

ALGORITHMIC THRESHOLDS FOR TENSOR PCA

BY GÉRARD BEN AROUS^{1,*}, REZA GHEISSARI^{1,†} AND AUKOSH JAGANNATH²

¹*Courant Institute, New York University, *benarous@cims.nyu.edu; †reza@cims.nyu.edu*

²*Department of Statistics and Actuarial Science, University of Waterloo, a.jagannath@uwaterloo.ca*

We study the algorithmic thresholds for principal component analysis of Gaussian k -tensors with a planted rank-one spike, via Langevin dynamics and gradient descent. In order to efficiently recover the spike from natural initializations, the signal-to-noise ratio must diverge in the dimension. Our proof shows that the mechanism for the success/failure of recovery is the strength of the “curvature” of the spike on the maximum entropy region of the initial data. To demonstrate this, we study the dynamics on a generalized family of high-dimensional landscapes with planted signals, containing the spiked tensor models as specific instances. We identify thresholds of signal-to-noise ratios above which order 1 time recovery succeeds; in the case of the spiked tensor model, these match the thresholds conjectured for algorithms such as approximate message passing. Below these thresholds, where the curvature of the signal on the maximal entropy region is weak, we show that recovery from certain natural initializations takes at least stretched exponential time. Our approach combines global regularity estimates for spin glasses with pointwise estimates to study the recovery problem by a perturbative approach.

1. Introduction. Optimization in high-dimensional landscapes can be computationally hard. This difficulty is often attributed to the topological complexity of the landscape. We show here that for planted signal recovery problems in high dimensions, there is another key obstruction to local optimization methods. Indeed, we find that a crucial factor in these settings is the competition between the strength of the signal and the entropy of the prior. We focus on a well-known optimization problem from high-dimensional statistics which is known to be NP hard [20], namely, maximum likelihood estimation for *tensor principal component analysis* (PCA) [33].

Suppose that we are given M i.i.d. observations, Y^ℓ , of a k -tensor of rank 1 which has been subject to Gaussian noise. That is,

$$Y^\ell = \sqrt{N}\lambda v^{\otimes k} + W^\ell,$$

where $v \in \mathbb{S}^{N-1}(1)$ is deterministic, W^ℓ are i.i.d. Gaussian k -tensors with $W_{i_1, \dots, i_k}^\ell \sim \mathcal{N}(0, 1)$ and $\lambda \geq 0$ is the *signal-to-noise ratio*. Our goal is to infer the “planted signal,” or “spike,” v , by maximum likelihood estimation.

Observe that maximum likelihood estimation for this problem boils down to optimizing an empirical risk of the form

$$(1.1) \quad \hat{R}(x) = \frac{1}{M} \sum_{\ell \leq M} (W^\ell, x^{\otimes k}) + \sqrt{N}\lambda(v, x)^k,$$

where (\cdot, \cdot) denotes the usual Euclidean inner product. Note further that in this setting, optimizing this risk is equivalent (in law) to optimizing the same risk for a single observation upon making the change $\lambda \mapsto \sqrt{M}\lambda$. We therefore restrict our analysis to the case $M = 1$.

Received August 2018; revised November 2019.

MSC2010 subject classifications. Primary 62F10, 62M05, 62F30; secondary 60H30, 65C05, 46N30, 82D30, 82C44.

Key words and phrases. Tensor PCA, tensor estimation, spiked tensor model, planted signal recovery, gradient descent, Langevin dynamics, spin glasses, free energy wells.

When $k = 2$, this is the well-known spiked matrix model [24]. In this setting it is known [30] that there is an order 1 critical signal-to-noise ratio, λ_c , such that below λ_c , it is information-theoretically impossible to detect the spike, and above λ_c , the maximum likelihood estimator is a distinguishing statistic. This transition is commonly referred to as the BBP transition [4]. In this setting the maximum likelihood estimator is the top eigenvector which can be computed in polynomial time by, for example, power iteration. Much more detailed information is known about this transition for spiked matrix models, including universality, fluctuations and large deviations. See, for example, [10, 12, 28] for a small sample of these works.

When $k \geq 3$, this is the spiked tensor model [33]. In this case there is a regime of signal-to-noise ratios for which it is information theoretically possible to recover the signal but for which there is no known algorithm to efficiently approximate it. This is called a *statistical-to-algorithmic gap*. In particular, it was shown in [29, 31, 33] that the minimal signal-to-noise ratio above which it is information-theoretically possible to detect the signal—called the *information-theoretic threshold*—is of order 1. See also [13, 27, 31, 32] for similar results with different priors. On the other hand, the minimal signal-to-noise ratio above which one can *efficiently* detect the signal—called the *algorithmic threshold*—has been proved or predicted to scale like $\lambda = N^\alpha$ for some $\alpha > 0$ for every studied algorithm. (By the correspondence mentioned above, the regime of diverging λ can be translated to the regime of λ of order 1 with a diverging number of observations $M = N^{\alpha/2}$, so that this regime is also of practical interest.) In [33], two local optimization methods, approximate message passing and tensor power iteration, were shown to have critical exponents α at most $(k - 1)/2$ with predicted thresholds at $\alpha = (k - 2)/2$. Semidefinite relaxation approaches have also been analyzed. Tensor unfolding was shown [33] to have a critical exponent of at most $\frac{\lfloor k/2 \rfloor - 1}{2}$ and, conjecturally, $(k - 2)/4$. It was also shown that the degree 4 sum-of-squares algorithm [22] and a related spectral algorithm [21] (in the case $k = 3$) have sharp critical thresholds of $(k - 2)/4$. See also [25] for a similar analysis in the case $k = 4$. We remark that statistical-to-algorithmic gaps, often diverging in the underlying dimension, have also been observed in myriad other problems of interest [1, 5, 6, 11, 15, 39].

Let us also discuss the complexity of the landscape $\hat{R}(x)$ given by (1.1). The complexity in the absence of a spike (the case where $\lambda = 0$) has been extensively studied [2, 3, 37]; see also [16] for a related line of work. When adding in the signal term so that $\lambda > 0$, it was proved [9] that the expected number of critical points of $\hat{R}(x)$, called the annealed complexity, is exponentially large in N and has a topological phase transition as one varies λ on the order 1 scale.

One might wonder why the statistical-to-algorithmic gap is diverging when $k \geq 3$. We investigate this issue for algorithms, which directly perform maximum likelihood estimation, by analyzing the behavior of a family of “plain vanilla” algorithms, called *Langevin dynamics*, as well as gradient descent. We find that, for natural initializations, the statistical-to-algorithmic gap for Langevin dynamics and gradient descent diverges like $\lambda \sim N^\alpha$. One may expect that this issue is due to the topological complexity of $\hat{R}(x)$. Our proof, however, suggests that this gap is actually due to the weakness of the signal in the region of maximal entropy for the uninformative prior.

To clarify this point, we study these dynamics on a more general family of random landscapes. For convenience, let us rescale our problem to be on $S^N = \mathbb{S}^{N-1}(\sqrt{N})$, the sphere in \mathbb{R}^N of radius \sqrt{N} . We consider a function $H : S^N \rightarrow \mathbb{R}$ of the form

$$(1.2) \quad H(x) = H_0(x) - N\lambda\phi(x),$$

where ϕ is a deterministic, nonlinear function and H_0 is a noise term. To put ourselves in a general setting, we only assume that H_0 is a mean-zero Gaussian process with a rotationally

invariant law that is well defined in all dimensions. That is, for every N , H_0 has covariance of the form

$$(1.3) \quad \text{Cov}(H_0(x), H_0(y)) = N\xi\left(\frac{(x, y)}{N}\right)$$

for some fixed function ξ , where (\cdot, \cdot) denotes the Euclidean inner product.¹ For simplicity, we take the function $\phi(x)$ to be a function of the inner product of x with some “unknown” vector $v \in \mathbb{R}^N$. As H_0 is isotropic, without loss of generality, we assume that $v = e_1$, the first canonical Euclidean basis vector, so that $\phi(x)$ is a function of

$$(1.4) \quad m_N(x) = \frac{(x, e_1)}{\sqrt{N}} = \frac{x_1}{\sqrt{N}},$$

which we call the *correlation*. In particular, we take ϕ of the form

$$(1.5) \quad \phi(x) = (m_N(x))^k,$$

where $k \geq 1$ is not necessarily integer. The case $\xi(t) = t^k$ and the integer $k \geq 2$ correspond to the setting of (1.1). The case where $\xi(t) = t^p$ corresponds to the $(p + k)$ -spin glass model from [18], whose topological phase transitions have been precisely analyzed in [34] via a computation of the quenched complexity using a novel replica Kac–Rice approach.

We analyze here the performance of Langevin dynamics and gradient descent in achieving order 1 correlation as one varies the initialization, the nonlinearity of the signal, k and the signal-to-noise ratio, λ . If $k > 2$, we find that the critical threshold for algorithmic recovery via Langevin dynamics diverges like $\lambda_{\text{alg}} \sim N^\alpha$, with $\alpha > 0$, for a natural class of initializations. On the other hand, we find that if $k < 2$, this algorithmic threshold is of order 1. In the former regime the second derivative of the signal is vanishing in the maximum entropy region of the uninformative prior, whereas in the latter it is diverging, matching the mechanism proposed above.

Our analysis has two main thrusts, efficient recovery above critical thresholds and refutation below them. In both of these settings, we find that the obstacle to recovering the signal e_1 via Langevin dynamics is escaping the equator, that is, the region where $m_N(x) = O(N^{-1/2})$, which corresponds to the maximum entropy region of the uninformative prior.

For the recovery problem we prove that as soon as $\alpha > \alpha_c(\infty) := (k - 2)/2$, Langevin dynamics and gradient descent solve the recovery problem in order 1 time when started uniformly at random; see Theorem 2.1. To isolate the importance of the initialization in this problem, we provide a hierarchy of sufficient conditions on the initial data that imply that Langevin dynamics with $\lambda = N^\alpha$ will efficiently solve the recovery problem down to a hierarchy of thresholds $\alpha_c(n)$, the lowest of these thresholds being $\alpha_c(\infty)$; see Section 2.2. In Section 2.3 we give examples of initial data that satisfy these conditions at different levels; the case of the volume measure is discussed in depth in Section 2.1. To prove these results, we build on the “bounding flows” strategy of [7]. In particular, we show that on $O(1)$ times, we can compare the evolution of the correlation, $m(X_t)$, to the gradient descent for the problem with no noise $H_0 = 0$. This follows by a stochastic Taylor expansion upon combining the Sobolev-type \mathcal{G} -norm estimates developed in [7] for spin glasses, with estimates on the regularity of the initial data developed here. This is discussed in more detail in Section 2.5.

For the refutation problem, the threshold $\alpha_c(\infty)$ also has a natural (heuristic) interpretation: if $\alpha < \alpha_c(\infty)$, then Langevin dynamics started uniformly at random would take exponential time to solve the recovery problem when given a pure signal, that is, $H_0(x) = 0$. We

¹It is classical [35] that the largest class of such ξ is of the form $\xi(t) = \sum_p a_p^2 t^p$ with $\xi(1 + \epsilon) < \infty$ for some $\epsilon > 0$.

conjecture that this is also the case when H_0 is added back, so that $\alpha_c(\infty)$ is sharp for uniform at random initialization. As added motivation for this conjecture, we prove that indeed $\alpha_c(\infty)$ is sharp for a natural Gibbs class of initial data: from such initializations, the dynamics efficiently recover the signal for $\alpha > \alpha_c(\infty)$ but take at least stretched exponential time to do so for $\alpha < \alpha_c(\infty)$. To prove this refutation theorem below $\alpha_c(\infty)$, we formalize the notion of *free energy wells* (see Definition 2.12), whose existence implies exponential lower bounds on the exit time of the well from a restriction of the corresponding Gibbs measure. We find that below the critical $\alpha_c(\infty)$ there is a free energy well around the equator and use this to deduce hardness of recovery for Gibbs initializations. For more on this, see Section 2.4.

2. Statements of main results. We focus on a canonical class of optimization algorithms called *Langevin dynamics* and *gradient descent* with *Hamiltonian* H . The Langevin dynamics interpolate between Brownian motion and gradient descent via a parameter $\beta > 0$, usually called the *inverse temperature*, with the case $\beta = 0$ corresponding to Brownian motion and the case $\beta = \infty$ corresponding to gradient descent. More precisely, for $0 \leq \beta < \infty$, let X_t^β solve the stochastic differential equation (SDE)

$$(2.1) \quad \begin{cases} dX_t^\beta = \sqrt{2} dB_t - \beta \nabla H(X_t^\beta) dt, \\ X_0^\beta = x, \end{cases}$$

where B_t is Brownian motion on S^N , ∇ denotes the covariant derivative on S^N , and H is called the *Hamiltonian* which is given here by (1.2). As H is C^1 , this martingale problem is well posed so that X_t^β is well defined [14, 36]. (When k is an integer, H is smooth so that one can solve this in the strong sense as well [23].) When $\beta = \infty$, let X_t^∞ denote the solution to the ODE

$$(2.2) \quad \begin{cases} dX_t^\infty = -\nabla H(X_t^\infty) dt, \\ X_0^\infty = x. \end{cases}$$

Note that on the complement of the set $m_N(x) = 0$, ∇H is locally Lipschitz so that X_t^∞ is locally well posed (in time) by the Picard–Lindelof theorem. Note that the infinitesimal generators of these processes are given by

$$(2.3) \quad L_\beta = \begin{cases} \Delta - \beta \langle \nabla H, \nabla \cdot \rangle & \beta \in (0, \infty), \\ -\langle \nabla H, \nabla \cdot \rangle & \beta = \infty, \end{cases}$$

where Δ is the Laplacian on S^N and $\langle \cdot, \cdot \rangle$ is the metric tensor. Whenever β is clear from context, we will write $X_t = X_t^\beta$ and $L = L_\beta$.

We aim to determine the minimal λ for which efficient recovery of the signal, e_1 , is possible via these dynamics and understand the role that the initialization plays. There are, of course, multiple notions of recovery. The main ones in which we are interested are *weak recovery* and *strong recovery*. For fixed ξ, β, k , sequence λ_N and sequence of initial data x_N , we say that the Langevin dynamics *weakly recovers* the signal in order 1 time if it attains order 1 correlation in $O(1)$ time with high probability. On the other hand, we say that the dynamics *strongly recovers* the signal in order 1 time if it attains $1 - o(1)$ correlation in $O(1)$ time with high probability. In the diverging signal-to-noise ratio regime, weak and strong recovery are equivalent (see Lemma 4.1). A more complete discussion of their relationship is provided in Section 4.

2.1. *Recovery initialized from the volume measure.* Perhaps the most natural initialization is a completely uninformative prior, that is, the (uniform) volume measure on S^N . This is particularly motivated from the algorithmic perspective as it is easy to sample from the volume measure on S^N in order 1 time. In order to focus on the key issues and deal with all k in a comprehensive manner, we restrict to the upper hemisphere: $\{x_1 > 0\}$. Of course, a point sampled from the volume measure on S^N is in the upper hemisphere with probability $\frac{1}{2}$.

We obtain the following recovery guarantees starting from the volume measure on the upper hemisphere.² Let dx be the volume measure on S^N , and let \mathbb{P} denote the law of the noise H_0 . In the following, denote the law of X_t^β started at x by Q_x , and in the case of gradient descent, referred to as $\beta = \infty$, interpret Q_x as a Dirac mass on the trajectory of X_t^∞ started from x .

THEOREM 2.1. *Fix any ξ , any $\beta \in (0, \infty) \cup \{\infty\}$ and any $k \in [1, \infty)$:*

1. *If $k > 2$ and $\lambda = N^\alpha$ for $\alpha > \frac{k-2}{2}$, then for every $\epsilon > 0$, there exists T_0 such that for every $T \geq T_0$,*

$$\lim_{N \rightarrow \infty} \int_{x_1 > 0} Q_x \left(\min_{t \in [T_0, T]} m_N(X_t^\beta) \geq 1 - \epsilon \right) dx = 1, \quad \mathbb{P}\text{-a.s.}$$

2. *If $k < 2$, $\beta < \infty$ and λ is a large enough constant, there exists T_0 such that for all $T \geq T_0$,*

$$\lim_{\epsilon \rightarrow 0} \lim_{N \rightarrow \infty} \int_{x_1 > 0} Q_x \left(\min_{t \in [T_0, T]} m_N(X_t^\beta) \geq \epsilon \right) dx = 1, \quad \mathbb{P}\text{-a.s.}$$

We first pause to comment on the special case of $\xi(t) = t^k$ with k integer, corresponding to maximum likelihood estimation for tensor PCA. The thresholds of Theorem 2.1 improve upon the rigorously known threshold for approximate message passing and tensor power iteration, and the $(k - 2)/2$ threshold matches the conjectured threshold for those algorithms.

These thresholds correspond to the signal-to-noise ratios for which the second derivative of $N\lambda\phi$ diverges at points of correlation $m_N(x) = \Theta(N^{-\frac{1}{2}})$ (the asymptotic support of dx). We predict that these thresholds are sharp for efficient algorithmic recovery of e_1 via local optimization so that, when the second derivative is $o(1)$ at these correlations, efficient recovery is not possible. When $H_0 \equiv 0$, it is easy to see that for $k < 2$ and λ order 1, or $k > 2$ and $\lambda = N^\alpha$ with $\alpha < \frac{k-2}{2}$, the Langevin dynamics takes an exponential time to correlate with e_1 with this initialization; it would be expected that this persists with the addition of noise; see the in-depth discussion in Section 2.5.

We are able to prove sharpness of these thresholds for a high-temperature Gibbs-type initialization that approximates the volume measure as $\beta \rightarrow 0$, that is, that below these thresholds, the corresponding dynamics fails to recover in stretched exponential time whereas above the thresholds the dynamics recover as above. In the next section we define general conditions, Condition 1 and Condition 2, on the initial data that guarantee recovery above these thresholds.

2.2. *General thresholds for recovery.* Theorem 2.1 is a particular case of more general recovery results that translate a natural hierarchy of conditions on the initial data to thresholds in λ above which we guarantee recovery of e_1 . Let $\mathcal{M}_1(S^N)$ denote the space of probability measures on S^N . A choice of initial data corresponds to a choice of measure $\mu_N \in \mathcal{M}_1(S^N)$.

²In the first preprint of this article, a suboptimal version of this result for initialization from the uniform measure was stated; the current formulation of Theorem 2.1 resolves what was there labeled as Conjecture 1.

Our main recovery guarantees apply to any initial data which satisfy the following two natural conditions.

The first condition is on the regularity of the initial data. For $\beta \in (0, \infty) \cup \{\infty\}$, define L_0 as

$$(2.4) \quad L_{0,\beta} = \begin{cases} \Delta - \beta \langle \nabla H_0, \nabla \cdot \rangle & \beta \in (0, \infty), \\ -\langle \nabla H_0, \nabla \cdot \rangle & \beta = \infty, \end{cases}$$

that is, the generator of Langevin dynamics and gradient descent with respect to H_0 . Correspondingly, let $e^{tL_{0,\beta}}$ refer to the semigroup induced by $L_{0,\beta}$. For every $\delta > 0$, $n \geq 1$ and N , let

$$E_{n,\delta,N}^\beta = \bigcap_{\ell=0}^{n-1} \{x : |L_{0,\beta}^\ell m_N(x)| \leq N^{-\frac{1}{2}+\delta}\},$$

and for every $\delta, T > 0$ and N , let

$$\tilde{E}_{T,\delta,N}^\beta = \left\{x : \sup_{t \leq T} |e^{tL_{0,\beta}} L_{0,\beta} m_N(x)| \leq N^{-\frac{1}{2}+\delta}\right\}.$$

Again, when understood from context, we drop the dependence on β in the notation.

DEFINITION 2.2. We say that a sequence of random probability measures $\mu_N \in \mathcal{M}_1(\mathcal{S}^N)$ satisfies *Condition 1 at level n* for inverse temperature $\beta \in (0, \infty) \cup \{\infty\}$ if for every $\delta > 0$,

$$(2.5) \quad \lim_{N \rightarrow \infty} \mu_N((E_{n,\delta,N})^c) = 0 \quad \mathbb{P}\text{-a.s.}$$

We say that a sequence of random probability measures $\mu_N \in \mathcal{M}_1(\mathcal{S}^N)$ weakly satisfies *Condition 1 at level ∞* for inverse temperature $\beta \in (0, \infty) \cup \{\infty\}$ if for every $\delta, T > 0$,

$$(2.6) \quad \lim_{N \rightarrow \infty} \mu_N((\tilde{E}_{T,\delta,N})^c) = 0 \quad \mathbb{P}\text{-a.s.}$$

Let us pause for a moment to interpret Condition 1 intuitively. Let us call Langevin dynamics or gradient descent with Hamiltonian H_0 the “pure noise” dynamics. Recall that $e^{tL_0} f(x)$ is the expected value of f with respect to the pure noise dynamics at time t started at x . The set $E_{T,\delta,N}$ is then the set of initializations for which the pure noise dynamics does not have an atypically strong push in the (opposite) direction of the spike before time T . Since the landscape is isotropic, one should expect that if one runs the pure noise dynamics started from any initial data that is agnostic to the spike, then correlations of observables with the spike will stay on their central limit theorem scales for a long time. Weakly satisfying Condition 1 at level ∞ states that the probability of such initializations with respect to the initial data tends to one. On the other hand, Condition 1 at level n says that the first n coefficients in the Taylor expansion of $e^{tL_0} m$ are all on this typical scale. Clearly, $E_{\infty,\delta,N}$ is contained in $\tilde{E}_{T,\delta',N}$ for every $\delta' > \delta$, hence the qualifier “weak.” The implication in the other direction does not obviously hold due to possible cancellations.

The second condition ensures that the initial correlation is on the typical $\Theta(N^{-\frac{1}{2}})$ scale so that the drift from gradient descent for the signal is not negligible at time zero.

DEFINITION 2.3. A sequence of random probability measures, $\mu_N \in \mathcal{M}_1(\mathcal{S}^N)$, satisfies Condition 2 if

$$\lim_{\epsilon \rightarrow 0} \overline{\lim}_{N \rightarrow \infty} \mu_N(x_1 < \epsilon) = 0 \quad \mathbb{P}\text{-a.s.}$$

A sequence of random probability measures $\mu_N \in \mathcal{M}_1(\mathcal{S}^N)$ satisfies Condition 2 if for every $\delta > 0$,

$$\lim_{N \rightarrow \infty} \mu_N(x_1 \leq N^{-\delta}) = 0 \quad \text{in prob.}$$

We emphasize that neither of these conditions involve the parameters k (the nonlinearity of the signal) or λ (the signal-to-noise ratio). The conditions can be shown to hold for various natural choices of initial data, such as the volume measure on the upper hemisphere, implying Theorem 2.1, as well as certain “high-temperature” Gibbs measures. For more on this, see Section 2.3.

Let us now turn to our main results under Condition 1 and Condition 2. We begin with the *supercritical* regime, $k > 2$, where one will need λ to diverge with N to efficiently recover, as the curvature of the signal in the region where $m_N(x) = \Theta(N^{-1/2})$ is negligible. For every $n \geq 1$, let

$$(2.7) \quad \alpha_c(n) = \frac{k-1}{2} - \frac{n-1}{2n} \quad \text{and} \quad \alpha_c(\infty) = \frac{k-2}{2}.$$

We then have the following result regarding strong recovery.

THEOREM 2.4. *Fix any ξ , any $\beta \in (0, \infty) \cup \{\infty\}$ and any $k > 2$. Let $\lambda = N^\alpha$, and consider a sequence of initializations $\mu_N \in \mathcal{M}_1(\mathcal{S}^N)$. If μ_N satisfies Condition 2, then we have the following:*

1. *If μ_N weakly satisfies Condition 1 at level ∞ for inverse temperature β , then, for every $\alpha > \alpha_c(\infty)$ and every $\epsilon > 0$, there exists a T_0 such that for every $T > T_0$,*

$$\lim_{N \rightarrow \infty} \int Q_x \left(\min_{t \in [T_0, T]} m_N(X_t^\beta) \geq 1 - \epsilon \right) d\mu_N(x) = 1 \quad \mathbb{P}\text{-a.s.}$$

1. *If μ_N satisfies Condition 1 at level n for inverse temperature β , then the above holds for every $\alpha > \alpha_c(n)$.*

If instead μ_N satisfies Condition 2, then the above convergence holds in probability.

The above theorem shows that, in the regime $k > 2$, we need λ to diverge for Langevin dynamics and gradient descent to recover the signal in order 1 time. Observe that for such k , the second derivative of $N\phi$ in the region $m_N(x) = O(N^{-1/2})$ is vanishing as $N \rightarrow \infty$. Let us now show, conversely, that in the *subcritical* regime $k < 2$, that is, the regime where the second derivative of $N\phi$ is diverging when $m_N(x) = \Theta(N^{-1/2})$, order 1 time weak recovery holds for large but finite signal-to-noise ratios. That is, the statistical-to-algorithmic gap (if one exists) is at most order 1 for $k < 2$. In this regime one cannot hope for strong recovery (see Remark 2.11). If we let

$$(2.8) \quad k_c(n) = 2 - \frac{1}{n} \quad \text{and} \quad k_c(\infty) = 2,$$

then we have the following weak recovery guarantee. In the following, in the cases $\beta = \infty$ and $1 < k < 2$, we interpret Q_x to be the zero measure (i.e., $Q_x = 0$) when $x_1 = 0$ to avoid issues related to well posedness. For more on this, see Remark 5.6.

THEOREM 2.5. *Fix any ξ and any $\beta \in (0, \infty) \cup \{\infty\}$. There exists $\lambda_0(\beta, \xi, k) > 0$ such that for all $\lambda > \lambda_0$ the following holds. If μ_N satisfies Condition 2, then we have the following.*

1. If μ_N weakly satisfies Condition 1 at level ∞ for inverse temperature β , then, for every $1 \leq k < k_c(\infty)$ and every $\eta > 0$, there exists $\epsilon > 0$ and $T_0 > 0$, such that for any $T > T_0$

$$(2.9) \quad \lim_{N \rightarrow \infty} \int \mathcal{Q}_x \left(\min_{t \in [T_0, T]} m_N(X_t^\beta) \geq \epsilon \right) d\mu_N(x) > 1 - \eta \quad \mathbb{P}\text{-a.s.}$$

2. If instead, $\mu_N \in \mathcal{M}_1(\mathcal{S}^N)$ satisfies Condition 1 at level n for inverse temperature β , then the same result holds for every $1 \leq k < k_c(n)$.

The main ideas behind Theorems 2.4 and 2.5 are essentially the same. We will explain the intuition behind their proofs presently (see Section 2.5 below). We end this section with the following remark on possible relaxations of Condition 1.

REMARK 2.6. One could not decrease the sets $E_{n,\delta,N}$ and $\tilde{E}_{T,\delta,N}$, since measures that don't contain information about the planted signal, for example, the volume measure on \mathcal{S}^N , have $m_N(x) = \Theta(N^{-1/2})$.

2.3. *Examples of initial data satisfying Condition 1 and Condition 2.* Let us now turn to some examples of initial data that satisfy the conditions of our theorems. When considering initial data for such problems there are a few natural choices:

EXAMPLE 1. Let us begin by observing that \mathbb{P} -a.s., any initial data which is concentrated on the region $m_N(x) = \Theta(N^{-1/2})$, for example, δ_x for any x having $m_N(x) = \Theta(N^{-1/2})$, satisfies Condition 1 at level 1 for every $\beta \in (0, \infty) \cup \{\infty\}$ and Condition 2 tautologically.

EXAMPLE 2. Initialization from the volume measure was discussed at length in Section 2.1. Theorem 2.1 is an immediate corollary of Theorems 2.4–2.5 combined with the following theorem:

THEOREM 2.7. *The normalized volume measure on $\mathcal{S}^N \cap \{x_1 > 0\}$ weakly satisfies Condition 1 at level ∞ at every $\beta \in (0, \infty) \cup \{\infty\}$ and satisfies Condition 2.*

At an intuitive level, Theorem 2.7 is saying that for the Langevin dynamics or gradient descent with respect to the pure-noise environment H_0 , initialized from the normalized volume measure, the observable $L_0 m$ does not leave its typical scale of $O(N^{-\frac{1}{2}})$ in order 1 timescales. This is very natural to expect as $L_0 m$ consists of the correlation with e_1 and the gradient of H_0 in the direction of e_1 , while the volume measure and H_0 are both isotropic and have no sense of the direction e_1 . While a similar intuition should hold for a broad set of initializations, the proof of Theorem 2.7, found in Section 6.1, relies crucially on the rotational invariance of the volume measure.³

EXAMPLE 3. Another natural class of initializations are Gibbs-type initializations in the “pure noise” environment. Let $d\pi_{0,\beta}(x) \propto \exp(-\beta H_0(x)) dx$ be the Gibbs measure on \mathcal{S}^N corresponding only to the noise H_0 at inverse temperature β , and let $\pi_{0,\beta}^+$ be $\pi_{0,\beta}$ conditioned on $\{x_1 > 0\}$, that is,

$$(2.10) \quad \pi_{0,\beta}^+ = \pi_{0,\beta}(\cdot \mid x_1 > 0).$$

One can show that such measures satisfy Condition 1 at level n for every n .

³In a previous version of this paper, the suboptimal version of Theorem 2.1 followed from checking Condition 1 at level 4 by hand in the absence of weak Condition 1 at level ∞ .

THEOREM 2.8. *Let ξ be even. There exists $\beta_0 > 0$ such that, for all $\beta < \beta_0$, the measure $\pi_{0,\beta}^+$ satisfies Condition 1 at level n for β , for every $n \geq 1$. Moreover, $\pi_{0,\beta}^+$ satisfies Condition 2'.*

COROLLARY 2.9. *Let ξ be even and $k > 2$. If $\beta < \beta_0$, for every $\beta > 0$, if $\lambda = N^\alpha$ with $\alpha > \alpha_c(\infty)$, the Langevin dynamics starting from $\pi_{0,\beta}^+$ strongly recovers the signal in order 1 time in \mathbb{P} -prob.*

Notice, in particular, that, like the volume measure, the measure $\pi_{0,\beta}$ is completely independent of the noise ϕ as well as the signal-to-noise ratio λ . Moreover, as $\beta \rightarrow 0$, the measure $\pi_{0,\beta}$ approximates the volume measure. We end this section with the following conjecture regarding the measure $\pi_{0,\beta}$. This result would imply an almost sure recovery result for $k > 2$ and $\alpha > \alpha_c(\infty)$ as well as the matching weak recovery result for $k < 2$. We also believe that it is of independent interest in the statistical physics community.

CONJECTURE. *For every $\beta > 0$, if $\gamma(\cdot)$ is the law of a standard Gaussian,*

$$\pi_{0,\beta}(x_1 \in \cdot) \xrightarrow{N \rightarrow \infty} \gamma(\cdot)$$

weakly as measures \mathbb{P} -a.s. In particular, $\pi_{0,\beta}^+$ satisfies Condition 2.

2.4. Refutation below $\alpha_c(\infty)$. The Gibbs initialization introduced in the previous section, $\pi_{0,\beta}^+$, is of further importance to us as it is an initialization for which we are able to prove the sharpness of the threshold $\alpha_c(\infty)$. For any $\epsilon > 0$, let τ_ϵ be the hitting time

$$\tau_\epsilon = \inf\{t > 0 : m_N(X_t) \geq \epsilon\},$$

and let $\pi_{0,\beta}^+$ be as in (2.10).

THEOREM 2.10. *Fix any ξ , any $\beta \in (0, \infty)$ and $k > 2$, and let $\lambda_N = N^\alpha$. If $\alpha < \alpha_c(\infty)$, there exists $c > 0$ such that for every $\epsilon > 0$ sufficiently small,*

$$\overline{\lim}_{N \rightarrow \infty} \frac{1}{N^{2\epsilon}} \log \int Q_x(\tau_{2N^{-\frac{1}{2}+\epsilon}} \leq e^{cN^\epsilon}) d\pi_{0,\beta}^+(x) < -c, \quad \mathbb{P}\text{-a.s.}$$

Motivated by the fact that in the limit as $\beta \rightarrow 0$, $\pi_{0,\beta}^+$ approximates the volume measure on $S^N \cap \{x_1 > 0\}$, we believe that a similar refutation result also holds for initialization from the volume measure whenever $\alpha < \alpha_c(\infty)$; this would make the thresholds $\alpha_c(\infty)$ and $k_c(\infty)$ sharp for initialization from the volume measure.

2.5. Ideas of proofs. We now sketch some of the key ideas underlying the above recovery and refutation results and their proofs.

2.5.1. Ideas of proofs of Theorem 2.4 and Theorem 2.5. We first discuss the intuition behind the proofs of Theorems 2.4–2.5. Our interest is in understanding the transition for signal recovery in short times. It turns out that the subcritical and supercritical problems are essentially the same. To see why, consider, for the moment, the recovery question for Langevin dynamics in the simpler setting where there is only a signal, $H = -N\lambda\phi$ and $H_0 \equiv 0$.

By rotation invariance, the question of escaping the equator for the problem with pure signal is effectively the same as studying the escape from the origin for a one-dimensional Langevin dynamics with Hamiltonian,

$$(2.11) \quad V(m) = \beta\lambda m^k + \frac{1}{2} \log(1 - m^2) \approx \beta\lambda m^k - \frac{1}{2} m^2,$$

in the small noise regime (noise of order $N^{-1/2}$); the term $\frac{1}{2} \log(1 - m^2)$ is the entropy of the set $\mathcal{S}^N \cap \{m(x) \approx m\}$. Evidently, this amounts to studying the ODE,

$$(2.12) \quad \dot{m} = \beta \lambda k m^{k-1} - m,$$

where $m(t) = m_N(X_t)$. The second term reverts to the origin. To escape, we must then hope that the first term dominates at the initial point. In particular, if m is positive and small and λ is large, one hopes to compare this ODE to the simpler system

$$(2.13) \quad \dot{m} \approx \beta \lambda k m^{k-1}.$$

In this setting one may then apply a standard comparison inequality (see Lemma 5.1) which compares solutions of this ODE to certain power laws.⁴ Evidently, the order of growth of the second derivative of m^k is the essential ingredient in resolving the tradeoff above. Indeed, under Condition 2, which places $m_N(0) = \Theta(\frac{1}{\sqrt{N}})$, if $k < 2$ and λ is order 1, then $m^{k-1} \gg m$, and, similarly, if $k > 2$, then $\lambda m^{k-1} \gg m$, provided λ grows sufficiently fast ($\alpha > \alpha_c(\infty)$). Notice also that if $\alpha < \alpha_c(\infty)$, then Langevin dynamics with $\beta < \infty$ from $m_N(0) = O(\frac{1}{\sqrt{N}})$ would not efficiently recover the signal even in this trivial pure spike problem.

When adding back H_0 , we consider the evolution equation for m given by

$$(2.14) \quad dm = (\beta \lambda k m^{k-1} (1 - m^2) + L_0 m) dt + dM_t^m,$$

where L_0 , given by (2.4), is the infinitesimal generator for Langevin dynamics with respect to H_0 , and M_t^m is a martingale. We will see that $M_t^m = O(\sqrt{t/N})$, so that, on short times, this is not far from the situation of (2.13). The remaining discrepancy, evidently, is to ensure that $|L_0 m|$ starts and remains smaller than λm^{k-1} . To this end, we use the \mathcal{G} -norm estimates from [7] to show that provided $L_0 m(X_0)$ is suitably localized, that is, provided Condition 1 holds at level n , then $L_0 m(X_t)$ remains localized on the relevant timescale needed to recover the signal above $\alpha_c(n)$ (see Theorem 5.3). The main result then follows by combining this localization, or its weaker version at level ∞ , with the comparison inequality of Lemma 5.1: this is developed in Section 5.

The proof of the recovery result for $\beta = \infty$ follows mutatis mutandis as that of $\beta < \infty$. For an explanation of the relevant modifications, see Section 5.4.

REMARK 2.11 (Strong recovery is impossible for finite λ). When λ is order 1, one cannot hope to obtain a strong recovery result. Indeed, if we start from any point sufficiently close to the north pole, correlation $m_N(x) \geq 1 - \epsilon$ for $\epsilon > 0$ sufficiently small, then $m_N(t)$ will decrease in correlation in order 1 time. To see this, we examine the drift in (2.14) and expand L_0 as

$$L_0 m = -m - \beta \langle \nabla H_0, \nabla m \rangle :$$

if ϵ is sufficiently small, then $-m$ will dominate $\beta \lambda k m^{k-1}$; furthermore, the maximum of $|\langle \nabla H_0, \nabla m \rangle|$ on the spherical cap $\{m_N(x) \geq 1 - \epsilon\}$ can be shown to scale down to zero as ϵ goes to zero. Putting these together with the types of arguments found in Section 4 would imply the desired.

⁴In the critical regime, $k = 2$, this is the classical Gronwall inequality which is, of course, not of power law type.

2.5.2. *Ideas of proofs of Theorem 2.10.* The underlying idea behind our refutation result is the presence of what we call a *free energy well* for the correlation which is defined as follows. Define the Gibbs measure for H by

$$d\pi_\beta(x) \propto \exp(-\beta H(x)) dx,$$

which is normalized to be a probability measure, where dx is the normalized volume measure on S^N . In the following, for a real number a , we let $B_\epsilon(a) = \{x : |x - a| < \epsilon\}$ denote the ball of radius ϵ around a . For any function $f : S^N \rightarrow \mathbb{R}$, we define the entropy

$$I_f(a; \epsilon) = -\log \pi(\{x : f(x) \in B_\epsilon(a)\}).$$

We can now define free energy wells for Lipschitz functions.

