

SHARP THRESHOLD SEQUENCE AND UNIVERSALITY FOR ISING PERCEPTRON MODELS

SHUTA NAKAJIMA* AND NIKE SUN^o

ABSTRACT. We study a family of Ising perceptron models with $\{0, 1\}$ -valued activation functions. This includes the classical half-space models, as well as some of the symmetric models considered in recent works. For each of these models we show that the free energy is self-averaging, there is a sharp threshold sequence, and the free energy is universal with respect to the disorder. A prior work of C. Xu (2019) used very different methods to show a sharp threshold sequence in the half-space Ising perceptron with Bernoulli disorder. Recent works of Perkins–Xu (2021) and Abbe–Li–Sly (2021) determined the sharp threshold and limiting free energy in a symmetric perceptron model. The results of this paper apply in more general settings, and are based on new “add one constraint” estimates extending Talagrand’s estimates for the half-space model (1999, 2011).

CONTENTS

1. Introduction	1
2. Talagrand’s results on the half-space perceptron	6
3. Adding a single constraint in perceptron models	14
4. Universality on average in perceptron models	24
5. Concentration, sharp threshold sequence, and universality	29
References	34

1. INTRODUCTION

The **perceptron** is a classical model in high-dimensional probability theory [Cov65, Wen62], which can be interpreted as a toy model of a simple neural network [Gar88]. In the simplest version, the **Ising (half-space) perceptron** refers to the random set S defined by intersecting the discrete cube $\{-1, +1\}^N$ with $M = N\alpha$ i.i.d. random half-spaces (formal definitions below). In the late 1980s it was shown that heuristic analytical tools of statistical physics (i.e., the “replica method” or “cavity method”) can be applied to derive precise predictions on the limiting behavior of the perceptron model [Gar88, GD88, KM89, Méz89]. In particular, it was predicted that the **free energy** $N^{-1} \log |S|$ of the model concentrates around an explicit constant (depending on the parameters of the model), and that there is a **sharp threshold** α_c such that $\mathbb{P}(|S| > 0)$ transitions from $1 - o_N(1)$ to $o_N(1)$ in an $o_N(1)$ window around α_c .

In contrast with what has been conjectured via physics heuristics, mathematical understanding of Ising perceptron models remains quite limited. For the half-space Ising perceptron, it follows by a trivial first moment calculation that $\mathbb{P}(|S| > 0) = o_N(1)$ for $\alpha = M/N$ large enough, but this bound is typically not tight. On the other hand, it was shown by [KR98] (see also [Tal99a]) that $\mathbb{P}(|S| > 0) = 1 - o_N(1)$ for $\alpha = M/N$ small enough. Moreover, for small enough α , the free energy concentrates around the value conjectured by physicists [Tal00, Tal11, BNSX21]. A result of [DS18] shows that, under an additional condition, the random set S is nonempty with *positive* probability for any α smaller than the critical value predicted by physicists [KM89]. These results leave open the question of whether the model has a sharp threshold, and whether the free energy concentrates for general α .

For the half-space Ising perceptron with **Bernoulli disorder** — meaning that the half-spaces point in directions chosen uniformly at random from $\{-1, +1\}^N$ — C. Xu applied Hatami’s pseudo-junta theorem [Hat12] to show that the model has a **sharp threshold sequence** [Xu21]. In a different direction, numerous recent works have

Date: April 8, 2022.

*Graduate School of Science and Technology, Meiji University, Tokyo. ^oDepartment of Mathematics, Massachusetts Institute of Technology.

studied **symmetric** perceptron models, which are significantly more tractable than the half-space models and can be analyzed for much finer properties [APZ19, PX21, ALS21b, ALS21a, GKPX22]. In particular, for a **symmetric interval** variant of the Ising perceptron (intersecting $\{-1, +1\}^N$ with i.i.d. random symmetric *slabs* rather than half-spaces), recent work rigorously pinpoints the limiting free energy and sharp threshold [PX21, ALS21b].

In this paper we study a family of Ising perceptron models with $\{0, 1\}$ -valued activation functions (defined more formally below). This includes the classical half-space models as well as some of the symmetric models considered in more recent works. Further, we allow the disorder random variables to come from any **subgaussian** distribution, in contrast with previous works which have assumed gaussian or Bernoulli disorder. For each of these models we show:

- There is a **sharp threshold sequence**, meaning that $\mathbb{P}(|S| > 0)$ transitions from $1 - o_N(1)$ to $o_N(1)$ in an $o_N(1)$ window of α values, although the location of the transition may depend on N .
- The model is **self-averaging**, meaning that the free energy $N^{-1} \log |S|$ concentrates around a deterministic value which may depend on N .
- The free energy and sharp threshold sequence are **universal** with respect to the disorder.

Our sharp threshold sequence result extends the main theorem of [Xu21], although by very different methods — our approach does not use Hatami’s theorem or other tools of boolean analysis.

For the half-space Ising perceptron model with gaussian or Bernoulli disorder, the sharp threshold sequence and self-averaging property appear to be relatively straightforward consequences of “add one constraint” estimates proved in [Tal99b, Tal11] (and discussed below). While Talagrand does indicate the implications of his estimates for self-averaging, to the best of our knowledge he does not explicitly address the sharp threshold problem; and one of the aims of this article is to clearly spell out this simple connection. The main technical contributions of this paper are new “add one constraint” estimates for general (not necessarily symmetric) interval perceptron models estimates with general (subgaussian) disorder. These results cannot be deduced by an easy generalization of Talagrand’s arguments, and we describe some of the main ideas in §1.1 below. We then apply our “add one constraint” estimates to deduce the sharp threshold sequence and free energy concentration results. Lastly we combine these with some central limit theorem considerations to deduce the universality results. A weaker version of our universality results (comparing gaussian with Bernoulli disorder, and with additional smoothness assumptions) appears in [Tal11, §9.9].

The remainder of this introductory section is organized as follows:

- In §1.1 we formally define the various perceptron models that we consider. We state our first set of results (Theorems 1.3–1.5) which control the effect of adding a single constraint in these perceptron models. We also describe some of the basic proof ideas.
- In §1.2 we state our second set of results (Theorems 1.7–1.11) which address self-averaging, sharp thresholds, and universality in the perceptron models. These may be viewed as consequences of the “add one constraint” estimates described in §1.1.
- In §1.3 we discuss the context given by the surrounding literature on perceptron models.

At the end of this section we give the outline for the rest of the paper.

1.1. “Add one constraint” estimates for perceptron models. In the above we used S to denote the Ising perceptron solution set, a random subset of $\{-1, +1\}^N$. In this paper we only consider the size of the random set, and we hereafter always denote $Z \equiv |S|$. We will prove results for the classical Ising perceptron model

$$Z \equiv Z_{M,N}(\kappa, \infty, \xi) \equiv \sum_{\sigma \in \{-1, +1\}^N} \prod_{k \leq M} \mathbf{1} \left\{ \frac{(\xi^k, \sigma)}{N^{1/2}} \geq \kappa \right\}, \quad (1.1)$$

which we will refer to as the **half-space (Ising) perceptron**. We also consider more general models of the form

$$Z_{M,N}(U; \xi) \equiv \sum_{\sigma \in \{-1, +1\}^N} \prod_{k \leq M} U \left(\frac{(\xi^k, \sigma)}{N^{1/2}} \right) \quad (1.2)$$

where $U : \mathbb{R} \rightarrow [0, 1]$ (measurable) is the **activation function**. We refer to this as the U -perceptron; most of this paper concerns the case that U is $\{0, 1\}$ -valued. We refer to the case $U(x) = \mathbf{1}\{x \in [a, b]\}$ as the **interval (Ising)**

perceptron. The **symmetric interval perceptron** mentioned above corresponds to the case $a = -b$. We assume that the ξ, ξ^k ($k \geq 1$) are i.i.d. random vectors satisfying the following:

Assumption 1.1. The ξ_i are i.i.d. random variables with mean zero, unit variance, such that

$$\mathbb{E} \exp(\lambda \xi_i) \leq \exp\left(\frac{\lambda^2 \nu}{2}\right)$$

for all $\lambda \in \mathbb{R}$ (that is, the ξ_i are subgaussian with variance proxy ν). This assumption includes the two most widely studied cases of the perceptron model: **gaussian disorder** (the ξ^k are standard gaussian vectors) and **Bernoulli disorder** (the ξ^k are sampled uniformly at random from $\{-1, +1\}^N$). Note that in general we must have $\nu \geq 1$, since the cumulant-generating function $\mathfrak{R}(\lambda) \equiv \log \mathbb{E} \exp(\lambda \xi_i)$ satisfies $\mathfrak{R}''(0) = \text{Var } \xi_i = 1$.

We first restate a result of Talagrand:

Theorem 1.2 (half-space perceptron, gaussian disorder [Tal99b, Tal11]). *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(N\delta)$, where δ is a small positive constant. If g is a standard gaussian random vector in $\mathbb{R}^N$, then*

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(g, \sigma)}{N^{1/2}} \geq \kappa \right\} \right| < \frac{1}{\exp(w)} \right) \leq \exp\left(-\frac{w}{C_\delta}\right).$$

provided $C_\delta \leq w \leq N/C_\delta$, where C_δ is a large finite constant that depends only on κ and δ .

Theorem 1.2 was in fact already proved by Talagrand in the more difficult setting of Bernoulli disorder (see [Tal99b, Propn. 2.3]). The result for the gaussian case follows by a simplification of the argument for the Bernoulli case, which we review in §2.1. Most of the argument for the gaussian case also appears in [Tal11, Ch. 9]. (The latter reference [Tal11] does not appear to treat the case of small δ , but this was already done previously in [Tal99b].)

As we discuss further in §2.2–2.3 below, the proof of Theorem 1.2 relies on two main ingredients:

- (i) a lower bound on the expected supremum of a stochastic process; and
- (ii) a concentration estimate for the supremum of a stochastic process.

In the gaussian setting these ingredients are supplied by the well-known **Sudakov minoration lower bound** and **Borell–TIS (Tsirelson–Ibragimov–Sudakov) concentration inequality**, restated in Lemmas 2.1 and 2.2 below. For more general distributions, however, both of these ingredients can be quite nontrivial.

By contrast, our next result gives a weaker bound than Theorem 1.2, but has the advantage that it uses only the simplest gaussian versions of the results mentioned above (Sudakov minoration and Borell–TIS). We transfer these bounds to general distributions by a rather weak form of the multivariate CLT. As a result, the bound of Theorem 1.3 is almost certainly suboptimal, but it has the advantage that it is relatively straightforward to derive, and applies for a general class of subgaussian distributions. To the best of our knowledge, using existing methods, the argument of Talagrand for Theorem 1.2 may only be extended to a limited subclass of subgaussian distributions (see §2.2–2.3 for the details).

Theorem 1.3 (half-space perceptron, general disorder). *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(N\delta)$, where δ is a small positive constant. Suppose ξ is a random vector in $\mathbb{R}^N$ satisfying Assumption 1.1. Then*

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \geq \kappa \right\} \right| < \frac{1}{\exp(w)} \right) \leq \exp\left(-\frac{w}{C_\delta}\right).$$

provided $C_\delta \leq w \leq N^{1/2}/C_\delta$, where C_δ is a large finite constant that depends only on κ, ν , and δ .

The next two results concern the U -perceptron (1.2), under the assumption that we have $U : \mathbb{R} \rightarrow \{0, 1\}$ with $U(x) \geq \mathbf{1}\{x \in [a, b]\}$ for some $-\infty < a < b < \infty$.

Theorem 1.4 (U -perceptron, gaussian disorder). *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(N\delta)$ where δ is a small positive constant. If g is a standard gaussian random vector in $\mathbb{R}^N$, then*

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(g, \sigma)}{N^{1/2}} \in [a, b] \right\} \right| < \frac{1}{\exp(w)} \right) \leq \exp\left(-\frac{w}{C_\delta}\right)$$

provided $C_\delta \leq w \leq N^{1/2}/C_\delta$, where C_δ is a large finite constant that depends only on a, b , and δ .

Theorem 1.5 (*U-perceptron, general disorder*). *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(N\delta)$ where δ is a small positive constant. Suppose ξ is a random vector in $\mathbb{R}^N$ satisfying Assumption 1.1. Then*

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \in [a, b] \right\} \right| < \frac{1}{\exp(w)} \right) \leq \exp\left(-\frac{w}{C_\delta}\right)$$

provided $C_\delta \leq w \leq N^{1/3}/C_\delta$, where C_δ is a large finite constant that depends only on a, b, v , and δ .

We now describe some of the basic ideas that appear in the proofs of Theorems 1.3–1.5. We repeatedly leverage a basic observation that comes from Talagrand’s proof of Theorem 1.2: if S is a large subset of $\{-1, +1\}^N$, then it must contain many elements that are well separated in Hamming distance. If g is a standard gaussian vector in $\mathbb{R}^N$, then the separation guarantee can be used to show that

$$\max \left\{ \frac{(g, \sigma)}{N^{1/2}} : \sigma \in S \right\}$$

is likely to be quite large. The details of this argument are reviewed in §2.1 below. We point out that it relies on theorems for gaussian processes, namely, concentration of the supremum and a Sudakov minoration lower bound on the expectation of the supremum. As we discuss in §2.2–2.3, these do not easily extend to more general subgaussian distributions.

In this paper we devise a more flexible version of Talagrand’s approach, as follows. First we decompose $[N] \equiv \{1, \dots, N\}$ into L blocks of size K each, where $KL = N$. Denote the blocks $I_1, \dots, I_L$. For any subset $I \subseteq [N]$ let $\sigma_I \equiv (\sigma_i)_{i \in I}$; we will say that two configurations $\sigma, \tau \in \{-1, +1\}^N$ are well-separated on I if σ_I and τ_I are not too close (see Definition 3.1 below). We show that if $S \subseteq \{-1, +1\}^N$ is large, then it must contain many elements that are well-separated on a positive fraction of the blocks — say, on the final $L\gamma$ blocks. We then consider the process

$$M_k(\sigma) = \frac{1}{N^{1/2}} \sum_{j=1}^k (\xi_{I_j}, \sigma_{I_j}) \quad (1.3)$$

for $\sigma \in S$, starting from $M_0(\sigma) = 0$ and with the goal of having say $M_L(\sigma) \in [a, b]$ (for the U -perceptron (1.2), assuming $U(x) \geq \mathbf{1}\{x \in [a, b]\}$). In the case that ξ is **gaussian**, we can apply Talagrand’s estimates for the half-space model (1.1) to control the increments of $M_k(\sigma)$ over the final $L\gamma$ steps, using the separation guarantees on the final $L\gamma$ blocks. With this approach we show that, with very good probability, S must contain many elements σ with $M_L(\sigma) \in [a, b]$.

In the case that ξ follows a more general **subgaussian** distribution, we follow a similar strategy of first obtaining estimates for the half-space model (1.1), then using a block decomposition to obtain estimates for the U -perceptron (1.2). However, as we already noted, it does not appear that Talagrand’s results for the half-space perceptron extend easily to the full class of subgaussian distributions. Instead, to prove the desired estimates in the half-space model, we again use the block decomposition approach, combined with a (rather weak) form of the multivariate central limit theorem which allows us to borrow the estimates from the gaussian setting. In particular, we restrict ourselves to quantities of the form

$$\max \left\{ \frac{(\xi, \sigma)}{N^{1/2}} : \sigma \in X \right\}$$

for subsets $X \subseteq \{-1, +1\}^N$ of *bounded* size, allowing us to apply central limit theorems only in bounded dimensions. This results in weaker bounds which can potentially be improved by appealing to stronger forms of the central limit theorem.

1.2. Self-averaging, sharp threshold sequence, and universality. We next describe the main results which we obtain as consequences of the “add one constraint” estimates presented in §1.1. To emphasize the dependence on the “add one constraint” estimates, we state some of our results for a more abstract model of the form

$$Z \equiv Z_{M,N} \equiv \sum_{\sigma \in \{-1, +1\}^N} \prod_{k \leq M} \Theta_k(\sigma), \quad (1.4)$$

where the Θ_k are a sequence of $[0, 1]$ -valued random functions, adapted to a filtration $(\mathcal{F}_k)_{k \geq 0}$, satisfying a fairly weak “add one constraint” estimate:

Assumption 1.6. For the model (1.4), let Z_{M+1} be the partition function that results from introducing one more factor Θ_{M+1} . Suppose that for all $\delta > 0$ small enough that on the event $Z_M \geq \exp(N\delta)$ we have

$$\mathbb{P}\left(\frac{Z_{M+1}}{Z_M} \leq \frac{1}{\exp(w)} \mid \mathcal{F}_M\right) \leq f_\delta(w)$$

for all $C_\delta \leq w \leq w_{\max}$, such that the function f_δ satisfies the bound

$$\int_{C_\delta}^{w_{\max}} 2w f_\delta(w) dw \leq C_{\delta,2}. \quad (1.5)$$

In the above, C_δ and $C_{\delta,2}$ are finite constants that depend only on the model and on δ , while $w_{\max}$ can depend on the model as well as on δ and N .

The half-space and U -perceptron models (1.1) and (1.2) are clearly a special case of the model (1.4). Therefore, Theorems 1.2–1.5 imply that the perceptron models (1.1) and (1.2) satisfy a stronger condition than Assumption 1.6; see Assumption 5.1 at the start of Section 5. However, we will show that the weaker Assumption 1.6 suffices for some similar results. Denote $\log_{N\delta} Z \equiv \max\{\log Z, N\delta\}$.¹

Theorem 1.7 (concentration of free energy). *For the model (1.4), under Assumption 1.6, we have*

$$\lim_{N \rightarrow \infty} \left| \log_{N\delta} Z - \mathbb{E} \log_{N\delta} Z \right| = 0$$

for any positive constant δ , where the limit holds in probability as $N \rightarrow \infty$.

Our next result concerns the behavior of the probability $\mathbb{P}(Z_M > 0)$ for the model (1.4). For this question, we may as well assume that the Θ_a are $\{0, 1\}$ -valued, since otherwise we can replace $\Theta_a(x)$ with $\mathbf{1}\{\Theta_a(x) > 0\}$. We introduce one more assumption below, which together with Assumption 1.6 ensures that $\mathbb{P}(Z_M > 0)$ has a transition in the proportional regime $M \asymp N$. This assumption is typically easily verifiable in practice:

Assumption 1.8. There is a positive constant c such that $\mathbb{E}(\Theta_{k+1} \mid \mathcal{F}_k) \leq \exp(-c)$ almost surely for all $k \geq 0$.

Theorem 1.9 (sharp threshold sequence). *For the model (1.4), if the Θ_a are $\{0, 1\}$ -valued and satisfy Assumptions 1.6 and 1.8, then there is a sharp threshold sequence: that is to say, there is a sequence $\alpha_N \asymp 1$ such that $\mathbb{P}(Z_{N\alpha_N} > 0)$ transitions from $1 - o_N(1)$ to $o_N(1)$ in an $o_N(1)$ window around α_N .*

In particular, by Theorems 1.2–1.5, the results of Theorems 1.7 and 1.9 apply to the perceptron models (1.1) and (1.2), assuming in the latter model that U is $\{0, 1\}$ -valued with $U(x) \geq \mathbf{1}\{x \in [a, b]\}$, and further assuming the ξ^k satisfy Assumption 1.1. For the sharp threshold question, we can replace $U(x)$ by $\bar{U}(x) = \mathbf{1}\{U(x) > 0\}$, so we can allow the original U to take all values in $[0, \infty]$. Thus Theorem 1.9 implies the following:

Corollary 1.10. *For a measurable function $U : \mathbb{R} \rightarrow [0, 1]$, let $\bar{U}(x) = \mathbf{1}\{U(x) > 0\}$. If $\bar{U}(x) \geq \mathbf{1}\{x \in [a, b]\}$ for some $-\infty < a < b < \infty$, and $\bar{U}$ is not almost everywhere one, then the results of Theorem 1.9 hold for the model (1.2), provided that the ξ^k satisfy Assumption 1.1.*

Proof. Since we are only interested in whether the partition function Z of (1.2) is nonnegative, we may assume without loss that $U = \bar{U}$. By the condition $U(x) \geq \mathbf{1}\{x \in [a, b]\}$, combined with Assumption 1.1 and Theorem 1.5, the model (1.2) satisfies Assumption 5.1. The condition that U is not almost everywhere one guarantees that Assumption 1.8 is also satisfied. The claim then directly follows from Theorem 1.9. $\square$

Finally, we can combine the above results with some central limit theorem estimates to obtain the following:

Theorem 1.11 (universality). *In the perceptron model (1.2), suppose U is $\{0, 1\}$ -valued and piecewise continuous, and neither identically one nor identically zero. Suppose the ξ^k are i.i.d. random vectors satisfying Assumption 1.1. Then the threshold sequence and free energy are universal with respect to the disorder.*

¹This differs slightly from the notation of [Tal11, §8.3] which has $\log_A x = \max\{\log x, -A\}$. We choose a notation which is more convenient for our choice of normalization.

1.3. Related work. As mentioned earlier, the perceptron model was analyzed in the (nonrigorous) statistical physics literature in the 1980s [Gar88, GD88, KM89, Méz89]. We refer to [Cov65, Gar88, MRSY19, MZZ21] for discussions of statistical motivations for the model.

There is an important variant of the perceptron model which is well understood at a rigorous level: the **spherical (half-space) perceptron**, defined by intersecting the sphere $N^{1/2}\mathbf{S}^{N-1}$ with i.i.d. random **spherical caps**. When the spherical caps are exactly *hemispheres* (pointing in uniformly random directions), the critical threshold $\alpha_c = 2$ was determined by combinatorial arguments [Cov65, Wen62]. Later work showed that as long as the spherical caps have *up to half* the volume of the entire sphere, the limiting free energy and critical threshold coincide with the physics predictions [ST03, Sto13]. These results rely crucially on the **convex** nature of the model. In the regime where the spherical caps have *more than half* the volume of the entire sphere, the problem is no longer convex, and the expected behavior is more complicated [FP16, FPS⁺17, AS20].

For **Ising** perceptron models, in comparison with the spherical versions, rather less has been rigorously proved. Many of the relevant works were already mentioned at the start of the section. There has been a very fruitful line of works investigating **symmetric** perceptron models, which are much more tractable because the **moment method** typically gives fairly sharp results [APZ19, PX21, ALS21b, ALS21a, GKPX22]. Several of these recent works have investigated algorithmic properties of the solution landscape, inspired in part by conjectures in the physics literature (see e.g. [BBC⁺16]). In particular, for a symmetric interval perceptron model — more precisely, the model (1.2) with $U(x) = \mathbf{1}\{x \in [-\kappa, \kappa]\}$ and gaussian or Bernoulli disorder ξ — the sharp threshold and free energy were determined by [PX21, ALS21b]. The symmetric interval model is also related to the discrepancy minimization problem; see e.g. [Spe85, Ban10, LM15, TMR20, ALS21c].

For non-symmetric Ising perceptron models, it is much more difficult to understand the typical size or behavior of the perceptron solution set S . The moment method does not give sharp results. Bounds on the critical window were given by [KR98, Tal99a]; and the free energy was computed for small α by [Tal00, Tal11, BNSX21]. A result of [DS18] shows that, under an additional condition, the random set S is nonempty with *positive* probability throughout the predicted regime; our result allows to improve this statement to high probability. For the half-space model with Bernoulli disorder, [Xu21] proved a sharp threshold sequence using a characterization of Hatami for low-influence boolean functions [Hat12]. It would be of interest to see if the methods of [Xu21] can be extended to more general perceptron models, in ways that do not require more precise estimates on these models [Xu22]. With respect to our current paper, the most closely related previous results are the estimates obtained by Talagrand for the half-space perceptron model [Tal99b, Tal11].

Organization. The remainder of this paper is organized as follows:

- In Section 2 we review Talagrand’s **proof of Theorem 1.2**, for the half-space perceptron (1.1) with gaussian disorder. We indicate the obstructions to extending this result to more general disorder distributions.
- In Section 3 we prove our main “add one constraint” estimates, Theorems 1.3–1.5, for the models (1.1) and (1.1) with general disorder.
- In Section 4 we give a weak form of the multivariate central limit theorem, and use it to prove an averaged universality statement, Theorem 4.1, for the perceptron models (1.1) and (1.2).
- In Section 5 we prove Theorems 1.7–1.11.

Acknowledgements. We wish to thank David Belius, Erwin Bolthausen, Ryoki Fukushima, Elchanan Mossel, Joe Neeman, and Changji Xu for many interesting conversations. S.N. is supported in part by SNSF grant 176918. N.S. is supported in part by NSF CAREER grant DMS-1940092 and NSF-Simons grant DMS-2031883.

2. TALAGRAND’S RESULTS ON THE HALF-SPACE PERCEPTRON

In this section we review Talagrand’s results on the half-space perceptron (1.1). We begin with the case of gaussian disorder, and then discuss the possibility of extending to more general distributions of ξ :

- In §2.1 we review Talagrand’s **proof of Theorem 1.2**. We highlight two key ingredients in the proof, Lemma 2.1 (Sudakov minoration) and Lemma 2.2 (Borell–TIS).
- In §2.2 we discuss extensions of Lemma 2.2 to more general distributions.

- In §2.3 we discuss extensions of Lemma 2.1 to more general distributions.

Throughout this paper we use C , $\bar{C}$, C' , and C'' to denote absolute constants. Following common convention, the value of the constant may change from one occurrence to the next, but in a way that does not depend on N or any of the parameters of the model. We use the different labels C , $\bar{C}$, C' , and C'' to avoid ambiguities when different constants interact in the same proof.

2.1. Intersection of cube and half-space with gaussian disorder. In this subsection we review some results from [Tal99b, Tal11], including the proof of Theorem 1.2. As commented above, the paper [Tal99b] considers the model (1.1) where the sum goes over $x \in \{-1, +1\}^N$, and the g^a are replaced with ξ^a which are i.i.d. uniform from $\{-1, +1\}^N$. Meanwhile [Tal11, Ch. 9] treats the gaussian case, but presents a simplified argument that does not appear to handle the small δ regime. We present a more complete summary of the gaussian case below. We begin by recalling two well-known results:

Lemma 2.1 (Sudakov minoration for gaussian processes). *Let $(u_i)_{i \leq n}$ be a centered gaussian process with $\mathbb{E}[(u_i)^2] = 1$ for all i , and $\mathbb{E}(u_i u_j) \leq 1 - \epsilon < 1$ for all $i \neq j$. Then*

$$\mathbb{E}\left(\max_{i \leq n} u_i\right) \geq \frac{(\epsilon \log n)^{1/2}}{2^{1/2}}.$$

Proof. See for example [BLM13, Thm. 13.4]. This is discussed further in §2.3 below. $\square$

Lemma 2.2 (Borell–TIS inequality). *Let $(u_i)_{i \leq n}$ be a centered gaussian process with $\mathbb{E}[(u_i)^2] = 1$ for all i . Let $u_{\max} \equiv \max_{i \leq n} u_i$. Then, for any $s \geq 0$, we have*

$$\max\left\{\mathbb{P}\left(u_{\max} - \mathbb{E}u_{\max} \geq s\right), \mathbb{P}\left(u_{\max} - \mathbb{E}u_{\max} \leq -s\right)\right\} \leq \frac{1}{\exp(s^2/2)}.$$

Proof. See for example [BLM13, Thm. 5.8]. The concentration for $u_{\max}$ is the consequence of a more general concentration result for Lipschitz functionals of gaussian processes; see [BLM13, Thm. 5.6]. This is discussed further in §2.3 below. $\square$

An immediate consequence of the two preceding lemmas is the following:

Corollary 2.3. *Let $(u_i)_{i \leq n}$ be a centered gaussian process with $\mathbb{E}[(u_i)^2] = 1$ for all i , and $\mathbb{E}(u_i u_j) \leq 1 - \epsilon < 1$ for all $i \neq j$. Let $u_{\max} \equiv \max_{i \leq n} u_i$. Then*

$$\mathbb{P}\left(u_{\max} \leq \frac{(\epsilon \log n)^{1/2}}{2}\right) \leq \frac{1}{n^{\epsilon/50}}.$$

Proof. It follows by Lemma 2.1 combined with Lemma 2.2 that

$$\mathbb{P}\left(u_{\max} \leq \frac{(\epsilon \log n)^{1/2}}{2}\right) \leq \mathbb{P}\left(u_{\max} - \mathbb{E}u_{\max} \leq -\left(\frac{1}{2^{1/2}} - \frac{1}{2}\right)(\epsilon \log n)^{1/2}\right) \leq \frac{1}{n^{\epsilon/50}},$$

as claimed. $\square$

We next record some basic notations which will be used throughout:

Definition 2.4. For $t = 1 - \epsilon$, denote the binary relative entropy function

$$k_2(t) \equiv H\left(\frac{1+t}{2} \middle| \frac{1}{2}\right) = \frac{1+t}{2} \log(1+t) + \frac{1-t}{2} \log(1-t) \equiv \log 2 - \psi_2(\epsilon), \quad (2.1)$$

so $0 \leq k_2(t) \leq \log 2$ with $k_2(0) = 0$ and $k_2(-1) = k_2(1) = \log 2$. It will also be useful to recall a simplified bound: for $0 \leq p \leq 1$ and $0 \leq tp \leq 1$ we have

$$H(tp | p) = tp \log t + (1 - tp) \log \frac{1 - tp}{1 - p} \geq tp \log t + (1 - tp) \log(1 - tp) \geq tp \log \frac{t}{e}, \quad (2.2)$$

where the last step uses that $(1 - x) \log(1 - x) \geq -x$ for all $x \leq 1$. (Note however that since the relative entropy is always nonnegative, the bound (2.2) is vacuous unless $t \geq e$.)

The following estimate says that in any large subset of $\{-1, +1\}^N$, most pairs of elements of that subset are separated in Hamming distance, where the separation guarantee depends on the size of the set:

Lemma 2.5 (rephrasing of [Tal99b, Lem. 2.2]). *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(N\delta)$. If μ denotes the uniform probability measure on S , then*

$$\mu^{\otimes 2} \left(\left| \frac{(\sigma^{(1)}, \sigma^{(2)})}{N} \right| \geq 1 - \epsilon \right) \leq \frac{2}{\exp(N\delta/2)},$$

for any $\epsilon \in (0, 1)$ with $\psi_2(\epsilon) \leq \delta/2$.²

Proof. Abbreviate $t \equiv 1 - \epsilon$. Conditioning on $\sigma^{(1)}$, we have

$$\begin{aligned} \mu \left(\left\{ \sigma^{(2)} : \left| \frac{(\sigma^{(1)}, \sigma^{(2)})}{N} \right| \geq t \right\} \right) &= \frac{1}{|S|} \left| \left\{ \sigma^{(2)} \in \{-1, +1\}^N : \left| \frac{(\sigma^{(1)}, \sigma^{(2)})}{N} \right| \geq t \right\} \right| \\ &\leq \frac{2^N}{\exp(N\delta)} \mathbb{P} \left(\left| \text{Bin} \left(N, \frac{1}{2} \right) - \frac{N}{2} \right| \geq \frac{Nt}{2} \right) \leq \frac{2^{N+1} \exp(-Nk_2(t))}{\exp(N\delta)} \leq \frac{2 \exp(N\psi_2(\epsilon))}{\exp(N\delta)}. \end{aligned}$$

The claim follows. $\square$

The next result is the main ingredient in Talagrand's proof of Theorem 1.2:

Proposition 2.6 (adaptation of [Tal99b, Propn. 2.1]). *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(N\delta)$ where δ is a positive constant. Let $\epsilon \in (0, 1)$ with $\psi_2(\epsilon) \leq \delta/2$. Then*

$$\mathbb{P} \left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(g, \sigma)}{N^{1/2}} \geq s\epsilon^{1/2} \right\} \right| \leq \frac{1}{\exp(C's^2)} \right) \leq \frac{1}{\exp(s^2\epsilon/C')},$$

provided $C' \leq s \leq (N\delta)^{1/2}/C'$, where C' is an absolute constant.

Proof. Let μ denote the uniform probability measure on S . The basic idea of the proof is the following. By Lemma 2.5, $\mu^{\otimes n}$ gives large weight to well-separated n -tuples of configurations. By Corollary 2.3, given a well-separated n -tuple, it is very unlikely that all n configurations violate the new constraint. This implies the desired bound. The details are as follows.

Step 1. First we show that $\mu^{\otimes n}$ gives large weight to well-separated n -tuples of configurations. Let $t = 1 - \epsilon > 0$ with $\psi_2(\epsilon) \leq \delta/2$ as in Lemma 2.5 above. For any positive integer n , let C_n denote the subset of configurations in $(\{-1, +1\}^N)^n$ that are pairwise well-separated:

$$C_n = \left\{ (\sigma^{(1)}, \dots, \sigma^{(n)}) \in (\{-1, +1\}^N)^n : \frac{(\sigma^{(i)}, \sigma^{(j)})}{N} \leq t \ \forall i \neq j \right\}.$$

It follows from Lemma 2.5 that $\mu^{\otimes n}(C_n)$ must be large: indeed, a union bound over all pairs $i < j$ gives

$$\mu^{\otimes n}(C_n) \geq 1 - \binom{n}{2} \mu^{\otimes 2} \left(\left\{ (\sigma^{(1)}, \sigma^{(2)}) : \left| \frac{(\sigma^{(1)}, \sigma^{(2)})}{N} \right| \geq t \right\} \right) \geq 1 - \binom{n}{2} \frac{2}{\exp(N\delta/2)} \geq \frac{3}{4}, \quad (2.3)$$

where the last inequality holds provided that

$$1 \leq n \leq \frac{\exp(N\delta/4)}{2}. \quad (2.4)$$

Step 2. Now let H_g denote the set of all σ such that $(g, \sigma)/N^{1/2} \geq (\epsilon \log n)^{1/2}/2$; we now lower bound the probability that $\mu(H_g)$ is too small. Let $D_{n,g}$ denote the subset of configurations in $(\{-1, +1\}^N)^n$ where all n points lie outside H_g :

$$D_{n,g} = \left\{ (\sigma^{(1)}, \dots, \sigma^{(n)}) \in (\{-1, +1\}^N)^n : \frac{(g, \sigma^{(i)})}{N^{1/2}} < \frac{(\epsilon \log n)^{1/2}}{2} \ \forall 1 \leq i \leq n \right\} = \left(\{-1, +1\}^N \setminus H_g \right)^n.$$

²Note this result is similar to but stronger than [Tal11, Lem. 9.2.1].

On the event that $\mu(H_g) \leq 1/(4n)$, the complement of $D_{n,g}$ has measure at most $1/4$ under $\mu^{\otimes n}$, so

$$\frac{3}{4} \stackrel{(2.3)}{\leq} \mu^{\otimes n}(C_n) \leq \mu^{\otimes n}(C_n \cap D_{n,g}) + \frac{1}{4}. \quad (2.5)$$

Rearranging the above gives $\mu^{\otimes n}(C_n \cap D_{n,g}) \geq 1/2$ on the event $\mu(H_g) \leq 1/(4n)$. It follows that

$$\mathbb{P}\left(\mu(H_g) \leq \frac{1}{4n}\right) \stackrel{(2.5)}{\leq} \mathbb{P}\left(\mu^{\otimes n}(C_n \cap D_{n,g}) \geq \frac{1}{2}\right) \leq 2\mathbb{E}\left[\mu^{\otimes n}(C_n \cap D_{n,g})\right] \leq \frac{2}{n^{\epsilon/50}},$$

where the intermediate step is by Markov's inequality, and the last step is by Corollary 2.3. Making a change of variables $s^2 = (\log n)/4$ gives the claim. $\square$

Proof of Theorem 1.2. If $C' \leq s \leq (N\delta)^{1/2}/C'$ and $s\epsilon^{1/2} \geq \kappa$, then Proposition 2.6 immediately implies

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(g, \sigma)}{N^{1/2}} \geq \kappa \right\} \right| \leq \frac{1}{\exp(C's^2)} \right) \leq \frac{1}{\exp(s^2\epsilon/C')}.$$

The claim follows. Note that if $\kappa \leq 0$ then the condition $s\epsilon^{1/2} \geq \kappa$ holds trivially for all $s \geq 0$, so the C_δ in the theorem statement can be taken independently of κ . By contrast, if $\kappa > 0$ then C_δ depends on κ as well. $\square$

For comparison, the next proposition is a variant of Proposition 2.6 which was also proved by Talagrand but with different methods. It gives similar although not exactly comparable results. In this paper we will only make use of Proposition 2.6, so Proposition 2.7 can be skipped by the reader.

Proposition 2.7 (adaptation of [Tal11, Thm. 8.2.4]). *Let $(u_i)_{i \leq n}$ be a centered gaussian process with $\mathbb{E}[(u_i)^2] = 1$ for all i , and $\mathbb{E}[u_i u_j] \leq 1 - \epsilon$ for all $i \neq j$. Then*

$$\mathbb{P}\left(\frac{1}{n} \left| \left\{ i \leq n : u_i \geq s\epsilon^{1/2} \right\} \right| < \frac{1}{\exp(\bar{C}s^2/\epsilon)} \right) \leq \frac{1}{\exp(s^2/2)}$$

provided $\bar{C} \leq s \leq (\log n)^{1/2}/\bar{C}$, where $\bar{C}$ is an absolute constant.

Proof. Fix $s > 0$ and consider the functions

$$F(u) \equiv \log \left(\sum_{i \leq n} \exp(su_i) \right), \quad F_2(u) \equiv \log \left(\sum_{i \leq n} \exp(2su_i) \right)$$

It follows from [Tal11, Propn. 8.2.2] (based on [Tal11, Lem. 8.2.1]) that if $(u_i)_{i \leq n}$ and $(v_i)_{i \leq n}$ are two centered gaussian processes with $\mathbb{E}[(u_i)^2] \geq \mathbb{E}[(v_i)^2]$ for all i , and $\mathbb{E}[u_i u_j] \leq \mathbb{E}[v_i v_j]$ for all $i \neq j$, then $\mathbb{E}F(u) \geq \mathbb{E}F(v)$. In the current setting we can take $v_i = (1 - \epsilon)^{1/2}z + \epsilon^{1/2}z_i$ where z, z_i are i.i.d. standard gaussian random variables. It follows by combining with [Tal11, Lem. 8.2.3] that

$$\mathbb{E}F(u) \geq \mathbb{E}F(v) = \log \left(\sum_{i \leq n} \exp(s\epsilon^{1/2}z_i) \right) \geq \log n + \frac{s^2\epsilon}{5},$$

where the last inequality holds provided $C \leq s\epsilon^{1/2} \leq (\log n)^{1/2}/C$, where C is an absolute constant. Moreover, gaussian concentration gives

$$\mathbb{P}\left(F(u) \leq \mathbb{E}F(u) - t\right) \leq \exp\left(-\frac{t^2}{4s^2}\right),$$

and combining the last two bounds implies

$$\mathbb{P}\left(F(u) \geq \log n + \frac{s^2\epsilon}{10}\right) \leq 1 - \exp\left(-\frac{s^2\epsilon}{400}\right).$$

On the other hand, it follows using Markov's inequality that

$$\mathbb{P}\left(F_2(u) \geq \log n + 3s^2\right) \leq \frac{\mathbb{E} \exp F_2(u)}{n \exp(3s^2)} = \frac{\exp(2s^2)}{\exp(3s^2)} = \frac{1}{\exp(s^2)}.$$

It follows by combining the last two bounds that the event

$$E_\star \equiv \left\{ F(u) \geq \log n + \frac{s^2\epsilon}{10}, F_2(u) \leq \log n + 3s^2 \right\}$$

has probability $\mathbb{P}(E_\star) \geq 1 - 2 \exp(-s^2\epsilon/400)$. Let $\mathbf{P}_n$ be the uniform probability measure on $\{1, \dots, n\}$, and on the probability space $([n], \mathbf{P}_n)$ consider the random variable $X : i \mapsto X(i) = \exp(su_i)$. Then $\mathbf{E}_n X = n^{-1} \exp F(u)$ and $\mathbf{E}_n(X^2) = n^{-1} \exp F_2(u)$, so the Paley–Zygmund inequality implies

$$\mathbf{P}_n\left(X \geq \frac{\mathbf{E}_n X}{2}\right) \geq \frac{(\mathbf{E}_n X)^2}{4\mathbf{E}_n(X^2)} = \frac{(n^{-1} \exp F(u))^2}{4n^{-1} \exp F_2(u)} \geq \frac{\exp(s^2\epsilon/5)}{4 \exp(3s^2)} \geq \frac{1}{\exp(3s^2)}, \quad (2.6)$$

where the second-to-last inequality holds on the event $E_\star$, and the last inequality holds using $s\epsilon^{1/2} \geq C$. Moreover,

$$\frac{\mathbf{E}_n X}{2} = \frac{\exp F(u)}{2n} \geq \frac{\exp(s^2\epsilon/10)}{2} \geq \exp\left(\frac{s^2\epsilon}{20}\right),$$

where the last inequality again uses $s\epsilon^{1/2} \geq C$. Thus, $X(i) = \exp(su_i) \geq \mathbf{E}_n X/2$ implies $u_i \geq s\epsilon/20$, so

$$\mathbf{P}_n\left(X \geq \frac{\mathbf{E}_n X}{2}\right) \leq \frac{1}{n} \left| \left\{ i \leq n : u_i \geq \frac{s\epsilon}{20} \right\} \right|.$$

Since the bound (2.6) holds on the event $E_\star$, it follows that

$$\begin{aligned} \mathbb{P}\left(\frac{1}{n} \left| \left\{ i \leq n : u_i \geq \frac{s\epsilon}{20} \right\} \right| < \frac{1}{\exp(3s^2)}\right) &\leq \mathbb{P}\left(\mathbf{P}_n\left(X \geq \frac{\mathbf{E}_n X}{2}\right) < \frac{1}{\exp(3s^2)}\right) \\ &\leq 1 - \mathbb{P}(E_\star) \leq 2 \exp\left(-\frac{s^2\epsilon}{400}\right) \leq \exp\left(-\frac{s^2\epsilon}{800}\right). \end{aligned}$$

Taking $s' = s\epsilon^{1/2}/20$ gives, for $C \leq 20s' \leq (\log n)^{1/2}/C$,

$$\mathbb{P}\left(\frac{1}{n} \left| \left\{ i \leq n : u_i \geq s'\epsilon^{1/2} \right\} \right| < \frac{1}{\exp(1200(s')^2/\epsilon)}\right) \leq \exp\left(-\frac{(s')^2}{2}\right).$$

The claimed bound follows by setting $\bar{C} = \max\{1200, 20C\}$. $\square$

2.2. Concentration of supremum of canonical processes. In the remainder of this section we discuss the possibility of extending the approach of Theorem 1.2 to more general distributions of the disorder ξ . The main conclusion of this discussion is that the proof of Theorem 1.2 can likely be extended to a limited class of distributions using results from the existing literature. By contrast, our main results follow a different (and perhaps simpler) approach, and apply to the wider class of all subgaussian distributions. However, we obtain somewhat weaker bounds than in Theorem 1.2 (e.g., Theorem 1.2 applies for $w \leq N/C_\delta$ while Theorem 1.3 applies for $w \leq N^{1/2}/C_\delta$).

Theorem 1.2 can be obtained as a consequence of either Proposition 2.6 or Proposition 2.7. Proposition 2.7 relies on a gaussian interpolation bound and appears difficult to extend more generally, so we turn to discussing the possibility of extending Proposition 2.6. The proposition relies on two essential ingredients:

- (i) The Sudakov minoration (Lemma 2.1), lower bounding the expected supremum of a gaussian process;
- (ii) The Borell–TIS inequality (Lemma 2.2), giving concentration of the supremum of a gaussian process.

We discuss each of these ingredients separately below, for a stochastic process

$$\left(\frac{(\xi, \sigma)}{N^{1/2}} : \sigma \in S\right), \quad (2.7)$$

where S is a subset of $\mathbb{R}^N$. The rescaling of S by $N^{1/2}$ is not essential; we chose it to maintain consistency with the scaling in the rest of this paper, where S is generally a subset of $\{-1, +1\}^N$. In the literature, the stochastic process (2.7) is sometimes termed the “canonical process” for ξ indexed by $S/N^{1/2}$. In this subsection we discuss concentration of the supremum of the process (point (ii) above); in §2.3 we discuss lower bounding the expected supremum (point (i) above).

For $S \subseteq \mathbb{R}^N$, let us consider the supremum of the canonical process (2.7) as a function of the disorder ξ :

$$f(\xi) \equiv f\left(\xi; \frac{S}{N^{1/2}}\right) \equiv \sup \left\{ \frac{(\xi, \sigma)}{N^{1/2}} : \sigma \in S \right\}. \quad (2.8)$$

If S is any subset of $N^{1/2}B_2$ where B_2 is the euclidean unit ball in $\mathbb{R}^N$, then $f : \mathbb{R}^N \rightarrow \mathbb{R}$ will be Lipschitz with respect to the euclidean norm on $\mathbb{R}^N$:

$$f(\xi') = \sup \left\{ \frac{(\xi, \sigma)}{N^{1/2}} + \frac{(\xi' - \xi, \sigma)}{N^{1/2}} : \sigma \in S \right\} \leq f(\xi) + \|\xi' - \xi\|_2 \quad (2.9)$$

for any $\xi, \xi' \in \mathbb{R}^N$. The Borell–TIS inequality gives concentration for any Lipschitz functional of a gaussian random variable, hence Lemma 2.2 (since the gaussian process $(u_i)_{i \leq n}$ can be expressed as $((g, v_i))_{i \leq n}$, where g is a standard gaussian vector and $(v_i)_{i \leq n}$ is a collection of unit vectors). We know of similar bounds in two other settings:

- Suppose ξ_i has density $\exp(-V(x)) dx$ where V is **uniformly convex** (i.e., the second derivative of V is uniformly bounded below by a positive constant). Then it follows from [Led01, Propn. 2.18] (see also [Ver18, Thm. 5.2.15]) that (2.9) satisfies a similar concentration bound as Lemma 2.2, for any $S \subseteq N^{1/2}B_2$.
- If the ξ_i are almost surely **bounded** random variables, then it can be deduced via the method of bounded differences (see e.g. [BLM13, Thm. 6.2]) that (2.9) satisfies a similar concentration bound as Lemma 2.2, again for any $S \subseteq N^{1/2}B_2$.

For subgaussian distributions, we do not have such strong guarantees, but we have the following lemma which gives a weaker result under the further assumption $S \subseteq \{-1, +1\}^N$ (rather than $S \subseteq N^{1/2}B_2$):

Lemma 2.8. *Suppose ξ is a random vector in $\mathbb{R}^N$ satisfying Assumption 1.1. If $S \subseteq \{-1, +1\}^N$, then*

$$\mathbb{P} \left(\left| f \left(\xi; \frac{S}{N^{1/2}} \right) - \mathbb{E} f \left(\xi; \frac{S}{N^{1/2}} \right) \right| \geq u \right) \leq \exp \left(- \min \left\{ \frac{u^2}{4A^2}, \frac{N^{1/2}u}{2A} \right\} \right)$$

for f as defined by (2.8), and A a constant depending only on ν .

Proof. Let $\mathcal{H}_i$ denote the σ -algebra generated by the first i coordinates $(\xi_1, \dots, \xi_i)$, and decompose

$$f(\xi) - \mathbb{E} f(\xi) = \sum_{i=1}^N (X_i - X_{i-1})$$

where $X_i \equiv \mathbb{E}(f(\xi) | \mathcal{H}_i)$. Let $\mathbb{E}_i$ be expectation over ξ_i only, and let $\mathbb{E}^i$ be expectation over all the $(\xi_j)_{j \geq i}$, so that we have $X_i = \mathbb{E}^{i+1} f(\xi)$ while $X_{i-1} = \mathbb{E}^i f(\xi) = \mathbb{E}_i \mathbb{E}^{i+1} f(\xi)$. We then have, using Jensen's inequality,

$$\begin{aligned} E_i(\lambda) &\equiv \mathbb{E} \left[\exp \left(\lambda |X_i - X_{i-1}| \right) \middle| \mathcal{H}_{i-1} \right] = \mathbb{E}_i \mathbb{E}^{i+1} \exp \left(\lambda \left| \mathbb{E}^{i+1} (f(\xi) - \mathbb{E}_i f(\xi)) \right| \right) \\ &\leq \mathbb{E}_i \mathbb{E}^{i+1} \exp \left(\lambda \left| f(\xi) - \mathbb{E}_i f(\xi) \right| \right), \end{aligned}$$

Let $\xi^{(i)}$ denote the vector that results from replacing the i -th coordinate of ξ with an independent copy ζ_i , and let $\mathbb{E}^{(i)}$ denote expectation over ζ_i only. Then applying Jensen's inequality again gives

$$E_i(\lambda) \leq \mathbb{E}_i \mathbb{E}^{i+1} \exp \left(\lambda \left| \mathbb{E}^{(i)} (f(\xi) - f(\xi^{(i)})) \right| \right) \leq \mathbb{E}_i \mathbb{E}^{i+1} \mathbb{E}^{(i)} \exp \left(\lambda \left| f(\xi) - f(\xi^{(i)}) \right| \right).$$

We then note, since we assume $S \subseteq \{-1, +1\}^N$, the bound (2.9) can be refined to

$$f(\xi^{(i)}) \leq f(\xi) + \max \left\{ \frac{(\xi' - \xi, \sigma)}{N^{1/2}} : \sigma \in \{-1, +1\}^N \right\} \leq f(\xi) + \frac{|\zeta_i - \xi_i|}{N^{1/2}}.$$

Since for any $x \in \mathbb{R}$ we have $\exp(|x|) \leq \exp(x) + \exp(-x)$, we can use Assumption 1.1 to bound

$$E_i(\lambda) \leq \mathbb{E} \exp \left(\frac{\lambda |\zeta_i - \xi_i|}{N^{1/2}} \right) \leq 2 \exp \left(\frac{\lambda^2 \nu}{N} \right).$$

In particular, taking $\lambda = N^{1/2}$ gives $E_i(N^{1/2}) \leq 2 \exp(\nu)$. Another application of Jensen's inequality gives that for A large enough (depending on ν only), we have

$$E_i \left(\frac{N^{1/2}}{A} \right) \leq \left(E_i(N^{1/2}) \right)^{1/A} \leq \left(2 \exp(\nu) \right)^{1/A} \leq 2.$$

It then follows by the martingale Bernstein inequality (see [Tal11, Thm. A.6.1]) that

$$\mathbb{P}\left(\left|f(\xi) - \mathbb{E}f(\xi)\right| \geq u\right) \leq \exp\left(-\min\left\{\frac{u^2}{4A^2}, \frac{N^{1/2}u}{2A}\right\}\right)$$

for all $u \geq 0$, as claimed. $\square$

We note that, for the purposes of Theorem 1.2, Lemma 2.8 can be substituted for Lemma 2.2 with similar results. Thus, in our setting, concentration of the supremum of (2.7) does not appear to be a major issue. Rather, as we next discuss, lowering bounding the expected supremum appears to be the main obstacle to extending the argument of Theorem 1.2 to more general distributions.

2.3. Sudakov minoration lower bound. We now discuss lower bounding the expected supremum

$$\mathbb{E}f\left(\xi, \frac{S}{N^{1/2}}\right) = \mathbb{E}\left[\sup\left\{\frac{(\xi, \sigma)}{N^{1/2}} : \sigma \in S\right\}\right], \quad (2.10)$$

for f as in (2.8) and $S \subseteq \mathbb{R}^N$. In the gaussian case, the desired lower bound is given by Lemma 2.1. In the **Bernoulli** case a comparable lower bound is given by [Tal93, Propn. 2.2]. Minoration bounds for more general processes do exist (e.g. [Tal93, Tal94, Lat97, Tal05]), but the dependence on the law of ξ_i is somewhat complicated. The most relevant results come from [Lat97] (see also [Tal05, Ch. 5]), for the situation that the ξ_i are symmetric random variables such that the tail probability function

$$\phi_\xi(x) \equiv \log \frac{1}{\mathbb{P}(|\xi_i| \geq x)} \quad (2.11)$$

is **convex**. For simplicity, we will restrict most of our discussion to the case that the distribution of ξ_i has density

$$\frac{\exp(-|x|^\alpha)}{z_\alpha} dx \quad (2.12)$$

with respect to Lebesgue measure, where $1 \leq \alpha < \infty$ and z_α denotes the normalizing constant. This is a special case of the setting of [Lat97], and was treated earlier by [Tal93, Tal94].

If ξ_i follows the distribution (2.12) with $1 \leq \alpha \leq 2$, or more generally if $\phi_\xi(x) \lesssim x^2$ for $x \geq 1$ (meaning that ξ_i has **heavier tails** than the gaussian distribution), then it follows from [Lat97, Thm. 1] (see also [Tal94, Thm. 1.2] and the discussion around [Tal05, Thm. 5.2.6]) that a similar lower bound as Lemma 2.1 holds. In more detail, for an N -dimensional gaussian vector g and for any $X \subseteq \mathbb{R}^N$, it was proved by [Tal87] that

$$\mathbb{E}\left[\sup\left\{(g, x) : x \in X\right\}\right] \asymp \gamma_2(X, \|\cdot\|_2) = \inf_{(X_n)} \sup_{t \in X} \sum_{n \geq 0} 2^{n/2} \min\left\{\|t - x\|_2 : x \in X_n\right\},$$

where X_n is a subset of X of cardinality at most 2^{2^n} (γ_2 is a standard notation in the topic of majorizing measures or generic chaining). If the ξ_i are symmetric random variables with $\phi_\xi(x) \lesssim x^2$ for $x \geq 1$, then we have

$$\mathbb{E}\left[\sup\left\{(\xi, x) : x \in X\right\}\right] \gtrsim \gamma_2(X, \|\cdot\|_2) \asymp \mathbb{E}\left[\sup\left\{(g, x) : x \in X\right\}\right]$$

for any $X \subseteq \mathbb{R}^N$ (cf. and [Tal05, Thm. 2.1.1 and display (5.44)]). If X is a subset of the unit sphere in $\mathbb{R}^N$ such that $(x, x') \leq 1 - \epsilon$ for all $x \neq x'$ in X , then the right-hand side of the above is lower bounded by Lemma 2.1.³

If ξ_i follows the distribution (2.12) with $\alpha \geq 2$ (so that ξ_i has **lighter tails** than the gaussian distribution), then characterizing the expected supremum is in general more complicated; see [Tal94, Thm. 1.3]. A lower bound that is easier to work with is given by [Tal94, Thm. 3.1] (see also [Lat97, Thm. 1]) which in fact applies for all $\alpha > 1$: writing $m \equiv m(S/N^{1/2})$ for the quantity in (2.10), and writing β for the dual (Hölder conjugate) of α , we have

$$m \gtrsim \log n \left(\frac{S}{N^{1/2}}, U_\beta(m)\right), \quad (2.13)$$

³The resulting lower bound on the left-hand side may not be tight when $\alpha < 2$, although this in itself does not appear to be an issue.

where $n(X, U)$ denotes the number of translates of U needed to cover X , and

$$U_\beta(u) \equiv \left\{ x : \sum_{i \geq 1} \eta_\beta(x_i) \leq u \right\}, \quad \eta_\beta(x) \equiv \begin{cases} x^2 & \text{for } |x| \leq 1, \\ |x|^\beta & \text{for } |x| \geq 1. \end{cases}$$

Adapting the argument following [Tal94, Thm. 3.2] yields the following bound, which (like Lemma 2.8) relies crucially on the restriction $S \subseteq \{-1, +1\}^N$:

Corollary 2.9. *Suppose $S \subseteq \{-1, +1\}^N$, and assume that for all $\sigma \neq \tau$ in S we have $|(\sigma, \tau)|/N \leq 1 - \epsilon$. Suppose ξ is an N -dimensional random vector with i.i.d. entries ξ_i distributed according to (2.12) with $\alpha > 1$. Then*

$$\mathbb{E} \left[\sup \left\{ \frac{(\xi, \sigma)}{N^{1/2}} : \sigma \in S \right\} \right] \gtrsim \epsilon (\log |S|)^{1/2}.$$

Proof. As above, we abbreviate $\mathbf{m} = \mathbf{m}(S/N^{1/2})$ for the quantity of interest. Let θ denote a positive constant. Applying (2.13) (from [Tal94, Thm. 3.1]) with θS in place of S gives

$$\theta \mathbf{m} \gtrsim \log n \left(\frac{\theta S}{N^{1/2}}, U_\beta(\theta \mathbf{m}) \right).$$

Suppose two points $x, y \in [-1/2, +1/2]^N$ are covered by the same translate of $U_\beta(\theta \mathbf{m})$. Then there must exist $z \in [-1/2, +1/2]^N$ such that both x and y belong to the translate of $U_\beta(\theta \mathbf{m})$ centered at z : that is,

$$\max \left\{ \sum_{i \leq N} \eta_\beta(x_i - z_i), \sum_{i \leq N} \eta_\beta(y_i - z_i) \right\} \leq \theta \mathbf{m}.$$

Since $|x_i - z_i| \leq 1$ and $|y_i - z_i| \leq 1$ for all $i \leq N$, the definition of η_β implies that both x and y must lie within euclidean distance $(\theta \mathbf{m})^{1/2}$ of z , and so x and y must also be covered by the same translate of $(\theta \mathbf{m})^{1/2} B_2$, where B_2 denotes the unit ball in euclidean norm. Thus, as long as $\theta/N^{1/2} \leq 1/2$, we have

$$\theta \mathbf{m} \gtrsim \log n \left(\frac{\theta S}{N^{1/2}}, (\theta \mathbf{m})^{1/2} B_2 \right) = \log n \left(\frac{S}{N^{1/2}}, \left(\frac{\mathbf{m}}{\theta} \right)^{1/2} B_2 \right).$$

If $\mathbf{m}/\epsilon \leq N^{1/2}/2$, then setting $\theta = \mathbf{m}/\epsilon$ gives

$$\frac{\mathbf{m}^2}{\epsilon} \geq \log n \left(\frac{S}{N^{1/2}}, \epsilon^{1/2} B_2 \right) = \log |S|,$$

since the separation assumption in the statement of the corollary implies that each translate of $\epsilon^{1/2} B_2$ can cover at most one element of S . Otherwise we must have $\mathbf{m} \geq \epsilon N^{1/2}/2 \gtrsim \epsilon (\log |S|)^{1/2}$, since $|S| \leq 2^N$. Combining these two cases gives the claimed lower bound. $\square$

For simplicity we have stated Corollary 2.9 for the distribution (2.12) with $\alpha > 1$, so that the bound can be obtained by applying [Tal94, Thm. 3.1]. For the more general setting where the function $\phi_\xi(x)$ of (2.11) is **convex**, a similar lower bound as Corollary 2.9 can be obtained by applying [Lat97, Thm. 1]. In summary, we have similar lower bounds for the canonical process (2.7) in the cases

- The ξ_i are gaussian (Lemma 2.1);
- The ξ_i are Bernoulli ([Tal93, Propn. 2.2]);
- The ξ_i are symmetric random variables such that the function $\phi_\xi(x)$ is convex ([Tal94, Thm. 3.1], [Lat97, Thm. 1], and Corollary 2.9).

We do not know of a similar lower bound for the more general class of subgaussian vectors.

Summary of conclusions of §2.2–2.3. The conclusion of the last two subsections is that, using only existing results in the literature, it may be possible to extend Theorem 1.2 to cover the case that the ξ_i are symmetric subgaussian random variables such that the function $\phi_\xi(x)$ of (2.11) is convex. One can adapt the argument of Theorem 1.2 with Corollary 2.9 in place of Lemma 2.1, and Lemma 2.8 in place of Lemma 2.2. To cover all subgaussian distributions, the main difficulty appears to be in obtaining a lower bound comparable to that of Lemma 2.8.

In the remainder of this paper we prove versions of Theorem 1.2 for general subgaussian disorder, using as input **only** the gaussian results Lemma 2.1 and Lemma 2.2. This results in weaker (and likely suboptimal) bounds. On the other hand, our argument is relatively simple, as it bypasses the more difficult Sudakov minoration bounds proved by [Tal94, Lat97].

3. ADDING A SINGLE CONSTRAINT IN PERCEPTRON MODELS

In this section we prove our main estimates for the perceptron model, Theorems 1.3–1.5.

- In §3.1 we prove a preliminary estimate (Proposition 3.2) towards our main results, which says that large subsets of $\{-1, +1\}^N$ must contain many pairs of configurations that are well-separated within blocks of the coordinates $[N]$.
- In §3.2 we prove Proposition 3.8, which is a version of Proposition 2.6 for general disorder. This leads to the [proof of Theorem 1.3](#).
- In §3.3 we apply Proposition 3.8 to give the [proof of Theorem 1.5](#). As we explain at the end of that subsection, a similar argument (using Proposition 2.6 in place of Proposition 3.8) yields the [proof of Theorem 1.4](#).

3.1. Extraction of subsets with separation on blocks. Recall that we decomposed $[N]$ into blocks $I_1, \dots, I_L$, each of size $K = N/L$. Given $S \subseteq \{-1, +1\}^N$ of size $|S| \geq \exp(8N\delta)$, our first goal is to extract $\Omega \subseteq S$ large, such that configurations in Ω have “good separation” on some of the blocks I_j . Formally:

Definition 3.1. For any subset of coordinates $I \subseteq [N]$ of size $|I| = K$, we say that two configurations $\sigma, \tau \in \{-1, +1\}^N$ are ϵ -**separated on I** if

$$\frac{|(\sigma_I, \tau_I)|}{K} \leq 1 - \epsilon.$$

Equivalently, the euclidean distance between σ_I and τ_I is at least $(2K\epsilon)^{1/2}$.

Note that in an extreme case we could have for instance

$$S = \left\{ \sigma \in \{-1, +1\}^N : \sigma_i = 1 \text{ for all } 1 \leq i \leq N \left(1 - \frac{8\delta}{\log 2}\right) \right\},$$

so in this case $|S| \geq \exp(8N\delta)$, but clearly any pair of elements in S will not be well-separated on most blocks I_j . Thus it will only be possible to guarantee good separation on a small fraction γ (depending on δ) of the blocks I_j . The main result of this subsection is the following:

Proposition 3.2. *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(8N\delta)$ where δ is a small positive constant. Suppose ϵ and γ are positive constants with $\psi_2(8\epsilon) \leq \delta$ and $\gamma \leq \delta/(2 \log 2)$. Suppose $C'/\gamma \leq L \leq N\delta/C'$, and divide $[N]$ into consecutive blocks $I_1, \dots, I_L$ of size $K = N/L$ each. Then there exists $J_\star \subseteq [L]$ with $|J_\star| = N\gamma$, and $\Omega \subseteq S$ with $|\Omega| \geq \exp(N\delta/L)$, such that all pairs $\sigma \neq \tau$ in Ω satisfy*

$$\frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} \leq 1 - \epsilon$$

for all $j \in J_\star$. That is to say, all pairs $\sigma \neq \tau$ in Ω are ϵ -separated on I_j for all $j \in J_\star$.

The [proof of Proposition 3.2](#) appears at the end of this subsection.

Lemma 3.3. *As in Proposition 3.2, assume that S is any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(8N\delta)$; and that we have $\psi_2(8\epsilon) \leq \delta$, $\gamma \leq \delta/(2 \log 2)$, and $C'/\gamma \leq L \leq N\delta/C'$. Then*

$$\mu^{\otimes 2} \left(\sum_{j \leq L} \mathbf{1} \left\{ \frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} \leq 1 - 8\epsilon \right\} < 2L\gamma \right) \leq \frac{1}{\exp(5N\delta)},$$

where μ is the uniform measure on S , and (σ, τ) is a pair of i.i.d. samples from μ .

Proof. Write $t \equiv 1 - 8\epsilon$. For $\sigma \in \{-1, +1\}^N$, we can bound

$$\begin{aligned} \mu_\sigma &\equiv \mu\left(\left\{\tau : \sum_{j \leq L} \mathbf{1}\left\{\frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} \leq 1 - 8\epsilon\right\} < 2L\gamma\right\}\right) \\ &\leq \frac{1}{|S|} \left|\left\{\tau \in \{-1, +1\}^N : \sum_{j \leq L} \mathbf{1}\left\{\frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} \leq 1 - 8\epsilon\right\} < 2L\gamma\right\}\right|. \end{aligned}$$

For any fixed σ , if τ is sampled uniformly at random from $\{-1, +1\}^N$, then the scalar product $(\sigma_{I_j}, \tau_{I_j})$ is equidistributed as $2(\text{Bin}(K, 1/2) - K/2)$, so the chance for $|(\sigma_{I_j}, \tau_{I_j})| \geq K(1 - 8\epsilon)$ is

$$\bar{p} \equiv \frac{1}{2^K} \left|\left\{\tau_{I_j} \in \{-1, +1\}^K : \frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} \leq t\right\}\right| = \mathbb{P}\left(\left|\text{Bin}\left(K, \frac{1}{2}\right) - K\right| \geq \frac{Kt}{2}\right) \leq \frac{2}{\exp(Kk_2(t))}$$

Moreover the $(\sigma_{I_j}, \tau_{I_j})$ are independent across the different j , so

$$\mu_\sigma \leq \frac{2^N}{\exp(8N\delta)} \mathbb{P}\left(\text{Bin}(L, \bar{p}) \geq L(1 - 2\gamma)\right) \leq \frac{2^N}{\exp(8N\delta)} \left(\frac{L}{2L\gamma}\right) \bar{p}^{L(1-2\gamma)}.$$

Combining with the preceding bound on $\bar{p}$, and recalling $k_2(t) = \log 2 - \psi_2(8\epsilon)$, we obtain

$$\begin{aligned} \mu_\sigma &\leq \frac{2^N \exp(LH(2\gamma))}{\exp(8N\delta)} \left(\frac{2}{\exp(Kk_2(t))}\right)^{L(1-2\gamma)} \\ &\leq \exp\left\{N\left(-8\delta + \psi_2(8\epsilon) + 2\gamma \log 2 - 2\gamma \psi_2(8\epsilon)\right) + L\left(H(2\gamma) + \log 2\right)\right\} \leq \frac{1}{\exp(5N\delta)}, \end{aligned}$$

where the last bound holds for parameters $\delta, \epsilon, \gamma, L$ as in the statement of the lemma. The result follows by averaging over $\sigma \in S$. $\square$

Corollary 3.4. *In the setting of Lemma 3.3 (and with all the same parameters), there must be a subset $S' \subseteq S$ with $|S'| \geq \exp(5N\delta/2)$ such that all pairs $\sigma \neq \tau$ in S' are well-separated in the sense that*

$$\sum_{j \leq L} \mathbf{1}\left\{\frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} \leq 1 - 8\epsilon\right\} \geq 2L\gamma. \quad (3.1)$$

Proof. Let $\sigma^{(1)}, \dots, \sigma^{(n)}$ be $n = \exp(5N\delta/2)$ i.i.d. samples from μ (the uniform measure on S). By Lemma 3.3, the probability that for any $1 \leq i < j \leq n$ the configurations σ^i, σ^j fail to be well-separated is upper bounded by

$$\binom{n}{2} \mu^{\otimes 2} \left(\sum_{j \leq L} \mathbf{1}\left\{\frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} \leq 1 - 8\epsilon\right\} < 2L\gamma \right) \leq \binom{n}{2} \frac{1}{\exp(5N\delta)} \leq \frac{1}{2},$$

where the last bound holds by the choice of n . Therefore the random set $\{\sigma^{(1)}, \dots, \sigma^{(n)}\}$ satisfies the required condition with probability at least $1/2$; and this implies the existence of the claimed set S' . $\square$

Lemma 3.5. *Let $S' \subseteq \{-1, +1\}^N$ with $|S'| \geq \exp(5N\delta/2)$, such that all pairs $\sigma \neq \tau$ in S' are well-separated in the sense of (3.1). If η satisfies $(2/|S'|)^{1/L} \leq \eta \leq 1/2$, then there exist $J_\circ \subseteq [L]$ with $|J_\circ| \geq 2L\gamma$, and $T \subseteq S'$ with $|T| \geq \eta^L |S'|$, such that we have*

$$\frac{1}{|T|} \left|\left\{\tau \in T : \frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} > 1 - \epsilon\right\}\right| \leq 2\eta. \quad (3.2)$$

for all $\sigma \in T$ and all $j \in J_\circ$.

Proof. We first give the construction of T and $J_\circ$. For $\sigma \in \{-1, +1\}^N$, $\Omega \subseteq \{-1, +1\}^N$, and $1 \leq j \leq L$, denote

$$T_j(\sigma, \Omega) \equiv \left\{\tau \in \Omega : \frac{(\sigma_{I_j}, \tau_{I_j})}{K} > 1 - \epsilon\right\}.$$

- *Initial step.* Let $\Omega^{(0)} \equiv S'$. If for all $\sigma \in \Omega^{(0)}$ and all $j \in [L]$ we have

$$\max \left\{ \frac{|T_j(\sigma, \Omega^{(0)})|}{|\Omega^{(0)}|}, \frac{|T_j(-\sigma, \Omega^{(0)})|}{|\Omega^{(0)}|} \right\} \leq \eta,$$

then we are done by simply taking $T = \Omega^{(0)}$ and $J_\circ = [L]$. If this is not the case, then we must have a configuration $\sigma^{(1)} \in \Omega^{(0)} \cup (-\Omega^{(0)})$ and an index $j^{(1)} \in [L]$ such that

$$\frac{|T_{j^{(1)}}(\sigma^{(1)}, \Omega^{(0)})|}{|\Omega^{(0)}|} > \eta.$$

Then set $\Omega^{(1)} = T_{j^{(1)}}(\sigma^{(1)}, \Omega^{(0)}) \subseteq \Omega^{(0)}$.

- *Inductive step.* Similarly, suppose inductively that we have constructed $\sigma^{(k)}, j^{(k)}, \Omega^{(k)}$. If for all $\sigma \in \Omega^{(k)}$ and all $j \in [L] \setminus \{j^{(1)}, \dots, j^{(k)}\}$ we have

$$\max \left\{ \frac{|T_j(\sigma, \Omega^{(k)})|}{|\Omega^{(k)}|}, \frac{|T_j(-\sigma, \Omega^{(k)})|}{|\Omega^{(k)}|} \right\} \leq \eta,$$

then we end the process by setting $T = \Omega^{(k)}$ and $J_\circ = [L] \setminus \{j^{(1)}, \dots, j^{(k)}\}$. If this is not the case, then we must have a configuration $\sigma^{(k+1)} \in \Omega^{(k)} \cup (-\Omega^{(k)})$ and an index $j^{(k+1)} \in [L] \setminus \{j^{(1)}, \dots, j^{(k)}\}$ such that

$$\frac{|T_{j^{(k+1)}}(\sigma^{(k+1)}, \Omega^{(k)})|}{|\Omega^{(k)}|} > \eta.$$

Then set $\Omega^{(k+1)} = T_{j^{(k+1)}}(\sigma^{(k+1)}, \Omega^{(k)}) \subseteq \Omega^{(k)}$.

We claim that the above construction ends with $T = \Omega^{(k')}$ for some $k' \leq \ell = L(1 - 2\gamma)$. Indeed, suppose for contradiction that it does not. For simplicity of notation, let us re-index such that $j^{(k)} = k$ for all k . Now consider the set $F = \Omega^{(\ell+1)} = T_{\ell+1}(\sigma^{(\ell+1)}, \Omega^{(\ell)})$ for $\ell = L(1 - 2\gamma)$. If $\tau \in F$, then for each $k \leq \ell + 1$ the configurations τ and $\sigma^{(k)}$ must be close on the block of coordinates I_k :

$$\frac{\|\tau_{I_k} - (\sigma^{(k)})_{I_k}\|^2}{K} = \frac{2K - 2(\tau_{I_k}, (\sigma^{(k)})_{I_k})}{K} \leq 2 - 2(1 - \epsilon) = 2\epsilon.$$

By the assumption $\eta^L |S'| \geq 2$ we must have

$$|F| = |\Omega^{(0)}| \prod_{k=1}^{\ell+1} \frac{|\Omega^{(k)}|}{|\Omega^{(k-1)}|} \geq |S'| \eta^{\ell+1} \geq |S'| \eta^L \geq 2,$$

so we can find a pair of distinct elements $\sigma \neq \tau$ in F . For this pair, the triangle inequality gives

$$\frac{\|\sigma_{I_k} - \tau_{I_k}\|^2}{K} \leq \left(2(2\epsilon)^{1/2}\right)^2 = 8\epsilon$$

for each $k \leq \ell + 1$. This contradicts the assumption that σ and τ must satisfy (3.1), since they are both elements of the original set S' . This verifies the claim that $T = \Omega^{(k')}$ for some $k' \leq \ell = L(1 - 2\gamma)$, and

$$\frac{|T|}{|S'|} = \prod_{k=1}^{k'} \frac{|\Omega^{(k)}|}{|\Omega^{(k-1)}|} \geq \eta^{k'} \geq \eta^\ell \geq \eta^L.$$

Finally, the required bound (3.2) holds because otherwise the above construction would not stop at $T = \Omega^{(k')}$. $\square$

Lemma 3.6. *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(8N\delta)$ where δ is a small positive constant. Suppose $\psi_2(8\epsilon) \leq \delta$, $\gamma \leq \delta/(2 \log 2)$, and $C'/\gamma \leq L \leq N\delta/C'$. Then there exist a subset $J_\circ \subseteq [L]$ with $|J_\circ| \geq 2L\gamma$, and a subset $T \subseteq S$ with $|T| \geq \exp(N\delta)$, such that for all $\sigma \in T$ we have*

$$\frac{1}{|T|} \left| \left\{ \tau \in T : \frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} > 1 - \epsilon \right\} \right| \leq 2 \exp \left(- \frac{9N\delta}{4L} \right)$$

for each $j \in J_\circ$.

Proof. Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(8N\delta)$. It follows from Corollary 3.4 that there must be a subset $S' \subseteq S$ with $|S'| \geq \exp(5N\delta/2)$ such that all pairs $\sigma \neq \tau$ in S' are well-separated in the sense of (3.1). It then follows from Lemma 3.5 that if $(2/|S'|)^{1/L} \leq \eta \leq 1/2$, then there exists $J_\circ \subseteq [L]$ with $|J_\circ| \geq 2L\gamma$, and $T \subseteq S'$ with $|T| \geq \eta^L |S'|$, such that we have

$$\frac{1}{|T|} \left| \left\{ \tau \in T : \frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} > 1 - \epsilon \right\} \right| \leq 2\eta$$

for all $\sigma \in T$ and all $j \in J_\circ$. In particular, we can choose $\eta = \exp(-9N\delta/(4L))$, so that

$$\eta^L = \exp\left(-\frac{9N\delta}{4}\right) \geq 2 \exp\left(-\frac{5N\delta}{2}\right) \geq \frac{2}{|S'|}$$

holds for all N large enough, as required. $\square$

Proof of Proposition 3.2. Let $T \subseteq S$ be as given by Lemma 3.6, and let μ be the uniform probability measure on T . Similarly to the proof of Corollary 3.4, let $(\sigma^{(1)}, \dots, \sigma^{(n)})$ be sampled according to $\mu^{\otimes n}$, i.e., uniformly at random from T^n . Let us say that a block I_k “fails” if there exists a pair of indices $1 \leq i < j \leq n$ such that the configurations $\sigma^{(i)}$ and $\sigma^{(j)}$ are not ϵ -separated on the block I_k . The probability that more than half of the blocks in $J_\circ$ fail can be upper bounded with Markov’s inequality:

$$\begin{aligned} \mu^{\otimes n} \left(\frac{1}{|J_\circ|} \sum_{j \in J_\circ} \mathbf{1}\{\text{block } I_j \text{ fails}\} \geq \frac{1}{2} \right) &\leq \frac{2}{|J_\circ|} \sum_{j \in J_\circ} \binom{n}{2} \mu^{\otimes 2} \left(\frac{|(\sigma_{I_j}, \tau_{I_j})|}{K} > 1 - \epsilon \right) \\ &\leq n^2 \exp\left(-\frac{9N\delta}{4L}\right) \leq \frac{1}{2}, \end{aligned}$$

where the last inequality holds for $n \leq \exp(N\delta/L)$. Therefore, with probability at least $1/2$, for the random set $\{\sigma^{(1)}, \dots, \sigma^{(n)}\}$ we can find $|J_\star| \geq L\gamma$ such that the required properties are satisfied. If we take $n = \exp(N\delta/L)$, then this implies the existence of the claimed set $\Omega \subseteq T \subseteq S$. $\square$

3.2. Intersection of cube and half-space with general disorder. The main results in this subsection are Proposition 3.8, which is a version of Proposition 2.6 for general disorder, and the proof of Theorem 1.3. The assumption of the following is based on the result of Proposition 3.2:

Proposition 3.7. *Suppose ξ is a random vector in $\mathbb{R}^N$ satisfying Assumption 1.1. Suppose δ is a small positive constant, and $\epsilon, \gamma \leq \delta$. Divide $[N]$ into consecutive blocks $I_1, \dots, I_L$ of size $K = N/L$ each. Suppose $\Omega \subseteq \{-1, +1\}^N$ with $|\Omega| \geq \exp(N\delta/L)$, such that all pairs $\sigma \neq \tau$ are ϵ -separated on I_j for all $j \in J_\star = [\ell + 1, L]$, where $\ell = L(1 - \gamma)$. Then we have the bound*

$$\mathbb{P} \left(\frac{1}{|\Omega|} \left| \left\{ \sigma \in \Omega : \frac{(\xi, \sigma)}{N^{1/2}} \geq \frac{\gamma(L\epsilon \log p)^{1/2}}{60} \right\} \right| < \frac{1}{2 \cdot (4p)^{L\gamma}} \right) \leq 4 \exp \left(-\frac{L\gamma^2 \epsilon \log p}{60\nu} \right),$$

provided $\log p \geq C'\nu/\epsilon$ and $C'/\gamma \leq L \leq N^{1/2}/(C' \log p)^{1/2}$.

Proof. Let p be a large integer. Let $\mathcal{F}_k$ be the σ -algebra generated by the vectors ξ_{I_j} for $1 \leq j \leq k$, so that $M_k(\sigma)$ (as defined by (1.3)) is measurable with respect to $\mathcal{F}_k$. Recall $\ell = L(1 - \gamma)$, and define

$$\Omega^{(\ell)} \equiv \left\{ \sigma \in \Omega : M_\ell(\sigma) \geq -\frac{\gamma(L\epsilon \log p)^{1/2}}{5} \right\}.$$

It follows using Markov’s inequality and the subgaussian tail bound that

$$\begin{aligned} \text{(I)} &\equiv \mathbb{P} \left(\frac{|\Omega^{(\ell)}|}{|\Omega|} \leq \frac{1}{2} \right) \leq \mathbb{P} \left(\frac{1}{|\Omega|} \left| \left\{ \sigma \in \Omega : M_\ell(\sigma) < -\frac{\gamma(L\epsilon \log p)^{1/2}}{5} \right\} \right| \geq \frac{1}{2} \right) \\ &\leq 2 \exp \left(-\frac{L\gamma^2 \epsilon \log p}{50\nu} \right). \end{aligned}$$

Next, for $\ell + 1 \leq k \leq L$, suppose inductively that $\Omega^{(k)}$ has been defined. Let

$$\Xi^{(k)} \equiv \left\{ \sigma \in \Omega^{(k)} : \frac{(\xi_{I_k}, \sigma_{I_k})}{K^{1/2}} \geq \frac{(\epsilon \log p)^{1/2}}{2} \right\} \subseteq \Omega^{(k)}.$$

We will show below that $\Xi^{(k)}$ occupies a constant fraction of $\Omega^{(k)}$ with good probability:

$$\mathbb{P}\left(\frac{|\Xi^{(k)}|}{|\Omega^{(k)}|} \leq \frac{1}{4p}\right) \leq \frac{4}{p^{\epsilon/32}}. \quad (3.3)$$

With this in mind, for $\ell + 1 \leq k \leq L$ we define the sets

$$\Omega^{(k+1)} = \begin{cases} \Xi^{(k)} & \text{if } |\Xi^{(k)}| \geq |\Omega^{(k)}|/(4p) \\ \Omega^{(k)} & \text{otherwise.} \end{cases}$$

The construction guarantees for all $\ell \leq k \leq L$ that

$$|\Omega^{(k)}| \geq |\Omega^{(L)}| \geq \frac{|\Omega|}{2} \left(\frac{1}{4p}\right)^{L\gamma} \geq \frac{1}{2} \exp\left(\frac{N\delta}{L} - L\gamma \log(4p)\right) \geq p, \quad (3.4)$$

where the last bound holds by the assumed upper bound on L . Now, turning to the proof of (3.3), from $\Omega^{(k)}$ let us extract disjoint subsets $X_1, \dots, X_m$ each of size p , whose union occupies more than half of $\Omega^{(k)}$ – the preceding bound guarantees that we can do this with $m \geq 1$, since $|\Omega^{(k)}| \geq p$. Combining Corollary 2.3 with the CLT estimate Corollary 4.4 (deferred to §4.1 below) gives

$$\begin{aligned} \mathbb{P}\left(\max_{\sigma \in X_a} \frac{(\xi, \sigma)}{N^{1/2}} \leq \frac{(\epsilon \log p)^{1/2}}{2}\right) &= \mathbb{E}\left[\prod_{\sigma \in X_a} \mathbf{1}\left\{\frac{(\xi, \sigma)}{N^{1/2}} \leq \frac{(\epsilon \log p)^{1/2}}{2}\right\}\right] \\ &\leq \mathbb{P}\left(\max_{\sigma \in X_a} \frac{(g, \sigma)}{N^{1/2}} \leq \frac{(\epsilon \log p)^{1/2}}{2}\right) + o_N(1) \leq \frac{2}{p^{\epsilon/50}} \end{aligned}$$

It then follows by the Markov inequality that

$$\mathbb{P}\left(\frac{1}{m} \sum_{a \leq m} \mathbf{1}\{\Xi^{(k)} \cap X_a = \emptyset\} \geq \frac{1}{2}\right) \leq \frac{4}{p^{\epsilon/50}}.$$

On the complementary event we have

$$|\Xi^{(k)}| \geq \sum_{a \leq m} \mathbf{1}\{\Xi^{(k)} \cap X_a \neq \emptyset\} \geq \frac{m}{2} \geq \frac{|\Omega^{(k)}|}{4p},$$

which proves (3.3). Next let $J_\bullet$ denote the subset of indices $k \in J_\star$ for which we have $\Omega^{(k)} = \Omega^{(k-1)}$. It follows from (3.3) that $|J_\bullet|$ is stochastically dominated by a binomial random variable with $L\gamma$ trials and success probability $4/p^{\epsilon/50}$. It follows by the Chernoff bound and (2.2) that

$$(II) \equiv \mathbb{P}\left(|J_\bullet| \geq \frac{L\gamma}{6\nu}\right) \leq \exp\left\{-L\gamma H\left(\frac{1}{6\nu} \middle| \frac{4}{p^{\epsilon/50}}\right)\right\} \leq \exp\left(-\frac{L\gamma\epsilon \log p}{6\nu \cdot 51}\right),$$

where the last bound holds because the restriction $\log p \geq C'\nu/\epsilon$ guarantees

$$\frac{t}{e} = \frac{1}{e} \cdot \frac{1}{6\nu} \cdot \frac{p^{\epsilon/50}}{4} \geq p^{\epsilon/51}.$$

By a union bound over all subsets $J \subseteq J_\star$, we have

$$\begin{aligned} (III) &\equiv \mathbb{P}\left(\sum_{j \in J} \frac{(\xi_{I_j}, \sigma_{I_j})}{N^{1/2}} \leq -\frac{\gamma(L\epsilon \log p)^{1/2}}{5} \text{ for any } J \subseteq J_\star\right) \\ &\leq 2^{L\gamma} \exp\left(-\frac{L\gamma\epsilon \log p}{50\nu}\right) \leq \exp\left(-\frac{L\gamma\epsilon \log p}{60\nu}\right), \end{aligned}$$

where the last bound again uses the restriction $\log p \geq C'\nu/\epsilon$. On the complement of the events bounded by (I), (II), and (III), we have for all $\sigma \in \Omega^{(L)}$ that

$$\begin{aligned} \frac{(\xi, \sigma)}{N^{1/2}} &= M_L(\sigma) = M_\ell(\sigma) + \sum_{j \in J} \frac{(\xi_{I_j}, \sigma_{I_j})}{N^{1/2}} + \sum_{j \in J_* \setminus J} \frac{(\xi_{I_j}, \sigma_{I_j})}{N^{1/2}} \\ &\geq -\frac{2\gamma(L\epsilon \log p)^{1/2}}{5} + L\gamma \left(1 - \frac{1}{6\nu}\right) \frac{(\epsilon \log p)^{1/2}}{2L^{1/2}} \geq \frac{\gamma(L\epsilon \log p)^{1/2}}{60}, \end{aligned}$$

where the last bound uses that we must have $\nu \geq 1$ (see Assumption 1.1). The claim then follows by recalling the lower bound on $|\Omega^{(L)}|$ from (3.4). $\square$

The next proposition is a version of Proposition 2.6 which applies in the case of general disorder. It has a somewhat worse ϵ -dependence Proposition 2.6. From our perspective the more important difference is that Proposition 3.8 applies for a more limited range of s than Proposition 2.6:

Proposition 3.8. *Let S be any subset of $\{-1, +1\}^N$ with $|S| \geq \exp(N\delta)$ where δ is a small positive constant. Let ϵ be a positive constant with $\psi_2(8\epsilon) \leq \delta/9$. Then*

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \geq s(\epsilon\nu)^{1/2} \right\} \right| < \exp\left(-\frac{C''\nu s^2}{2\epsilon}\right) \right) \leq \frac{4}{\exp(s^2\epsilon)}.$$

provided $C'' \leq s \leq (N\epsilon^3/\nu)^{1/4}/C''$, where C'' is a large absolute constant.

Proof. Since we can relabel the constants later, we assume $|S| \geq \exp(9N\delta)$ and $\psi_2(8\epsilon) \leq \delta$. Recalling the statement of Proposition 3.2, let $\gamma = \epsilon$. (For δ small enough, the assumption $\psi_2(8\epsilon) \leq \delta$ and the choice $\gamma = \epsilon$ together guarantee that the condition $\gamma \leq \delta/(2\log 2)$ from Proposition 3.2 is satisfied.) It follows from Proposition 3.2 that there exists $J_1 \subseteq [L]$ with $|J_1| = L\gamma$, and $\Omega_1 \subseteq S$ with $|\Omega_1| = \exp(N\delta/L)$, such that all pairs $\sigma \neq \tau$ in Ω_1 are ϵ -separated on I_j for all $j \in J_1$. We can apply Proposition 3.2 again on $S \setminus \Omega_1$, and so on, to extract disjoint subsets $\Omega_1, \dots, \Omega_m \subseteq S$, up to the first m such that

$$\sum_{a \leq m} |\Omega_a| \geq \frac{|S|}{2}.$$

Then for each a we will have $|\Omega_a| = \exp(N\delta/L)$, and all pairs $\sigma \neq \tau$ in Ω_a will be ϵ -separated on I_j for all $j \in J_a$, where $J_a \subseteq [L]$ with $|J_a| = L\gamma$. Recalling the statement of Proposition 3.7, let $\log p = C'\nu/\epsilon$, and define

$$\Xi \equiv \left\{ \sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \geq \frac{\gamma(L\epsilon \log p)^{1/2}}{60} \right\}, \quad \Xi_a \equiv \Xi \cap \Omega_a.$$

By Markov's inequality and Proposition 3.7, as long as $C'/\gamma \leq L \leq N^{1/2}/(C' \log p)^{1/2}$, we have

$$\mathbb{P}\left(\frac{1}{m} \sum_{a \leq m} \mathbf{1}\left\{ \frac{|\Xi_a|}{|\Omega_a|} \leq \frac{1}{2 \cdot (4p)^{L\gamma}} \right\} \geq \frac{1}{2} \right) \leq 4 \exp\left(-\frac{L\gamma^2 \epsilon \log p}{60\nu}\right).$$

On the complementary event we must have

$$|\Xi| \geq \sum_{a \leq m} |\Xi_a| \geq \frac{m|\Omega_1|}{2} \cdot \frac{1}{2 \cdot (4p)^{L\gamma}} \geq \frac{|S|}{8 \cdot (4p)^{L\gamma}},$$

so we have shown that, for any L satisfying $C'/\gamma \leq L \leq N^{1/2}/(C' \log p)^{1/2}$, we have

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \geq \frac{\gamma(L\epsilon \log p)^{1/2}}{60} \right\} \right| < \frac{1}{8 \cdot (4p)^{L\gamma}} \right) \leq 4 \exp\left(-\frac{L\gamma^2 \epsilon \log p}{60\nu}\right).$$

Recalling that $\log p = C'\nu/\epsilon$, the above implies

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \geq \frac{(C'L\nu)^{1/2}\gamma}{60} \right\} \right| < \frac{1}{8} \exp\left(-\frac{2C'L\nu\gamma}{\epsilon}\right) \right) \leq 4 \exp\left(-\frac{C'L\gamma^2}{60}\right).$$

Making the change of variables $s = (C'L\gamma)^{1/2}/60$ gives

$$\mathbb{P}\left(\frac{1}{|S|}\left|\left\{\sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \geq s(\gamma\gamma)^{1/2}\right\}\right| < \frac{1}{8} \exp\left(-\frac{(2 \cdot 60^2)\gamma s^2}{\epsilon}\right)\right) \leq \frac{4}{\exp(60\gamma s^2)}.$$

and the conclusion follows by recalling that we took $\gamma = \epsilon$. Note that the bounds on s in the statement of the result guarantee that L satisfies the requirements $C'/\gamma \leq L \leq N^{1/2}/(C' \log p)^{1/2}$. $\square$

Proof of Theorem 1.3. Let ϵ be a positive constant with $\psi_2(8\epsilon) \leq \delta/9$. It follows directly from Proposition 3.8 that

$$\mathbb{P}\left(\frac{1}{|S|}\left|\left\{\sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \geq \kappa\right\}\right| < \exp\left(-\frac{C''\gamma s^2}{2\epsilon}\right)\right) \leq \frac{4}{\exp(s^2\epsilon)},$$

provided that s satisfies (with C'' a large absolute constant)

$$\max\left\{C'', \frac{\kappa_+}{(\epsilon\gamma)^{1/2}}\right\} \leq s \leq \frac{1}{C''}\left(\frac{N\epsilon^3}{\gamma}\right)^{1/4}.$$

The claim follows by fixing ϵ to satisfy $\psi_2(8\epsilon) = \delta/9$, and making the change of variables $w = C''\gamma s^2/(2\epsilon)$. $\square$

3.3. Intersection of cube and slab with general disorder. The proof of Theorem 1.5 appears at the end of this subsection. The next two lemmas are motivated by the following considerations. In view of Proposition 3.2, suppose we have a large subset $\Omega \subseteq \{-1, +1\}^N$ such that all pairs $\sigma \neq \tau$ in Ω are well-separated on the blocks I_j for $j \in J_*$. By re-indexing, we may suppose without loss that $J_* = [L] \setminus [\ell]$ for $\ell = L(1-\gamma)$. For $\sigma \in \Omega$, recall the process defined by (1.3), and suppose we have $|M_\ell(\sigma)| \leq s$ for s large. Then, in order to achieve the desired outcome $M_L(\sigma) \in [a, b]$, the process $M_k(\sigma)$ may need to traverse a distance at most $2s$ over the final $L\gamma$ blocks, $k \in [\ell+1, L]$. This suggests that we consider the events

$$\left\{\frac{(g_{I_k}, \sigma_{I_k})}{N^{1/2}} \geq \frac{2s}{L\gamma}\right\}, \quad \left\{-\frac{(g_{I_k}, \sigma_{I_k})}{N^{1/2}} \geq \frac{2s}{L\gamma}\right\}$$

so that the process can traverse enough distance on each block. On the other hand, if $M_{k-1}(\sigma) \in [a, b]$, then we can ensure $M_k(\sigma) \in [a, b]$ by requiring one of the two events

$$\left\{0 \leq \frac{(g_{I_k}, \sigma_{I_k})}{N^{1/2}} \leq \frac{b-a}{2}\right\}, \quad \left\{0 \leq -\frac{(g_{I_k}, \sigma_{I_k})}{N^{1/2}} \leq \frac{b-a}{2}\right\}.$$

This leads to the statement of the following:

Lemma 3.9. *Let $\chi \in (0, 1]$ be a constant. Let S be any subset of $\{-1, +1\}^K$ with $|S| \geq \exp(K\delta)$ where δ is a positive constant. Let ϵ be a positive constant with $\psi_2(\epsilon) \leq \delta/2$.*

$$\mathbb{P}\left(\frac{1}{|S|}\left|\left\{\sigma \in S : \frac{(g, \sigma)}{N^{1/2}} \in \left[\frac{r\epsilon^{1/2}}{L^{1/2}}, \chi\right]\right\}\right| \leq \frac{1}{\exp(2C's^2)}\right) \leq \frac{2}{\exp(s^2\epsilon/C')},$$

provided that the parameters satisfy the bounds

$$C' \leq r \leq s \leq \min\left\{\frac{L^{1/2}\chi}{(6C')^{1/2}}, \frac{(K\delta)^{1/2}}{C'}\right\}.$$

Proof. It follows directly from Proposition 2.6 that for $C' \leq r \leq s \leq (K\delta)^{1/2}/C'$ we have

$$\mathbb{P}\left(\frac{1}{|S|}\left|\left\{\sigma \in S : \frac{(g, \sigma)}{N^{1/2}} \geq \frac{r\epsilon^{1/2}}{L^{1/2}}\right\}\right| \leq \frac{1}{\exp(C's^2)}\right) \leq \frac{1}{\exp(s^2\epsilon/C')},$$

On the other hand, it following using Markov's inequality and the gaussian tail bound that

$$\mathbb{P}\left(\frac{1}{|S|}\left|\left\{\sigma \in S : \frac{(g, \sigma)}{N^{1/2}} \geq \chi\right\}\right| \geq \frac{1}{\exp(2C's^2)}\right) \leq \frac{\exp(2C's^2)}{\exp(\chi^2 L/2)} \leq \frac{1}{\exp(C's^2)},$$

where the last bound holds provided that $s^2 \leq L\chi^2/(6C')$. Combining these bounds gives the claim. $\square$

The following is a version of Lemma 3.9 for general disorder. It has a worse ϵ -dependence and applies in a more limited range of s , as a consequence of applying Proposition 3.8 in place of Proposition 2.6.

Lemma 3.10. *Let $\chi \in (0, 1]$ be a constant. Let S be any subset of $\{-1, +1\}^K$ with $|S| \geq \exp(K\delta)$ where δ is a small positive constant. Let ϵ be a positive constant with $\psi_2(8\epsilon) \leq \delta/9$. Then*

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \in \left[\frac{r(\epsilon\nu)^{1/2}}{L^{1/2}}, \chi \right] \right\} \right| < \exp\left(-\frac{C''\nu s^2}{\epsilon}\right)\right) \leq \frac{5}{\exp(s^2\epsilon)},$$

provided that the parameters satisfy the bounds

$$C'' \leq r \leq s \leq \min \left\{ \frac{(L\epsilon)^{1/2}\chi}{2(C'')^{1/2}\nu}, \frac{1}{C''} \left(\frac{K\epsilon^3}{\nu} \right)^{1/4} \right\}.$$

Proof. It follows directly from Proposition 3.8 that for $C'' \leq r \leq s \leq (K\epsilon^3/\nu)^{1/4}/C''$ we have

$$\mathbb{P}\left(\frac{1}{|S|} \left| \left\{ \sigma \in S : \frac{(\xi, \sigma)}{N^{1/2}} \geq \frac{r(\epsilon\nu)^{1/2}}{L^{1/2}} \right\} \right| < \exp\left(-\frac{C''\nu s^2}{2\epsilon}\right)\right) \leq \frac{4}{\exp(s^2\epsilon)}.$$

On the other hand, it following using Markov's inequality and the subgaussian tail bound that

$$\mathbb{P}\left(\frac{1}{n} \left| \left\{ i \leq n : \frac{u_i}{L^{1/2}} \geq \chi \right\} \right| \geq \frac{1}{\exp(C''\nu s^2/\epsilon)}\right) \leq \frac{\exp(C''\nu s^2/\epsilon)}{\exp(\chi^2 L/(2\nu))} \leq \exp\left(-\frac{C''\nu s^2}{\epsilon}\right),$$

where the last bound holds provided that $s^2 \leq L\chi^2\epsilon/(4C''\nu^2)$. Combining these bounds gives the claim. $\square$

The assumption of the next proposition is based on the result of Proposition 3.2:

Proposition 3.11. *Let $\chi \equiv \min\{1, (b-a)/2\}$. Define*

$$\bar{R}^2 \equiv \frac{\bar{r}}{\gamma} \equiv \frac{\max\{|a|, |b|, C'\}}{\gamma}. \quad (3.5)$$

As before, divide $[N]$ into consecutive blocks $I_1, \dots, I_L$ of size $K = N/L$ each, such that L satisfies

$$\frac{4C''\nu^2\bar{r}}{(\chi\epsilon)^2} \leq L \leq \frac{(N\epsilon)^{1/3}\nu}{C''\chi^2}. \quad (3.6)$$

Now suppose $\Omega \subseteq \{-1, +1\}^N$ with $|\Omega| \geq \exp(2N\delta/L)$, such that all pairs $\sigma \neq \tau$ are ϵ -separated on I_j for all for all $j \in J_\star = [\ell + 1, L]$, where $\ell = L(1 - \gamma)$. Then we have the bound

$$\mathbb{P}\left(\frac{1}{|\Omega|} \left| \left\{ \sigma \in \Omega : \frac{(\xi, \sigma)}{N^{1/2}} \in [a, b] \right\} \right| \leq \frac{1}{4} \exp\left(-\frac{8C''\nu\bar{r}L}{\epsilon}\right)\right) \leq \exp\left(-\frac{L\epsilon^2}{2}\right),$$

provided that $\gamma = \epsilon \leq \delta$, and δ is small enough.

Proof. We emphasize that all the randomness is in the N -dimensional random vector ξ . