gamma_* = \beta^2(1 - \text{mmse}(\gamma_* + t)). \quad (4.21)$$

The solution to this equation is unique for all $t > 0$.

(d) The function $(\beta, t) \mapsto \gamma_*(\beta, t)$ is differentiable for $t > 0$.

(e) For all $\beta < 1$ and $t > 0$,

$$1 - \beta^{2k} \leq \frac{\gamma_k(\beta, t)}{\gamma_*(\beta, t)} \leq 1. \quad (4.22)$$

(f) For $\beta < 1$ and $T > 0$, there exist constants $c(\beta, T), C(\beta, T) \in (0, \infty)$ such that, for all $t \in (0, T]$,

$$c(\beta, T) \leq \frac{\gamma_*(\beta, t)}{t} \leq C(\beta, T). \quad (4.23)$$

(g) For $\beta < 1$ and any $t_1, t_2 \in (0, \infty)$,

$$\gamma_*(\beta, t_1) - \gamma_*(\beta, t_2) \leq \frac{\beta^2}{1 - \beta^2} |t_1 - t_2|. \quad (4.24)$$

Proof. Lemma 6.1 in [DAM17] proves that $\gamma \mapsto \text{mmse}(\gamma)$ is differentiable, strictly decreasing, and convex in $\gamma \in \mathbb{R}_{\geq 0}$ (note that the statement of that Lemma does not claim differentiability, but this is actually proved there by a simple application of dominated convergence). This proves point (a).

Point (b) follows by a direct calculation, cf. [DAM17]. Indeed, by Stein's lemma (Gaussian integration by parts), with $Z = \gamma + W\sqrt{\gamma}$,

$$\begin{aligned} -\text{mmse}'(\gamma) &= \frac{d}{d\gamma} \mathbb{E}[\tanh(\gamma + W\sqrt{\gamma})^2] \\ &= \mathbb{E}[2 \tanh(Z) \tanh'(Z) + \tanh'(Z)^2 + \tanh(Z) \tanh''(Z)] \end{aligned}$$

Evaluating at $\gamma = 0$ shows

$$\text{mmse}'(0) = -1.$$

Also, dominated convergence yields the desired limit values.

Point (c), namely existence and uniqueness of solutions of Eq. (4.21) follows from the above monotonicity and convexity properties. Point (d) follows from the implicit function theorem.

We are left with the task of proving (4.22), (4.23) and (4.24), which are not given in [DAM17].

Define

$$f_t(\gamma) \equiv \beta^2(1 - \text{mmse}(\gamma + t))$$

so that $f_t(\gamma_k(\beta, t)) = \gamma_{k+1}(\beta, t)$. By point (b), $f_t(0) \geq 0$. By point (a), $f_t(\cdot)$ is increasing and concave. Combined with the computation above, we conclude that $f_t'(\gamma) \in [0, \beta^2]$ for all $\gamma \geq 0$. By the mean value theorem, it follows that for $\gamma < \gamma_*$,

$$0 \leq \gamma_*(\beta, t) - f_t(\gamma) = f(\gamma_*(\beta, t)) - f_t(\gamma) \leq \beta^2(\gamma_*(\beta, t) - \gamma). \quad (4.25)$$

Setting $\gamma = \gamma_j(\beta, t)$, we obtain

$$0 \leq \frac{\gamma_*(\beta, t) - \gamma_{j+1}(\beta, t)}{\gamma_*(\beta, t) - \gamma_j(\beta, t)} \leq \beta^2.$$

Multiplying for $j \in \{0, \dots, k-1\}$, we find

$$0 \leq \frac{\gamma_*(\beta, t) - \gamma_k(\beta, t)}{\gamma_*(\beta, t)} \leq \beta^{2k},$$

or,

$$\frac{\gamma_k(\beta, t)}{\gamma_*(\beta, t)} \in [1 - \beta^{2k}, 1],$$

which proves (4.22).

To prove (4.23), note that we just showed

$$\frac{\gamma_1(\beta, t)}{\gamma_*(\beta, t)} \in [1 - \beta^2, 1].$$

Therefore it suffices to show that

$$c(\beta, T) \leq \frac{\gamma_1(\beta, t)}{t} \leq C(\beta, T), \quad t \in (0, T]. \quad (4.26)$$

By definition, $\gamma_1(\beta, t) = \beta^2(1 - \text{mmse}(t))$. Thus (4.26) follows from the fact that $\text{mmse}(0) = 1$, $\text{mmse}'(0) = -1$, and $\text{mmse} : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ is convex and strictly decreasing. In fact we have $\gamma_1(\beta, T)/T \leq \gamma_1(\beta, t)/t \leq \beta^2$ for all $t \in (0, T]$.

Finally, we prove (4.24). Since $|\text{mmse}'(t)| \leq 1$ for all $t \geq 0$ we find that for $t_1, t_2 \geq 0$,

$$\begin{aligned} |\gamma_*(\beta, t_1) - \gamma_*(\beta, t_2)| &= \beta^2 |\text{mmse}(\gamma_*(\beta, t_1) + t_1) - \text{mmse}(\gamma_*(\beta, t_2) + t_2)| \\ &\leq \beta^2 |\gamma_*(\beta, t_1) - \gamma_*(\beta, t_2)| + \beta^2 |t_1 - t_2|. \end{aligned}$$

Rearranging, we obtain

$$\frac{|\gamma_*(\beta, t_1) - \gamma_*(\beta, t_2)|}{|t_1 - t_2|} \leq \frac{\beta^2}{1 - \beta^2}.$$

□

For $(\mathbf{A}, \mathbf{y}) \sim \mathbb{P}$ and $\mathbf{x} \sim \mu_{\mathbf{A}, \mathbf{y}(t)}$, define

$$\text{MSE}_{\text{AMP}}(k; \beta, t) = \text{p-lim}_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \|\mathbf{x} - \widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}(t))\|_2^2, \quad \widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}(t)) := \text{AMP}(\mathbf{A}, \mathbf{y}(t); k), \quad (4.27)$$

where the limit is guaranteed to exist by Proposition 4.3.

Lemma 4.6. *We have*

$$\text{MSE}_{\text{AMP}}(k; \beta, t) = 1 - \frac{\gamma_{k+1}(\beta, t)}{\beta^2}.$$

In particular,

$$\lim_{k \rightarrow \infty} \text{MSE}_{\text{AMP}}(k; \beta, t) = 1 - \frac{\gamma_*(\beta, t)}{\beta^2}.$$

Proof. By state evolution

$$\begin{aligned}
\text{MSE}_{\text{AMP}}(k; \beta, t) &= \text{p-lim}_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \|\widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}(t)) - \mathbf{x}\|_2^2 \\
&= \mathbb{E} [(\tanh(\gamma_k X + \sigma_k W + Y) - X)^2] \\
&= \mathbb{E} [(\tanh(\tilde{\gamma}_k X + \tilde{\sigma}_k W) - X)^2] \\
&= 1 - 2 \mathbb{E}[\tanh(\tilde{\gamma}_k X + \tilde{\sigma}_k W)X] + \mathbb{E}[\tanh(\tilde{\gamma}_k X + \tilde{\sigma}_k W)^2] \\
&= 1 - 2\gamma_{k+1}/\beta^2 + \sigma_{k+1}^2/\beta^2 \\
&= 1 - \gamma_{k+1}/\beta^2,
\end{aligned}$$

where the last line follows from Proposition 4.4. \square

We next show that, for any $t > 0$, the mean square error achieved by AMP is the same as the Bayes optimal error, i.e., the mean squared error achieved by the posterior expectation $\mathbf{m}(\mathbf{A}, \mathbf{y}(t))$.

Proposition 4.7. *Fix $\beta < 1$ and $t \geq 0$. We have*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} [\|\mathbf{x} - \mathbf{m}(\mathbf{A}, \mathbf{y}(t))\|_2^2] = \frac{\gamma_*(\beta, t)}{\beta^2}. \quad (4.28)$$

Proof. The proof is an adaptation from [DAM17], which we will present succinctly.

Let $I(X; Y)$ denote the mutual information between random variables X, Y on the same probability space. Letting $X \sim \text{Unif}(\{-1, +1\})$ independent of $W \sim \mathbf{N}(0, 1)$, define the function

$$l(\gamma) := I(X; \gamma X + \sqrt{\gamma}W) \quad (4.29)$$

$$= \gamma - \mathbb{E} \log \cosh(\gamma + \sqrt{\gamma}W). \quad (4.30)$$

We also define the function

$$\Psi(\gamma; \beta, t) := \frac{\beta^2}{4} + \frac{\gamma^2}{4\beta^2} - \frac{\gamma}{2} + l(\gamma + t). \quad (4.31)$$

As in [DAM17], it is easy to check that $\partial_\gamma \Psi(\gamma_*(\beta, t); \beta, t) = 0$ and, using the continuity of $(\beta, t) \mapsto \gamma_*(\beta, t)$,

$$\frac{d}{d(\beta^2)} \Psi(\gamma_*(\beta, t); \beta, t) = \frac{1}{4} \left(1 - \frac{\gamma_*(\beta, t)^2}{\beta^4} \right), \quad (4.32)$$

$$\frac{d}{dt} \Psi(\gamma_*(\beta, t); \beta, t) = \frac{1}{2} \left(1 - \frac{\gamma_*(\beta, t)}{\beta^2} \right). \quad (4.33)$$

We further note that by the de Bruijn identity (also known as I-MMSE relation [GSV05])

$$\frac{d}{d(\beta^2)} I(\mathbf{x}; \mathbf{A}(\beta), \mathbf{y}(t)) = \frac{1}{4n} \mathbb{E} \left[\|\mathbf{x}\mathbf{x}^\top - \mathbb{E}\{\mathbf{x}\mathbf{x}^\top | \mathbf{A}(\beta), \mathbf{y}(t)\}\|_F^2 \right], \quad (4.34)$$

$$\frac{d}{dt} I(\mathbf{x}; \mathbf{A}(\beta), \mathbf{y}(t)) = \frac{1}{2} \mathbb{E} \left[\|\mathbf{x} - \mathbb{E}\{\mathbf{x} | \mathbf{A}(\beta), \mathbf{y}(t)\}\|_2^2 \right]. \quad (4.35)$$

Here we write $\mathbf{A} = \mathbf{A}(\beta)$ to emphasize the dependence upon β . Using Eqs. (4.32) and (4.34), we have

$$\begin{aligned}
\log 2 - \mathfrak{l}(t) &= \lim_{\beta \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{n} [I(\mathbf{x}; \mathbf{A}(\beta), \mathbf{y}(t)) - I(\mathbf{x}; \mathbf{A}(0), \mathbf{y}(t))] \\
&= \lim_{n \rightarrow \infty} \int_0^\infty \frac{1}{4n} \mathbb{E} \left[\|\mathbf{x}\mathbf{x}^\top - \mathbb{E}\{\mathbf{x}\mathbf{x}^\top | \mathbf{A}(\beta), \mathbf{y}(t)\}\|_F^2 \right] d\beta^2 \\
&\leq \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \int_0^\infty \frac{1}{4n} \mathbb{E} \left[\|\mathbf{x}\mathbf{x}^\top - \widehat{\mathbf{m}}^k(\mathbf{A}(\beta), \mathbf{y}(t))\widehat{\mathbf{m}}^k(\mathbf{A}(\beta), \mathbf{y}(t))^\top\|_F^2 \right] d\beta^2 \\
&= \lim_{k \rightarrow \infty} \int_0^\infty \frac{1}{4} \left(1 - \frac{\gamma_k(\beta, t)^2}{\beta^4}\right) d\beta^2 \\
&= \int_0^\infty \frac{1}{4} \left(1 - \frac{\gamma_*(\beta, t)^2}{\beta^4}\right) d\beta^2 \\
&= \lim_{\beta \rightarrow \infty} [\Psi(\gamma_*(\beta, t); \beta, t) - \Psi(\gamma_*(0, t); 0, t)].
\end{aligned}$$

(The exchanges of limits are justified by dominated convergence.)

Finally, a direct calculation reveals that $\lim_{\beta \rightarrow \infty} [\Psi(\gamma_*(\beta, t); \beta, t) - \Psi(\gamma_*(0, t); 0, t)] = \log(2) - \mathfrak{l}(t)$ and therefore equality holds at each of the steps above. We deduce that $\lim_{n \rightarrow \infty} n^{-1} I(\mathbf{x}; \mathbf{A}(\beta), \mathbf{y}(t)) = \Psi(\gamma_*(\beta, t); \beta, t)$.

Using this fact, together with Eqs. (4.33), (4.35) and the fact that the right hand sides of these equations are monotone decreasing in t , we get that the following holds for almost every $t > 0$:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\|\mathbf{x} - \mathbb{E}\{\mathbf{x} | \mathbf{A}(\beta), \mathbf{y}(t)\}\|_2^2 \right] = 1 - \frac{\gamma_*(\beta, t)}{\beta^2}. \quad (4.36)$$

This coincides with the claim (4.28), and actually holds for every $t > 0$ since the right-hand side of Eq. (4.28) is continuous in $t > 0$ by Lemma 4.5. \square

It follows that AMP approximately computes the posterior mean $\mathbf{m}(\mathbf{A}, \mathbf{y}(t))$ in the following sense.

Proposition 4.8. *Fix $\beta < 1$, $T > 0$ and let $t \in (0, T]$. Recalling that $\widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}(t)) := \text{AMP}(\mathbf{A}, \mathbf{y}(t); k)$ denotes the AMP estimate after k iterations, and that \mathbf{z}^k is defined by Eq. (4.15), we have*

$$\lim_{k \rightarrow \infty} \sup_{t \in (0, T)} \text{p-lim}_{n \rightarrow \infty} \frac{\|\mathbf{m}(\mathbf{A}, \mathbf{y}(t)) - \widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}(t))\|_2}{\|\mathbf{m}(\mathbf{A}, \mathbf{y}(t))\|_2} = 0. \quad (4.37)$$

Moreover

$$\lim_{k \rightarrow \infty} \sup_{t \in (0, T)} \text{p-lim}_{n \rightarrow \infty} \frac{\|\mathbf{z}^{k+1} - \mathbf{z}^k\|}{\|\mathbf{z}^k\|} = 0. \quad (4.38)$$

Remark 4.2. A somewhat similar result has recently been proved by Chen and Tang [CT21] where the external field vector $\mathbf{y}(t)$ is replaced by a multiple of the all-ones vector $h\mathbf{1}$, for any pair (β, h) for which a certain condition of uniform concentration of the overlap between two independent draws from the measure $\mu_{\mathbf{A}, h\mathbf{1}}$ holds. In our setting, we are concerned with a different family of external fields, namely the ones generated by the stochastic localization process (1.4). The argument, which proceeds via the planted model, does not require the uniform concentration condition.

Proof. Throughout this proof we write \mathbf{y} instead of $\mathbf{y}(t)$ for ease of notation. To show Eq. (4.37), observe that the bias-variance decomposition yields (recalling the definition $\text{MSE}_{\text{AMP}}(\cdot)$ in Eq. (4.27))

$$\text{MSE}_{\text{AMP}}(k; \beta, t) = \text{p-lim}_{n \rightarrow \infty} \left\{ \frac{1}{n} \mathbb{E} \left[\|\widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}) - \mathbf{m}(\mathbf{A}, \mathbf{y})\|_2^2 \right] + \frac{1}{n} \mathbb{E} \left[\|\mathbf{x} - \mathbf{m}(\mathbf{A}, \mathbf{y})\|_2^2 \right] \right\}.$$

Using Lemma 4.6 for the left-hand side and Proposition 4.7 for the second step the second term on the right-hand side, we get

$$\text{p-lim}_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\|\widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}) - \mathbf{m}(\mathbf{A}, \mathbf{y})\|_2^2 \right] = \frac{\gamma_*(\beta, t) - \gamma_{k+1}(\beta, t)}{\beta^2}. \quad (4.39)$$

Claim (4.37) now follows by combining Eq. (4.39) with Eqs. (4.22) and (4.23) of Lemma 4.5.

Finally, Eq. (4.38) is an immediate consequence of Proposition 4.3 and Proposition 4.4. Indeed, by Proposition 4.3, we have

$$\text{p-lim}_{n \rightarrow \infty} \frac{1}{n} \|\mathbf{z}^k\|_2^2 = \mathbb{E} [(\gamma_k X + G_k + Y)^2] = (\gamma_k + t)^2 + \gamma_k + t, \quad (4.40)$$

$$\text{p-lim}_{n \rightarrow \infty} \frac{1}{n} \|\mathbf{z}^{k+1} - \mathbf{z}^k\|_2^2 = \mathbb{E} [((\gamma_{k+1} - \gamma_k)X + G_{k+1} - G_k)^2] \quad (4.41)$$

$$= (\gamma_{k+1} - \gamma_k)^2 + (\Sigma_{k+1, k+1} - 2\Sigma_{k, k+1} + \Sigma_{k, k}) \quad (4.42)$$

$$= (\gamma_{k+1} - \gamma_k)^2 + (\gamma_{k+1} - \gamma_k), \quad (4.43)$$

where in the last step we used Proposition 4.4. We therefore obtained we have

$$\text{p-lim}_{n \rightarrow \infty} \frac{\|\mathbf{z}^{k+1} - \mathbf{z}^k\|_2^2}{\|\mathbf{z}^k\|_2^2} = \frac{(\gamma_{k+1} - \gamma_k)^2 + (\gamma_{k+1} - \gamma_k)}{(\gamma_k + t)^2 + \gamma_k + t}. \quad (4.44)$$

Hence Eq. (4.38) also follows from Eq. (4.22). \square

We conclude this subsection with a lemma controlling the regularity of the posterior path $t \mapsto \mathbf{m}(\mathbf{A}, \mathbf{y}(t))$, which will be useful later.

Lemma 4.9. *Fix $\beta < 1$ and $0 \leq t_1 < t_2 \leq T$. Then*

$$\text{p-lim}_{n \rightarrow \infty} \sup_{t \in [t_1, t_2]} \frac{1}{n} \|\mathbf{m}(\mathbf{A}, \mathbf{y}(t)) - \mathbf{m}(\mathbf{A}, \mathbf{y}(t_1))\|_2^2 = \text{p-lim}_{n \rightarrow \infty} \frac{1}{n} \|\mathbf{m}(\mathbf{A}, \mathbf{y}(t_2)) - \mathbf{m}(\mathbf{A}, \mathbf{y}(t_1))\|_2^2 \quad (4.45)$$

$$= \frac{\gamma_*(\beta, t_2) - \gamma_*(\beta, t_1)}{\beta^2}. \quad (4.46)$$

Proof. We will exploit the fact that $(\mathbf{m}(\mathbf{A}, \mathbf{y}(t)))_{t \geq 0}$ is a martingale, as a consequence of Lemma 3.1 (with $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ given by $\varphi(\mathbf{x}) = \mathbf{x}$).

Using Proposition 4.7, we obtain, for any $t_1 < t_2$

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E} \left[\|\mathbf{m}(\mathbf{A}, \mathbf{y}(t_2)) - \mathbf{m}(\mathbf{A}, \mathbf{y}(t_1))\|_2^2 \right] &= \text{p-lim}_{n \rightarrow \infty} \frac{1}{n} \left\{ \mathbb{E} \left[\|\mathbf{x} - \mathbf{m}(\mathbf{A}, \mathbf{y}(t_1))\|_2^2 \right] - \mathbb{E} \left[\|\mathbf{x} - \mathbf{m}(\mathbf{A}, \mathbf{y}(t_1))\|_2^2 \right] \right\} \\ &= \frac{\gamma_*(\beta, t_2) - \gamma_*(\beta, t_1)}{\beta^2}, \end{aligned}$$

where the first equality uses the fact that $\mathbb{E}\{\mathbf{m}(\mathbf{A}, \mathbf{y}(t_2)) | \mathbf{A}, \mathbf{y}(t_1)\} = \mathbf{m}(\mathbf{A}, \mathbf{y}(t_1))$. By Lemma 4.8, we have, with high probability, $\|\mathbf{m}(\mathbf{A}, \mathbf{y}(t_i)) - \widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}(t_i))\|_2^2/n \leq \varepsilon_k$, for some deterministic constants ε_k so that $\varepsilon_k \rightarrow 0$ as $k \rightarrow \infty$. As a consequence

$$\text{p-lim}_{n \rightarrow \infty} \frac{1}{n} \|\mathbf{m}(\mathbf{A}, \mathbf{y}(t_2)) - \mathbf{m}(\mathbf{A}, \mathbf{y}(t_1))\|_2^2 = \frac{\gamma_*(\beta, t_2) - \gamma_*(\beta, t_1)}{\beta^2}. \quad (4.47)$$

Now, since $t \rightarrow \mathbf{m}_{\mathbf{A}, \mathbf{y}(t)}$ is a bounded martingale, it follows that, for any fixed constant c , the process

$$Y_{n,t} := |M_{n,t} - c|, \quad \text{where} \quad M_{n,t} := \frac{1}{\sqrt{n}} \|\mathbf{m}(\mathbf{A}, \mathbf{y}(t)) - \mathbf{m}(\mathbf{A}, \mathbf{y}(t_1))\|_2, \quad (4.48)$$

is a positive bounded submartingale for $t \geq t_1$. Therefore by Doob's maximal inequality [Dur19],

$$\mathbb{P} \left(\sup_{t \in [t_1, t_2]} Y_{n,t} \geq a \right) \leq \frac{1}{a} \mathbb{E} [Y_{n,t_2}] \leq \frac{1}{a} \mathbb{E} [Y_{n,t_2}^2]^{1/2}, \quad (4.49)$$

for any $a > 0$. We choose $c = \sqrt{\gamma_*(\beta, t_2) - \gamma_*(\beta, t_1)}/\beta$. By (4.47), we have

$$\text{p-lim}_{n \rightarrow \infty} M_{n,t_2}^2 = \frac{\gamma_*(t_2) - \gamma_*(t_1)}{\beta^2} = c^2,$$

and therefore, since $M_{n,t}$ is bounded, for any fixed $a > 0$

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [t_1, t_2]} M_{n,t} \geq c + a \right) &\leq \lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [t_1, t_2]} Y_{n,t} \geq a \right) \\ &\leq \frac{1}{a} \lim_{n \rightarrow \infty} \mathbb{E} [(M_{n,t_2} - c)^2]^{1/2} = 0. \end{aligned}$$

Together with Eq. (4.47), this yields

$$\text{p-lim}_{n \rightarrow \infty} \sup_{t \in [t_1, t_2]} M_{n,t}^2 = \frac{\gamma_*(t_2) - \gamma_*(t_1)}{\beta^2},$$

which coincides with the claim (4.46). \square

4.3 Natural Gradient Descent

Algorithm 3: NATURAL GRADIENT DESCENT ON $\mathcal{F}_{\text{TAP}}(\cdot; \mathbf{y}, q)$

Input: Initialization $\mathbf{u}^0 \in \mathbb{R}^n$, data $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\widehat{\mathbf{y}} \in \mathbb{R}^n$, step size $\eta > 0$, $q \in (0, 1)$, integer $K > 0$.

- 1 $\widehat{\mathbf{m}}^{+,0} = \tanh(\mathbf{u}^0)$.
- 2 **for** $k = 0, \dots, K - 1$ **do**
- 3 $\mathbf{u}^{k+1} \leftarrow \mathbf{u}^k - \eta \cdot \nabla \mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^{+,k}; \mathbf{y}, q)$,
- 4 $\widehat{\mathbf{m}}^{+,k+1} = \tanh(\mathbf{u}^{+,k+1})$,
- 5 **end**
- 6 **return** $\widehat{\mathbf{m}}^{+,K}$

The main objective of this section is to show that $\mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q)$ behaves well for $q = q_*(\beta, t)$ and for \mathbf{m} in a neighborhood of $\widehat{\mathbf{m}}^{K_{\text{AMP}}}$. Namely it has a unique local minimum $\mathbf{m}_* = \mathbf{m}_*(\mathbf{A}, \mathbf{y})$ in such a neighborhood, and NGD approximates \mathbf{m}_* well for large number of iterations K . Crucially, the map $\mathbf{y} \mapsto \mathbf{m}_*$ will be Lipschitz. For reference, we reproduce the NGD algorithm as Algorithm 3. This corresponds to lines 6-11 of Algorithm 1.

Lemma 4.10. Let $\beta < \frac{1}{2}$, $c \in (0, 1 - 2\beta)$, and $T > 0$ be fixed. Then there exists $\varepsilon_0 = \varepsilon_0(\beta, T)$ such that, for all $\varepsilon \in (0, \varepsilon_0)$ there exists $K_{\text{AMP}} = K_{\text{AMP}}(\beta, T, \varepsilon)$ and $\rho_0 = \rho_0(\beta, T, \varepsilon)$ such that for all $\rho \in (0, \rho_0)$ there exists $K_{\text{NGD}} = K_{\text{NGD}}(\beta, T, \varepsilon, \rho)$, such that the following holds.

Let $\widehat{\mathbf{m}}^{\text{AMP}} = \text{AMP}(\mathbf{A}, \mathbf{y}(t); K_{\text{AMP}})$ be the output of the AMP after K_{AMP} iterations, when applied to $\mathbf{y}(t)$. Fix $K \geq K_{\text{AMP}}$. With probability $1 - o_n(1)$ over $(\mathbf{A}, \mathbf{y}) \sim \mathbb{P}$, for all $t \in (0, T]$ and all $\widehat{\mathbf{y}} \in B(\mathbf{y}(t), c\sqrt{\varepsilon tn}/4)$, setting $q_* := q_*(\beta, t)$:

1. The function

$$\mathbf{m} \mapsto \mathcal{F}_{\text{TAP}}(\mathbf{m}; \widehat{\mathbf{y}}, q_*)$$

restricted to $B(\widehat{\mathbf{m}}^{\text{AMP}}, \sqrt{\varepsilon tn}) \cap (-1, 1)^n$ has a unique stationary point

$$\mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}) \in B(\widehat{\mathbf{m}}^{\text{AMP}}, \sqrt{\varepsilon tn}/2) \cap (-1, 1)^n$$

which is also a local minimum. In the case $\widehat{\mathbf{y}} = \mathbf{y}(t)$, $\mathbf{m}_*(\mathbf{A}, \mathbf{y}(t))$ also satisfies

$$\mathbf{m}_*(\mathbf{A}, \mathbf{y}) \in B(\widehat{\mathbf{m}}^{k'}, \sqrt{\varepsilon tn}/2) \cap (-1, 1)^n$$

for all $k' \in [K_{\text{AMP}}, K]$, where $\widehat{\mathbf{m}}^{k'} = \text{AMP}(\mathbf{A}, \mathbf{y}(t); k')$.

2. The stationary point $\mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}})$ satisfies (recall that $\mathbf{m}(\mathbf{A}, \mathbf{y})$ denotes the mean of the Gibbs measure)

$$\|\mathbf{m}(\mathbf{A}, \mathbf{y}) - \mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}})\|_2 \leq \rho\sqrt{tn}.$$

3. The stationary point \mathbf{m}_* satisfies the following Lipschitz property for all $\widehat{\mathbf{y}}, \widehat{\mathbf{y}}' \in B(\mathbf{y}(t), c\sqrt{\varepsilon tn}/4)$:

$$\|\mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}) - \mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}')\| \leq c^{-1}\|\widehat{\mathbf{y}} - \widehat{\mathbf{y}}'\|. \quad (4.50)$$

4. There exists a learning rate $\eta = \eta(\beta, T, \varepsilon)$ such that the following holds. Let $\widehat{\mathbf{m}}^{\text{NGD}}(\mathbf{A}, \widehat{\mathbf{y}})$ be the output of NGD (Algorithm 3), when run for K_{NGD} iterations with parameter q_* , $\widehat{\mathbf{y}}$, η . Assume that the initialization \mathbf{u}^0 satisfies

$$\|\mathbf{u}^0 - \text{arctanh}(\widehat{\mathbf{m}}^{\text{AMP}})\| \leq \frac{c\sqrt{\varepsilon tn}}{200}. \quad (4.51)$$

Then the algorithm output satisfies

$$\|\widehat{\mathbf{m}}^{\text{NGD}}(\mathbf{A}, \widehat{\mathbf{y}}) - \mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}})\| \leq \rho\sqrt{tn}. \quad (4.52)$$

The proof of this lemma is deferred to the appendix. Here we will prove the two key elements: first that $\widehat{\mathbf{m}}^{\text{AMP}}$ is an approximate stationary point of $\mathcal{F}_{\text{TAP}}(\cdot; \mathbf{y}(t), q_*)$ (Lemma 4.11), and second that $\mathcal{F}_{\text{TAP}}(\cdot; \widehat{\mathbf{y}}, q_*)$ is strongly convex in a neighborhood of $\widehat{\mathbf{m}}^{\text{AMP}}$ (Lemma 4.12). Let us point out that, in the local convexity guarantee, it is important that the neighborhood has radius $\Theta(\sqrt{tn})$ as $t \rightarrow 0$.

We recall below the expressions for the gradient and Hessian of $\mathcal{F}_{\text{TAP}}(\cdot; \mathbf{y}, q)$ at $\mathbf{m} \in (-1, 1)^n$:

$$\nabla \mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q) = -\beta \mathbf{A} \mathbf{m} - \mathbf{y} + \text{arctanh}(\mathbf{m}) + \beta^2 (1 - q) \mathbf{m} \quad (4.53)$$

$$\nabla^2 \mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q) = -\beta \mathbf{A} + \mathbf{D}(\mathbf{m}) + \beta^2 (1 - q) \mathbf{I}_n, \quad \mathbf{D}(\mathbf{m}) := \text{diag}(\{(1 - m_i^2)^{-1}\}_{i \leq n}). \quad (4.54)$$

In (4.53), arctanh is applied coordinate-wise to $\mathbf{m} \in (-1, 1)^n$.

For $t > 0, k \geq 0$ we let $\widehat{\mathbf{m}}^k = \text{AMP}(\mathbf{A}, \mathbf{y}(t); k)$ and define the quantities

$$q_k(\beta, t) := \frac{\gamma_{k+1}(\beta, t)}{\beta^2}, \quad q_*(\beta, t) := \frac{\gamma_*(\beta, t)}{\beta^2}. \quad (4.55)$$

Note that, by Lemma 4.6, we have

$$q_k(\beta, t) = \text{p-lim}_{n \rightarrow \infty} \frac{\|\widehat{\mathbf{m}}^k\|^2}{n}, \quad q_*(\beta, t) = \lim_{k \rightarrow \infty} q_k(\beta, t). \quad (4.56)$$

We will use the bounds (4.22), (4.23) in Lemma 4.5 several times below, which ensures that $(q_k(\beta, t)/t) \in [c, C]$ holds for constants $c, C > 0$ independent of $t \in (0, T]$ and $k \geq 1$.

Lemma 4.11. *Let $\widehat{\mathbf{m}}^k = \widehat{\mathbf{m}}^k(\mathbf{A}, \mathbf{y}(t))$ denote the AMP iterates on input $\mathbf{A}, \mathbf{y}(t)$. Then for any $T > 0$,*

$$\lim_{k \rightarrow \infty} \sup_{t \in (0, T]} \sup_{q \in [q_k(\beta, t), q_*(\beta, t)]} \text{p-lim}_{n \rightarrow \infty} \frac{\|\nabla_{\mathcal{F}_{\text{TAP}}}(\widehat{\mathbf{m}}^k; \mathbf{y}(t), q)\|}{\sqrt{tn}} = 0.$$

Proof. As in Algorithm 1, let

$$\mathbf{z}^{k+1} = \text{arctanh}(\widehat{\mathbf{m}}^{k+1}) = \beta \mathbf{A} \widehat{\mathbf{m}}^k + \mathbf{y} - \beta^2 \left(1 - \frac{1}{n} \|\widehat{\mathbf{m}}^k\|^2\right) \widehat{\mathbf{m}}^{k-1}.$$

Let $q \in [q_k(\beta, t), q_*(\beta, t)]$. Combining the above with Eqs. (4.53) and (4.56) yields

$$\begin{aligned} \frac{1}{\sqrt{n}} \|\nabla_{\mathcal{F}_{\text{TAP}}}(\text{AMP}(\mathbf{A}, \mathbf{y}; k); \mathbf{y}, q)\| &= \frac{1}{\sqrt{n}} \left\| -\beta \mathbf{A} \widehat{\mathbf{m}}^k - \mathbf{y} + \text{arctanh}(\widehat{\mathbf{m}}^k) + \beta^2 (1 - q) \widehat{\mathbf{m}}^k \right\| \\ &= \frac{1}{\sqrt{n}} \left\| \mathbf{z}^k - \beta \mathbf{A} \widehat{\mathbf{m}}^k - \mathbf{y} + \beta^2 (1 - q) \widehat{\mathbf{m}}^k \right\| \\ &\leq \frac{1}{\sqrt{n}} \|\mathbf{z}^{k+1} - \mathbf{z}^k\| + \frac{1}{\sqrt{n}} \left\| \mathbf{z}^{k+1} - \beta \mathbf{A} \widehat{\mathbf{m}}^k - \mathbf{y} + \beta^2 (1 - q) \widehat{\mathbf{m}}^k \right\| \\ &= \frac{1}{\sqrt{n}} \|\mathbf{z}^{k+1} - \mathbf{z}^k\| + \frac{\beta^2}{\sqrt{n}} \left\| (1 - \|\widehat{\mathbf{m}}^k\|^2/n) \widehat{\mathbf{m}}^{k-1} - (1 - q) \widehat{\mathbf{m}}^k \right\| \\ &\leq \frac{1}{\sqrt{n}} \|\mathbf{z}^{k+1} - \mathbf{z}^k\| + \frac{\beta^2}{\sqrt{n}} \|\widehat{\mathbf{m}}^{k-1} - \widehat{\mathbf{m}}^k\| \\ &\quad + \beta^2 (q_*(\beta, t) - q_k(\beta, t)) + o_{n, \mathbb{P}}(1). \end{aligned}$$

Here $o_{n, \mathbb{P}}(1)$ denotes terms which converge to 0 in probability as $n \rightarrow \infty$. By (4.38), (4.56) and the bound $(q_k(\beta, t)/t) \in [c, C]$

$$\lim_{k \rightarrow \infty} \sup_{t \in (0, T]} \text{p-lim}_{n \rightarrow \infty} \frac{\|\mathbf{z}^{k+1} - \mathbf{z}^k\|}{\sqrt{tn}} = 0.$$

Moreover, $\|\widehat{\mathbf{m}}^{k-1} - \widehat{\mathbf{m}}^k\| \leq \|\mathbf{z}^{k-1} - \mathbf{z}^k\|$ since the function $x \mapsto \tanh(x)$ is 1-Lipschitz. Finally (4.22) and (4.23) of Lemma 4.5 imply

$$\lim_{k \rightarrow \infty} \sup_{t \in (0, T]} \frac{q_*(\beta, t) - q_k(\beta, t)}{\sqrt{t}} = 0.$$

Combining the above statements concludes the proof. \square

We next control on the Hessian $\nabla^2 \mathcal{F}_{\text{TAP}}(\cdot; \mathbf{y}, q)$. As anticipated in Remark 2.1, this is the only part of our proof that requires $\beta < 1/2$ instead of $\beta < 1$.

Lemma 4.12. *Let $\beta > 0$, $\mathbf{y} \in \mathbb{R}^n$ and $q \in [0, 1]$. Then for all $\mathbf{m} \in (-1, 1)^n$,*

$$(1 - \beta \|\mathbf{A}\|_{\text{op}}) \mathbf{D}(\mathbf{m}) \preceq \nabla^2 \mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q) \preceq (1 + \beta^2 + \beta \|\mathbf{A}\|_{\text{op}}) \mathbf{D}(\mathbf{m}). \quad (4.57)$$

In particular if $\beta \leq \frac{1}{2} - c$, for $c > 0$, then with probability $1 - o_n(1)$, for all $\mathbf{m} \in (-1, 1)^n$,

$$c \mathbf{D}(\mathbf{m}) \preceq \nabla^2 \mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q) \preceq 2 \mathbf{D}(\mathbf{m}). \quad (4.58)$$

Proof. The upper and lower bounds in Eq. (4.57) are obtained from (4.54) using the fact that $\mathbf{D}(\mathbf{m}) \succeq \mathbf{I}_n$ for all $\mathbf{m} \in (-1, 1)^n$. Further, we use the fact that $\|\mathbf{A}\|_{\text{op}} \leq 2 + o_n(1)$ with probability $1 - o_n(1)$. Therefore, Eq. (4.58) follows from the assumption $\beta \leq \frac{1}{2} - c$. \square

As mentioned above, our convergence analysis of NGD, and proof of Lemma 4.10 are given in appendix A. The key insight is that the main iterative step in line 3 of Algorithm 3 can be expressed as a version of mirror descent. Define the concave function $h(\mathbf{m}) = \sum_{i=1}^n h(m_i)$ for $\mathbf{m} \in (-1, 1)^n$ (recall that $h(x) := -((1+x)/2) \log((1+x)/2) - ((1-x)/2) \log((1-x)/2)$). Following [LFN18], we define for $\mathbf{m}, \mathbf{n} \in (-1, 1)^n$ the Bregman divergence

$$D_{-h}(\mathbf{m}, \mathbf{n}) = -h(\mathbf{m}) + h(\mathbf{n}) + \langle \nabla h(\mathbf{n}), \mathbf{m} - \mathbf{n} \rangle. \quad (4.59)$$

Then with $L = 1/\eta$, the update in line 3 admits the alternate description

$$\widehat{\mathbf{m}}^{+,k+1} = \arg \min_{\mathbf{x} \in (-1, 1)^n} \langle \nabla \mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^{+,k}; \mathbf{y}, q), \mathbf{x} - \widehat{\mathbf{m}}^{+,k} \rangle + L \cdot D_{-h}(\mathbf{x}, \widehat{\mathbf{m}}^{+,k}). \quad (4.60)$$

We will use this description to prove convergence.

Remark 4.3. If the Hessian $\nabla^2 \mathcal{F}_{\text{TAP}}$ were bounded above and below by constant multiples of the identity matrix instead of $\mathbf{D}(\mathbf{m})$, then we could use simple gradient descent instead of NGD in Algorithm 1. This would also simplify the proof. However, $\nabla^2 \mathcal{F}_{\text{TAP}}$ is not bounded above near the boundaries of $(-1, +1)^n$. The use of NGD to minimize TAP free energy was introduced in [CFM21], which however considered a different regime in the planted model.

Remark 4.4. Our proof of Lemma 4.10 does not require $\nabla^2 \mathcal{F}_{\text{TAP}}$ to be globally convex. Instead, we only use the fact that, with probability $1 - o_n(1)$,

$$\nabla^2 \mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}, q) \succeq c \mathbf{D}(\mathbf{m}), \quad \forall \mathbf{m} \in B(\widehat{\mathbf{m}}^{\text{AMP}}, \sqrt{\varepsilon t n}) \cap (-1, 1).$$

For $\beta \in [1/2, 1)$ we expect only this weaker guarantee to hold. We believe the technique of [CFM21] could be used to prove such local strong convexity in the full regime $\beta \in [0, 1)$.

4.4 Continuous limit and proof of Theorem 2.1

We fix (β, T) and choose constants $K_{\text{AMP}} = K_{\text{AMP}}(\beta, T, \varepsilon)$, $\rho_0 = \rho_0(\beta, T, \varepsilon, K_{\text{AMP}})$, $\rho \in (0, \rho_0)$ and $K_{\text{NGD}} = K_{\text{NGD}}(\beta, T, \varepsilon, \rho)$ so that Lemma 4.10 holds.

We couple the discretized process $(\widehat{\mathbf{y}}_\ell)_{\ell \geq 0}$ defined in Eq. (2.6) (line 6 of Algorithm 2) to the continuous time process $(\mathbf{y}(t))_{t \in \mathbb{R}_{\geq 0}}$ (cf. Eq. (4.8)) via the driving noise, as follows:

$$\mathbf{w}_{\ell+1} = \frac{1}{\sqrt{\delta}} \int_{\ell\delta}^{(\ell+1)\delta} d\mathbf{B}(t). \quad (4.61)$$

We denote by $\widehat{\mathbf{m}}(\mathbf{A}, \mathbf{y})$ the output of the mean estimation algorithm 1 on input \mathbf{A}, \mathbf{y} . By Lemma 4.10, which ensures that, for any $t \in (0, T]$, with probability $1 - o_n(1)$,

$$\|\widehat{\mathbf{m}}(\mathbf{A}, \mathbf{y}(t)) - \mathbf{m}_*(\mathbf{A}, \mathbf{y}(t); q_*(\beta, t))\| \leq \rho\sqrt{tn}. \quad (4.62)$$

Here and below we note explicitly the dependence of \mathbf{m}_* on t via q_* . The next lemma provides a crude estimate on the Lipschitz continuity of AMP with respect to its input.

Lemma 4.13. *Recall that $\text{AMP}(\mathbf{A}, \mathbf{y}; k) \in \mathbb{R}^n$ denotes the output of the AMP algorithm on input (\mathbf{A}, \mathbf{y}) , after k iterations, cf. Eq. (2.2). If $\|\mathbf{A}\|_{\text{op}} \leq 3$, then, for any $\mathbf{y}, \widehat{\mathbf{y}} \in \mathbb{R}^n$,*

$$\|\arctanh(\text{AMP}(\mathbf{A}, \mathbf{y}; k)) - \arctanh(\text{AMP}(\mathbf{A}, \widehat{\mathbf{y}}; k))\|_2 \leq k6^k \|\mathbf{y} - \widehat{\mathbf{y}}\|_2. \quad (4.63)$$

Proof. For $0 \leq j \leq k$, set:

$$\begin{aligned} \mathbf{m}^j &= \text{AMP}(\mathbf{A}, \mathbf{y}; j), & \mathbf{z}^j &= \arctanh(\mathbf{m}^j), & \mathbf{b}_j &= \frac{\beta^2}{n} \sum_{i=1}^n (1 - \tanh^2(z_i^j)), \\ \widehat{\mathbf{m}}^j &= \text{AMP}(\mathbf{A}, \widehat{\mathbf{y}}; j), & \widehat{\mathbf{z}}^j &= \arctanh(\widehat{\mathbf{m}}^j), & \widehat{\mathbf{b}}_j &= \frac{\beta^2}{n} \sum_{i=1}^n (1 - \tanh^2(\widehat{z}_i^j)). \end{aligned}$$

Using the AMP update equation (line 4 of Algorithm 1) and the fact that $\tanh(\cdot)$ is 1-Lipschitz, we obtain

$$\begin{aligned} \|\mathbf{z}^{j+1} - \widehat{\mathbf{z}}^{j+1}\| &\leq \|\beta \mathbf{A}(\mathbf{m}^j - \widehat{\mathbf{m}}^j)\| + \|\mathbf{y} - \widehat{\mathbf{y}}\| + \|\mathbf{b}_j \mathbf{m}^{j-1} - \mathbf{b}_j \widehat{\mathbf{m}}^{j-1}\| + \|\mathbf{b}_j \widehat{\mathbf{m}}^{j-1} - \widehat{\mathbf{b}}_j \widehat{\mathbf{m}}^{j-1}\| \\ &\leq 3\beta \|\mathbf{z}^j - \widehat{\mathbf{z}}^j\| + \|\mathbf{y} - \widehat{\mathbf{y}}\| + \mathbf{b}_j \|\mathbf{z}^{j-1} - \widehat{\mathbf{z}}^{j-1}\| + |\mathbf{b}_j - \widehat{\mathbf{b}}_j| \sqrt{n}. \end{aligned}$$

Note that $|1 - \tanh^2(x)| \leq 1$ for all $x \in \mathbb{R}$ and $|\mathbf{b}_j| \leq \beta^2$. Setting $E_j = \max_{i \leq j} \|\mathbf{z}^{i+1} - \widehat{\mathbf{z}}^{i+1}\|$, we find

$$\begin{aligned} E_{j+1} &\leq (3\beta^2 + 3\beta)E_j + \|\mathbf{y} - \widehat{\mathbf{y}}\| \\ &\leq 6E_j + \|\mathbf{y} - \widehat{\mathbf{y}}\|. \end{aligned}$$

It follows by induction that

$$E_j \leq j6^j \|\mathbf{y} - \widehat{\mathbf{y}}\|.$$

Setting $j = k$ concludes the proof. \square

Define the random approximation errors

$$A_\ell := \frac{1}{\sqrt{n}} \|\widehat{\mathbf{y}}_\ell - \mathbf{y}(\ell\delta)\|, \quad (4.64)$$

$$B_\ell := \frac{1}{\sqrt{n}} \|\widehat{\mathbf{m}}(\mathbf{A}, \widehat{\mathbf{y}}_\ell) - \mathbf{m}(\mathbf{A}, \mathbf{y}(\ell\delta))\|. \quad (4.65)$$

Note that $A_0 = B_0 = 0$. In the next lemma we bound the above quantities:

Lemma 4.14. *For $\beta < 1/2$ and $T > 0$, there exists a constant $C = C(\beta) < \infty$, and a deterministic non-negative sequence $\xi(n)$ with $\lim_{n \rightarrow \infty} \xi(n) = 0$ such that the following holds with probability $1 - o_n(1)$. For every $\ell \geq 0$, $\delta \in (0, 1)$ such that $\ell\delta \leq T$,*

$$A_\ell \leq C e^{C\ell\delta} \ell\delta (\rho\sqrt{\ell\delta} + \sqrt{\delta}) + \xi(n), \quad (4.66)$$

$$B_\ell \leq C e^{C\ell\delta} \ell\delta (\rho\sqrt{\ell\delta} + \sqrt{\delta}) + C\rho\sqrt{\ell\delta} + \xi(n). \quad (4.67)$$

Proof. Throughout the proof, we denote by $\xi(n)$ a deterministic non-negative sequence $\xi(n)$ with $\lim_{n \rightarrow \infty} \xi(n) = 0$, which can change from line to line. Also, C will denote a generic constant that may depend on β, T, K_{AMP} .

The proof proceeds by induction on ℓ . As the base case is trivial, we assume the result holds for all $j \leq \ell$ and we prove it for $\ell + 1$. We first claim that with probability $1 - o_n(1)$,

$$A_{\ell+1} \leq A_\ell + \delta B_\ell + C\delta^{3/2}. \quad (4.68)$$

Indeed, using (4.61) we find

$$\begin{aligned} A_{\ell+1} - A_\ell &\leq n^{-1/2} \int_{\ell\delta}^{(\ell+1)\delta} \|\widehat{\mathbf{m}}(\mathbf{A}, \widehat{\mathbf{y}}_\ell) - \mathbf{m}(\mathbf{A}, \mathbf{y}(t))\| dt \\ &\leq \delta n^{-1/2} \left(\|\widehat{\mathbf{m}}(\mathbf{A}, \widehat{\mathbf{y}}_\ell) - \mathbf{m}(\mathbf{A}, \mathbf{y}(\ell\delta))\| + \sup_{t \in [\ell\delta, (\ell+1)\delta]} \|\mathbf{m}(\mathbf{A}, \mathbf{y}(t)) - \mathbf{m}(\mathbf{A}, \mathbf{y}(\ell\delta))\| \right) \\ &\leq \delta B_\ell + \delta n^{-1/2} \cdot \sup_{t \in [\ell\delta, (\ell+1)\delta]} \|\mathbf{m}(\mathbf{A}, \mathbf{y}(t)) - \mathbf{m}(\mathbf{A}, \mathbf{y}(\ell\delta))\| \\ &\leq \delta B_\ell + C(\beta)\delta^{3/2} + \xi(n), \end{aligned}$$

where the last line holds with high probability by Lemma 4.9 and Eq. (4.24) of Lemma 4.5. Using this bound together with the inductive hypothesis on A_ℓ and B_ℓ , we obtain

$$\begin{aligned} A_{\ell+1} &\leq C e^{C(\ell+1)\delta} \ell\delta(\rho\sqrt{\ell\delta} + \sqrt{\delta}) + C\rho\delta\sqrt{\ell\delta} + C\delta^{3/2} + \xi(n) \\ &\leq C e^{C(\ell+1)\delta} (\ell+1)\delta(\rho + \sqrt{\delta}) + \xi(n). \end{aligned}$$

This implies Eq. (4.66) for $\ell + 1$.

We next show that Eq. (4.67) holds with ℓ replaced by $\ell + 1$. By the bound (4.66) for $\ell + 1$, taking $\delta \leq \delta(\beta, \varepsilon, K_{\text{AMP}}, T)$ and $\rho \in (0, \rho_0)$ $\rho = \rho(\beta, \varepsilon, K_{\text{AMP}}, T)$ ensures that

$$A_{\ell+1} \leq \frac{c\sqrt{\varepsilon\ell\delta}}{200K_{\text{AMP}}6^{K_{\text{AMP}}}},$$

where ε can be chosen an arbitrarily small constant. So by Lemma 4.13, we have with probability $1 - o_n(1)$,

$$\begin{aligned} \left\| \arctanh(\text{AMP}(\mathbf{A}, \mathbf{y}((\ell+1)\delta); K_{\text{AMP}})) - \arctanh(\text{AMP}(\mathbf{A}, \widehat{\mathbf{y}}_{\ell+1}; K_{\text{AMP}})) \right\|_2 &\leq K_{\text{AMP}}6^{K_{\text{AMP}}} A_{\ell+1} \sqrt{n} \\ &\leq \frac{c\sqrt{\varepsilon\ell\delta n}}{200}. \end{aligned}$$

By choosing $\varepsilon \leq \varepsilon_0(\beta, T)$, we obtain that Lemma 4.10, part 4 applies. We thus find

$$\|\widehat{\mathbf{m}}(\mathbf{A}, \widehat{\mathbf{y}}_{\ell+1}) - \mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}_{\ell+1})\| \leq \rho\sqrt{\ell\delta n}.$$

Using parts 3 and 2 respectively of Lemma 4.10 on the other terms below, by triangle inequality we obtain (writing for simplicity $q_\ell := q_*(\beta, \ell\delta)$)

$$\begin{aligned} \|\widehat{\mathbf{m}}(\mathbf{A}, \widehat{\mathbf{y}}_{\ell+1}) - \mathbf{m}(\mathbf{A}, \mathbf{y}((\ell+1)\delta))\| &\leq \|\widehat{\mathbf{m}}(\mathbf{A}, \widehat{\mathbf{y}}_{\ell+1}) - \mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}_{\ell+1}; q_{\ell+1})\| \\ &\quad + \|\mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}_{\ell+1}; q_{\ell+1}) - \mathbf{m}_*(\mathbf{A}, \mathbf{y}((\ell+1)\delta); q_{\ell+1})\| \\ &\quad + \|\mathbf{m}_*(\mathbf{A}, \mathbf{y}((\ell+1)\delta); q_{\ell+1}) - \mathbf{m}(\mathbf{A}, \mathbf{y}((\ell+1)\delta))\| \\ &\leq (\rho\sqrt{\ell\delta} + c^{-1}A_{\ell+1} + \rho\sqrt{\ell\delta} + \xi(n))\sqrt{n}. \end{aligned} \quad (4.69)$$

In other words with probability $1 - o_n(1)$,

$$B_{\ell+1} \leq c^{-1}A_{\ell+1} + 2\rho\sqrt{\ell\delta} + \xi(n).$$

Using this together with the bound (4.66) for $\ell + 1$ verifies the inductive step for (4.67) and concludes the proof. \square

Finally we show that standard randomized rounding is continuous in $W_{2,n}$.

Lemma 4.15. *Suppose probability distributions μ_1, μ_2 on $[-1, 1]^n$ are given. Sample $\mathbf{m}_1 \sim \mu_1$ and $\mathbf{m}_2 \sim \mu_2$ and let $\mathbf{x}_1, \mathbf{x}_2 \in \{-1, +1\}^n$ be standard randomized roundings, respectively of \mathbf{m}_1 and \mathbf{m}_2 . (Namely, the coordinates of \mathbf{x}_i are conditionally independent given \mathbf{m}_i , with $\mathbb{E}[\mathbf{x}_i | \mathbf{m}_i] = \mathbf{m}_i$.) Then*

$$W_{2,n}(\mathcal{L}(\mathbf{x}_1), \mathcal{L}(\mathbf{x}_2)) \leq 2\sqrt{W_{2,n}(\mu_1, \mu_2)}.$$

Proof. Let $(\mathbf{m}_1, \mathbf{m}_2)$ be distributed according to a $W_{2,n}$ -optimal coupling between μ_1, μ_2 . Couple the roundings $\mathbf{x}_1, \mathbf{x}_2$ by choosing i.i.d. uniform random variables $u_i \sim \text{Unif}([0, 1])$ for $i \in [n]$, and for $(i, j) \in [n] \times \{1, 2\}$ setting

$$(\mathbf{x}_j)_i = \begin{cases} +1, & \text{if } u \leq \frac{1+(\mathbf{m}_j)_i}{2}, \\ -1, & \text{else.} \end{cases}$$

Then it is not difficult to see that

$$\begin{aligned} \frac{1}{n} \mathbb{E} [\|\mathbf{x}_1 - \mathbf{x}_2\|^2 | (\mathbf{m}_1, \mathbf{m}_2)] &= \frac{2}{n} \sum_{i=1}^n |(\mathbf{m}_1)_i - (\mathbf{m}_2)_i| \\ &\leq 2\sqrt{\frac{1}{n} \|\mathbf{m}_1 - \mathbf{m}_2\|^2}. \end{aligned}$$

Averaging over the choice of $(\mathbf{m}_1, \mathbf{m}_2)$ implies the result. \square

Proof of Theorem 2.1. Set $\ell = L = T/\delta$ and $\rho = \sqrt{\delta}$ in Eq. (4.67). With all laws $\mathcal{L}(\cdot)$ conditional on \mathbf{A} below, we find

$$\begin{aligned} \mathbb{E} W_{2,n}(\mu_{\mathbf{A}}, \mathcal{L}(\widehat{\mathbf{m}}(\mathbf{A}, \widehat{\mathbf{y}}_L))) &\leq \mathbb{E} W_{2,n}(\mu_{\mathbf{A}}, \mathcal{L}(\mathbf{m}(\mathbf{A}, \mathbf{y}(T)))) + \mathbb{E} W_{2,n}(\mathcal{L}(\mathbf{m}(\mathbf{A}, \mathbf{y}(T))), \mathcal{L}(\widehat{\mathbf{m}}(\mathbf{A}, \widehat{\mathbf{y}}_L))) \\ &\leq T^{-1/2} + C(\beta, T)\sqrt{\delta} + o_n(1). \end{aligned}$$

Here the first term was bounded by Eq. (3.4) in Section 3 and the second by Eq. (4.67). Taking T sufficiently large, δ sufficiently small, and n sufficiently large, we may obtain

$$\mathbb{E} W_{2,n}(\mu_{\mathbf{A}}, \mathcal{L}(\widehat{\mathbf{m}}_{\text{NGD}}(\mathbf{A}, \widehat{\mathbf{y}}_L))) \leq \frac{\varepsilon^2}{4}$$

for any desired $\varepsilon > 0$. Applying Lemma 4.15 shows that

$$\mathbb{E} W_{2,n}(\mu_{\mathbf{A}}, \mathbf{x}^{\text{alg}}) \leq \varepsilon.$$

The Markov inequality now implies that (2.7) holds with probability $1 - o_n(1)$ as desired. \square

5 Algorithmic stability and disorder chaos

In this section we prove Theorem 2.3 establishing that our sampling algorithm, Algorithm 2 is stable. Next, we prove that the Sherrington-Kirkpatrick measure $\mu_{\mathbf{A},\beta}$ exhibits W_2 -disorder chaos for $\beta > 1$, proving Theorem 2.5 and deduce that no stable algorithm can sample in normalized W_2 distance for $\beta > 1$, see Theorem 2.6.

5.1 Algorithmic stability: Proof of Theorem 2.3

Recall Definition 2.2, defining sampling algorithms as measurable functions $\text{ALG}_n : (\mathbf{A}, \beta, \omega) \mapsto \text{ALG}_n(\mathbf{A}, \beta, \omega) \in [-1, 1]^n$ where $\beta \geq 0$ and ω is an independent random variable taking values in some probability space.

Remark 5.1. In light of Lemma 4.15, we can always turn a stable sampling algorithm ALG with codomain $[-1, 1]^n$ into a stable sampling algorithm with binary output:

$$\widetilde{\text{ALG}}_n(\mathbf{A}, \beta, \tilde{\omega}) \in \{-1, +1\}^n.$$

Indeed this is achieved by standard randomized rounding, i.e., drawing a (conditionally independent) random binary value with mean $(\widetilde{\text{ALG}}(\mathbf{A}, \beta, \tilde{\omega}))_i$ for each coordinate $1 \leq i \leq n$.

Recall the definition of the interpolating family $(\mathbf{A}_s)_{s \in [0,1]}$ whereby $\mathbf{A}_0, \mathbf{A}_1 \sim \text{GOE}(n)$ i.i.d. and

$$\mathbf{A}_s = \sqrt{1-s^2} \mathbf{A}_0 + s \mathbf{A}_1, \quad s \in [0, 1], \quad (5.1)$$

We take $\mu_{\mathbf{A}_s, \beta}(\mathbf{x}) \propto \exp\{(\beta/2)\langle \mathbf{x}, \mathbf{A}_s \mathbf{x} \rangle\}$ to be the corresponding Gibbs measure.

We start with the following simple estimate.

Lemma 5.1. *There exists an absolute constant $C > 0$ such that*

$$\inf_{s \in (0,1)} \mathbb{P}\left(\|\mathbf{A}_0 \mathbf{u} - \mathbf{A}_s \mathbf{v}\| \leq C(\|\mathbf{u} - \mathbf{v}\| + s\sqrt{n}), \quad \forall \mathbf{u}, \mathbf{v} \in [-1, 1]^n\right) = 1 - o_n(1). \quad (5.2)$$

Proof. We write

$$\begin{aligned} \|\mathbf{A}_0 \mathbf{u} - \mathbf{A}_s \mathbf{v}\| &\leq \|\mathbf{A}_0 \mathbf{u} - \mathbf{A}_0 \mathbf{v}\| + \|\mathbf{A}_0 \mathbf{v} - \mathbf{A}_s \mathbf{v}\| \\ &\leq \|\mathbf{A}_0\|_{\text{op}} \|\mathbf{u} - \mathbf{v}\| + \|(1 - \sqrt{1-s^2})\mathbf{A}_0 - s\mathbf{A}_1\|_{\text{op}} \|\mathbf{v}\|. \end{aligned}$$

We note that $(1 - \sqrt{1-s^2})\mathbf{A}_0 - s\mathbf{A}_1 \stackrel{d}{=} \sqrt{2(1 - \sqrt{1-s^2})}\mathbf{A}_0$ and $\sqrt{2(1 - \sqrt{1-s^2})} \sim s$ for small s and this quantity is bounded above by a constant for any $s \in [0, 1]$. The result follows since $\|\mathbf{A}_0\|_{\text{op}} \leq 2.1$ with probability $1 - o_n(1)$. \square

Proposition 5.2. *Suppose an algorithm ALG is given by an iterative procedure*

$$\begin{aligned} \mathbf{z}^{k+1} &= G_k((\mathbf{z}^j, \beta \mathbf{A} \mathbf{m}^j, \mathbf{A} \mathbf{m}^j, \beta^2 \mathbf{m}^j, \mathbf{w}^j)_{0 \leq j \leq k}), \quad 0 \leq k \leq K-1, \\ \mathbf{m}^k &= \rho_k(\mathbf{z}^k), \quad 0 \leq k \leq K-1, \\ \text{ALG}_n(\mathbf{A}, \beta, \omega) &:= \mathbf{m}^K, \end{aligned}$$

where the sequence $\omega = (\mathbf{w}^0, \dots, \mathbf{w}^{K-1}) \in (\mathbb{R}^n)^K$, the initialization $\mathbf{z}^0 \in \mathbb{R}^n$, and \mathbf{A} are mutually independent, and the functions $G_k : (\mathbb{R}^n)^{5k+5} \rightarrow \mathbb{R}^n$ and $\rho_k : \mathbb{R}^n \rightarrow [-1, 1]^n$ are L_0 -Lipschitz for $L_0 \geq 0$ an n -independent constant. Then ALG is both disorder-stable and temperature-stable.

Proof. Let us generate iterates $\mathbf{z}^k = \mathbf{z}^k(\mathbf{A}_0, \beta)$ and $\tilde{\mathbf{z}}^k = \mathbf{z}^k(\mathbf{A}_s, \tilde{\beta})$ for $0 \leq k \leq K$ using the same initialization $\mathbf{z}^0 = \tilde{\mathbf{z}}^0$ and external randomness $\omega = (\mathbf{w}^0, \dots, \mathbf{w}^{K-1})$, but with different Hamiltonians and inverse temperatures. Similarly let $\mathbf{m}^k = \rho_k(\mathbf{z}^k)$ and $\tilde{\mathbf{m}}^k = \rho_k(\tilde{\mathbf{z}}^k)$. We will allow C to vary from line to line in the proof below.

First by Lemma 5.1, with probability $1 - o_n(1)$,

$$\begin{aligned} \|\beta \mathbf{A}_0 \mathbf{m}^k - \tilde{\beta} \mathbf{A}_s \tilde{\mathbf{m}}^k\| &\leq \|\beta \mathbf{A}_0 \mathbf{m}^k - \beta \mathbf{A}_s \tilde{\mathbf{m}}^k\| + \|\beta \mathbf{A}_s \tilde{\mathbf{m}}^k - \tilde{\beta} \mathbf{A}_s \tilde{\mathbf{m}}^k\| \\ &\leq C\beta \|\mathbf{m}^k - \tilde{\mathbf{m}}^k\| + C\beta s\sqrt{n} + |\beta - \tilde{\beta}| \cdot \|\mathbf{A}_s \tilde{\mathbf{m}}^k\| \\ &\leq C(\|\mathbf{m}^k - \tilde{\mathbf{m}}^k\| + s\sqrt{n} + |\beta - \tilde{\beta}|\sqrt{n}). \end{aligned}$$

Similarly as long as $\tilde{\beta} \leq 2\beta$ so that $|\beta^2 - \tilde{\beta}^2| \leq 3\beta|\beta - \tilde{\beta}|$, we have

$$\begin{aligned} \|\beta^2 \mathbf{m}^k - \tilde{\beta}^2 \tilde{\mathbf{m}}^k\| &\leq \|\beta^2 \mathbf{m}^k - \beta^2 \tilde{\mathbf{m}}^k\| + \|\beta^2 \tilde{\mathbf{m}}^k - \tilde{\beta}^2 \tilde{\mathbf{m}}^k\| \\ &\leq \beta^2 \|\mathbf{m}^k - \tilde{\mathbf{m}}^k\| + 3\beta|\beta - \tilde{\beta}|\sqrt{n}. \end{aligned}$$

It follows that the error sequence

$$A_k = \frac{1}{\sqrt{n}} \max_{j \leq k} \|\mathbf{z}^{j+1}(\mathbf{A}_0, \beta) - \mathbf{z}^{j+1}(\mathbf{A}_s, \tilde{\beta})\|$$

satisfies with probability $1 - o_n(1)$ the recursion

$$\begin{aligned} A_{k+1} &\leq L_0 k^{1/2} C (A_k + s + |\beta - \tilde{\beta}|), \\ A_0 &= 0, \end{aligned}$$

for a suitable $C = C(\beta)$. It follows that with probability $1 - o_n(1)$,

$$A_K \leq \sum_{k=1}^K (L_0 k^{1/2} C)^k (s + |\beta - \tilde{\beta}|) \leq K (L_0 K C)^K (s + |\beta - \tilde{\beta}|). \quad (5.3)$$

Since $\|\mathbf{m}^K(\mathbf{A}_0) - \mathbf{m}^K(\mathbf{A}_s)\| \leq 2\sqrt{n}$ almost surely, we obtain for any $\eta > 0$

$$n^{-1} \mathbb{E} \left[\|\mathbf{m}^K(\mathbf{A}_0) - \mathbf{m}^K(\mathbf{A}_s)\|^2 \right] \leq (L_0 K (L_0 K C)^K (s + |\beta - \tilde{\beta}|))^2 + \eta$$

if $n \geq n_0(\eta)$ is large enough so that Eq. (5.3) holds with probability at least $1 - \frac{\eta}{4}$. The stability of the algorithm follows. \square

Proof of Theorem 2.3. We show that Algorithm 2 with n -independent parameters $(\beta, \eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta)$ is of the form in Proposition 5.2 for a constant $L_0 = L_0(\beta, \eta, K_{\text{AMP}}, K_{\text{NGD}}, L, \delta)$. Indeed note that the algorithm goes through L iterations, indexed by $\ell \in \{0, \dots, L-1\}$.

During each of these iterations, two loops are run (here we modify the notation introduced in Algorithm 1 and Algorithm 2, to account for the dependence on ℓ , and to get closer to the notation of Proposition 5.2):

1. The AMP loop, whereby, for $k = 0, \dots, K_{\text{AMP}} - 1$,

$$\widehat{\mathbf{m}}^{\ell, k} = \tanh(\mathbf{z}^{\ell, k}), \quad \mathbf{b}(\widehat{\mathbf{m}}^{\ell, k}) = \frac{\beta^2}{n} \sum_{i=1}^n \tanh'(z_i^{\ell, k}), \quad (5.4)$$

$$\mathbf{z}^{\ell, k+1} = \beta \mathbf{A} \widehat{\mathbf{m}}^{\ell, k} + \widehat{\mathbf{y}}_\ell - \mathbf{b}(\mathbf{z}^{\ell, k}) \widehat{\mathbf{m}}^{\ell, k-1}. \quad (5.5)$$

(Here $\tanh'(x)$ denotes the first derivative of $\tanh(x)$.)

2. The NGD loop, whereby, for $k = K_{\text{AMP}}, \dots, K_{\text{AMP}} + K_{\text{NGD}} - 1$, setting $q_\ell = q_{K_{\text{AMP}}}(\beta, t = \ell\delta)$,

$$\widehat{\mathbf{m}}^{\ell,k} = \tanh(\mathbf{z}^{\ell,k}), \quad (5.6)$$

$$\mathbf{z}^{\ell,k+1} = \mathbf{z}^{\ell,k} + \eta[\beta \mathbf{A} \widehat{\mathbf{m}}^{\ell,k} + \mathbf{y}_\ell - \mathbf{z}^{\ell,k} - \beta^2 (1 - q_\ell) \mathbf{m}^{\ell,k}]. \quad (5.7)$$

Further, recalling line 6 of Algorithm 2, $\widehat{\mathbf{y}}_\ell$ is updated via

$$\widehat{\mathbf{y}}_{\ell+1} = \widehat{\mathbf{y}}_\ell + \widehat{\mathbf{m}}^{\ell, K_{\text{AMP}} + K_{\text{NGD}}} \delta + \sqrt{\delta} \mathbf{w}_{\ell+1}. \quad (5.8)$$

These updates take the same form as in Proposition 5.2, with iterations indexed by (ℓ, k) , $\omega = (\mathbf{w}_\ell)_{\ell \leq L}$, $\rho_{\ell,k}(\mathbf{z}) = \tanh(\mathbf{z})$ for all ℓ, k , and

$$G_{\ell,k} \left((\mathbf{z}^{\ell',j}, \beta \mathbf{A} \widehat{\mathbf{m}}^{\ell',j}, \mathbf{A} \widehat{\mathbf{m}}^{\ell',j}, \beta^2 \widehat{\mathbf{m}}^{\ell',j}, \mathbf{w}_{\ell'})_{\ell',j} \right) = \beta \mathbf{A} \widehat{\mathbf{m}}^{\ell,k} + \widehat{\mathbf{y}}_\ell - \mathbf{b}(\mathbf{z}^{\ell,k}) \widehat{\mathbf{m}}^{\ell,k-1}, \quad 0 \leq k \leq K_{\text{AMP}} - 1, \quad (5.9)$$

$$\begin{aligned} G_{\ell,k} \left((\mathbf{z}^{\ell',j}, \beta \mathbf{A} \widehat{\mathbf{m}}^{\ell',j}, \mathbf{A} \widehat{\mathbf{m}}^{\ell',j}, \beta^2 \widehat{\mathbf{m}}^{\ell',j}, \mathbf{w}_{\ell'})_{\ell',j} \right) &= \\ &= \mathbf{z}^{\ell,k} + \eta[\beta \mathbf{A} \widehat{\mathbf{m}}^{\ell,k} + \mathbf{y}_\ell - \mathbf{z}^{\ell,k} - \beta^2 (1 - q_\ell) \mathbf{m}^{\ell,k}], \quad K_{\text{AMP}} \leq k \leq K_{\text{AMP}} + K_{\text{NGD}} - 1. \end{aligned} \quad (5.10)$$

Notice that these functions depend on previous iterates both explicitly, as noted, and implicitly through $\widehat{\mathbf{y}}_\ell$. By summing up Eq. (5.8), we obtain

$$\widehat{\mathbf{y}}_\ell = \sum_{j=0}^{\ell-1} \widehat{\mathbf{m}}^{j, K_{\text{AMP}} + K_{\text{NGD}}} \delta + \sqrt{\delta} \sum_{j=1}^{\ell} \mathbf{w}_{\ell+1}, \quad (5.11)$$

which is Lipschitz in the previous iterates $(\mathbf{m}^{j,k})_{j \leq \ell-1, k < K_{\text{AMP}} + K_{\text{NGD}}}$. Since both (5.9) and (5.10) depend linearly on $\widehat{\mathbf{y}}_\ell$ (with n -independent coefficients), it is sufficient to consider the explicit dependence on previous iterates of $G_{\ell,k}$. Namely, it is sufficient to control the Lipschitz modulus of the following functions

$$\widetilde{G}_{\ell,k} \left(\mathbf{z}^{\ell,k}, \beta \mathbf{A} \widehat{\mathbf{m}}^{\ell,k}, \widehat{\mathbf{m}}^{\ell,k-1} \right) = \beta \mathbf{A} \widehat{\mathbf{m}}^{\ell,k} - \mathbf{b}(\mathbf{z}^{\ell,k}) \widehat{\mathbf{m}}^{\ell,k-1}, \quad k < K_{\text{AMP}} \quad (5.12)$$

$$\widetilde{G}_{\ell,k} \left(\mathbf{z}^{\ell,k}, \beta \mathbf{A} \widehat{\mathbf{m}}^{\ell,k}, \beta^2 \widehat{\mathbf{m}}^{\ell,k} \right) = \mathbf{z}^{\ell,k} + \eta[\beta \mathbf{A} \widehat{\mathbf{m}}^{\ell,k} - \mathbf{z}^{\ell,k} - \beta^2 (1 - q_\ell) \widehat{\mathbf{m}}^{\ell,k}], \quad k > K_{\text{AMP}}. \quad (5.13)$$

Consider first Eq. (5.12). Since $|\tanh''(x)| \leq 2$ for all $x \in \mathbb{R}$, it follows that

$$\|\mathbf{b}(\mathbf{z}) - \mathbf{b}(\widetilde{\mathbf{z}})\| \leq \frac{2\beta^2}{n} \sum_{i=1}^n |z_i - \widetilde{z}_i| \leq \frac{2\beta^2}{\sqrt{n}} \|\mathbf{z} - \widetilde{\mathbf{z}}\|_2.$$

Therefore, that for any $(\mathbf{u}, \mathbf{v}, \beta, \widetilde{\mathbf{u}}, \widetilde{\mathbf{v}}, \widetilde{\beta})$ (noting explicitly the dependence of \mathbf{b} upon β):

$$\begin{aligned} \|\mathbf{b}_\beta(\mathbf{u}) \tanh(\mathbf{v}) - \mathbf{b}_{\widetilde{\beta}}(\widetilde{\mathbf{u}}) \tanh(\widetilde{\mathbf{v}})\| &\leq \|\mathbf{b}_\beta(\mathbf{u}) \tanh(\mathbf{v}) - \mathbf{b}_\beta(\widetilde{\mathbf{u}}) \tanh(\mathbf{v})\| + \|\mathbf{b}_\beta(\widetilde{\mathbf{u}}) \tanh(\mathbf{v}) - \mathbf{b}_{\widetilde{\beta}}(\widetilde{\mathbf{u}}) \tanh(\widetilde{\mathbf{v}})\| \\ &\leq \frac{2\beta^2}{\sqrt{n}} \|\mathbf{u} - \widetilde{\mathbf{u}}\| \cdot \|\tanh(\mathbf{v})\| + \left(\frac{1}{n} \sum_{i=1}^n \tanh'(\widetilde{u}_i) \right) \|\beta^2 \tanh(\mathbf{v}) - \widetilde{\beta}^2 \tanh(\widetilde{\mathbf{v}})\| \\ &\leq 2\beta^2 \|\mathbf{u} - \widetilde{\mathbf{u}}\| + \|\beta^2 \tanh(\mathbf{v}) - \widetilde{\beta}^2 \tanh(\widetilde{\mathbf{v}})\|. \end{aligned}$$

Using this bound implies that the function \widetilde{G} of Eq. (5.12) satisfies the Lipschitz assumption of Proposition 5.2.

Consider next Eq. (5.13). Since this function is linear in its arguments, with coefficients independent of n , it follows that it satisfies Lipschitz assumption of Proposition 5.2. This completes the proof. \square

5.2 Hardness for stable algorithms: Proof of Theorems 2.5 and 2.6

Before proving Theorem 2.5 and Theorem 2.6 we recall a known result about disorder chaos, already stated in Eq. (2.11). Draw $\mathbf{x}^0 \sim \mu_{\mathbf{A},\beta}$ independently of $\mathbf{x}^s \sim \mu_{\mathbf{A}_s,\beta}$, and denote by $\mu_{\mathbf{A},\beta}^{(0,s)} := \mu_{\mathbf{A},\beta} \otimes \mu_{\mathbf{A}_s,\beta}$ their joint distribution. Then [Cha14, Theorem 1.11] implies that, for all $\beta \in (0, \infty)$,

$$\lim_{s \rightarrow 0} \lim_{n \rightarrow \infty} \mathbb{E} \mu_{\mathbf{A},\beta}^{(0,s)} \left\{ \left(\frac{1}{n} \langle \mathbf{x}^0, \mathbf{x}^s \rangle \right)^2 \right\} = 0. \quad (5.14)$$

The following simple estimate will be used in our proof.

Lemma 5.3. *Recall that $\mathcal{P}(\{-1, +1\}^n)$ denotes the space of probability distributions over $\{-1, +1\}^n$, and let the function $f : \mathcal{P}(\{-1, +1\}^n)^2 \rightarrow \mathbb{R}$ be defined as*

$$f(\mu, \mu') = \mathbb{E}_{(\mathbf{x}, \mathbf{x}') \sim \mu \otimes \mu'} \left\{ \frac{1}{n} |\langle \mathbf{x}, \mathbf{x}' \rangle| \right\}. \quad (5.15)$$

Then, for all $\mu_1, \mu_2, \nu_1, \nu_2 \in \mathcal{P}(\{-1, +1\}^n)$, we have

$$|f(\mu_1, \nu_1) - f(\mu_2, \nu_2)| \leq W_{2,n}(\mu_1, \mu_2) + W_{2,n}(\nu_1, \nu_2).$$

Proof. Let the vector pairs $(\mathbf{x}^{\mu_1}, \mathbf{x}^{\mu_2})$ and $(\mathbf{x}^{\nu_1}, \mathbf{x}^{\nu_2})$ be independently drawn from the optimal $W_{2,n}$ -couplings of the pairs (μ_1, μ_2) and (ν_1, ν_2) , respectively. Then we have:

$$\begin{aligned} \left| \mathbb{E} \{ |\langle \mathbf{x}^{\mu_1}, \mathbf{x}^{\nu_1} \rangle| \} - \mathbb{E} \{ |\langle \mathbf{x}^{\mu_2}, \mathbf{x}^{\nu_2} \rangle| \} \right| &\leq \left| \mathbb{E} \{ |\langle \mathbf{x}^{\mu_1}, \mathbf{x}^{\nu_1} \rangle| - |\langle \mathbf{x}^{\mu_2}, \mathbf{x}^{\nu_1} \rangle| \} \right| + \left| \mathbb{E} \{ |\langle \mathbf{x}^{\mu_2}, \mathbf{x}^{\nu_1} \rangle| - |\langle \mathbf{x}^{\mu_2}, \mathbf{x}^{\nu_2} \rangle| \} \right| \\ &\leq \sqrt{n} \left(\mathbb{E} \|\mathbf{x}^{\mu_1} - \mathbf{x}^{\mu_2}\| + \mathbb{E} \|\mathbf{x}^{\nu_1} - \mathbf{x}^{\nu_2}\| \right) \\ &\leq \sqrt{n} \left(\mathbb{E} \left[\|\mathbf{x}^{\mu_1} - \mathbf{x}^{\mu_2}\|^2 \right]^{1/2} + \mathbb{E} \left[\|\mathbf{x}^{\nu_1} - \mathbf{x}^{\nu_2}\|^2 \right]^{1/2} \right), \end{aligned}$$

where the second inequality follows from the fact that $\mathbf{x} \mapsto |\langle \mathbf{v}, \mathbf{x} \rangle|$ is Lipschitz continuous with Lipschitz constant $\|\mathbf{v}\|_2$. \square

We are now in position to prove Theorem 2.5.

Proof of Theorem 2.5. Using the notations of the last lemma Eq. (5.14) implies that for all $s \in (0, 1]$,

$$\lim_{n \rightarrow \infty} \mathbb{E} f(\mu_{\mathbf{A}_s,\beta}, \mu_{\mathbf{A}_0,\beta}) = 0. \quad (5.16)$$

Therefore, Theorem 2.5 follows from Lemma 5.3 if we can show that $f(\mu_{\mathbf{A}_0,\beta}, \mu_{\mathbf{A}_0,\beta})$ remains bounded away from zero. This is in turn a well-known consequence of the Parisi formula, as we recall below.

Define the free energy density of the SK model as

$$F_n(\beta) = \frac{1}{n} \mathbb{E} \log \left\{ \sum_{\mathbf{x} \in \{-1, +1\}^n} e^{\beta \langle \mathbf{x}, \mathbf{A} \mathbf{x} \rangle / 2} \right\}. \quad (5.17)$$

The free energy F_n is convex in β and one obtains by Gaussian integration parts that

$$\frac{d}{d\beta} F_n(\beta) = \frac{\beta}{2} \left(1 - \mathbb{E} \mu_{\mathbf{A}_0,\beta} \otimes \mu_{\mathbf{A}_0,\beta} \left\{ \left(\frac{1}{n} \langle \mathbf{x}_1, \mathbf{x}_2 \rangle \right)^2 \right\} \right). \quad (5.18)$$

Moreover, the limit of $F_n(\beta)$ for large n is known to exist for all $\beta > 0$ and its value is given by the Parisi formula [Tal06b]:

$$\lim_{n \rightarrow \infty} F_n(\beta) = \inf_{\zeta \in \mathcal{P}([0,1])} \mathbf{P}_\beta(\zeta), \quad (5.19)$$

where $\mathcal{P}([0,1])$ denotes the set of Borel probability measures supported on $[0,1]$, and \mathbf{P}_β is the Parisi functional at inverse temperature β ; see for instance [Tal06b] or [Pan13, Chapter 3] for definitions.

The following properties are known:

1. A unique minimizer $\zeta_\beta^* \in \mathcal{P}([0,1])$ of \mathbf{P}_β exists for all β [AC15].
2. If $\beta > 1$, then ζ_β^* is not an atom on 0: $\zeta_\beta^* \neq \delta_0$. This follows from Toninelli's theorem [Ton02] that $\limsup_{n \rightarrow \infty} F_n(\beta) \leq \log 2 + \beta^2/4 - \varepsilon(\beta)$ for some continuous $\varepsilon(\beta)$, with $\varepsilon(\beta) > 0$ when $\beta > 1$.
3. The function $\beta \mapsto \mathbf{P}_\beta(\zeta_\beta^*)$ is convex and differentiable at all $\beta > 0$, and

$$\frac{d}{d\beta} \mathbf{P}_\beta(\zeta_\beta^*) = \frac{\beta}{2} \left(1 - \int q^2 \zeta_\beta^*(dq) \right). \quad (5.20)$$

See for instance [Pan13, Theorem 3.7] or [Tal06a, Theorem 1.2] for a proof.

The convexity of F_n implies that for almost all $\beta > 0$, $\lim_{n \rightarrow \infty} F'_n(\beta) = \frac{d}{d\beta} \mathbf{P}_\beta(\zeta_\beta^*)$. Using Eq. (5.18) and Eq. (5.20) we obtain

$$\lim_{n \rightarrow \infty} \frac{\beta}{2} \left(1 - \mathbb{E} \mu_{\mathbf{A}_0, \beta} \otimes \mu_{\mathbf{A}_0, \beta} \left\{ \left(\frac{1}{n} \langle \mathbf{x}_1, \mathbf{x}_2 \rangle \right)^2 \right\} \right) = \frac{\beta}{2} \left(1 - \int q^2 \zeta_\beta^*(dq) \right) < \frac{\beta}{2} - \varepsilon(\beta), \quad (5.21)$$

where the last inequality holds for almost all $\beta > 1$ by Property 2 above. Since the both sides are non-decreasing and the right hand side is continuous, the inequality holds for all β . This is equivalent to

$$\lim_{n \rightarrow \infty} \mathbb{E} f(\mu_{\mathbf{A}_0, \beta}, \mu_{\mathbf{A}_0, \beta}) > 0. \quad (5.22)$$

Now, using Eq. (5.16) and Eq. (5.22), together with the continuity of f (Lemma 5.3) implies the claim of the theorem. \square

We next prove that Theorem 2.6 is an immediate consequence of Theorem 2.5.

Proof of Theorem 2.6. Fix $s \in (0, 1)$ and $\mu_{\mathbf{A}_s, \beta}^{\text{alg}}$ be the law of $\text{ALG}_n(\mathbf{A}_s, \beta, \omega)$ conditional on \mathbf{A}_s . By the triangle inequality,

$$W_{2,n}(\mu_{\mathbf{A}_s, \beta, s}, \mu_{\mathbf{A}_0, \beta}) \leq W_{2,n}(\mu_{\mathbf{A}_s, \beta}, \mu_{\mathbf{A}_s, \beta}^{\text{alg}}) + W_{2,n}(\mu_{\mathbf{A}_s, \beta}^{\text{alg}}, \mu_{\mathbf{A}_0, \beta}^{\text{alg}}) + W_{2,n}(\mu_{\mathbf{A}_0, \beta}^{\text{alg}}, \mu_{\mathbf{A}_0, \beta, 0}).$$

Taking expectations over \mathbf{A} and \mathbf{A}_s , we have $\mathbb{E} [W_{2,n}(\mu_{\mathbf{A}_s, \beta}, \mu_{\mathbf{A}_s, \beta}^{\text{alg}})] = \mathbb{E} [W_{2,n}(\mu_{\mathbf{A}_0, \beta}^{\text{alg}}, \mu_{\mathbf{A}_0, \beta})]$. Further, by stability of the algorithm, $\mathbb{E} [W_{2,n}(\mu_{\mathbf{A}_s, \beta}^{\text{alg}}, \mu_{\mathbf{A}_0, \beta}^{\text{alg}})] \rightarrow 0$ when $n \rightarrow \infty$ followed by $s \rightarrow 0$. Therefore, using Theorem 2.5 and choosing s sufficiently small, we obtain

$$\liminf_{n \rightarrow \infty} \mathbb{E} [W_{2,n}(\mu_{\mathbf{A}_0, \beta}^{\text{alg}}, \mu_{\mathbf{A}_0, \beta})] \geq W_* > 0.$$

\square

References

- [AC15] Antonio Auffinger and Wei-Kuo Chen, *The Parisi formula has a unique minimizer*, Communications in Mathematical Physics **335** (2015), no. 3, 1429–1444.
- [AH87] Michael Aizenman and Richard Holley, *Rapid convergence to equilibrium of stochastic Ising models in the Dobrushin Shlosman regime*, Percolation theory and ergodic theory of infinite particle systems, Springer, 1987, pp. 1–11.
- [AJK⁺21] Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong, *Entropic Independence I: Modified Log-Sobolev Inequalities for Fractionally Log-Concave Distributions and High-Temperature Ising Models*, arXiv preprint arXiv:2106.04105 (2021).
- [ALR87] Michael Aizenman, Joel L Lebowitz, and David Ruelle, *Some rigorous results on the Sherrington–Kirkpatrick spin glass model*, Communications in Mathematical Physics **112** (1987), no. 1, 3–20.
- [BAJ18] Gérard Ben Arous and Aukosh Jagannath, *Spectral gap estimates in mean field spin glasses*, Communications in Mathematical Physics **361** (2018), no. 1, 1–52.
- [BB19] Roland Bauerschmidt and Thierry Bodineau, *A very simple proof of the LSI for high temperature spin systems*, Journal of Functional Analysis **276** (2019), no. 8, 2582–2588.
- [BD11] Joseph Blitzstein and Persi Diaconis, *A sequential importance sampling algorithm for generating random graphs with prescribed degrees*, Internet Mathematics **6** (2011), no. 4, 489–522.
- [BH21] Guy Bresler and Brice Huang, *The Algorithmic Phase Transition of Random k -SAT for Low Degree Polynomials*, 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS), 2021, pp. 298–309.
- [BM11] Mohsen Bayati and Andrea Montanari, *The dynamics of message passing on dense graphs, with applications to compressed sensing*, IEEE Transactions on Information Theory **57** (2011), no. 2, 764–785.
- [CDHL05] Yuguo Chen, Persi Diaconis, Susan P Holmes, and Jun S Liu, *Sequential Monte Carlo methods for statistical analysis of tables*, Journal of the American Statistical Association **100** (2005), no. 469, 109–120.
- [CE22] Yuansi Chen and Ronen Eldan, *Localization schemes: A framework for proving mixing bounds for Markov chains*, arXiv preprint arXiv:2203.04163 (2022).
- [CFM21] Michael Celentano, Zhou Fan, and Song Mei, *Local convexity of the TAP free energy and AMP convergence for \mathbf{Z}_2 -synchronization*, arXiv preprint arXiv:2106.11428 (2021).
- [CGPR19] Wei-Kuo Chen, David Gamarnik, Dmitry Panchenko, and Mustazee Rahman, *Suboptimality of local algorithms for a class of max-cut problems*, The Annals of Probability **47** (2019), no. 3, 1587–1618.
- [Cha09] Sourav Chatterjee, *Disorder chaos and multiple valleys in spin glasses*, arXiv preprint arXiv:0907.3381 (2009).
- [Cha14] ———, *Superconcentration and related topics*, vol. 15, Springer, 2014.

- [Che13] Wei-Kuo Chen, *Disorder chaos in the Sherrington–Kirkpatrick model with external field*, The Annals of Probability **41** (2013), no. 5, 3345–3391.
- [CHHS15] Wei-Kuo Chen, Hsi-Wei Hsieh, Chii-Ruey Hwang, and Yuan-Chung Sheu, *Disorder chaos in the spherical mean-field model*, Journal of Statistical Physics **160** (2015), no. 2, 417–429.
- [CP18] Wei-Kuo Chen and Dmitry Panchenko, *Disorder chaos in some diluted spin glass models*, The Annals of Applied Probability **28** (2018), no. 3, 1356–1378.
- [CT21] Wei-Kuo Chen and Si Tang, *On Convergence of the Cavity and Bolthausen’s TAP Iterations to the Local Magnetization*, Communications in Mathematical Physics **386** (2021), no. 2, 1209–1242.
- [DAM17] Yash Deshpande, Emmanuel Abbe, and Andrea Montanari, *Asymptotic mutual information for the balanced binary stochastic block model*, Information and Inference: A Journal of the IMA **6** (2017), no. 2, 125–170.
- [Dur19] Rick Durrett, *Probability: Theory and Examples*, vol. 49, Cambridge university press, 2019.
- [EKZ21] Ronen Eldan, Frederic Koehler, and Ofer Zeitouni, *A spectral condition for spectral gap: fast mixing in high-temperature Ising models*, Probability Theory and Related Fields (2021), 1–17.
- [Eld20] Ronen Eldan, *Taming correlations through entropy-efficient measure decompositions with applications to mean-field approximation*, Probability Theory and Related Fields **176** (2020), no. 3, 737–755.
- [ES22] Ronen Eldan and Omer Shamir, *Log concavity and concentration of Lipschitz functions on the Boolean hypercube*, Journal of Functional Analysis (2022), 109392.
- [GJ19] Reza Gheissari and Aukosh Jagannath, *On the spectral gap of spherical spin glass dynamics*, Annales de l’Institut Henri Poincaré, Probabilités et Statistiques, vol. 55, Institut Henri Poincaré, 2019, pp. 756–776.
- [GJ21] David Gamarnik and Aukosh Jagannath, *The overlap gap property and approximate message passing algorithms for p -spin models*, The Annals of Probability **49** (2021), no. 1, 180–205.
- [GJW20] David Gamarnik, Aukosh Jagannath, and Alexander S. Wein, *Low-degree hardness of random optimization problems*, Proceedings of 61st FOCS, IEEE, 2020, pp. 131–140.
- [GJW21] ———, *Circuit lower bounds for the p -spin optimization problem*, arXiv preprint arXiv:2109.01342 (2021).
- [GK21] David Gamarnik and Eren C. Kızıldağ, *Algorithmic obstructions in the random number partitioning problem*, arXiv preprint arXiv:2103.01369 (2021).
- [GS14] David Gamarnik and Madhu Sudan, *Limits of local algorithms over sparse random graphs*, Proceedings of the 5th conference on Innovations in theoretical computer science, ACM, 2014, pp. 369–376.
- [GS17] ———, *Performance of sequential local algorithms for the random NAE- K -sat problem*, SIAM Journal on Computing **46** (2017), no. 2, 590–619.

- [GSV05] Dongning Guo, Shlomo Shamai, and Sergio Verdú, *Mutual information and minimum mean-square error in gaussian channels*, IEEE Trans. Inform. Theory **51** (2005), 1261–1282.
- [HS21] Brice Huang and Mark Sellke, *Tight Lipschitz Hardness for Optimizing Mean Field Spin Glasses*, arXiv preprint arXiv:2110.07847 (2021).
- [JM13] Adel Javanmard and Andrea Montanari, *State evolution for general approximate message passing algorithms, with applications to spatial coupling*, Information and Inference: A Journal of the IMA **2** (2013), no. 2, 115–144.
- [JVV86] Mark R Jerrum, Leslie G Valiant, and Vijay V Vazirani, *Random generation of combinatorial structures from a uniform distribution*, Theoretical computer science **43** (1986), 169–188.
- [LFN18] Haihao Lu, Robert M Freund, and Yurii Nesterov, *Relatively smooth convex optimization by first-order methods, and applications*, SIAM Journal on Optimization **28** (2018), no. 1, 333–354.
- [LS77] Robert Shevilevich Liptser and Al’bert Nikolaevich Shiriaev, *Statistics of random processes: General theory*, vol. 394, Springer, 1977.
- [MPV87] Marc Mézard, Giorgio Parisi, and Miguel A. Virasoro, *Spin glass theory and beyond*, World Scientific, 1987.
- [MV21] Andrea Montanari and Ramji Venkataramanan, *Estimation of low-rank matrices via approximate message passing*, The Annals of Statistics **49** (2021), no. 1, 321–345.
- [NSZ22] Danny Nam, Allan Sly, and Lingfu Zhang, *Ising model on trees and factors of IID*, Communications in Mathematical Physics (2022), 1–38.
- [Pan13] Dmitry Panchenko, *The Sherrington-Kirkpatrick model*, Springer Science & Business Media, 2013.
- [RV17] Mustazee Rahman and Bálint Virág, *Local algorithms for independent sets are half-optimal*, The Annals of Probability **45** (2017), no. 3, 1543–1577.
- [SZ81] Haim Sompolinsky and Annette Zippelius, *Dynamic theory of the spin-glass phase*, Physical Review Letters **47** (1981), no. 5, 359.
- [Tal06a] Michel Talagrand, *Parisi measures*, Journal of Functional Analysis **231** (2006), no. 2, 269–286.
- [Tal06b] ———, *The Parisi formula*, Annals of Mathematics (2006), 221–263.
- [Tal10] Michel Talagrand, *Mean field models for spin glasses: Volume i*, Springer-Verlag, Berlin, 2010.
- [Ton02] Fabio Lucio Toninelli, *About the Almeida-Thouless transition line in the Sherrington-Kirkpatrick mean-field spin glass model*, EPL (Europhysics Letters) **60** (2002), no. 5, 764.
- [VdV98] Aad W Van der Vaart, *Asymptotic statistics*, vol. 3, Cambridge university press, 1998.
- [Wei22] Alexander S Wein, *Optimal low-degree hardness of maximum independent set*, Mathematical Statistics and Learning (2022).

A Convergence analysis of Natural Gradient Descent

The main objective of this appendix is to prove Lemma 4.10, which we will do in Section A.2, after some technical preparations in Section A.1.

A.1 Technical preliminaries

Definition A.1. Let $Q \subseteq (-1, 1)^n$ be a convex set. We say that a twice differentiable function $\mathcal{F} : Q \rightarrow \mathbb{R}$ is relatively c -strongly convex if it satisfies

$$\nabla^2 \mathcal{F}(\mathbf{m}) \succeq c\mathbf{D}(\mathbf{m}) \quad \forall \mathbf{m} \in Q. \quad (\text{A.1})$$

We say it is relatively C -smooth if it satisfies

$$\nabla^2 \mathcal{F}(\mathbf{m}) \preceq C\mathbf{D}(\mathbf{m}) \quad \forall \mathbf{m} \in Q. \quad (\text{A.2})$$

As $\mathbf{D}(\mathbf{m}) = \nabla^2(-h(\mathbf{m})) \succeq \mathbf{I}_n$, it follows that (A.1) implies ordinary c -strong convexity in Euclidean norm. The next proposition connects relative strong convexity with the Bregman divergence introduced in Eq. 4.59.

Proposition A.2 (Proposition 1.1 in [LFN18]). A twice differentiable function $\mathcal{F} : Q \rightarrow \mathbb{R}$ is relatively c -strongly convex if and only if

$$\mathcal{F}(\mathbf{m}) \geq \mathcal{F}(\mathbf{n}) + \langle \nabla \mathcal{F}(\mathbf{n}), \mathbf{m} - \mathbf{n} \rangle + cD_{-h}(\mathbf{m}, \mathbf{n}), \quad \forall \mathbf{m}, \mathbf{n} \in Q. \quad (\text{A.3})$$

Lemma A.3. For $\mathbf{m}, \mathbf{n} \in (-1, 1)^n$,

$$D_{-h}(\mathbf{m}, \mathbf{n}) \geq \frac{\|\mathbf{m} - \mathbf{n}\|_2^2}{2}, \quad (\text{A.4})$$

$$D_{-h}(\mathbf{m}, \mathbf{n}) \leq 10n \left(1 + \frac{\|\text{arctanh}(\mathbf{n})\|_2}{\sqrt{n}} \right), \quad (\text{A.5})$$

$$D_{-h}(\mathbf{m}, \mathbf{n}) \leq \|\text{arctanh}(\mathbf{m}) - \text{arctanh}(\mathbf{n})\|_2^2. \quad (\text{A.6})$$

Proof. Observe that $h''(x) = -1/(1-x^2) \leq -1$ for all $x \in (-1, 1)$ with equality if and only if $x = 0$. Therefore

$$\begin{aligned} D_{-h}(\mathbf{m}, \mathbf{n}) &= \sum_{i=1}^n \int_{m_i}^{n_i} (x - m_i)(-h''(x)) \, dx \\ &= \sum_{i=1}^n \frac{(n_i - m_i)^2}{2}. \end{aligned}$$

This proves Eq. (A.4).

Next, Eq. (A.5) follows from Eq. (4.59) and the fact that the binary entropy $h : \mathbb{R} \rightarrow \mathbb{R}$ is uniformly bounded.

Finally Eq. (A.6) follows from

$$\begin{aligned} D_{-h}(\mathbf{m}, \mathbf{n}) &\leq \langle \nabla h(\mathbf{n}) - \nabla h(\mathbf{m}), \mathbf{m} - \mathbf{n} \rangle \\ &= \langle \text{arctanh}(\mathbf{m}) - \text{arctanh}(\mathbf{n}), \mathbf{m} - \mathbf{n} \rangle \\ &\leq \|\text{arctanh}(\mathbf{m}) - \text{arctanh}(\mathbf{n})\|_2^2. \end{aligned}$$

Here in the last step we used that $\tanh(\cdot)$ is 1-Lipschitz. □

Lemma A.4. *If $\mathcal{F} : Q \rightarrow \mathbb{R}$ is relatively c -strongly convex for some convex set $Q \subseteq (-1, 1)^n$, and $\nabla \mathcal{F}(\mathbf{m}_*) = 0$ for $\mathbf{m}_* \in Q$, it follows that*

$$\mathcal{F}(\mathbf{m}) - \mathcal{F}(\mathbf{m}_*) \geq \frac{c \|\mathbf{m} - \mathbf{m}_*\|_2^2}{2}.$$

for all $\mathbf{m} \in Q$.

Proof. Using (A.3) and (A.4), and observing that $\nabla \mathcal{F}(\mathbf{m}_*) = 0$, we obtain

$$\frac{\mathcal{F}(\mathbf{m}) - \mathcal{F}(\mathbf{m}_*)}{\|\mathbf{m} - \mathbf{m}_*\|_2^2} \geq \frac{\mathcal{F}(\mathbf{m}) - \mathcal{F}(\mathbf{m}_*)}{2 \cdot D_{-h}(\mathbf{m}, \mathbf{m}_*)} \geq \frac{c}{2}.$$

□

Lemma A.5. *Suppose $\mathcal{F} : Q_* \rightarrow \mathbb{R}$ is c -strongly convex in the convex set $Q_* := B(\rho) \cap (-1, 1)^n$. If $\mathbf{x}_* \in \partial Q_*$, $x_{*,k} = +1$ (respectively, $x_{*,k} = -1$) and $|x_j| < 1$ for all $j \in [n] \setminus \{k\}$, then $\lim_{t \rightarrow 0^+} \partial_{x_k} \mathcal{F}(\mathbf{x}_* - t\mathbf{e}_k) = +\infty$ (respectively $\lim_{t \rightarrow 0^+} \partial_{x_k} \mathcal{F}(\mathbf{x}_* + t\mathbf{e}_k) = -\infty$).*

Proof. Consider the case $x_k = +1$ (as the case $x_k = -1$ follows by symmetry.) Then there exists $t_0 > 0$ such that $\mathbf{x}_* - t\mathbf{e}_k \in Q_*$ for all $t \in (0, t_0]$. Let $\mathbf{x}(s) := \mathbf{x}_* - (t_0 - s)\mathbf{e}_k$, $s \in [0, t_0]$. Then

$$\begin{aligned} \partial_{x_k} F(\mathbf{x}(s)) &= \partial_{x_k} F(\mathbf{x}(0)) + \int_0^s \partial_{x_k}^2 F(\mathbf{x}(u)) \, du \\ &= \partial_{x_k} F(\mathbf{x}(0)) + \int_0^s \langle \mathbf{e}_k, \nabla^2 F(\mathbf{x}(u)) \mathbf{e}_k \rangle \, du \\ &\geq \partial_{x_k} F(\mathbf{x}(0)) + c \int_0^s (1 - x_k(u)^2)^{-1} \, du \\ &\geq \partial_{x_k} F(\mathbf{x}(0)) + c \int_0^s (1 - (1 - t_0 + u)^2)^{-1} \, du, . \end{aligned}$$

The last integral diverges as $s \uparrow t_0$, thus proving the claim. □

Lemma A.6. *Suppose $\mathcal{F} : Q \rightarrow \mathbb{R}$ is c -strongly convex for a convex set $Q \subseteq (-1, 1)^n$. Moreover suppose that*

$$\|\nabla \mathcal{F}(\mathbf{m})\| \leq c\sqrt{\varepsilon n}$$

for some $\mathbf{m} \in Q$ with

$$B(\mathbf{m}, 2\sqrt{\varepsilon n}) \cap (-1, 1)^n \subseteq Q.$$

Then there exists a unique $\mathbf{m}_* \in B(\mathbf{m}, 2\sqrt{\varepsilon n}) \cap (-1, 1)^n$ satisfying $\nabla \mathcal{F}(\mathbf{m}_*) = 0$, which is in fact a global minimizer of \mathcal{F} on Q . Moreover

$$\mathcal{F}(\mathbf{m}) - \mathcal{F}(\mathbf{m}_*) \leq 2c\varepsilon n. \tag{A.7}$$

Proof. Let $Q_{\leq} := \{\mathbf{x} \in Q : \mathcal{F}(\mathbf{x}) \leq \mathcal{F}(\mathbf{m})\}$. Then, for any $\mathbf{x} \in Q_0$, we have

$$\begin{aligned} 0 &\geq \mathcal{F}(\mathbf{x}) - \mathcal{F}(\mathbf{m}) \\ &\geq -c\sqrt{\varepsilon n} \|\mathbf{x} - \mathbf{m}\|_2 + cD_{-h}(\mathbf{x}; \mathbf{m}) \\ &\geq -c\sqrt{\varepsilon n} \|\mathbf{x} - \mathbf{m}\|_2 + \frac{c}{2} \|\mathbf{x} - \mathbf{m}\|_2^2. \end{aligned}$$

Hence $Q_{\leq} \subseteq Q_* := B(\mathbf{m}, \sqrt{\varepsilon n}) \cap (-1, 1)^n$, $Q_* \subseteq Q$. By continuity three cases are possible: (i) The minimum of \mathcal{F} is achieved in the interior of Q_{\leq} ; (ii) The minimum is achieved along a sequence $(\mathbf{x}_i)_{i \geq 0}$, $\|\mathbf{x}_i\|_{\infty} \rightarrow 1$; (iii) the minimum is achieved at $\mathbf{m}_* \neq \mathbf{m}$ such that $\mathcal{F}(\mathbf{m}_*) = \mathcal{F}(\mathbf{m})$. Case (iii) cannot hold by strong convexity, and case (ii) cannot hold by Lemma A.5.

Uniqueness of \mathbf{m}_* follows by strong convexity, and $\nabla \mathcal{F}(\mathbf{m}_*) = 0$ by differentiability. Finally

$$\mathcal{F}(\mathbf{m}) - \mathcal{F}(\mathbf{m}_*) \leq \|\nabla \mathcal{F}(\mathbf{m})\| \cdot \|\mathbf{m} - \mathbf{m}_*\| \leq 2c\varepsilon n.$$

□

Lemma A.7. *Suppose $\mathcal{F} : Q \rightarrow \mathbb{R}$ is relatively c -strongly convex. Let \mathbf{m}_* be a local minimum of \mathcal{F} belonging to the interior of Q , and suppose that $B(\mathbf{m}_*, 2\sqrt{\varepsilon n}) \cap (-1, 1)^n \subseteq Q$. Consider for $\mathbf{y} \in \mathbb{R}^n$ the function*

$$\mathcal{F}_{\mathbf{y}}(\mathbf{m}) = \mathcal{F}(\mathbf{m}) - \langle \mathbf{y}, \mathbf{m} \rangle.$$

Then $\mathcal{F}_{\mathbf{y}}$ is relatively c -strongly convex on Q for any $\mathbf{y} \in \mathbb{R}^n$. If $\|\mathbf{y}\| \leq (c/2)\sqrt{\varepsilon n}$, then $\mathcal{F}_{\mathbf{y}}$ has a unique stationary point and minimizer $\mathbf{m}_(\mathbf{y}) \in Q$. Moreover if $\|\mathbf{y}\|, \|\hat{\mathbf{y}}\| \leq \frac{c\sqrt{\varepsilon n}}{2}$ then*

$$\|\mathbf{m}_*(\mathbf{y}) - \mathbf{m}_*(\hat{\mathbf{y}})\| \leq \frac{\|\mathbf{y} - \hat{\mathbf{y}}\|}{c}. \quad (\text{A.8})$$

Proof. The relative c -strong convexity of $\mathcal{F}_{\mathbf{y}}$ is clear as the Hessian of $\mathcal{F}_{\mathbf{y}}$ does not depend on \mathbf{y} . For $\|\mathbf{y}\| \leq (c/2)\sqrt{\varepsilon n}$, because

$$\|\nabla \mathcal{F}_{\mathbf{y}}(\mathbf{m}_*)\| = \|\mathbf{y}\| \leq \frac{c\sqrt{\varepsilon n}}{2} \quad \text{and} \quad B(\mathbf{m}_*, \sqrt{\varepsilon n}) \cap (-1, 1)^n \subseteq Q,$$

Lemma A.6 implies the existence of a unique minimizer

$$\mathbf{m}_*(\mathbf{y}) \in B(\mathbf{m}_*, \sqrt{\varepsilon n}) \cap (-1, 1)^n \subseteq Q.$$

If $\|\hat{\mathbf{y}}\| \leq (c/2)\sqrt{\varepsilon n}$ also holds, $\mathcal{F}_{\hat{\mathbf{y}}}$ is c -strongly convex on

$$B(\mathbf{m}_*(\hat{\mathbf{y}}), \sqrt{\varepsilon n}) \cap (-1, 1)^n \subseteq B(\mathbf{m}_*, 2\sqrt{\varepsilon n}) \cap (-1, 1)^n \subseteq Q.$$

Moreover since $\|\mathbf{y} - \hat{\mathbf{y}}\| \leq c\sqrt{\varepsilon n}$, we obtain

$$\|\nabla \mathcal{F}_{\hat{\mathbf{y}}}(\mathbf{m}_*(\mathbf{y}))\| = \|\mathbf{y} - \hat{\mathbf{y}}\| = c\sqrt{\varepsilon' n},$$

for $\varepsilon' = \frac{\|\mathbf{y} - \hat{\mathbf{y}}\|^2}{c^2 n} \leq \varepsilon$. Therefore the conditions of Lemma A.6 are satisfied with $(\mathcal{F}_{\hat{\mathbf{y}}}, \mathbf{m}_*(\mathbf{y}), \varepsilon')$ in place of $(\mathcal{F}, \mathbf{m}, \varepsilon)$. Equation (A.8) now follows since

$$\|\mathbf{m}_*(\mathbf{y}) - \mathbf{m}_*(\hat{\mathbf{y}})\| \leq \sqrt{\varepsilon' n} = \frac{\|\mathbf{y} - \hat{\mathbf{y}}\|}{c}.$$

□

We now analyze the convergence of Algorithm 3 from a good initialization.

Lemma A.8. Suppose $\mathcal{F}(\cdot) = \mathcal{F}_{\text{TAP}}(\cdot; \mathbf{y}, q_K(\beta, t))$ has a local minimum at \mathbf{m}_* and is relatively c -strongly-convex on $B(\mathbf{m}_*, \sqrt{\varepsilon n}) \cap (-1, 1)^n$, and also C -relatively smooth on $(-1, 1)^n$. Suppose

$$\widehat{\mathbf{m}}^0 \in B(\mathbf{m}_*, \sqrt{\varepsilon n}) \cap (-1, 1)^n \quad (\text{A.9})$$

satisfies

$$\mathcal{F}(\widehat{\mathbf{m}}^0) < \mathcal{F}(\mathbf{m}_*) + \frac{c\varepsilon n}{8}. \quad (\text{A.10})$$

Then there exist constants $\eta_0, C' > 0$ depending only on (C, c, ε) such that the following holds. If Algorithm 3 is initialized at $\widehat{\mathbf{m}}^0$ with learning rate $\eta = 1/L \in (0, \eta_0)$, then, for every $K \geq 1$

$$\mathcal{F}(\widehat{\mathbf{m}}^K) \leq \mathcal{F}(\mathbf{m}_*) + C'n \left(1 + \frac{\|\text{arctanh}(\widehat{\mathbf{m}}^0)\|_2}{\sqrt{n}} \right) (1 - c\eta)^K, \quad (\text{A.11})$$

$$\|\widehat{\mathbf{m}}^K - \mathbf{m}_*\|_2 \leq C'\sqrt{n} \left(1 + \frac{\|\text{arctanh}(\widehat{\mathbf{m}}^0)\|_2}{\sqrt{n}} \right) (1 - c\eta)^{K/2}. \quad (\text{A.12})$$

Proof. Recall Eq. (4.60), which we copy here for the reader's convenience:

$$\widehat{\mathbf{m}}^{i+1} = \arg \min_{\mathbf{x} \in (-1, 1)^n} \langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \mathbf{x} - \widehat{\mathbf{m}}^i \rangle + L \cdot D_{-h}(\mathbf{x}, \widehat{\mathbf{m}}^i). \quad (\text{A.13})$$

If $\eta_0 \leq \frac{1}{2C}$ then [LFN18, Lemma 3.1] applied to the linear (hence convex) function $\langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \cdot \rangle$ states that for all $\mathbf{m} \in (-1, 1)^n$,

$$\langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \widehat{\mathbf{m}}^{i+1} \rangle + LD_{-h}(\widehat{\mathbf{m}}^{i+1}, \widehat{\mathbf{m}}^i) + LD_{-h}(\mathbf{m}, \widehat{\mathbf{m}}^{i+1}) \leq \langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \mathbf{m} \rangle + LD_{-h}(\mathbf{m}, \widehat{\mathbf{m}}^i). \quad (\text{A.14})$$

Moreover the global relative smoothness shown in (4.57) implies that for $\mathbf{m}, \mathbf{m}' \in (-1, 1)^n$,

$$\mathcal{F}(\mathbf{m}) \leq \mathcal{F}(\mathbf{m}') + \langle \nabla \mathcal{F}(\mathbf{m}'), \mathbf{m} - \mathbf{m}' \rangle + C \cdot D_{-h}(\mathbf{m}, \mathbf{m}'). \quad (\text{A.15})$$

Combining Eqs. (A.14) and (A.15) yields

$$\begin{aligned} \mathcal{F}(\widehat{\mathbf{m}}^{i+1}) &\leq \mathcal{F}(\widehat{\mathbf{m}}^i) + \langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i \rangle + LD_{-h}(\widehat{\mathbf{m}}^{i+1}, \widehat{\mathbf{m}}^i) \\ &\leq \mathcal{F}(\widehat{\mathbf{m}}^i) + \langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \mathbf{m} - \widehat{\mathbf{m}}^i \rangle + LD_{-h}(\mathbf{m}, \widehat{\mathbf{m}}^i) - LD_{-h}(\mathbf{m}, \widehat{\mathbf{m}}^{i+1}). \end{aligned} \quad (\text{A.16})$$

Setting $\mathbf{m} = \widehat{\mathbf{m}}^i$, we find

$$\mathcal{F}(\widehat{\mathbf{m}}^{i+1}) \leq \mathcal{F}(\widehat{\mathbf{m}}^i), \quad \forall i \in [K].$$

We next prove by induction that for each $i \geq 1$,

$$\mathcal{F}(\widehat{\mathbf{m}}^i) < \mathcal{F}(\mathbf{m}_*) + \frac{c\varepsilon n}{8}, \quad \|\widehat{\mathbf{m}}^i - \mathbf{m}_*\| < \sqrt{\varepsilon n}. \quad (\text{A.17})$$

The base case $i = 0$ holds by assumption. Suppose (A.17) holds for i . It follows that

$$\mathcal{F}(\widehat{\mathbf{m}}^{i+1}) \leq \mathcal{F}(\widehat{\mathbf{m}}^i) \leq \mathcal{F}(\mathbf{m}_*) + \frac{c\varepsilon n}{8}.$$

In fact, local c -strong convexity

$$\nabla^2 \mathcal{F}(\mathbf{m}) \succeq c\mathbf{D}(\mathbf{m}) \succeq c\mathbf{I}_n, \quad \mathbf{m} \in B(\mathbf{m}_*, \sqrt{\varepsilon n}) \cap (-1, 1)^n$$

implies $\widehat{\mathbf{m}}^i$ is even closer to \mathbf{m}_* than required by (A.17):

$$\|\widehat{\mathbf{m}}^i - \mathbf{m}_*\|_2 \leq \sqrt{\frac{\mathcal{F}(\widehat{\mathbf{m}}^i) - \mathcal{F}(\mathbf{m}_*)}{c}} \leq \frac{\sqrt{\varepsilon n}}{2}.$$

Next we bound the movement from a single NGD step. Comparing values of (A.13) at $\widehat{\mathbf{m}}^i$ and the minimizer $\widehat{\mathbf{m}}^{i+1}$ implies

$$\langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i \rangle + LD_{-h}(\widehat{\mathbf{m}}^{i+1}, \widehat{\mathbf{m}}^i) \leq 0. \quad (\text{A.18})$$

From definition of Bregman divergence and the fact that (on the high probability event $\|\mathbf{A}\|_{\text{op}} \leq 3$) $\|\nabla \mathcal{F} + \nabla h\|_2 \leq C\sqrt{n}$ (thanks to the special form of $\mathcal{F}(\cdot) = \mathcal{F}_{\text{TAP}}(\cdot; \mathbf{y}, q_K(\beta, t))$,

$$\begin{aligned} |\langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i \rangle + D_{-h}(\widehat{\mathbf{m}}^{i+1}, \widehat{\mathbf{m}}^i)| &= |\langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i) + \nabla h(\widehat{\mathbf{m}}^i), \widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i \rangle - h(\widehat{\mathbf{m}}^{i+1}) + h(\widehat{\mathbf{m}}^i)| \\ &\leq C_1 n \left(1 + \frac{\|\widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i\|}{\sqrt{n}} \right). \end{aligned}$$

Moreover assuming $L > 1$, (A.4) implies

$$(L-1)D_{-h}(\widehat{\mathbf{m}}^{i+1}, \widehat{\mathbf{m}}^i) \geq \frac{L-1}{2} \|\widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i\|^2.$$

Substituting the previous two displays into (A.18) yields

$$0 \geq \frac{L-1}{2} \|\widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i\|^2 - C_2 \sqrt{n} \|\widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i\|_2 - C_2 n$$

and so

$$\|\widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i\|_2 \leq \frac{C_3 \sqrt{n}}{\sqrt{L-1}}.$$

Taking L large enough, it follows that

$$\|\widehat{\mathbf{m}}^{i+1} - \mathbf{m}_*\| \leq \|\widehat{\mathbf{m}}^{i+1} - \widehat{\mathbf{m}}^i\|_2 + \|\widehat{\mathbf{m}}^i - \mathbf{m}_*\|_2 \leq \sqrt{\varepsilon n}.$$

This completes the inductive proof of Eq. (A.17), which we now use to analyze convergence of Algorithm 3. Indeed from the first part of (A.17), the local relative strong convexity of \mathcal{F} implies

$$\mathcal{F}(\widehat{\mathbf{m}}^i) + \langle \nabla \mathcal{F}(\widehat{\mathbf{m}}^i), \mathbf{m}_* - \widehat{\mathbf{m}}^i \rangle \leq \mathcal{F}(\mathbf{m}_*) - cD_{-h}(\mathbf{m}_*, \widehat{\mathbf{m}}^i), \quad \forall i \in [K].$$

Setting $\mathbf{m} = \mathbf{m}_*$ in (A.16) and combining yields

$$\mathcal{F}(\widehat{\mathbf{m}}^{i+1}) \leq \mathcal{F}(\mathbf{m}_*) + (L-c)D_{-h}(\mathbf{m}_*, \widehat{\mathbf{m}}^i) - LD_{-h}(\mathbf{m}_*, \widehat{\mathbf{m}}^{i+1}).$$

Multiplying by $\left(\frac{L}{L-c}\right)^{i+1}$ and summing over i gives

$$\sum_{i=0}^{K-1} \left(\frac{L}{L-c}\right)^{i+1} \mathcal{F}(\widehat{\mathbf{m}}^{i+1}) \leq \sum_{i=0}^{K-1} \left(\frac{L}{L-c}\right)^{i+1} \mathcal{F}(\mathbf{m}_*) + LD_{-h}(\mathbf{m}_*, \widehat{\mathbf{m}}^0).$$

Since the values $\mathcal{F}(\widehat{\mathbf{m}}^i)$ are decreasing, we find

$$\begin{aligned} \mathcal{F}(\widehat{\mathbf{m}}^K) &\leq \mathcal{F}(\mathbf{m}_*) + L \left(\sum_{i=0}^{K-1} \left(\frac{L}{L-c}\right)^{i+1} \right)^{-1} D_{-h}(\mathbf{m}_*, \widehat{\mathbf{m}}^0) \\ &\leq \mathcal{F}(\mathbf{m}_*) + L(1-c\eta)^K D_{-h}(\mathbf{m}_*, \widehat{\mathbf{m}}^0). \end{aligned}$$

Using Eq. (A.5) together with the last display proves Eq. (A.11).

It was shown above by induction that $\widehat{\mathbf{m}}^K$ is in a c -strongly convex neighborhood of \mathbf{m}_* . Using strong convexity in Euclidean norm yields

$$\|\widehat{\mathbf{m}}^k - \mathbf{m}_*\| \leq \sqrt{\frac{\mathcal{F}(\widehat{\mathbf{m}}^K) - \mathcal{F}(\mathbf{m}_*)}{c}}$$

and so (A.12) follows as well. \square

Lemma A.9. Assume $\|\mathbf{A}\|_{op} \leq 3$. For any $\mathbf{m}, \mathbf{n} \in (-1, 1)^n$, and $\mathbf{y}, \widehat{\mathbf{y}} \in \mathbb{R}^n$, and $q \in [0, 1]$:

$$\|\nabla \mathcal{F}_{\text{TAP}}(\mathbf{m}, \mathbf{y}, q) - \nabla \mathcal{F}_{\text{TAP}}(\mathbf{n}, \widehat{\mathbf{y}}, q)\| \leq (4\beta^2 + 4)\|\text{arctanh}(\mathbf{m}) - \text{arctanh}(\mathbf{n})\| + \|\mathbf{y} - \widehat{\mathbf{y}}\|. \quad (\text{A.19})$$

Proof. The inequality (A.19) follows with the smaller constant factor $\beta^2 + 3\beta + 1 \leq 4\beta^2 + 4$ using (4.53) and the fact that $\tanh(\cdot)$ is 1-Lipschitz. \square

A.2 Proof of Lemma 4.10

We split the proof into four parts.

Proof of Lemma 4.10, Part 1. Fix $c = (1/4) - (\beta/2) > 0$. Lemma 4.11 implies that for $K_{\text{AMP}} = K_{\text{AMP}}(\beta, T, \varepsilon)$ sufficiently large, we have with probability $1 - o_n(1)$

$$\begin{aligned} \|\nabla \mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^{\text{AMP}}; \mathbf{y}, q_*)\| &\leq \frac{c\sqrt{\varepsilon tn}}{4}, \\ \widehat{\mathbf{m}}^{\text{AMP}} &:= \text{AMP}(\mathbf{A}, \mathbf{y}(t); K_{\text{AMP}}), \quad q_* := q_*(\beta, t). \end{aligned} \quad (\text{A.20})$$

Therefore, if $\|\mathbf{y}(t) - \widehat{\mathbf{y}}\| \leq (c\sqrt{\varepsilon tn})/4$ then

$$\|\nabla \mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^{\text{AMP}}; \widehat{\mathbf{y}}, q_*)\| \leq \|\nabla \mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^{\text{AMP}}; \mathbf{y}(t), q_*)\| + \|\mathbf{y} - \widehat{\mathbf{y}}\| \leq \frac{c}{2}\sqrt{\varepsilon t}$$

Moreover Lemma 4.12 implies that there exist $\varepsilon_0, c > 0$ such that for all $\varepsilon \in (0, \varepsilon_0)$,

$$\nabla^2 \mathcal{F}_{\text{TAP}}(\mathbf{m}; \widehat{\mathbf{y}}, q_*) = \nabla^2 \mathcal{F}_{\text{TAP}}(\mathbf{m}; \mathbf{y}(t), q_*) \succeq c\mathbf{D}(\mathbf{m}), \quad \forall \mathbf{m} \in B\left(\widehat{\mathbf{m}}^{\text{AMP}}, \sqrt{\varepsilon tn}\right) \cap (-1, 1)^n.$$

Using $\varepsilon t/4$ in place of ε in Lemma A.6, it follows that there exists a local minimum

$$\mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}; q_*) \in B\left(\widehat{\mathbf{m}}^{\text{AMP}}, \frac{\sqrt{\varepsilon tn}}{2}\right) \cap (-1, 1)^n$$

of $\mathcal{F}_{\text{TAP}}(\cdot, \widehat{\mathbf{y}}; q_*)$ which is also the unique stationary point in $B\left(\widehat{\mathbf{m}}^{\text{AMP}}, (1/2)\sqrt{\varepsilon tn}\right) \cap (-1, 1)^n$.

We next claim that, for any $K > K_{\text{AMP}}$, with probability $1 - o_n(1)$, this local minimum is also the unique stationary point in $B\left(\text{AMP}(\mathbf{A}, \mathbf{y}(t); k), (1/2)\sqrt{\varepsilon tn}\right) \cap (-1, 1)^n$. Indeed for K_{AMP} sufficiently large (writing for simplicity $\mathbf{y} = \mathbf{y}(t)$):

$$\begin{aligned} \text{p-lim}_{n \rightarrow \infty} \sup_{k_1, k_2 \in [K_{\text{AMP}}, K]} \|\text{AMP}(\mathbf{A}, \mathbf{y}; k_1) - \text{AMP}(\mathbf{A}, \mathbf{y}; k_2)\|^2 &= \sup_{k_1, k_2 \in [k_{\text{alg}}, K]} \text{p-lim}_{n \rightarrow \infty} \|\text{AMP}_\beta(\mathbf{A}, \mathbf{y}; k_1) - \text{AMP}_\beta(\mathbf{A}, \mathbf{y}; k_2)\|^2 \\ &\leq n \cdot \sup_{k_1, k_2 \geq K_{\text{AMP}}} |q_{k_1}(\beta, t) - q_{k_2}(\beta, t)|. \end{aligned}$$

From Eq. (4.22), by eventually increasing K_{AMP} , we have

$$\sup_{k_1, k_2 \geq K_{\text{AMP}}} |q_{k_1}(\beta, t) - q_{k_2}(\beta, t)| \leq \frac{\varepsilon t}{16}.$$

For such K_{AMP} , with probability $1 - o_n(1)$, all $k \in [K_{\text{AMP}}, K]$ satisfy

$$\begin{aligned} \|\mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_{K_{\text{AMP}}}) - \text{AMP}(\mathbf{A}, \mathbf{y}; k)\| &\leq \|\mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_{\text{AMP}}) - \text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}})\| \\ &\quad + \|\text{AMP}(\mathbf{A}, \mathbf{y}; k) - \text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}})\| \\ &\leq \frac{\sqrt{\varepsilon t n}}{2} + \sqrt{\frac{\varepsilon t n}{4}} \\ &\leq \frac{3}{4} \sqrt{\varepsilon t n}. \end{aligned}$$

Let

$$S(k, \rho) := B(\text{AMP}_\beta(\mathbf{A}, \mathbf{y}; k), \rho) \cap (-1, 1)^n, \quad \rho_{n,t} := \sqrt{\varepsilon n t}$$

Recall that $\mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_*)$ is the unique stationary point of $\mathcal{F}_{\text{TAP}}(\cdot; \mathbf{y}, q_*)$ in $S(K_{\text{AMP}}, \rho_{n,t})$. By the above, it is also a stationary point in $S(k, \rho_{n,t})$, for $k \in [K_{\text{AMP}}, K]$. Repeating the same argument as before, there is only one stationary point inside $S(k, \rho_{n,t})$, hence this must coincide with $\mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_*)$. \square

Proof of Lemma 4.10, Part 2. Because K_{AMP} is large depending on δ_0 , Lemma 4.11 implies that with probability $1 - o_n(1)$,

$$\|\nabla \mathcal{F}_{\text{TAP}}(\text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}}), \mathbf{y}; q_*)\| \leq \frac{c\delta_0 \sqrt{t n}}{4}.$$

Using $\frac{\delta_0 \sqrt{t}}{4}$ in place of ε in Lemma A.6, it follows that the local minimizer $\mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_*)$ of $\mathcal{F}_{\text{TAP}}(\cdot; \mathbf{y}, q_*)$ satisfies

$$\|\text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}}) - \mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_*)\| \leq \frac{\delta_0 \sqrt{t n}}{2}.$$

Since K is sufficiently large depending on δ_0 , Lemma implies that with probability $1 - o_n(1)$,

$$\|\mathbf{m}(\mathbf{A}, \mathbf{y}) - \text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}})\| \leq \frac{\delta_0 \sqrt{t n}}{2}.$$

Combining, we obtain that with probability $1 - o_n(1)$,

$$\begin{aligned} \|\mathbf{m}(\mathbf{A}, \mathbf{y}) - \mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_*)\| &\leq \|\mathbf{m}(\mathbf{A}, \mathbf{y}) - \text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}})\| + \|\text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}}) - \mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_*)\| \\ &\leq \delta_0 \sqrt{t n}. \end{aligned}$$

\square

Proof of Lemma 4.10, Part 3. The result is immediate from (A.8). \square

Proof of Lemma 4.10, Part 4. We apply Lemma A.8 with $\mathcal{F}(\cdot) = \mathcal{F}_{\text{TAP}}(\cdot; \hat{\mathbf{y}}, q_*)$ and $\mathbf{m}_* = \mathbf{m}_*(\mathbf{A}, \hat{\mathbf{y}}; q_*)$ (with $q_* = q_*(\beta, t)$). We need to check that assumptions (A.9), (A.10) of Lemma A.8 hold for $\widehat{\mathbf{m}}^0 = \tanh(\mathbf{u}^0)$ with \mathbf{u}^0 satisfying Eq. (4.51).

To check assumption (A.9), we take K_{AMP} sufficiently large and δ_0 sufficiently small, obtaining

$$\begin{aligned} \|\widehat{\mathbf{m}}^0 - \mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}; q_*)\| &\leq \|\widehat{\mathbf{m}}^0 - \text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}})\| + \|\text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}}) - \mathbf{m}(\mathbf{A}, \mathbf{y})\| \\ &\quad + \|\mathbf{m}(\mathbf{A}, \mathbf{y}) - \mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_*)\| + \|\mathbf{m}_*(\mathbf{A}, \mathbf{y}; q_*) - \mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}; q_*)\| \\ &\stackrel{(a)}{\leq} \frac{c\sqrt{\varepsilon tn}}{96(\beta^2 + 1)} + \frac{1}{100}\sqrt{\varepsilon tn} + \delta_0\sqrt{tn} + \frac{\|\mathbf{y} - \widehat{\mathbf{y}}\|}{c} \\ &\leq \frac{\sqrt{\varepsilon tn}}{3} \end{aligned}$$

where inequality (a) holds with probability $1 - o_n(1)$. In the last step we used $c \leq 1$.

To check Eq. (A.10), we use (A.19) we find that with probability $1 - o_n(1)$,

$$\begin{aligned} \|\nabla \mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^0; \widehat{\mathbf{y}}, q_*)\| &\leq \|\nabla \mathcal{F}_{\text{TAP}}(\text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}}); \mathbf{y}, q_*)\| + \|\mathbf{y} - \widehat{\mathbf{y}}\| \\ &\quad + (4\beta^2 + 4)\|\text{arctanh}(\widehat{\mathbf{m}}^0) - \text{arctanh}(\text{AMP}(\mathbf{A}, \widehat{\mathbf{y}}; K_{\text{AMP}}))\| \\ &\leq \|\nabla \mathcal{F}_{\text{TAP}}(\text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}}); \mathbf{y}, q_*)\| + \frac{c\sqrt{\varepsilon tn}}{24} + \frac{c\sqrt{\varepsilon tn}}{4}. \end{aligned}$$

Combining with Eq. (A.20), we find that with probability $1 - o_n(1)$,

$$\|\nabla \mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^0; \widehat{\mathbf{y}}, q_*)\| \leq \frac{c\sqrt{\varepsilon tn}}{6}.$$

Finally, we apply Lemma A.6 with $\frac{\varepsilon t}{9}$ in place of ε , to get

$$\mathcal{F}_{\text{TAP}}(\widehat{\mathbf{m}}^0; \widehat{\mathbf{y}}, q_*) \leq \mathcal{F}_{\text{TAP}}(\mathbf{m}_*(\mathbf{A}, \widehat{\mathbf{y}}; q_*); \widehat{\mathbf{y}}, q_*) + \frac{nc\varepsilon t}{9}.$$

Lemma A.8 now applies for η_0 sufficiently small. Moreover, with probability $1 - o_n(1)$ the initialization \mathbf{x}^0 satisfies

$$\begin{aligned} \|\text{arctanh}(\widehat{\mathbf{m}}^0)\| &\leq \|\text{arctanh}(\widehat{\mathbf{m}}^0) - \text{arctanh}(\text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}}))\| + \|\text{arctanh}(\text{AMP}(\mathbf{A}, \mathbf{y}; K_{\text{AMP}}))\| \\ &\leq \frac{c\sqrt{\varepsilon tn}}{96(\beta^2 + 1)} + \sqrt{3(\gamma_*(\beta, t) + t)}\sqrt{n} \\ &\leq C(\beta, c, T)\sqrt{tn}. \end{aligned}$$

Thus, (A.12) implies (4.52) for a sufficiently large number K_{NGD} of natural gradient iterations. \square