DEFINITION 2.12. A Lipschitz function, $f : S^N \rightarrow \mathbb{R}$, has an ϵ -free energy well of height h in $[a, b]$ if the following holds: there exists $c \in (a, b)$ and $\eta > 0$ such that $B_\epsilon(a) \cap B_\epsilon(b) \cap B_\eta(c) = \emptyset$ and

$$\min\{I_f(a, \epsilon), I_f(b, \epsilon)\} - I_f(c, \eta) \geq h.$$

Such free energy wells are the exit time analog of the of free energy barriers formalized in [8] for spectral gap estimates. We show in Theorem 7.4 that free energy wells confine the dynamics on timescales that are exponential in the height, h , when started from this Gibbs measure π_β restricted to the well. We then show that, for $\alpha < \alpha_c(\infty)$, there is a free energy well for the correlation, namely, the function $f(x) = (x, e_1)$ has a free energy well of height N^ϵ in $[-N^\epsilon, N^\epsilon]$ (see Proposition 7.1). Theorem 2.10 then follows by combining this with the facts that π_β and $\pi_{0,\beta}$ are comparable when restricted to this band of correlations, and $\pi_{0,\beta}$ is asymptotically supported in this region.

We conclude with a remark regarding exceptional points which facilitate recovery at order $1/\lambda$.

REMARK 2.13 (Equatorial passes). In light of the free energy well for the correlation, one might hope to prove an even stronger refutation theorem. It may be tempting to believe that when $k > 2$ and λ is order 1, the Langevin dynamics cannot recover the signal in subexponential times, uniformly, over *all* initial X_0 with $m_N(X_0) = O(N^{-\frac{1}{2}})$. Indeed, this is the case for the simpler ‘‘pure signal’’ problem where $H_0 = 0$. As a consequence of (2.14) and lower bounds on $\|\partial_1 H_0\|_\infty$ (see, e.g., [2]), however, this guess does not hold. One can show the following: for every ξ , k and $\lambda > \lambda_0(\xi, k)$, there exist initial data X_0 such that $m_N(X_0) = O(N^{-\frac{1}{2}})$ and such that Langevin dynamics started from X_0 weakly recovers in order 1 time.

3. Preliminaries: Regularity theory and stochastic analysis in high dimensions.

Throughout the paper we will make frequent use of certain uniform Sobolev-type estimates for H_0 , developed in the context of spin-glass dynamics, as well as properties of solutions to certain Langevin-type stochastic differential equations. We recall these results in this section. In what follows, for functions f, g we say that $f \lesssim g$ if there is a constant $C > 0$ such that $f \leq Cg$.

3.1. *Regularity theory of spin glasses and the \mathcal{G} -norm.* As is often the case in such problems, it will be important to understand the regularity of the related Hamiltonians. It turns out that, in high-dimensional analysis problems, one needs to define these Sobolev spaces carefully, as the scaling of the norms in the dimension is often crucial to the problem at hand.

With this in mind, let us recall the \mathcal{G}^k -norm, which provides a (topologically) equivalent norm on the usual Sobolev space, $W^{k,\infty}$, but which is better suited to high-dimensional problems as well as the related \mathcal{G} -norm regularity of H_0 , established in [7] which will be crucial to our analysis.

DEFINITION 3.1. A function $f : \mathcal{S}^N \rightarrow \mathbb{R}$ is in the space $\mathcal{G}_K^k(\mathcal{S}^N)$, if

$$\|f\|_{\mathcal{G}_K^k} := \sum_{0 \leq \ell \leq k} K^{\ell/2} \|\nabla^\ell f\|_{\text{op}} < \infty.$$

Here, $\|\nabla^k f\|_{\text{op}}(x)$ denotes the natural operator norm when $\nabla^k f$ is viewed as a k -form acting on the k -fold product of the tangent space $T_x \mathcal{S}^N$. Throughout the paper, unless otherwise specified, $\|\nabla^k f\|$ will denote this norm.

REMARK 3.2. By the equivalence of norms on finite dimensional vector spaces, for fixed N , this space is equivalent to the canonical Sobolev space $W^{k,\infty}$, which is defined using Frobenius norms. We use the operator norm instead for the following reason. For $n \geq 2$, we need to bound the operator norms of random tensors. It is well known that, for such random tensors, there is a marked difference in the scaling of the Frobenius and operator norms in the dimension (see, e.g., [38]).

We let $\mathcal{G}^n(\mathcal{S}^N) = \mathcal{G}_N^n(\mathcal{S}^N)$ denote the special case $K = N$ which is chosen precisely such that the scaling in N of the \mathcal{G}^n -norm is independent of n . Namely, in [7] it was shown that, for every n , the \mathcal{G}^n -norm of H_0/N is order 1 in N . Recall $\xi(t) = \sum a_p t^p$ from (1.3) and that $\xi(1 + \epsilon) < \infty$ implies that $H_0 \in C^\infty(\mathcal{S}^N)$.

THEOREM 3.3 ([7], Theorem 3.3). For every n , there exist $K(\xi, n)$, $c(\xi, n) > 0$ such that H_0/N is in \mathcal{G}^n uniformly in N with high probability: for every $r > 0$,

$$(3.1) \quad \mathbb{P}(\|H_0\|_{\mathcal{G}^n} \geq KN + r) \lesssim \exp(-cr^2/N).$$

REMARK 3.4. The result was stated there for H_0 with only one nonzero a_p , that is, the p -spin model; however, as observed in [7], Remark 3.4, it easily extends to this setting by Borell’s inequality and the fact that, in that case, the corresponding $K(p, n)$ is of at most polynomial growth in p .

As further motivation for the definition of the norm \mathcal{G} , specifically with $K = N$, we note here the following easy observations which are useful in bounding the regularity of observables with respect to Langevin-type operators; we call these the *ladder relations* for \mathcal{G} :

LEMMA 3.5 (Ladder relations). For every $n \geq 1$, there exists $c(n)$ such that for every N ,

$$(3.2) \quad \|\Delta\|_{\mathcal{G}^n \rightarrow \mathcal{G}^{n-2}} \leq 1, \quad \text{and} \quad \|\langle \nabla g, \nabla \cdot \rangle\|_{\mathcal{G}^n \rightarrow \mathcal{G}^{n-1}} \leq \frac{c}{N} \|g\|_{\mathcal{G}^n}.$$

In particular, if $L = \Delta + \langle \nabla g_1, \nabla \cdot \rangle$ and $A = \langle \nabla g_2, \nabla \cdot \rangle$ for some g_i that satisfy $\|g_1\|_{\mathcal{G}^{2\ell}} \leq c_1 N$ and $\|g_2\|_{\mathcal{G}^1} \leq c_2 N$, then there are $\tilde{c}_1(n, c_1) > 0$ and $\tilde{c}_2(n, c_1, c_2) > 0$ such that

$$(3.3) \quad \|L^\ell\|_{\mathcal{G}^{n+2} \rightarrow \mathcal{G}^n} \leq \tilde{c}_1, \quad \text{and} \quad \|AL^\ell\|_{\mathcal{G}^{2\ell+1} \rightarrow L^\infty} \leq \tilde{c}_2.$$

PROOF. The first result in (3.2) follows from the fact that traces commute with covariant derivatives. Indeed, for f smooth observe that

$$|\nabla^k \Delta f| = |\text{tr}(\nabla^{k+2} f)| \leq N |\nabla^{k+2} f|$$

so that $\|\Delta f\|_{\mathcal{G}^n} \leq \|f\|_{\mathcal{G}^{n+2}}$.

To see the second inequality in (3.2), observe that for every $h \in \mathcal{G}^n$,

$$\begin{aligned} \|\langle \nabla g, \nabla f \rangle\|_{\mathcal{G}^{n-1}} &= \sum_{0 \leq k \leq n-1} N^{k/2} \|\nabla^k \langle \nabla g, \nabla h \rangle|_{op}\|_{L^\infty} \\ &\leq \sum_{0 \leq k \leq n-1} N^{k/2} \sum_{0 \leq \ell \leq k} \binom{k}{\ell} \|\langle \nabla^{\ell+1} g, \nabla^{k-\ell+1} h \rangle|_{op}\|_{L^\infty} \\ &\leq \sum_{0 \leq \ell \leq k \leq n-1} N^{k/2} \binom{k}{\ell} \frac{\|g\|_{\mathcal{G}^n}}{N^{(\ell+1)/2}} \frac{\|h\|_{\mathcal{G}^n}}{N^{(k-\ell+1)/2}} \leq \frac{c(n)}{N} \|g\|_{\mathcal{G}^n} \|h\|_{\mathcal{G}^n}, \end{aligned}$$

where the third inequality follows by Cauchy–Schwarz and the definition of \mathcal{G}^n . Equation (3.3) then follows directly from (3.2). \square

As a result of the above and explicit calculation, one also sees the following:

COROLLARY 3.6. *For every n , there exist $K(\xi, n)$ and $c(\xi, n) > 0$ such that for every $r > 0$,*

$$\begin{aligned} \mathbb{P}(\|\Delta H_0\|_{\mathcal{G}^n} \geq KN + r) &\leq \exp(-cr^2/N) \quad \text{and} \\ \mathbb{P}(\|\nabla H_0\|^2_{\mathcal{G}^n} \geq KN + r) &\leq \exp(-cr/N). \end{aligned}$$

Finally, an explicit computation also shows that $\phi = \phi_k$ always lives in the space \mathcal{G}^m for any m . In particular, for every k and every n , there exists $K(n, k) > 0$ such that

$$(3.4) \quad \|m_N\|_{\mathcal{G}^n} \leq n \quad \text{and} \quad \|\phi\|_{\mathcal{G}^n} \leq K.$$

3.2. The Langevin operator and existence of the martingale solution. Let us now recall some elementary results from stochastic analysis. For a function $g(x)$, we always let $g(t) = g(X_t)$ denote its evolution under the Langevin dynamics (2.1). We also let M_t^g denote the martingale part of this evolution,

$$M_t^g = g(t) - g(0) - \int_0^t Lg(s) ds.$$

Observe that M_t^g is well defined as the martingale problem for L given by (2.3) is well posed. Moreover, its quadratic variation satisfies $[M_t^g] \leq t \|\nabla g\|_{L^\infty}^2 \leq \frac{t}{N} \|g\|_{\mathcal{G}^1}^2$.

Let us now recall the following elementary estimate. Suppose that g is smooth and $\|g\|_{\mathcal{G}^1} \leq K$, then the exponential moment of M_t^g satisfies $E_{Q_x}[\exp(\lambda M_t^g)] \leq \exp(K\lambda^2 t / (2N))$. Thus, by Doob’s maximal inequality and optimization in λ yields

$$(3.5) \quad \sup_x Q_x \left(\sup_{t \leq T} |M_t^g| \geq \epsilon \right) \lesssim \exp\left(-\frac{N\epsilon^2}{2K^2 T}\right).$$

As we will frequently use the following estimate, we note that, in the case that $m(x) = x_1/\sqrt{N}$, one has by (3.4) and Doob’s maximal inequality, that there is a universal K such that for every $\gamma, T > 0$, and every N ,

$$(3.6) \quad \sup_x Q_x \left(\sup_{t \leq T} |M_t^m| \geq \frac{\gamma}{\sqrt{N}} \right) \leq K \exp\left(-\frac{\gamma^2}{KT}\right).$$

We also define here the following notation which is used throughout the paper. Let \mathcal{F}_1 be given by

$$(3.7) \quad \mathcal{F}_1 = Lm = \beta\lambda km^{k-1}(1 - m^2) + L_{0,\beta}m,$$

and recall that $L_{0,\beta}m = \Delta m - \beta \langle \nabla H_0, \nabla m \rangle = -m - \beta \langle \nabla H_0, \nabla m \rangle$.

4. On weak and strong recovery. As mentioned in the [Introduction](#), there are two main notions of recovery that we study in this paper, weak and strong recovery. In this section we discuss the relationship between these. We prove that weak recovery implies strong recovery in the diverging λ regime. We then show that, depending on the rate of divergence in λ , there is a certain related radius of correlations, r_N , which is $o(1)$, such that one can weakly recover in order 1 time from *every* initial point with correlation greater than r_N . This reduces the difficulty of proving our recovery theorems to showing that the dynamics “escapes the equator.” We end the section observing the stability of weak recovery.

4.1. *Weak recovery implies strong recovery.* We show that, as long as λ is diverging, weak recovery with Langevin dynamics implies strong recovery. In the following we let $\Lambda > 0$ be such that $\|H_0\|_{\mathcal{G}^1} \leq \Lambda N$ eventually \mathbb{P} -almost surely. Recall that such a Λ exists by [Theorem 3.3](#). In this section, for any $\theta \in [-1, 1]$, we let τ_θ denote the first hitting time for the set

$$\{x \in \mathcal{S}^N : m_N(x) = \theta\}.$$

LEMMA 4.1. *Fix k and $\beta > 0$. For every $\epsilon > 0$ and every sequence $\lambda_N \rightarrow \infty$, there exists $T_0 > 0$ such that for all $T \geq T_0$,*

$$\lim_{N \rightarrow \infty} \inf_{x: m(x) \geq \epsilon} Q_x \left(\inf_{t \in [T_0, T]} m(t) \geq (1 - \epsilon) \right) = 1 \quad \mathbb{P}\text{-a.s.}$$

PROOF. Fix $\epsilon > 0$, and suppose that λ_N is any diverging sequence. Let $\tau_{\epsilon/2}$ be the first hitting time of $\{m(x) = \epsilon/2\}$. We wish to show that there exists T_0 such that, uniformly over all $m_N(X_0) \in (\epsilon, 1 - \epsilon/2)$, we have that $\tau_{1-\epsilon/2} \leq T_0$ and for every T , $\tau_{\epsilon/2} \geq T$ with Q_x -probability $1 - o(1)$. For such initial data, eventually, \mathbb{P} -almost surely, \mathcal{F}_1 from (3.7) satisfies

$$\mathcal{F}_1(t) \geq 2^{-k+1} \beta \lambda k \epsilon^{k-1} (\epsilon - \epsilon^2/4) - \|L_{0,\beta} m\|_{L^\infty} \geq 2^{-k+1} \beta \lambda k \epsilon^{k-1} (\epsilon - \epsilon^2/4) - 1 - \beta \Lambda,$$

for all $t \leq \tau_{\epsilon/2} \wedge \tau_{1-\epsilon/2}$. Thus, for any $\epsilon, k, \beta > 0$, there is a $\lambda_0 > 0$ such that for every $\lambda > \lambda_0$, $\mathcal{F}_1(t) \geq c_\lambda > 0$ for all $t \leq \tau_{\epsilon/2} \wedge \tau_{1-\epsilon/2}$. Applying (3.6) with $\gamma = \sqrt{N}\epsilon/2$, we see that, for every T ,

$$\inf_{x: m(x) \in (\epsilon, 1-\epsilon/2)} Q_x \left(m(t) \geq \frac{\epsilon}{2} + c_\lambda t \text{ for all } t \leq \tau_{1-\epsilon/2} \wedge T \right) \geq 1 - K e^{-N\epsilon^2/KT},$$

for some universal $K > 0$. As a consequence, for, say, $T_0(\epsilon, \beta, \lambda) = c_\lambda^{-1} > 0$, we have that $\tau_{1-\epsilon/2} \leq T_0$. By similar reasoning, for every T ,

$$\inf_{x: m(x) \geq 1-\epsilon/2} Q_x (m(t) \geq 1 - \epsilon \text{ for all } t \leq T) \geq 1 - K e^{-N\epsilon^2/KT}.$$

The strong Markov property and a union bound over the above two estimates then implies the desired. \square

4.2. *Weak recovery from microscopic scales.* By a similar argument to the preceding, one can show that, in this regime, weak recovery occurs as soon as X_t has crossed a certain microscopic correlation. More precisely, we obtain the following:

THEOREM 4.2. *Fix $k, \beta > 0$ and $\lambda_N \rightarrow \infty$. There exist $\gamma_0, C > 0$ such that if $r_N \rightarrow 0$ satisfies*

$$\lambda_N r_N^{k-1} = \gamma \quad \text{and} \quad r_N \geq \frac{C}{N^{1-\delta}},$$

for some $\delta > 0$ and some $\gamma \geq \gamma_0$, then for every $\eta > 0$, there is a $T_0 > 0$ such that for all T ,

$$\lim_{N \rightarrow \infty} \inf_{x: m_N(x) \geq r_N} Q_x \left(\inf_{t \in [T_0, T]} m_N(t) \geq (1 - \eta) \right) = 1 \quad \mathbb{P}\text{-a.s.}$$

The proof of Theorem 4.2 follows from Lemma 4.1 and the following lemma showing that the initial boost when $m_N(X_0) \geq r_N$ is sufficiently large to ensure weak recovery.

LEMMA 4.3. Fix k, β . Suppose that $r_N \rightarrow 0$ with $\lambda_N r_N^{k-1} = \gamma > 0$ and that γ, θ , and r_N satisfy

$$-\theta + \frac{k}{2^{k-1}} \gamma \beta (1 - \theta^2) - \beta \Lambda =: c_\gamma > \theta, \quad \text{and} \quad c_\gamma r_N \geq \frac{1}{N^{1-\delta}}$$

for some $\delta \in (0, \frac{1}{2})$. Then, \mathbb{P} -almost surely,

$$\lim_{N \rightarrow \infty} \inf_{x: m(x) \in (r_N, \theta)} Q_x(\tau_\theta \leq 1) = 1.$$

PROOF. Let $\tau_{r_N/2}$ be the first hitting time of the set $\{x : m(x) = r_N/2\}$. We wish to show that for all X_0 with $m(X_0) \in (r_N, \theta)$, we have $\tau_{r_N/2} \geq \tau_\theta$ and moreover, $\tau_\theta \leq 1$. We first claim that $\tau_{r_N/2} \geq \tau_\theta \wedge 1$. To see this, observe that eventually \mathbb{P} -almost surely, \mathcal{F}_1 from (3.7) satisfies

$$\mathcal{F}_1(t) \geq -\theta + \frac{k}{2^{k-1}} \gamma \cdot \beta (1 - \theta^2) - \beta \Lambda = c_\gamma > 0,$$

for $t \leq \tau_{r_N/2} \wedge \tau_\theta$. By (3.6), applied with $\gamma = \sqrt{N} r_N/2$, it follows that for some universal $K > 0$,

$$(4.1) \quad \inf_{x: m(x) \in (r_N, \theta)} Q_x \left(m(t) \geq \frac{r_N}{2} + c_\gamma t \text{ for all } t \leq \tau_{r_N/2} \wedge \tau_\theta \wedge 1 \right) \geq 1 - K e^{-N^{2\delta}/4K}.$$

Since r_N is positive and $c_\gamma > \theta$, we deduce that

$$\inf_{x: m(x) \in (r_N, \theta)} Q_x(\tau_\theta \leq \tau_{r_N/2} \wedge 1) \geq 1 - K e^{-N^{2\delta}/4K}$$

which implies the desired. \square

PROOF OF THEOREM 4.2. Let $\theta > 0$ be sufficiently small. We will first show that from any initial data satisfying $m_N(X_0) \in (r_N, 1 - \theta)$, we have $\tau_{1-\theta} \leq T_0$ for some T_0 with high Q_x -probability. Suppose first that $r_N \leq m_N(0) < \theta$. In this case, Lemma 4.3 implies that \mathbb{P} -a.s.,

$$\lim_{N \rightarrow \infty} \inf_{x: m(x) \in (r_N, \theta)} Q_x(\tau_\theta \leq 1) = 1.$$

By the strong Markov property for X_t , it remains to consider the case that $m_N(0) \geq \theta$ for some $\theta > 0$ sufficiently small. By Lemma 4.1 we see that for every $\theta > 0$, there exists T_0 such that for every $T \geq T_0$, \mathbb{P} -a.s.,

$$\lim_{N \rightarrow \infty} \inf_{x: m(x) \geq \theta} Q_x \left(\inf_{t \in [T_0, T]} m(t) \geq (1 - \theta) \right) = 1,$$

yielding the result. \square

4.3. *Stability of weak recovery.* Lemma 4.1 showed that, for λ_N diverging and every $\epsilon > 0$, if $m(t)$ ever exceeds $1 - \epsilon$, then it will remain above $1 - 2\epsilon$ for all sufficiently large (but order 1) times. Here, we show an analogous result in the weak recovery regime when λ is order 1 in N . These results can be used to establish “certificates” for recovery via the Langevin dynamics.

LEMMA 4.4. *Fix $k, \beta > 0$. For every $\epsilon \in (0, \frac{1}{2})$, there exists a $\lambda_0 > 0$ such that for all $\lambda \geq \lambda_0$ and every T, \mathbb{P} -almost surely,*

$$\lim_{N \rightarrow \infty} \sup_{x: m(x) \geq 2\epsilon} Q_x \left(\inf_{t \in [0, T]} m_N(t) \leq \epsilon \right) = 0.$$

PROOF. The proof follows by an analogous strategy Lemma 4.1. Fix $k, \beta > 0, \epsilon > 0$ and any $T > 0$. Let τ_ϵ and $\tau_{1-\epsilon}$ be the first hitting times of $\{x : m(x) = \epsilon\}$ and $\{x : m(x) = 1 - \epsilon\}$, respectively, and notice that for every x such that $m_N(x) \in (\epsilon, 1 - \epsilon)$,

$$\mathcal{F}_1(x) \geq -1 + k\beta\lambda\epsilon^{k-1}(2\epsilon - \epsilon^2) - \beta\Lambda.$$

Clearly, there exists a λ_0 sufficiently large such that for all $\lambda > \lambda_0$, the above is positive. Then, by (3.5), there exists $K > 0$ such that for every s ,

$$\sup_x Q_x \left(\sup_{t \leq T} |M_t^m| \geq s \right) \leq K e^{-Ns^2/KT}.$$

Setting $s = \epsilon/2$ and using the fact that $Lm(x) = \mathcal{F}_1(x)$, we see that, for every T , we obtain the desired inequality. \square

5. Recovery for Langevin dynamics under Conditions 1 and 2. We turn now to the proof of the main results of this paper, namely, those from Section 2.2. We first recall an elementary comparison inequality which will be at the heart of our comparison between $m_N(X_t)$ and the related gradient flow. We then provide a stochastic Taylor-type bound that allows us to propagate the regularity of the initial data on order 1 time scales. We then end this section with the proofs of the main theorems.

Observe the following elementary comparison inequality:

LEMMA 5.1. *Let $\gamma > 0$ with $\gamma \neq 1, c > 0$, and $f \in C_{loc}([0, T])$ with $f(0) > 0$. Suppose there exists T such that f satisfies the integral inequality,*

$$f(t) \geq a + \int_0^t c f^\gamma(s) ds,$$

for every $t \leq T$ and some $a > 0$. Then, for $t \geq 0$ satisfying $(\gamma - 1)ca^{\gamma-1}t < 1$, we have

$$f(t) \geq a(1 - (\gamma - 1)ca^{\gamma-1}t)^{-\frac{1}{\gamma-1}}.$$

If the integral inequality holds in reverse, then the corresponding upper bound holds. If $\gamma > 1$, then $T \leq t_{**}$, where $t_{**} = [(\gamma - 1)ca^{\gamma-1}]^{-1}$ is called the blow-up time.

REMARK 5.2. Observe that the case $\gamma = 1$ is excluded. This case is the well-known Gronwall’s inequality. In our setting this corresponds to the critical regime where $k = 2$.

PROOF. We prove the lower bound, as the upper bound is identical. Furthermore, it suffices to take $f \in C_{loc}^1([0, T])$, as otherwise it suffices to bound $g(t) = f(0) + \int_0^t c f^\gamma ds$. Since $\gamma > 0, F(x) = x^\gamma$ is locally Lipschitz on (ϵ, ∞) for any $\epsilon > 0$. Thus, the equation

$$\begin{cases} \dot{h} = ch^\gamma, \\ h(0) = a \end{cases}$$

has as unique solution

$$h(t) = f(0)(1 - (\gamma - 1)cf^{\gamma-1}(0)t)^{-\frac{1}{\gamma-1}}$$

until a blow up time t_* given by the solution to $(\gamma - 1)cf^{\gamma-1}(0)t_* = 1$. We use the convention that if $t_* < 0$, then we take $t_* = \infty$, and the solution is global-in-time. Thus, since $f(0) \geq h(0)$, f satisfies $f(t) \geq h(t)$ as desired. \square

5.1. *A growth estimate for Langevin dynamics under gradient-type perturbations.* In this section we seek to estimate the growth of well-behaved observables under the evolution of some Markov process whose infinitesimal generator is a perturbation of Langevin dynamics for a sufficiently regular Hamiltonian. For every $\delta > 0$, we let $B_\delta = [-N^{-\frac{1}{2}+\delta}, N^{-\frac{1}{2}+\delta}]$.

THEOREM 5.3. *Let $E \subset \mathcal{S}^N$, L be the infinitesimal generator of an Ito process X_t , f be smooth and $x_0 \in E$ with exit time τ_{E^c} . Suppose that these satisfy the following for some $n \geq 1$:*

1. *L is a differential operator of the form $L = L_0 + a(x)A$ where:*
 - (a) *A is of a gradient type: $A = \langle \nabla \psi, \nabla \cdot \rangle$ for some $\psi \in C^\infty$ with $\|\psi\|_{\mathcal{G}^{2n}} \leq c_0N$,*
 - (b) *$a \in C(\mathcal{S}^N)$, with $\|a\|_\infty \leq c_1$,*
 - (c) *$L_0 = \Delta + \langle \nabla U, \nabla \cdot \rangle$ for some $U \in C^\infty$ with $\|U\|_{\mathcal{G}^{2n}} \leq c_2N$.*
2. *f is smooth with $\|f\|_{\mathcal{G}^{2n}} \leq c_3$.*
3. *There exist $\delta > 0$ such that the initial x_0 satisfies $L_0^\ell f(x_0) \in B_\delta$ for every $0 \leq \ell \leq n - 1$.*
4. *There is an $\epsilon \in (0, 1)$ and a t_0 , possibly depending on ϵ , such that for any $t \leq \tau_{E^c} \wedge t_0$,*

$$\int_0^t |a(X_s)| ds \leq \epsilon |a(X_t)|.$$

Then, there exists $K > 0$ depending only on c_i and δ such that, for every T_0 ,

$$(5.1) \quad |f(X_t)| \leq K \left(\frac{N^\delta}{\sqrt{N}} \sum_{\ell=0}^{n-1} (1 + t^\ell) + t^n + \frac{1}{1 - \epsilon} \int_0^t |a(X_s)| ds \right)$$

for all $t \leq \tau_{E^c} \wedge t_0 \wedge T_0$ with Q_{x_0} -probability $1 - O(\exp(-N^{2\delta}/(KT_0)))$.

If instead they satisfy the above with item (3) replaced by

- (3') *There exists $t_1, \delta > 0$ such that the initial x_0 satisfies $e^{tL_0} f(x_0) \in B_\delta$ for every $t < t_1$,*

then the same (5.1) would hold for every for all $t \leq \tau_{E^c} \wedge t_0 \wedge t_1 \wedge T_0$.

PROOF. For any function g , let $g(t) = g(X_t)$. We begin by claiming that f has expansion

$$(5.2) \quad \begin{aligned} f(t) = & f(0) + M_t^f + \sum_{\ell=1}^{n-1} \int_0^t \dots \int_0^{t_{\ell-1}} L_0^\ell f(0) + M_{t_\ell}^{L_0^\ell f} dt_\ell \dots dt_1 \\ & + \int_0^t \dots \int_0^{t_{n-1}} L_0^n f(t_n) dt_n \dots dt_1 \\ & + \sum_{\ell=1}^n \int_0^t \dots \int_0^{t_{\ell-1}} a(t_\ell) A L_0^{\ell-1} f(t_\ell) dt_\ell \dots dt_1. \end{aligned}$$

The proof is by induction. The base case, $n = 1$, is simply the definition of M_t^f . Assume that the result holds in the n th case. Then, for the $n + 1$ st case we may expand the second-to-last term as

$$\begin{aligned} & \int_0^t \cdots \int_0^{t_{n-1}} L_0^n f(t_n) dt_n \cdots dt_1 \\ &= \int_0^t \cdots \int_0^{t_{n-1}} L_0^n f(0) + M_{t_n}^{L_0^n f} dt_n \cdots dt_1 \\ & \quad + \int_0^t \cdots \int_0^{t_n} L_0^{n+1} f(t_{n+1}) + a(t_{n+1})AL_0^n f(t_{n+1}) dt_{n+1} \cdots dt_1 \end{aligned}$$

by the definition of $M^{L_0^n f}$ and the splitting $L = L_0 + a(x)A$. Combining the terms yields the desired expression, by induction.

To obtain (5.1) under assumptions (1)–(4), bound the absolute values, term-by-term, in (5.2). We first observe that by the second assumption and the ladder relations (3.3),

$$\|L_0^\ell f\|_{\mathcal{G}^{2n-2\ell}} \leq \|f\|_{\mathcal{G}^{2n}} \leq c_3.$$

In particular, $\|\nabla L_0^\ell f\|_\infty \leq c_3/\sqrt{N}$. Thus, by (3.5), with \mathcal{Q}_{x_0} -probability $1 - O(\exp(-N^{2\delta}/(KT_0)))$,

$$\sup_{t \leq T_0} |M_t^{L_0^\ell f}| \leq \frac{N^\delta}{\sqrt{N}}.$$

Meanwhile, by the third assumption we can bound $|L_0^\ell f(0)|$ for every $0 \leq \ell \leq n - 1$ by $N^{-\frac{1}{2}+\delta}$. Integrating these two inequalities implies that the first line of (5.2) is upper bounded in absolute value by the first sum in (5.1). The second term in (5.1) bounds the integral of $L_0^n f(t_n)$ by the ladder relation (3.2). For the last term in (5.1), note that, by Lemma 3.5, $|AL_0^\ell f| \leq c$ for some $c > 0$; thus, the bound follows by applying the fourth assumption to obtain

$$\int_0^t \cdots \int_0^{t_{\ell-1}} |a(X_{t_\ell})| dt_\ell \cdots dt_1 \leq \epsilon^\ell |a(X_t)|.$$

To replace assumption (3) by (3'), the same argument applies except we bound the third term in the first line of (5.2) instead as follows. By Taylor expanding and applying (3.2) again,

$$\left| e^{tL_0} f(x_0) - \sum_{\ell=0}^{n-1} L_0^\ell f(x_0) \frac{t^\ell}{\ell!} \right| \leq \|L_0^n f\|_\infty t^n \leq c_3 t^n.$$

Combining this with assumption (3'), gives the desired bound. \square

5.2. *Proof of Theorem 2.5 for finite β .* Fix $n \geq 1$ and $T_0 > 0$. For any $\epsilon > 0$, let τ_ϵ denote the hitting time of $\{x : m(x) \geq \epsilon\}$. For a fixed μ_N , every $\epsilon, \gamma > 0$, and every sequence $A_N \subset \mathcal{S}^N$,

$$\begin{aligned} (5.3) \quad & \int \mathcal{Q}_x(\tau_\epsilon \geq T_0) d\mu_N \\ & \leq \mu_N(A_N^c) + \mu_N(x_1 < \gamma) + \int \mathcal{Q}_x(\tau_\epsilon \geq T_0) \mathbf{1}(A_N \cap \{x_1 \geq \gamma\}) d\mu_N. \end{aligned}$$

If we take the limit superior in N , the first term vanishes almost surely by Condition 1 at level n upon choosing $A_N = E_{n,\delta,N}$ and by the weak form of Condition 1 at level ∞ , upon choosing $A_N = \tilde{E}_{T,\delta,N}$ for any order 1 T . For any fixed $\eta > 0$, it follows by Condition 2

that there is a γ sufficiently small such that the second term is less than $\eta/2$. Finally, by the following theorem, for any such choice of γ sufficiently small there is an $\epsilon_0 > 0$ such that for every $\epsilon < \epsilon_0$, the third term is less than $\eta/2$ as well. Thus, the process reaches $m_N(X_t) \geq 2\epsilon$ in time T_0 with probability at least $1 - \eta$. To conclude the proof, note that by Lemma 4.4, after $\tau_{2\epsilon}$, $m_N(X_t)$ remains above ϵ . \square

THEOREM 5.4. *Let $k < 2$. For any $\epsilon > 0$, let τ_ϵ denote the hitting time of the set $\{x : m(x) \geq \epsilon\}$. For every $\beta > 0$ and $n \geq 1$, there exist $\lambda_0, \epsilon_0, c, K > 0$ such that for every $\lambda > \lambda_0$, every $\gamma \in (0, 1)$ and every sequence $\mu_N \in \mathcal{M}_1(S^N)$, the following holds:*

- if $k < k_c(n)$ and $\epsilon < \epsilon_0$, then for every $\delta > 0$ sufficiently small,

$$\int Q_x(\tau_\epsilon \geq c\lambda^{-1}\epsilon^{2-k})\mathbf{1}(E_{n,\delta,N} \cap \{x_1 \geq \gamma\}) d\mu_N(x) \leq K \exp\left(-\frac{\lambda\gamma^2}{K\epsilon^{2-k}}\right),$$

eventually \mathbb{P} -almost surely.

- if $k < k_c(\infty)$ and $\epsilon < \epsilon_0$, then for every $T > c\lambda^{-1}\epsilon^{2-k}$ and sufficiently small $\delta > 0$,

$$\int Q_x(\tau_\epsilon \geq c\lambda^{-1}\epsilon^{2-k})\mathbf{1}(\tilde{E}_{T,\delta,N} \cap \{x_1 \geq \gamma\}) d\mu_N(x) \leq K \exp\left(-\frac{\lambda\gamma^2}{K\epsilon^{2-k}}\right),$$

eventually \mathbb{P} -almost surely.

PROOF. The proofs of the two statements are identical up to minor modifications. We will focus on the proof of the first and explain the modifications necessary for the second throughout the proof wherever the two proofs differ. Throughout the following, for the second statement we let n denote the smallest n such that $k < k_c(n)$. In particular, if $k = k_c(\ell)$, then $n = \ell + 1$.

Let $\delta > 0$, later to be chosen sufficiently small, and let $A' = A'(\gamma, \delta)$ denote the event that the initial data, $x \sim \mu_N$, is in $E_{n,\delta,N} \cap \{x_1 \geq \gamma\}$ (resp., in $\tilde{E}_{T,\delta,N} \cap \{x_1 \geq \gamma\}$). Let $\theta = 2 - k$, so that, by assumption, $\theta > \frac{1}{n}$. Finally, without loss of generality take $\epsilon < 1/2$.

Let \mathcal{T}_L be the first hitting time of the bad set

$$\left\{x : |L_0m(x)| > \frac{1}{2}\beta k\lambda m^{k-1}(x)\right\}.$$

On the event A' , by continuity of X_t , $\mathcal{T}_L > 0$. Furthermore, by (3.6) it follows that there is a K_0 such that for every γ and T ,

$$Q_x\left(\sup_{t \leq T} |M_t^m| > \frac{\gamma}{2\sqrt{N}}\right) \leq K_0 \exp\left(-\frac{\gamma^2}{K_0T}\right).$$

For T_0 positive, to be chosen later (see (5.6) proportionally to ϵ^{2-k}/λ , let

$$A = A(\gamma, \delta, T_0) = A'(\gamma, \delta) \cap \left\{\sup_{t \leq T_0} |M_t^m| \leq \frac{\gamma}{2\sqrt{N}}\right\}.$$

For the remainder of the proof, we restrict our attention to the event A . By definitions of \mathcal{F}_1 (3.7) and \mathcal{T}_L , for every $t \leq \mathcal{T}_L \wedge \tau_\epsilon$,

$$\frac{1}{2}\lambda\beta km^{1-\theta}(t) \leq \mathcal{F}_1(t) \leq \frac{3}{2}\lambda\beta km^{1-\theta}(t).$$

Then, $m(t)$ satisfies the integral inequality

$$(5.4) \quad \frac{1}{2}m(0) + \int_0^t \lambda c_2 m^{1-\theta}(s) ds \leq m(t) \leq \frac{3}{2}m(0) + \int_0^t \lambda c_1 m^{1-\theta}(s) ds,$$

for $t \leq \tau_\epsilon \wedge \mathcal{T}_L \wedge T_0$, uniformly in λ , for some $c_1, c_2 > 0$ which depend on β and k . Thus, by Lemma 5.1 we have that

$$(5.5) \quad g_2(t) := \left[\left(\frac{1}{2}m(0) \right)^\theta + \theta\lambda c_2 t \right]^{1/\theta} \leq m(t) \leq \left[\left(\frac{3}{2}m(0) \right)^\theta + \theta\lambda c_1 t \right]^{1/\theta} =: g_1(t),$$

for all $t \leq \mathcal{T}_L \wedge \tau_\epsilon \wedge T_0$.

As a consequence of this comparison inequality, it suffices to show that $\mathcal{T}_L > T_0$, where we choose T_0 to solve the equation

$$(5.6) \quad \epsilon = \frac{1}{2}(\theta c_2 \lambda T_0)^{1/\theta}, \quad \text{so that } T_0 = \frac{(2\epsilon)^{2-k}}{(2-k)c_2 \lambda}.$$

To this end, let us first check that $f(t) = L_0 m(t)$ satisfies the conditions of Theorem 5.3. Indeed, if we let $U = H_0$, $\psi(x) = \sqrt{N}x_1$, and $a(x) = \beta k \lambda m^{k-1}(x)$, then item (1) is satisfied eventually \mathbb{P} -a.s. for every $n \geq 1$ by Theorem 3.3 and (3.4). Item (2) is satisfied by (3.2) combined with (3.4). Item (3) (resp., (3') with the choice $t_1 = T$) follows by assumption on the initial data, namely, the event $A'(\gamma, \delta)$. It remains to check the fourth condition. Note that by the integral inequality, (5.4), and the comparison (5.5), whenever $t \leq \mathcal{T}_L \wedge \tau_\epsilon \wedge T_0$,

$$\int_0^t |a(X_s)| ds \leq \frac{1}{c_2} m(t) \leq \frac{1}{c_2} \left[\left(\frac{3}{2}m(0) \right)^\theta + \theta\lambda c_1 t \right]^{1/\theta}.$$

Observe now that, if we let $T_1 = (\theta\lambda c_2)^{-1} (c_2/(2c_1))^{-1/\theta}$, then for every $t < T_1$, as long as $\beta k \lambda \geq 1$,

$$\frac{1}{c_2} \left(\theta\lambda c_1 t + \left(\frac{3}{2}m(0) \right)^\theta \right)^{1/\theta} < \frac{\beta k \lambda}{2} \left(\theta\lambda c_2 t + \left(\frac{1}{2}m(0) \right)^\theta \right)^{1-\frac{\theta}{2}} \leq \frac{1}{2} a(t),$$

where the second inequality follows from (5.5). Thus, the fourth condition is satisfied for every $t \leq \mathcal{T}_L \wedge \tau_\epsilon \wedge T_0 \wedge T_1$.

As a consequence, applying Theorem 5.3 on the event A , for every $\delta > 0$

$$(5.7) \quad |f(t)| \leq K \left(\frac{N^\delta}{\sqrt{N}} \sum_{\ell=0}^{n-1} (1+t^\ell) + t^n + 2 \int_0^t |a(X_s)| ds \right),$$

for $t \leq \tau_\epsilon \wedge \mathcal{T}_L \wedge T_0 \wedge T_1$ with \mathcal{Q}_x -probability at least $1 - 2K \exp(-N^{2\delta}/KT_0)$, for some $K > 0$, depending only on k, n and δ .

With this in hand, we aim to show the desired lower bound on \mathcal{T}_L , namely, that $\mathcal{T}_L \geq 2\epsilon^\theta / (\theta c_2 \lambda)$. It suffices to show that for all $t \leq T_0$, each term in (5.7) is bounded above by $\beta k \lambda m^{k-1}(t) / (4n + 2)$. Begin by observing that for every $0 \leq s \leq n - 1$, for large enough N ,

$$\frac{KN^\delta}{\sqrt{N}} (1+t^s) \leq \frac{\beta k \lambda}{4n+2} m^{1-\theta}(0) < \frac{\beta k \lambda}{4n+2} g_2^{1-\theta}(t),$$

provided $1+t^s = o(N^{\frac{\theta}{2}-\delta})$, which certainly holds for all $t \leq T_0 \wedge T_1$, provided δ is small enough since T_0, T_1 are order 1 in N . Similarly, there exists an $\epsilon_0 > 0$ (depending only on n, θ, K, c_2 , and not λ) such that if $\lambda \geq 1$, then for every $\epsilon < \epsilon_0$ and for every $t \leq \tau_\epsilon$,

$$2K \int_0^t |a(s)| ds \leq 2K \cdot \frac{\beta k}{c_2} m(t) \leq \frac{\beta k \lambda}{4n+2} m^{k-1}(t).$$

Finally, $h(t) = K t^n$ satisfies

$$\begin{aligned} \frac{d^\ell}{dt^\ell} h(0) &= 0 \leq \frac{d^\ell}{dt^\ell} g_2(0) \quad \text{for } \ell \leq n-1, \quad \text{and,} \\ \frac{d^n}{dt^n} h(t) &= K \cdot n! < C \lambda^{n+1} \left(\theta\lambda c_2 t + \left(\frac{1}{2}m(0) \right)^\theta \right)^{(1-n\theta)/\theta} = \frac{d^n}{dt^n} \left(\frac{\beta k \lambda}{4n+2} g_2(t) \right), \end{aligned}$$

for some C which depends on $\beta, k, \theta, n,$ and $c_2,$ provided

$$t \leq \frac{1}{2\theta\lambda c_2} \left(\frac{C\lambda^{n+1}}{n!K} \right)^{\theta/(n\theta-1)} = T_2,$$

since $n\theta > 1$ and $m(0) = o(1).$ Thus if we let $t_1 = T_1 \wedge T_2,$ we see that on the event that A holds, $t_1 \wedge \tau_\epsilon \wedge T_0 \leq \mathcal{T}_L.$ For λ sufficiently large, depending only on $\theta, \beta,$ we see that $t_1 = T_1 < T_2$ and that $T_1 = c(\theta, \beta, k)/\lambda$ for some such $c.$ Consequently, there is an $\epsilon_0(\theta, \beta, k)$ (independent of λ) such that if $\epsilon \leq \epsilon_0$ and T_0 satisfies (5.6), then $T_0 \leq T_1$ and as a result, $\tau_\epsilon \leq T_0$ as desired. \square

5.3. *Proof of Theorem 2.4 for finite $\beta.$* Fix $n \geq 1$ and $C, T > 0.$ For any sequence $r_N,$ let τ_{Cr_N} is the hitting time for the set $\{x : m(x) \geq Cr_N\}.$ For any sequence $\mu_N \in \mathcal{M}_1(\mathcal{S}^N),$ any sequence of events $A_N \subset \mathcal{S}^N$ and any sequence r_N it follows that

$$\begin{aligned} \int Q_x(\tau_{Cr_N} \geq T_0) d\mu_N(x) &\leq \mu_N(A_N^c) + \mu_N(x_1 < N^{-\delta}) \\ &+ \int Q_x(\tau_{Cr_N} \geq T_0) \mathbf{1}(A_N \cap \{x_1 \geq N^{-\delta}\}) d\mu_N(x). \end{aligned}$$

We take the limit superior in N on both sides and bound these terms one-by-one.

If μ_N satisfies Condition 1 at level $n,$ then the first term goes to zero \mathbb{P} -a.s. by taking $A_N = E_{n,\delta,N}.$ Otherwise, if μ_N weakly satisfies Condition 1 at level $\infty,$ then the same is true upon taking $A_N = \tilde{E}_{T,\delta,N}.$ If Condition 2 holds, then the second term goes to zero \mathbb{P} -a.s. as well. Otherwise, if Condition 2' holds, then it goes to zero in probability. The third term goes to zero \mathbb{P} -a.s. by the following theorem upon taking $r_N = N^{-\frac{\alpha}{k-1}}$ and $C > 0.$ The result will then follow upon taking $C > 0$ sufficiently large so that we may apply Theorem 4.2 and applying the strong Markov property. \square

THEOREM 5.5. *Let $k > 2, \lambda = N^\alpha, r_N = N^{-\frac{\alpha}{k-1}}$ and $C, \beta > 0.$ Take a sequence $\mu_N \in \mathcal{M}_1(\mathcal{S}^N),$ and let τ_{Cr_N} denote the hitting time of the set $\{x : m(x) \geq Cr_N\}.$ The following hold:*

- For every $n \geq 1,$ if $\alpha > \alpha_c(n),$ then for δ sufficiently small,

$$\int Q_x(\tau_{Cr_N} > 1) \mathbf{1}(E_{n,\delta,N} \cap \{x_1 \geq N^{-\delta}\}) d\mu_N(x) \leq \exp(-cN^\delta),$$

eventually, \mathbb{P} -almost surely.

- If $\alpha > \alpha_0,$ then for δ sufficiently small and $T > 1,$

$$\int Q_x(\tau_{Cr_N} > 1) \mathbf{1}(\tilde{E}_{T,\delta,N} \cap \{x_1 \geq N^{-\delta}\}) d\mu_N(x) \leq \exp(-cN^\delta),$$

eventually, \mathbb{P} -almost surely.

PROOF. The proofs of the two statements are identical up to minor modifications. We will focus on the proof of the first and explain the modifications necessary for the second in parentheses throughout the proof wherever the two proofs differ. Throughout the following for the second statement, we let n denote the smallest n such that $\alpha > \alpha_c(n).$ In particular, if $\alpha = \alpha_c(k),$ then $n = k + 1.$

Let $A' = A'(\delta);$ denote the event that the initial data X_0 is in $E_{n,\delta,N} \cap \{x_1 \geq N^{-\delta}\}$ (resp., in $\tilde{E}_{T,\delta,N} \cap \{x_1 \geq N^{-\delta}\}.$) Let \mathcal{T}_L denote the hitting time of the bad set

$$\left\{ x : |L_0 m(x)| > \frac{1}{2} \beta k \lambda m^{k-1} \right\}.$$

By our assumptions and continuity of $X_t,$ on A' both $\mathcal{T}_L > 0.$

By (3.6) it follows that, for every $\iota > 0$ and $T < N^{-\iota}$,

$$(5.8) \quad \inf_x Q_x \left(\sup_{t \leq T} |M_t^m| \leq N^{-\iota/4} N^{-1/2} \right) \geq 1 - K \exp(-N^{\iota/2}/K).$$

Recalling (3.7) and setting $\iota = 5\delta$, it follows that on the intersection of this event and A' , $m_N(t)$ satisfies the integral inequality

$$(5.9) \quad \frac{1}{2}m(0) + \int_0^t \mathcal{F}_1(s) ds \leq m(t) \leq \frac{3}{2}m(0) + \int_0^t \mathcal{F}_1(s) ds,$$

for all $t < N^{-5\delta}$, eventually, \mathbb{P} -a.s. Let us call the intersection of these events $A = A(\delta)$. In the following we restrict our attention to this event.

By definition of \mathcal{T}_L , there are c_1, c_2 positive which depend only on β and k , such that

$$(5.10) \quad \lambda c_2 m^{k-1}(t) \leq \mathcal{F}_1(t) \leq \lambda c_1 m^{k-1}(t),$$

for $t \leq \mathcal{T}_L$. By (5.9), (5.10) and Lemma 5.1, we then obtain upper and lower bounds on $m(t)$ of the form

$$(5.11) \quad g_2(t) \leq m(t) \leq g_1(t),$$

for any $t \leq \mathcal{T}_L \wedge N^{-5\delta}$, where

$$g_i(t) = a_{N,i} \left(1 - \frac{t}{t_{**}^i} \right)^{-\frac{1}{k-2}}.$$

Here, $a_{N,1} = N^{-\frac{1}{2}+\delta}$, $a_{N,2} = \frac{1}{2}N^{-\frac{1}{2}-\delta}$ and g_i have blow-up times

$$t_{**}^i = [(k-2)\lambda c_i a_{N,i}^{k-2}]^{-1} \quad \text{for } i = 1, 2.$$

Since $k > 2$, these blow-up times satisfy $t_{**}^1 \leq t_{**}^2 < \infty$. In particular, it must be that

$$\mathcal{T}_L \wedge N^{-5\delta} \leq t_{**}^2$$

since $m \leq 1$. Furthermore, if we let $\theta = (k-1)/2 - \alpha$, then, since $t_{**}^2 = O(N^{-\frac{1}{2}+\theta+\delta(k-2)})$ and $\theta < \frac{1}{2}$, we have that $t_{**}^2 < N^{-5\delta}$ provided $\delta < \delta_0(k, \alpha)$ for some $\delta_0(k, \alpha)$.

With the above in hand, we aim to show that $\tau_{Cr_N} \leq \mathcal{T}_L$ on A . Indeed, if this were the case, then we would have, on N , the desired

$$\tau_{Cr_N} \leq t_{**}^2 < 1.$$

Suppose, by way of contradiction, that $\tau_{Cr_N} > \mathcal{T}_L$. It would suffice to show that

$$t_* = (1 - (2C)^{-(k-2)} N^{-(k-2)(\frac{\theta}{k-1}+\delta)}) t_{**}^2 < \mathcal{T}_L,$$

as we would then arrive at a contradiction since, by design, $g_2(t_*) = Cr_N$.

To this end, let us observe the following. First, there is a $\delta_0(k, \alpha)$ such that for every $\epsilon > 0$, $\delta < \delta_0$ and N sufficiently large, on the event $A = A(\delta)$,

$$(5.12) \quad |m(t)| \leq \epsilon \beta k \lambda m^{k-1}(t) \quad \text{for all } t \leq \mathcal{T}_L,$$

since, by (5.11), $m(t)$ is lower bounded by an increasing function for which the inequality holds at time zero.

Second, let us observe that $f(t) = L_0 m$ satisfies the conditions of Theorem 5.3. Indeed, if we let $U = H_0$, $\psi(x) = \sqrt{N}x_1$ and $a(x) = \beta k \lambda m^{k-1}(x)$, then item (1) is satisfied, eventually, \mathbb{P} -a.s. for any $n \geq 1$ by Theorem 3.3. Item (2) for f follows by Theorem 3.3 again, eventually,

\mathbb{P} -a.s. for any $n \geq 1$. Item (3) (respectively, item (3') with $t_1 = T$) follows, by definition, on the event A' . To see the fourth condition, observe that by the lower bound in (5.9) and (5.12),

$$\int_0^t |a(s)| ds \leq \frac{1}{c_2} m(t) \leq \frac{1}{2} \beta k \lambda m^{k-1}(t)$$

for $t \leq \mathcal{T}_L$ provided N is sufficiently large, yielding the fourth condition. Thus, by Theorem 5.3 (choosing $T_0 = 1 \gg t_{**}^2$) there exists $K > 0$, such that for some $c > 0$,

$$(5.13) \quad |f(t)| \leq K \left(\frac{N^\delta}{\sqrt{N}} \sum_{\ell=0}^{n-1} (1 + t^\ell) + t^n + 2 \int_0^t |a(X_s)| ds \right),$$

for $t \leq \mathcal{T}_L \wedge 1$ (in the second case, one should also take the minimum with T), on A , with Q_x -probability $1 - O(\exp(-cN^\delta))$, eventually, \mathbb{P} -a.s. provided $\delta < \delta_0(k, \alpha)$. (We note here that for the second case this inequality still holds at $t > 0$.) Now, in order to show that $\mathcal{T}_L > t_*$, thereby concluding the proof, it suffices to show that every term in the sum (5.13) is less than $\beta k \lambda m^{k-1}(t)/(4n + 4)$ for $t \leq \mathcal{T}_L$.

To this end, suppose first that for some $t \leq \mathcal{T}_L$, some $s \leq n - 1$ and some $K > 0$,

$$\beta k \lambda m(t)^{k-1} \leq (4n + 4) K \frac{N^\delta}{\sqrt{N}} (1 + t^s).$$

Because $g_2(t) \leq m(t)$ by (5.11) and g_2 is increasing, it would follow that

$$N^{-\theta - \delta(k-1)} = \lambda g_2(0)^{k-1} \leq \lambda g_2(t)^{k-1} \leq \frac{2(4n + 4)K}{\beta k} N^{-\frac{1}{2} + \delta},$$

where we use here that $\mathcal{T}_L \leq N^{-5\delta} < 1$. If we choose δ sufficiently small, this is impossible for N sufficiently large. Suppose instead that for some $t \leq \mathcal{T}_L$ and some $K > 0$,

$$\frac{\beta k}{2} \lambda m(t)^{k-1} \leq 2K(4n + 4) \int_0^t |a(s)| ds.$$

Since $t \leq \mathcal{T}_L$, the comparisons (5.9) and (5.10) imply that

$$2 \int_0^t |a(s)| ds = 2 \int_0^t \beta k \lambda m^{k-1}(s) ds \leq \frac{2}{c_1} m(t).$$

Thus, for some $t \leq \mathcal{T}_L$, it would have to be that

$$\beta k \lambda m^{k-2}(t) \leq \frac{2(4n + 4)}{c_1} K.$$

Applying the comparison inequality, (5.11), again yields that on $t \leq t_{**}^2$,

$$\beta k \lambda g_2(t)^{k-2} \leq \frac{2(4n + 4)}{c_1} K.$$

Since $g_2(t)$ is increasing, this would imply that

$$\beta k N^{\frac{1}{2} - \theta - \delta(k-2)} = \beta k \lambda g_2(0)^{k-2} \leq \frac{2(4n + 4)}{c_1} K.$$

As $\theta < \frac{1}{2}$, this cannot happen either for δ sufficiently small (depending on α and k).

Thus, the only remaining way for $\mathcal{T}_L < t_*$ would be that, for some $t < t_*$,

$$G(t) := \beta k \lambda g_2(t)^{k-1} \leq 2(4n + 4) K t^n =: h(t).$$

Observe that for N sufficiently large and $t \leq t_*$,

$$\begin{aligned} \frac{d^\ell}{dt^\ell} h(0) &= 0 < \frac{d^\ell}{dt^\ell} G(0) \quad \text{for } \ell \leq n - 1, \quad \text{and} \\ \frac{d^n}{dt^n} h(t) &= 2(4n + 4)K \cdot n! \leq \frac{d^n}{dt^n} G(t), \end{aligned}$$

where the second inequality follows by an explicit calculation and the fact that $\theta < \frac{1}{2} \frac{n-1}{n}$ as $\alpha > \alpha_c(n)$, as long as δ is sufficiently small. This again yields a contradiction.

Thus, as long as δ was sufficiently small and N sufficiently large, then on the event A , $t_* \leq \mathcal{T}_L$ and, in particular, $\tau_{c_{r_N}} \leq \mathcal{T}_L \leq 1$ with Q_x -probability $1 - O(\exp(-cN^\delta))$ as desired. \square

5.4. *Proofs of Theorem 2.4–2.5 for $\beta = \infty$.* In this section we demonstrate that the above proofs go through with the appropriate modifications in the setting of gradient descent instead of Langevin dynamics. Before we begin this proof we note the following useful remark regarding the local well posedness of X_t^∞ under Condition 2.

REMARK 5.6. We note that in the case $1 < k < 2$ for gradient descent, the equation for X_t^∞ is ill posed on \mathcal{S}^N , as ∇H is only $(k - 1)$ -Hölder for such k . That being said, on $\{x_1 \geq N^{-\delta}\}$, ∇H is locally Lipschitz so that it is locally well defined by the Picard–Lindelöf theorem wherever we work. (In fact, by virtue of the proof one also obtains global in time existence.) Consequently, by the assumption that Condition 2 holds, this subtlety can be avoided since we have interpreted Q_x to be zero on the set $x_1 = 0$ so that the inequality (5.3) still holds.

Recall that in this setup we replace the infinitesimal generator L with $L_\infty = -\langle \nabla H, \nabla \cdot \rangle$ and replace L_0 with $L_{0,\infty} = -\langle \nabla H_0, \nabla \cdot \rangle$. As a consequence, the drift \mathcal{F}_1 for $m_N(x)$ is replaced by

$$\tilde{\mathcal{F}}_1 = L_\infty m = \lambda k m^{k-1} (1 - m^2) + L_{0,\infty} m.$$

Since $L_{0,\infty} m \geq L_{0,1} m$, as long as $m_N(x) \geq 0$, we see that $\tilde{\mathcal{F}}_1$ will satisfy the same lower bounds satisfied by \mathcal{F}_1 when $\beta = 1$ and the proofs of Lemma 4.1 and Theorem 4.2 go through (but now with no martingale terms to control), and we have that, eventually, \mathbb{P} -a.s., for every $\epsilon > 0$, if λ, r_N are taken as in Theorem 4.2, there exists $T_0 > 0$ such that

$$(5.14) \quad \inf_{X_0^\infty: m_N(X_0^\infty) \geq r_N} \inf_{t \geq T_0} m_N(X_t^\infty) \geq (1 - \epsilon)$$

for all N large enough. We then notice that the stochastic Taylor expansion in Theorem 5.3 is only simplified if the Ito process X_t is replaced by the gradient descent process X_t^∞ , where L_0 is replaced by $L_{0,\infty}$; in particular, the bound (5.1) holds deterministically for $|f(X_t^\infty)|$. From there it is evident that following the proof of Theorem 5.4 and 5.5, making the appropriate modifications and omitting the L^∞ bounds on M_t^m , will yield its desired analog for X_t^∞ . Combined with (5.14) and the assumptions on μ_N satisfying Condition 1 at level n , or weakly satisfying Condition 1 at level ∞ , at inverse temperature $\beta = \infty$ and Condition 2 (resp., Condition 2') then allows one to conclude the proof of Theorems 2.4–2.5.

6. Checking Conditions 1 and 2: Regularity of initial data. In this section we provide some natural examples of initial data satisfying Condition 1 at different levels as well as Condition 2 or Condition 2'. Both of these examples are completely independent of the planted

signal, so they can be viewed as uninformative while facilitating recovery of the planted signal. The first of these examples, the volume measure on S^N , is handled in Section 6.1, which we show weakly satisfies Condition 1 at level ∞ for every $\beta \in (0, \infty) \cup \{\infty\}$. The second of these examples, the high-temperature Gibbs measure for H_0 , is handled in Section 6.2 and shown to satisfy Condition 1 at level n for every n . Both of these imply recovery thresholds of exactly $\alpha_c(\infty) = (k - 2)/2$ and $k_c(\infty) = 2$.

6.1. *Regularity of the initial data under the volume measure.* In this section we show Theorem 2.7 holds. Let us begin first by recalling the following result which usually goes by the name of the Poincaré lemma (see, e.g., [38]). This lemma and related concentration and anticoncentration estimates will appear frequently in the following. Their proofs are standard and follow from explicit computation of volumes of spherical caps. We summarize them here.

LEMMA 6.1. *The normalized volume measure dx or $d \text{vol}$ on S^N satisfies the following:*

- (Poincaré lemma) if X is drawn from dx , then

$$(6.1) \quad \sqrt{N}m_N(X) \xrightarrow{(d)} Z,$$

where Z is a standard Gaussian random variable.

- (Concentration) There exists a universal constant $C > 0$ such that for every $t > 0$,

$$\text{vol}(|m_N(X)| > t) \leq C \exp(-t^2 N/C).$$

- (Anticoncentration) There exists a universal $C > 0$ such that for every $\epsilon > 0$,

$$(6.2) \quad \text{vol}(|m_N(X)| \leq N^{-\frac{1}{2}-\epsilon}) \leq CN^{-\epsilon}.$$

We now wish to check that the normalized volume measure vol^+ on $S^N \cap \{x_1 > 0\}$ weakly satisfies Condition 1 at level ∞ . Specifically, we wish to prove the following concentration estimate:

THEOREM 6.2. *For every $\beta \in (0, \infty) \cup \{\infty\}$ and $T > 0$, there exists $C > 0$ so that for every $\delta > 0$,*

$$\text{vol}^+\left(\sup_{t \leq T} |e^{tL_{0,\beta}} L_{0,\beta} m_N(X)| \geq N^{-\frac{1}{2}+\delta}\right) \leq Ce^{-N^\delta/C}, \quad \text{eventually, } \mathbb{P}\text{-a.s.}$$

Notice, first of all, that it suffices to prove Theorem 6.2 under the law $d \text{vol}$ instead of $d \text{vol}^+$ as $\text{vol}^+(A) \leq 2 \text{vol}(A)$ for every set $A \subset S^N$. Though it should be the case that any reasonable initialization that is independent of H_0 and the direction e_1 weakly satisfies Condition 1 at level ∞ , our proof relies in an essential way on the rotational invariance of the uniform measure on S^N .

PROOF OF THEOREM 2.7. The fact that the normalized volume measure vol on $S^N \cap \{x_1 \geq 0\}$ satisfies Condition 2 follows from Lemma 6.1, namely, (6.1). For every $\beta \in (0, \infty) \cup \{\infty\}$, the fact that the normalized volume measure weakly satisfies Condition 1 at level ∞ at inverse temperature β follows from Theorem 6.2. \square

Toward the proof of Theorem 6.2, it will help to treat the gradient descent case of $\beta = \infty$ distinctly from the $\beta \in (0, \infty)$ situation where there is an additional martingale noise driven by the Brownian motion under Q_x . Toward this, let \hat{X}_t denote the Langevin dynamics generated by $L_{0,\beta}$ when $\beta \in (0, \infty)$, and let \hat{Y}_t be the gradient descent flow generated by $L_{0,\infty}$.

We will need several preliminary estimates on the gradients $\nabla H_0(\hat{X}_t)$ and $\nabla H_0(\hat{Y}_t)$. It will be important to ensure that for most initial conditions and realizations of H , the norm of the gradient stays bounded away from zero, so that $\frac{\nabla H_0(\hat{X}_t)}{|\nabla H_0(\hat{X}_t)|}$ and $\frac{\nabla H_0(\hat{Y}_t)}{|\nabla H_0(\hat{Y}_t)|}$ make sense and their regularity in t can be controlled by $\|H_0\|_{\mathcal{G}^k}$. In this section we let $\mu_N = dx$ denote the uniform measure on S^N .

LEMMA 6.3. *There exists $c, \Lambda > 0$ (independent of N) such that with \mathbb{P} -probability $1 - O(e^{-cN})$,*

$$|\nabla H_0(\hat{Y}_t)| \geq |\nabla H_0(\hat{Y}_0)|e^{-\Lambda t} \quad \text{for every } \hat{Y}_0 \text{ and every } t > 0.$$

In particular, if $\hat{Y}_0 \sim \mu_N$, then with $\mu_N \otimes \mathbb{P}$ -probability $1 - O(e^{-cN})$, we have for every $C > 0$,

$$|\nabla H_0(\hat{Y}_t)| \geq c\sqrt{N}e^{-\Lambda t} \quad \text{for all } t > 0, \quad \text{and} \quad \inf_{t \leq \frac{C}{\Lambda} \log N} |\nabla H_0(\hat{Y}_t)| \geq cN^{-C+\frac{1}{2}}.$$

PROOF. By Theorem 3.3, there exists $c, \Lambda(\xi) > 0$ such that with \mathbb{P} -probability $1 - O(e^{-cN})$, we have $\|\nabla^2 H_0(x)|_{\text{op}}\|_\infty < \Lambda$. Differentiating in time, we have

$$\frac{d}{dt} |\nabla H_0(\hat{Y}_t)|^2 = -2\nabla^2 H_0(\nabla H_0, \nabla H_0)(\hat{Y}_t) \geq -2\Lambda |\nabla H_0(\hat{Y}_t)|^2,$$

so that by Gronwall’s inequality we obtain

$$|\nabla H_0(\hat{Y}_t)|^2 \geq |\nabla H_0(\hat{Y}_0)|^2 e^{-2\Lambda t},$$

as desired. The second claim follows from the fact that there exist $c_1, c_2 > 0$ such that

$$(6.3) \quad \mu_N \otimes \mathbb{P}(|\nabla H_0(\hat{Y}_0)| \leq c_1\sqrt{N}) \leq e^{-c_2N}.$$

Specifically, by the rotational invariance of μ_N and the law of $H_0(x)$, for any fixed $x \in S^N$,

$$\mu_N \otimes \mathbb{P}(|\nabla H_0(\hat{Y}_0)| < c_1\sqrt{N}) = \mathbb{P}(|\nabla H_0(x)| < c_1\sqrt{N}) < e^{-c_2N}$$

for some $c_1, c_2 > 0$, as $\nabla H_0(x)$ is distributed as an i.i.d. Gaussian vector in dimension N with entries of order one variance $\sigma_p > 0$. \square

REMARK 6.4. One could prove a corresponding bound for the Langevin dynamics by controlling the effect of the noise in the dynamics via standard martingale estimates. For our purposes, toward Theorem 6.2, we only need a lower bound on $|\nabla H_0(\hat{X}_t)|$ for order one timescales, and therefore it suffices to use the bounds of the prequel [7]: one can read off from (6.3) and Theorems 1.2 and 1.4 of [7] that for every β and $T > 0$, there exists $c > 0$ such that

$$(6.4) \quad \mu_N \otimes \mathbb{P} \otimes Q_x \left(\inf_{t \in [0, T]} |\nabla H_0(\hat{X}_t)| < c\sqrt{N} \right) \leq 1 - O(e^{-cN}).$$

We also rely on the rotational invariance of the gradient when $\hat{Y}_0 \sim \mu_N$. Note that as a corollary of the preceding, $\nabla H_0(\hat{Y}_t)/|\nabla H_0(\hat{Y}_t)|$ is well defined modulo a measure zero set of $\mu_N \otimes \mathbb{P}$ and likewise for \hat{X}_t under $\mu_N \otimes \mathbb{P} \otimes Q_x$.

LEMMA 6.5. *For every $t \in [0, \infty)$, if $\hat{Y}_0 \sim \mu$, then the laws of $\frac{1}{\sqrt{N}}\hat{Y}_t$ and $\frac{\nabla H_0(\hat{Y}_t)}{|\nabla H_0(\hat{Y}_t)|}$ under $\mu \otimes \mathbb{P}$ are distributed as $\text{Unif}(S^{N-1}(1))$. Similarly, if $\hat{X}_0 \sim \mu$, and \hat{X}_t is the Langevin dynamics for H_0 at inverse temperature β , then, for every $t \in [0, \infty)$, the laws of $\frac{1}{\sqrt{N}}\hat{X}_t$ and $\frac{\nabla H_0(\hat{X}_t)}{|\nabla H_0(\hat{X}_t)|}$ under $\mu \otimes \mathbb{P} \otimes Q_x$ are distributed as $\text{Unif}(S^{N-1}(1))$.*

PROOF. It suffices for us to show that the law of \hat{Y}_t and $\frac{\nabla H_0(\hat{Y}_t)}{|\nabla H_0(\hat{Y}_t)|}$ are invariant under rotations. To this end, fix a rotation $\mathcal{O} \in O(N)$. Let

$$\tilde{Y}_0 = \mathcal{O}\hat{Y}_0 \quad \text{and} \quad \tilde{H}_0(x) = H_0(\mathcal{O}^{-1}x),$$

and let \tilde{Y}_t be the gradient descent on \tilde{H}_0 started from \tilde{Y}_0 . Since \tilde{Y}_0 is equidistributed with \hat{Y}_0 and $(\tilde{H}_0(x))_x$ is equidistributed with $(H_0(x))_x$, $\nabla \tilde{H}_0(\tilde{Y}_t)$ is equal in distribution to $\nabla H_0(\hat{Y}_t)$ for all $t \geq 0$. Since $\tilde{Y}_t = \mathcal{O}\hat{Y}_t$, we then see that \hat{Y}_t is equal in distribution to $\mathcal{O}\tilde{Y}_t$ as desired. Plugging in, it also implies that

$$\nabla \tilde{H}_0(\tilde{Y}_t) = \mathcal{O}\nabla H_0(\mathcal{O}^{-1}\tilde{Y}_t) = \mathcal{O}\nabla H_0(\hat{Y}_t).$$

From this we deduce that for every $t \in [0, \infty)$ and every rotation \mathcal{O} ,

$$\nabla H_0(\hat{Y}_t) \stackrel{(d)}{=} \mathcal{O}\nabla H_0(\hat{Y}_t).$$

Since $|\nabla H_0(\hat{Y}_t)| > 0$ holds $\mu_N \otimes \mathbb{P}$ -almost surely, it follows that the distribution of $\frac{\nabla H_0(\hat{Y}_t)}{|\nabla H_0(\hat{Y}_t)|}$ is uniform on the $(N - 1)$ -dimensional unit sphere.

In the case where \hat{X}_t and \tilde{X}_t are the Langevin dynamics, a similar argument applies upon coupling the driving Brownian motions as $\tilde{B}_t = \mathcal{O}B_t$, and then the facts that \tilde{X}_t is equidistributed to \hat{X}_t and $\nabla \tilde{H}_0(\tilde{X}_t)$ is equidistributed to $\nabla H_0(\hat{X}_t)$ hold as before, yielding the desired. \square

With the above two lemmas in hand, we are now in position to prove that the volume measure weakly satisfies Condition 1 at level ∞ . In particular, it holds on timescales that are $O(\log N)$ for gradient descent and, due to the sub-optimality of Remark 6.4, the Langevin dynamics for timescales that are $O(1)$. We begin with the proof for the gradient descent, then we modify the proof to carry it over to the Langevin dynamics.

PROOF OF THEOREM 6.2 AT $\beta = \infty$. By definition of the semigroup

$$e^{tL_0}L_0m(\hat{Y}_0) = L_0m(\hat{Y}_t),$$

and therefore it suffices for us to study $L_0m(\hat{Y}_t)$. Now, by definition,

$$L_0m = -\langle \nabla H_0, \nabla m \rangle = -\frac{1}{\sqrt{N}}\langle \nabla H_0, \nabla x_1 \rangle.$$

For each $r \in [-1, 1]$, define $\text{Cap}(e_1, r) := \{y \in \mathbb{S}^{N-1}(1) : \langle y, e_1 \rangle \geq r\}$. Taking $r := N^{-\frac{1}{2}+\delta}$, by Lemma 6.5, and (6.1) of Lemma 6.1, at every fixed time $t > 0$,

$$(6.5) \quad \mu \otimes \mathbb{P}\left(\left\langle \frac{\nabla H_0(\hat{Y}_t)}{|\nabla H_0(\hat{Y}_t)|}, e_1 \right\rangle \geq N^{-\frac{1}{2}+\delta}\right) = \text{vol}(\text{Cap}(e_1, N^{-\frac{1}{2}+\delta})) \leq e^{-cN^{2\delta}}.$$

Our goal is to union bound over $T = C \log N$ many discrete times to show that, with high probability, $L_0m(X_t)$ is small for all $t \leq T$. However, since this is a continuous time process, we define a series of bad events, on the complement of which we can push our union bound through. For constants Λ and $\epsilon > 0$, let

$$B_1 := \{\|H_0\|_{\mathcal{G}^2} \geq \Lambda N\}, \quad B_2 := \{X_0 : |\nabla H_0(X_0)| \leq \epsilon\sqrt{N}\}.$$

On the event B_1^c , for every $t \geq s \geq 0$,

$$|\nabla H_0(\hat{Y}_t) - \nabla H_0(\hat{Y}_s)| \leq \int_s^t |\nabla^2 H_0(\nabla H_0, \cdot)(u)| du \leq \Lambda^2\sqrt{N}(t - s).$$

As a consequence, under B_2^c , for every $C > 0$ and $t \leq C \log N$, by Lemma 6.3,

$$\begin{aligned} \left| \frac{\nabla H_0(\hat{Y}_t)}{|\nabla H_0(\hat{Y}_t)|} - \frac{\nabla H_0(\hat{Y}_s)}{|\nabla H_0(\hat{Y}_s)|} \right| &\leq \epsilon^{-2} N^{2C-1} \left| |\nabla H_0(\hat{Y}_s)| \nabla H_0(\hat{Y}_t) - |\nabla H_0(\hat{Y}_t)| \nabla H_0(\hat{Y}_s) \right| \\ &\leq \epsilon^{-2} N^{2C-1} (|\nabla H_0(\hat{Y}_s)| |\nabla H_0(\hat{Y}_t) - \nabla H_0(\hat{Y}_s)| \\ &\quad + |\nabla H_0(\hat{Y}_s)| |\nabla H_0(\hat{Y}_s)| - |\nabla H_0(\hat{Y}_t)|), \end{aligned}$$

which in turn, additionally under B_1^c , is at most

$$(6.6) \quad \epsilon^{-2} N^{2C-1} \Lambda^2 \sqrt{N} (t-s) \|\nabla H_0\|_\infty \leq \epsilon^{-2} \Lambda^3 N^{2C} (t-s).$$

Now, fix any $T \leq C \log N$, and partition the interval $[0, T]$ into $\lceil TN^{4C} \rceil$ intervals of size $\frac{1}{N^{4C}}$, $t_1 := 0, \dots, t_{\lceil TN^{4C} \rceil} := T$. Since $r = N^{-\frac{1}{2} + \delta}$,

$$d\left(\text{Cap}\left(e_1, \frac{r}{2}\right), \text{Cap}(e_1, r)^c\right) \geq \frac{1}{2} N^{-\frac{1}{2} + \delta},$$

whence in order for $\frac{\nabla H_0(\hat{Y}_t)}{|\nabla H_0(\hat{Y}_t)|} \in \text{Cap}(e_1, \frac{r}{2})$ for some $t \leq T$, by (6.6) either it must have been in $\text{Cap}(e_1, r)$ for one of $\{t_1, \dots, t_{\lceil TN^{4C} \rceil}\}$ or one of B_1 or B_2 must have occurred. Therefore, by a union bound and (6.5), we have

$$\mu \otimes \mathbb{P}\left(\sup_{t \leq T} \left\langle \frac{\nabla H_0(\hat{Y}_t)}{|\nabla H_0(\hat{Y}_t)|}, e_1 \right\rangle \geq \frac{1}{2} N^{-\frac{1}{2} + \delta}\right) \leq \lceil TN^{4C} \rceil e^{-cN^{2\delta}} + \mathbb{P}(B_1) + \mu \otimes \mathbb{P}(B_2).$$

By Theorem 3.3 and Borrell's inequality, an upper bound of $O(e^{-cN})$ on $\mathbb{P}(B_1)$ follows for Λ large enough. The matching bound on $\mu \otimes \mathbb{P}(B_2)$ follows from (6.3) for any $\epsilon < c_1$ there. \square

PROOF OF THEOREM 6.2 FOR $\beta \in (0, \infty)$. Let E_{Q_x} denote expectation with respect to the law of \hat{X}_t , the Langevin dynamics on H_0 . Recall that, by definition of the Markov semigroup e^{tL_0} ,

$$e^{tL_0} L_0 m(x) = E_{Q_x}[L_0 m(\hat{X}_t)].$$

Let us break up $L_0 m$ into its two constituent parts,

$$L_0 m = \left[-\frac{1}{2} \Delta - \langle \nabla H_0, \nabla \cdot \rangle \right] m = -\frac{1}{2} m - \frac{1}{\sqrt{N}} \langle \nabla H_0, \nabla x_1 \rangle.$$

(Note that since ∇H_0 is in TS^N and that for $x \in S^N$ we have $\nabla x_1 = e_1 - (x_1/\sqrt{N})x$, it follows that $\langle \nabla H_0, \nabla x_1 \rangle = \langle \nabla H_0, e_1 \rangle$.) By Lemma 6.5 and Lemma 6.1, the quantities $m(\hat{X}_t)$ and $\langle \frac{\nabla H_0(\hat{X}_t)}{|\nabla H_0(\hat{X}_t)|}, e_1 \rangle$ satisfied the sub-Gaussian tail bounds that for every $r > 0$,

$$\mu \otimes \mathbb{P} \otimes Q_x(|m(\hat{X}_t)| \geq r) = \text{vol}(\text{Cap}(e_1, r)) \leq C e^{-r^2 N/C} \quad \text{and}$$

$$\mu \otimes \mathbb{P} \otimes Q_x\left(\left| \left\langle \frac{\nabla H_0(\hat{X}_t)}{|\nabla H_0(\hat{X}_t)|}, e_1 \right\rangle \right| \geq r\right) = \text{vol}(\text{Cap}(e_1, r)) \leq C e^{-r^2 N/C}.$$

Integrating out the Brownian motion, it is straightforwardly seen that the sub-Gaussian tail bounds transfer to $E_{Q_x}[m(\hat{X}_t)]$ and $E_{Q_x}[\langle \frac{\nabla H_0(\hat{X}_t)}{|\nabla H_0(\hat{X}_t)|}, e_1 \rangle]$ under $\mu \otimes \mathbb{P}$.

For constants Λ, T and $\epsilon > 0$, define the bad events,

$$B_1 := \{\|H_0\|_{\mathcal{G}^3} \geq \Lambda \sqrt{N}\}, \quad B_2 := \left\{x : \inf_{t \leq T} E_{Q_x} |\nabla H_0(\hat{X}_t)| \leq \epsilon \sqrt{N}\right\}.$$

We can now express on $B_1^c \cap B_2^c$, by Lemma 3.5,

$$\begin{aligned} |E_{Q_x}[m(\hat{X}_t) - m(\hat{X}_s)]| &\leq 2 \left| E_{Q_x} \left[\int_s^t L_0 m(\hat{X}_u) du \right] \right| + 2E_{Q_x}[|\langle B_t - B_s, e_1 \rangle|] \\ &\leq 2(t-s)\|L_0 m\|_\infty + \sqrt{t-s} \leq 2(t-s)(1 + \Lambda) + 2\sqrt{t-s}. \end{aligned}$$

At the same time, by the Cauchy–Schwarz inequality,

$$\begin{aligned} &\left| E_{Q_x} \left[\left\langle \frac{\nabla H_0(\hat{X}_t)}{|\nabla H_0(\hat{X}_t)|}, e_1 \right\rangle - \left\langle \frac{\nabla H_0(\hat{X}_s)}{|\nabla H_0(\hat{X}_s)|}, e_1 \right\rangle \right] \right| \\ &= \left| E_{Q_x} \left[\frac{\langle \nabla H_0(\hat{X}_t) | \nabla H_0(\hat{X}_s) \rangle - |\nabla H_0(\hat{X}_t)| |\nabla H_0(\hat{X}_s)|}{|\nabla H_0(\hat{X}_t)| |\nabla H_0(\hat{X}_s)|} \right] \right| \\ &\leq E_{Q_x} \left[\frac{|\langle \nabla H_0(\hat{X}_t), e_1 \rangle - \langle \nabla H_0(\hat{X}_s), e_1 \rangle|}{|\nabla H_0(\hat{X}_t)|} + \frac{||\nabla H_0(\hat{X}_t)| - |\nabla H_0(\hat{X}_s)||}{|\nabla H_0(\hat{X}_t)|} \right] \\ &\leq 2 \sqrt{\frac{E_{Q_x} |\langle \nabla H_0(\hat{X}_t), e_1 \rangle - \langle \nabla H_0(\hat{X}_s), e_1 \rangle|^2 \vee E_{Q_x} ||\nabla H_0(\hat{X}_t)| - |\nabla H_0(\hat{X}_s)||^2}{E_{Q_x} |\nabla H_0(\hat{X}_t)|^2}}. \end{aligned}$$

On the event $B_1^c \cap B_2^c$, this is at most

$$\begin{aligned} &\epsilon^{-1} N^{-\frac{1}{2}} \{ E_{Q_x} [|\langle \nabla H_0(\hat{X}_t), e_1 \rangle - \langle \nabla H_0(\hat{X}_s), e_1 \rangle|]^2 \vee E_{Q_x} [||\nabla H_0(\hat{X}_t)| - |\nabla H_0(\hat{X}_s)||]^2 \}^{1/2} \\ &\leq 2\epsilon^{-1} N^{-\frac{1}{2}} [(t-s)(\|L_0 \langle \nabla H_0, e_1 \rangle\|_\infty \vee \|L_0 |\nabla H_0|\|_\infty) + \sqrt{(t-s)} \|\nabla^2 H_0\|_\infty] \\ &\leq K \epsilon^{-1} N^{-\frac{1}{2}} [(t-s)\Lambda + \sqrt{t-s}\Lambda], \end{aligned}$$

for some constant $K > 0$ by (6.4). Fixing any finite T and partitioning $[0, T]$ in to $\lceil TN^4 \rceil$ intervals of width N^{-4} and arguing as in the proof of the gradient descent estimate, we see that in order for $e^{tL_0} L_0 m$ to have been at least $\frac{1}{2} N^{-\frac{1}{2}+\delta}$, either B_1 or B_2 occurred, or one of $|E_{Q_x}[m(\hat{X}_t)]| \geq \frac{1}{4} N^{-\frac{1}{2}+\delta}$ or $|E_{Q_x}[\langle \frac{\nabla H_0(\hat{X}_t)}{|\nabla H_0(\hat{X}_t)|}, e_1 \rangle]| \geq \frac{1}{4} N^{-\frac{1}{2}+\delta}$ for some $t \in \{0, N^{-4}, \dots, T\}$. By a union bound and the above,

$$\mu \otimes \mathbb{P} \left(\sup_{t \leq T} |e^{tL_0} L_0 m(x)| \geq N^{-\frac{1}{2}+\delta} \right) \leq \lceil TN^4 \rceil e^{-cN^{2\delta}} + \mathbb{P}(B_1) + \mu \otimes \mathbb{P}(B_2).$$

By Theorem 3.3 and Borrell’s inequality, an upper bound of $O(e^{-cN})$ on $\mathbb{P}(B_1)$ follows for Λ large enough. The matching bound on $\mu \otimes \mathbb{P}(B_2)$ follows from (6.4) for any $\epsilon < c$ there. □

6.2. *Regularity of the initial data under the high-temperature Gibbs measure.* Recall from the Introduction that $d\pi_{0,\beta}(x) \propto e^{-\beta H_0(x)} dx$ denotes the Gibbs measure corresponding to the pure noise Hamiltonian, $H_0(x)$, at inverse temperature $\beta > 0$, and $\pi_{0,\beta}^+$ is the same Gibbs measure conditioned on $\{x_1 \geq 0\}$. Observe the following consequence of (6.2) and the isotropy of H_0 :

LEMMA 6.6. *For every ξ , every $\beta > 0$ and every $\delta > 0$,*

$$\pi_{0,\beta}^+(x_1 \leq N^{-\delta}) \rightarrow 0 \quad \text{in prob.}$$

In particular, $\pi_{0,\beta}^+$ satisfies Condition 2' for every $\delta > 0$.

PROOF. By isotropy of the law of H_0 and (6.2), we have

$$\mathbb{E}[\pi_{0,\beta}^+(x_1 \leq N^{-\delta})] = \text{vol}(|x_1| \leq N^{-\delta}) \leq CN^{-\delta/2},$$

which after an application of Markov’s inequality, implies the desired. \square

We now turn to the main estimate in this section.

LEMMA 6.7. *For every ξ , there exists a $\beta_0 > 0$ such that, eventually, \mathbb{P} -almost surely, for every $\beta < \beta_0$, the measure $\pi_{0,\beta}$ satisfies Condition 1 at level n at inverse temperature β , for every n . Consequently, if ξ is even, the measure $\pi_{0,\beta}^+$ satisfies Condition 1 at level n for every n as well.*

PROOF. We wish to show that there exists $K(\xi, \beta) > 0$ such that for every $\ell \geq 0$, eventually, \mathbb{P} -a.s.,

$$(6.7) \quad \pi_{0,\beta}(|L_0^\ell m_N(x)| \geq N^{-\frac{1}{2}+\delta}) \leq Ke^{-N^{2\delta}/K},$$

as long as β is sufficiently small. Recall that there is an important implicit dependence on β in $L_0 = L_{0,\beta}$. A union bound over the first n such events would then imply $\pi_{0,\beta}$ satisfies Condition 1 at level n . For the case $\ell = 0$, it suffices to check

$$\mathbb{E}[\pi_{0,\beta}(|x_1| \geq N^\delta)] \leq e^{-N^{2\delta}}.$$

This holds for every $\beta > 0$ from the fact that $\mathbb{E}[\pi_{0,\beta}(A)] = \text{vol}(A)$ and the concentration of (X, e_1) under the volume measure dx . This then implies (6.7) with $\ell = 0$ by Markov’s inequality.

We now proceed to the cases $\ell \geq 1$. Recall that for every β , the generator $L_0 = L_{0,\beta}$ is essentially selfadjoint in $C^\infty(\mathcal{S}^N) \subset L^2(d\pi_{0,\beta})$.

Furthermore, it was shown in [17] that there exists $\beta_0 > 0$ such that for every $\beta < \beta_0$, the measure $\pi_{0,\beta}$ satisfies a log-Sobolev inequality with constant $c(\xi, \beta) > 0$, eventually, \mathbb{P} -almost surely.

By Herbst’s argument [26], it then follows that for every smooth, 1-Lipschitz function F , we have

$$\pi_{0,\beta}(|F(x) - \pi_{0,\beta}[F]| \geq r) \lesssim e^{-r^2/(2\sqrt{c})},$$

where $\pi_{0,\beta}[F]$ is the expectation of $F(x)$ under $\pi_{0,\beta}$. Now, fix any $\ell \geq 1$, and let $F(x) = \sqrt{N}L_0^\ell m_N(x)$: by essential selfadjointness of $L_{0,\beta}$ with respect to $\pi_{0,\beta}$ and smoothness of $m_N(x)$, we see that

$$\int (L_0^\ell m)(x) d\pi_{0,\beta}(x) = \langle 1, L_0^\ell m \rangle_{L^2(\pi_{0,\beta})} = \langle L_0^\ell 1, m \rangle_{L^2(\pi_{0,\beta})} = 0.$$

At the same time, we claim that

$$\|L_0^\ell m_N\|_{\text{Lip}} \leq \|\nabla L_0^\ell m_N\|_\infty \leq \frac{C_{\xi,\beta,\ell}}{\sqrt{N}},$$

for some $C_{\xi,\beta,\ell} > 0$. Recall from (3.4) that $\|m_N\|_{\mathcal{G}^k} \leq k$ for all k . On the other hand, by (3.1)–(3.3), for every n and ℓ there exists $C_{\xi,\beta,\ell} > 0$ such that, eventually, \mathbb{P} -almost surely,

$$\|L_0^\ell\|_{\mathcal{G}^n \rightarrow \mathcal{G}^{n-2}} \leq C_{\xi,\beta,\ell}.$$

Thus, for every k , eventually, \mathbb{P} -almost surely we have $\|L_0^\ell m_N\|_{\mathcal{G}^k} \leq C_{\xi,\beta,\ell}$ for some other $C_{\xi,\beta,\ell}$, and it then follows that $\|\nabla L_0^\ell m_N\|_\infty \leq \frac{C_{\xi,\beta,\ell}}{\sqrt{N}}$. Combining these, we see that for every $\ell \geq 1$,

$$\pi_{0,\beta}(|L_0^\ell m(x)| \geq rN^{-\frac{1}{2}}) \lesssim e^{-r^2/2},$$

eventually, \mathbb{P} -a.s., which implies the desired concentration estimate when plugging in $r = N^\delta$. \square

7. Free energy wells and obstructions to recovery. In Section 7.1 we formalize the notion of free energy wells and obtain an elementary exit time lower bound that is exponential in the height of the well, started from the natural Gibbs initialization in the well. Recall from Remark 2.13 that imposing such an initialization is necessary to have such an exit time lower bound.

We then show that the measure $d\pi \propto \exp(-\beta H) dx$ has a free energy well around the equator.

PROPOSITION 7.1. *Fix ξ, k and $\beta > 0$. If $\alpha < \frac{k-2}{2}$, then for $\epsilon \in (0, \frac{1}{2})$ sufficiently small, eventually, \mathbb{P} -a.s., the function $f(x) = (x, e_1)$ has a $(\frac{1}{2}N^\epsilon)$ -free energy well of height $N^{2\epsilon}$ at $[-\frac{3}{2}N^\epsilon, \frac{3}{2}N^\epsilon]$.*

We end the section by combining this with a crude comparison between π and π_0 near the equator to deduce Theorem 2.10.

7.1. Free energy wells and hitting time lower bounds. The existence of free energy wells is closely related to the behavior of both equilibrium and off-equilibrium Langevin dynamics. In particular, free energy barriers govern the eigenvalues of the Dirichlet problems on domains and, in turn, the exit times of sets from their interiors as well as spectral gaps. Importantly, though, these relationships are, in a sense, with respect to worst possible initial starts or with respect to Gibbs, typical initializations.

For the recovery problem in our setting, what will be relevant are free energy barriers between the equator, or volume-typical latitudes, and the signal.

In this section we will work in full generality of Langevin dynamics on S^N with respect to some Hamiltonian V_N . Let $f : S^N \rightarrow \mathbb{R}$ be some smooth, say $C^\infty(S^N)$, observable. The Langevin dynamics has invariant measure which we denote (abusing notation) $d\pi(x) \propto \exp(-\beta V_N)$. For the purposes of the recovery problem, one should have in mind the example $V_N = H$ and $f(x) = (x, e_1)$.

For any function $f : S^N \rightarrow \mathbb{R}$, we can define the following “rate function:”

$$I_f(a; \epsilon) = -\log \pi(\{x : f(x) \in B_\epsilon(a)\}).$$

DEFINITION 7.2. The function f has an ϵ -free energy well of height h in $[a, b]$ if the following holds: there exists $c \in (a, b)$ and a pair ϵ, η such that $B_\epsilon(a), B_\epsilon(b), B_\eta(c)$ are disjoint, and

$$\min\{I_f(a, \epsilon), I_f(b, \epsilon)\} - I_f(c, \eta) \geq h.$$

REMARK 7.3. It will sometimes be useful to consider one-sided energy wells. To do so, we may take either $a = -\infty$ or $b = \infty$ in which case, for every ϵ , we would set

$$I_f(\pm\infty, \epsilon) = \infty.$$

THEOREM 7.4. *There is a universal constant $C > 0$ such that the following holds. Suppose that f is a smooth function with no critical values in an open neighbourhood of $[a, b]$. If f has an ϵ -free energy well of height h in $[a, b]$ with*

$$\frac{\epsilon^2}{h} \leq \frac{C \|\nabla f\|_{L^\infty(A_\epsilon)}^2}{N + \sqrt{N}\beta \|\nabla V\|_{L^\infty(A_\epsilon)}},$$

where $A_\epsilon = f^{-1}([a, b])$, then the exit time of A_ϵ , denoted τ_{A_ϵ} , satisfies

$$\int Q_x(\tau_{A_\epsilon} \leq T) d\pi(x | A_\epsilon) \lesssim (1 + \epsilon^{-4} h T \|\nabla f\|_{L^\infty(A_\epsilon)}^4) \exp(-h).$$

PROOF. For ease of notation, let $A_\epsilon = f^{-1}([a, b])$, $A = f^{-1}([a, b] \setminus (B_\epsilon(a) \cup B_\epsilon(b)))$ and $B_\epsilon = A_\epsilon \setminus A$. Since $[a, b]$ consists of only regular points, $f^{-1}([a, b])$ is a submanifold with smooth boundary by the preimage theorem [19]. Therefore, Langevin dynamics reflected at the boundary of A_ϵ is well defined in the sense of the martingale problem [36]. Call this process \bar{X}_t , and denote its law by \bar{Q}_x . Recall that it is reversible, with stationary measure

$$\nu = \pi(\cdot | A_\epsilon).$$

By definition, X_t and \bar{X}_t are equal in law conditionally on $t \leq \tau_{\partial A_\epsilon}$. As a result,

$$\int \mathcal{Q}_x(\tau_{\partial A_\epsilon} < T) d\nu(x) = \int \bar{\mathcal{Q}}_x(\tau_{\partial A_\epsilon} < T) d\nu(x).$$

Let us estimate the right-hand side above. Let s_0 be a positive number, to be chosen later, and let $t_i = (i \cdot s_0) \wedge T$ for each $i = 0, \dots, \lceil T/s_0 \rceil$. Then,

$$\begin{aligned} \int_A \bar{\mathcal{Q}}_x(\tau_{\partial A_\epsilon} < T) d\nu(x) &\leq \int_{A_\epsilon} \bar{\mathcal{Q}}_x(\exists i : \bar{X}_{t_i} \in B_\epsilon) d\nu(x) \\ &\quad + \int_A \bar{\mathcal{Q}}_x(\tau_{\partial A_\epsilon} < T, \bar{X}_{t_i \wedge \tau_{\partial A_\epsilon}} \in A \forall i) d\nu(x). \end{aligned}$$

Call the first term above I and the second term II , and bound them separately. By a union bound and reversibility with respect to ν ,

$$I \leq \frac{T}{s_0} \nu(B_\epsilon).$$

On the other hand, by the Markov property, the equality in law of $X_{t \wedge \partial A_\epsilon}$ and $\bar{X}_{t \wedge \partial A_\epsilon}$ and a union bound,

$$\begin{aligned} II &\leq \int_A \bar{\mathcal{Q}}_x(\exists i : t_i < \tau_{\partial A_\epsilon} < t_{i+1}, \bar{X}_{t_i} \in A) d\nu(x) \\ &= \int_A \mathcal{Q}_x(\exists i : t_i < \tau_{\partial A_\epsilon} < t_{i+1}, X_{t_i} \in A) d\nu(x) \\ &\leq \frac{T}{s_0} \sup_x \mathcal{Q}_x \left(\sup_{t \leq s_0} d(\bar{X}_t, x) \geq \epsilon / \|\nabla f\|_\infty \right). \end{aligned}$$

To bound this last term, observe that for any $x_0 \in \mathcal{S}^N$, the function $g(x) = |x - x_0|^2$ is smooth and, by an explicit calculation,

$$\|\Delta g\|_{L^\infty(\mathcal{S}^N)} \leq CN \quad \text{and} \quad \|\nabla g\|_{L^\infty(\mathcal{S}^N)} \leq C\sqrt{N},$$

for some universal constant $C > 0$. In particular,

$$\|Lg\|_{L^\infty(A_\epsilon)} \leq C(N + \beta\sqrt{N}\|\nabla V\|_{L^\infty(A_\epsilon)}).$$

Thus, provided $C(N + \beta\sqrt{N}\|\nabla V\|_{L^\infty(A_\epsilon)})s_0 \leq \epsilon^2 / (2\|\nabla f\|_{L^\infty(A_\epsilon)}^2)$, we obtain

$$\begin{aligned} \mathcal{Q}_x \left(\sup_{t \leq s_0} d(X_t, x) \geq \frac{\epsilon}{\|\nabla f\|_{L^\infty(A_\epsilon)}} \right) &\leq \mathcal{Q}_x \left(\sup_{t \leq s_0} |M_t^g| \geq \frac{\epsilon^2}{2\|\nabla f\|_{L^\infty(A_\epsilon)}^2} \right) \\ &\leq K \exp \left(-\frac{\epsilon^4}{K\|\nabla f\|_{L^\infty(A_\epsilon)}^4 s_0} \right), \end{aligned}$$

for some universal $K > 0$ by Doob’s maximal inequality (3.5). If we choose

$$s_0 = \frac{1}{Kh} \left(\frac{\epsilon^2}{\|\nabla f\|_{L^\infty(A_\epsilon)}^2} \right)^2,$$

then, by assumption, s_0 satisfies this inequality from which it follows that the right-hand side of the preceding display is bounded by $K \exp(-h)$.

By the assumption of f having an ϵ -free energy well of height h , we know that $\nu(B_\epsilon) \leq \exp(-h)$. Combing these estimates yields

$$\int Q_x(\tau_{\partial A_\epsilon} < T) d\nu \lesssim \left(1 + \frac{T}{s_0}\right) \exp(-h)$$

which yields the desired bound, after plugging in for s_0 . \square

7.2. Spin-glass free energies. We now turn to the specific case of an equatorial free energy well for π_β . Define the following *restricted spin-glass free energy*: for a Borel set $A \subset \mathcal{S}^N$ define

$$F_0(A) = F_{0,N}(A; \beta) = \frac{1}{N} \log \int_A e^{-\beta H_0(x)} dx.$$

Similarly, define the restricted free energy

$$F(A) = \frac{1}{N} \log \int_A e^{-\beta H(x)} dx,$$

and define $F(A)$ analogously. Finally, let

$$(7.1) \quad A_{\delta_N} = \{x : m_N(x) \in (-\delta_N, \delta_N)\}$$

denote the equatorial band of correlation at most δ_N . We begin by showing the following estimate for the spin-glass free energy.

LEMMA 7.5. *Fix any ξ , and $\beta, \epsilon > 0$. Letting $\delta_N = N^{-\frac{1}{2}+\epsilon}$, there exists $C > 0$ such that, eventually, \mathbb{P} -almost surely,*

$$F_0(A_{\delta_N}) - F_0(A_{2\delta_N} \setminus A_{\delta_N}) \geq CN^{-1+2\epsilon}.$$

PROOF. We will in fact show the stronger

$$(7.2) \quad F_0(A_{\delta_N}) - F_0(A_{\delta_N}^c) \geq CN^{-1+2\epsilon},$$

which would imply the desired, since $A_{2\delta_N} \setminus A_{\delta_N} \subset A_{\delta_N}^c$ and thus $F_0(A_{2\delta_N} \setminus A_{\delta_N}) \leq F_0(A_{\delta_N}^c)$. Since for any A ,

$$F_0(A) - F_0(A^c) = \frac{1}{N} \log \frac{Z_0(A)}{Z_0(A^c)} = \frac{1}{N} \log \left(\frac{1}{\pi_{0,\beta}(A^c)} - 1 \right),$$

it suffices to show that \mathbb{P} -eventually almost surely,

$$\pi_{0,\beta}(A_{\delta_N/2}^c) \lesssim \exp(-cN^{2\epsilon})$$

for some $c(\xi, \beta) > 0$.

In order to upper bound $\pi_{0,\beta}(A_{\delta_N}^c)$, notice that for any fixed $A \subset \mathcal{S}^N$,

$$\mathbb{E}[\pi_{0,\beta}(x \in A)] = \text{vol}(A),$$

by isotropy of the law of H_0 . On the other hand, the concentration estimate in Lemma 6.1 shows that $\text{Vol}(A_{\delta_N}) \leq e^{-N^{2\epsilon}/2}$ so that, by Markov's inequality,

$$\mathbb{P}(\pi_{0,\beta}(x \in A_{\delta_N}^c) \geq e^{-\frac{N^{2\epsilon}}{4}}) \leq e^{-\frac{N^{2\epsilon}}{4}}.$$

Combining the above, we obtain the desired inequality. \square

PROOF OF PROPOSITION 7.1. Let $\delta = \delta_N = N^{-\frac{1}{2}+\epsilon}$ for some ϵ to be chosen sufficiently small. Define A_δ as in (7.1), and let $B_\delta = A_{2\delta} - A_\delta$. First of all,

$$\begin{aligned} F(A_\delta) &= \frac{1}{N} \log \int_{A_\delta} e^{-\beta(H_0(x) - \lambda N \phi(x))} dx \\ &\geq \frac{1}{N} \log \int_{A_\delta} e^{-\beta H_0(x) - \lambda N \delta_N^k} dx = F_0(A_{\delta/2}) - \beta \lambda \delta_N^k, \end{aligned}$$

while, on the other hand,

$$F(B_\delta) \leq F_0(B_\delta) + \beta \lambda \delta_N^k.$$

Thus, by Lemma 7.5 and the above inequalities, \mathbb{P} -eventually almost surely,

$$I_f\left(\frac{3}{2}N^\epsilon; \frac{1}{2}N^\epsilon\right) - I_f(0; N^\epsilon) = F(B_\delta) - F(A_\delta) \leq -\frac{C}{N^{1-2\epsilon}} + 2\beta \lambda \delta_N^k,$$

for some $C(\xi, \beta, k) > 0$. But by the assumption that $\alpha < (k - 2)/2$, if we choose $\epsilon > 0$ sufficiently small, namely $0 < \epsilon < \frac{1}{k}(\frac{k-2}{2} - \alpha)$, we see that the above satisfies

$$I_f\left(\frac{3}{2}N^\epsilon; \frac{1}{2}N^\epsilon\right) - I_f(0; N^\epsilon) < -\frac{C'}{N^{1-2\epsilon}}$$

for some $C'(\xi, \beta, k) > 0$. When k is even, by symmetry, we similarly have

$$I_f\left(-\frac{3}{2}N^\epsilon; \frac{1}{2}N^\epsilon\right) - I_f(0; N^\epsilon) < -\frac{C'}{N^{1-2\epsilon}},$$

and when k is odd, we would have

$$I_f\left(-\frac{3}{2}N^\epsilon; \frac{1}{2}N^\epsilon\right) - I_f(0; N^\epsilon) \leq -\frac{C}{N^{1-2\epsilon}} - \beta \lambda \delta_N^k$$

which is negative regardless of the choice of α . By Definition 7.2, this concludes the proof. \square

7.3. Proof of Theorem 2.10. We begin by combining Proposition 7.1 and Theorem 7.4 to obtain an exit time lower bound for the Gibbs measure π restricted to the equatorial free energy well. Fix $\beta > 0$, and take $V_N = H_N, \epsilon_N = \frac{1}{2}N^\epsilon, [a, b] = [-\frac{3}{2}N^\epsilon, \frac{3}{2}N^\epsilon]$ and $h_N = N^{2\epsilon}$ for any ϵ sufficiently small. By (3.4), $\|\nabla f\|_\infty \leq 1$, and so we end up with

$$(7.3) \quad \int \mathcal{Q}_x(\tau_{|x_1| \geq 2N^\epsilon} \leq T) d\pi\left(x \mid |x_1| \leq \frac{3}{2}N^\epsilon\right) \lesssim T \cdot \exp(-N^{2\epsilon}).$$

We now wish to boost this to an estimate on initialization from $\pi_0 = \pi_{0,\beta}$. As before, let

$$\nu = \pi\left(\cdot \mid |x_1| \leq \frac{3}{2}N^\epsilon\right), \quad \text{and} \quad \nu_0 = \pi_0\left(\cdot \mid |x_1| \leq \frac{3}{2}N^\epsilon\right).$$

We begin by bounding the Radon–Nikodym derivative

$$\exp\left(-\beta\left(\frac{3}{2}\right)^k N^{-\frac{k-2}{2}+\alpha+\epsilon k}\right) \leq \max_{x: |x_1| \leq \frac{3}{2}N^\epsilon} \frac{d\nu(x)}{d\nu_0(x)} \leq \exp\left(\beta\left(\frac{3}{2}\right)^k N^{-\frac{k-2}{2}+\alpha+\epsilon k}\right),$$

but, as soon as ϵ is sufficiently small, the power of N is negative since $\alpha < \alpha_c(\infty)$ and thus, for N sufficiently large, the Radon–Nikodym derivative satisfies $1/2 \leq \|d\nu/d\nu_0\|_\infty \leq 2$. As a result,

$$\int \mathcal{Q}_x(\tau_{|x_1| \geq 2N^\epsilon} \leq T) d\nu_0(x) \lesssim 2T \exp(-N^{2\epsilon}).$$

Now, notice that for some universal $C > 0$,

$$\pi_0\left(|x_1| \geq \frac{3}{2}N^\epsilon\right) \leq C \exp(-N^{2\epsilon}/C);$$

this follows as in the $\ell = 0$ case of π_0 satisfying Condition 1 at level n , namely, by concentration of x_1 under dx , and Markov's inequality. Thus, writing

$$\int \mathcal{Q}_x(\tau_{|x_1| \geq 2N^\epsilon} \leq T) d\pi_0(x) \leq \pi_0\left(|x_1| \geq \frac{3}{2}N^\epsilon\right) + 2 \int \mathcal{Q}_x(\tau_{|x_1| \geq 2N^\epsilon} \leq T) d\nu_0(x)$$

and plugging in the two inequalities above, we obtain the desired inequality by choosing $T = e^{cN^\epsilon}$ for $c > 0$ sufficiently small.

Acknowledgments. We thank the anonymous referee for helpful comments and suggestions. We also thank Giulio Biroli, Chiara Cammarota, Florent Krzakala, Lenka Zdeborova, Afonso Bandeira and David Gamarnik for inspiring discussions. This research was conducted while A.J. was supported by NSF OISE-1604232. R.G. thanks the Miller Institute for Basic Research in Science for its support.

REFERENCES

- [1] ABBE, E. and SANDON, C. (2018). Proof of the achievability conjectures for the general stochastic block model. *Comm. Pure Appl. Math.* **71** 1334–1406. MR3812075 <https://doi.org/10.1002/cpa.21719>
- [2] AUFFINGER, A. and BEN AROUS, G. (2013). Complexity of random smooth functions on the high-dimensional sphere. *Ann. Probab.* **41** 4214–4247. MR3161473 <https://doi.org/10.1214/13-AOP862>
- [3] AUFFINGER, A., BEN AROUS, G. and ČERNÝ, J. (2013). Random matrices and complexity of spin glasses. *Comm. Pure Appl. Math.* **66** 165–201. MR2999295 <https://doi.org/10.1002/cpa.21422>
- [4] BAIK, J., BEN AROUS, G. and PÉCHÉ, S. (2005). Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Ann. Probab.* **33** 1643–1697. MR2165575 <https://doi.org/10.1214/009117905000000233>
- [5] BARAK, B., HOPKINS, S. B., KELNER, J., KOTHARI, P., MOITRA, A. and POTECHIN, A. (2016). A nearly tight sum-of-squares lower bound for the planted clique problem. In *57th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2016* 428–437. IEEE Computer Soc., Los Alamitos, CA. MR3631005
- [6] BARAK, B. and MOITRA, A. (2016). Noisy tensor completion via the sum-of-squares hierarchy. In *Conference on Learning Theory* 417–445.
- [7] BEN AROUS, G., GHEISSARI, R. and JAGANNATH, A. (2020). Bounding flows for spherical spin glass dynamics. *Comm. Math. Phys.* **373** 1011–1048. <https://doi.org/10.1007/s00220-019-03649-4>
- [8] BEN AROUS, G. and JAGANNATH, A. (2018). Spectral gap estimates in mean field spin glasses. *Comm. Math. Phys.* **361** 1–52. MR3825934 <https://doi.org/10.1007/s00220-018-3152-6>
- [9] BEN AROUS, G., MEI, S., MONTANARI, A. and NICA, M. (2019). The landscape of the spiked tensor model. *Comm. Pure Appl. Math.* **72** 2282–2330. MR4011861 <https://doi.org/10.1002/cpa.21861>
- [10] BENAYCH-GEORGES, F., GUIONNET, A. and MAIDA, M. (2012). Large deviations of the extreme eigenvalues of random deformations of matrices. *Probab. Theory Related Fields* **154** 703–751. MR3000560 <https://doi.org/10.1007/s00440-011-0382-3>
- [11] BERTHET, Q. and RIGOLLET, P. (2013). Optimal detection of sparse principal components in high dimension. *Ann. Statist.* **41** 1780–1815. MR3127849 <https://doi.org/10.1214/13-AOS1127>
- [12] CAPITAINE, M., DONATI-MARTIN, C. and FÉRAL, D. (2009). The largest eigenvalues of finite rank deformation of large Wigner matrices: Convergence and nonuniversality of the fluctuations. *Ann. Probab.* **37** 1–47. MR2489158 <https://doi.org/10.1214/08-AOP394>
- [13] CHEN, W.-K. (2019). Phase transition in the spiked random tensor with Rademacher prior. *Ann. Statist.* **47** 2734–2756. MR3988771 <https://doi.org/10.1214/18-AOS1763>
- [14] ETHIER, S. N. and KURTZ, T. G. (1986). *Markov Processes: Characterization and Convergence*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. Wiley, New York. MR0838085 <https://doi.org/10.1002/9780470316658>
- [15] GAMARNIK, D. and ZADIK, I. (2017). High dimensional regression with binary coefficients. Estimating squared error and a phase transition. In *Conference on Learning Theory* 948–953.

- [16] GE, R. and MA, T. (2017). On the optimization landscape of tensor decompositions. In *Advances in Neural Information Processing Systems* 30 (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, eds.) 3653–3663. Curran Associates, Red Hook.
- [17] GHEISSARI, R. and JAGANNATH, A. (2019). On the spectral gap of spherical spin glass dynamics. *Ann. Inst. Henri Poincaré Probab. Stat.* **55** 756–776. MR3949952 <https://doi.org/10.1214/18-aihp897>
- [18] GILLIN, P. and SHERRINGTON, D. (2000). $p > 2$ spin glasses with first-order ferromagnetic transitions. *J. Phys. A* **33** 3081–3091. MR1766894 <https://doi.org/10.1088/0305-4470/33/16/302>
- [19] GUILLEMIN, V. and POLLACK, A. (2010). *Differential Topology*. AMS Chelsea Publishing, Providence, RI. MR2680546 <https://doi.org/10.1090/chel/370>
- [20] HILLAR, C. J. and LIM, L.-H. (2013). Most tensor problems are NP-hard. *J. ACM* **60** Art. 45, 39. MR3144915 <https://doi.org/10.1145/2512329>
- [21] HOPKINS, S. B., SHI, J., SCHRAMM, T. and STEURER, D. (2016). Fast spectral algorithms from sum-of-squares proofs: Tensor decomposition and planted sparse vectors. In *STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing* 178–191. ACM, New York. MR3536564
- [22] HOPKINS, S. B., SHI, J. and STEURER, D. (2015). Tensor principal component analysis via sum-of-square proofs. In *Conference on Learning Theory* 956–1006.
- [23] HSU, E. P. (2002). *Stochastic Analysis on Manifolds. Graduate Studies in Mathematics* **38**. Amer. Math. Soc., Providence, RI. MR1882015 <https://doi.org/10.1090/gsm/038>
- [24] JOHNSTONE, I. (2000). *On the Distribution of the Largest Principal Component*. Department of Statistics, Stanford Univ., Stanford.
- [25] KIM, C., BANDEIRA, A. S. and GOEMANS, M. X. (2017). Community detection in hypergraphs, spiked tensor models, and sum-of-squares. In *Sampling Theory and Applications (SampTA), 2017 International Conference on* 124–128. IEEE, Los Alamitos.
- [26] LEDOUX, M. (2001). *The Concentration of Measure Phenomenon. Mathematical Surveys and Monographs* **89**. Amer. Math. Soc., Providence, RI. MR1849347
- [27] LESIEUR, T., KRZAKALA, F. and ZDEBOROVÁ, L. (2017). Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications. *J. Stat. Mech. Theory Exp.* **7** 073403, 86. MR3683819 <https://doi.org/10.1088/1742-5468/aa7284>
- [28] MAÏDA, M. (2007). Large deviations for the largest eigenvalue of rank one deformations of Gaussian ensembles. *Electron. J. Probab.* **12** 1131–1150. MR2336602 <https://doi.org/10.1214/EJP.v12-438>
- [29] MONTANARI, A., REICHMAN, D. and ZEITOUNI, O. (2017). On the limitation of spectral methods: From the Gaussian hidden clique problem to rank one perturbations of Gaussian tensors. *IEEE Trans. Inf. Theory* **63** 1572–1579. MR3625981 <https://doi.org/10.1109/TIT.2016.2637959>
- [30] PÉCHÉ, S. (2006). The largest eigenvalue of small rank perturbations of Hermitian random matrices. *Probab. Theory Related Fields* **134** 127–173. MR2221787 <https://doi.org/10.1007/s00440-005-0466-z>
- [31] PERRY, A., WEIN, A. S. and BANDEIRA, A. S. (2016). Statistical limits of spiked tensor models. Preprint. Available at [arXiv:1612.07728](https://arxiv.org/abs/1612.07728).
- [32] PERRY, A., WEIN, A. S., BANDEIRA, A. S. and MOITRA, A. (2018). Optimality and sub-optimality of PCA I: Spiked random matrix models. *Ann. Statist.* **46** 2416–2451. MR3845022 <https://doi.org/10.1214/17-AOS1625>
- [33] RICHARD, E. and MONTANARI, A. (2014). A statistical model for tensor PCA. In *Advances in Neural Information Processing Systems* 2897–2905.
- [34] ROS, V., BEN AROUS, G., BIROLI, G. and CAMMAROTA, C. (2018). Complex energy landscapes in spiked-tensor and simple glassy models: Ruggedness, arrangements of local minima and phase transitions. Preprint. Available at <https://arxiv.org/abs/1804.02686>.
- [35] SCHOENBERG, I. J. (1942). Positive definite functions on spheres. *Duke Math. J.* **9** 96–108. MR0005922
- [36] STROOCK, D. W. and VARADHAN, S. R. S. (2006). *Multidimensional Diffusion Processes. Classics in Mathematics*. Springer, Berlin. MR2190038
- [37] SUBAG, E. (2017). The complexity of spherical p -spin models—a second moment approach. *Ann. Probab.* **45** 3385–3450. MR3706746 <https://doi.org/10.1214/16-AOP1139>
- [38] VERSHYNIN, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge Series in Statistical and Probabilistic Mathematics* **47**. Cambridge Univ. Press, Cambridge. MR3837109 <https://doi.org/10.1017/9781108231596>
- [39] ZDEBOROVÁ, L. and KRZAKALA, F. (2016). Statistical physics of inference: Thresholds and algorithms. *Adv. Phys.* **65** 453–552.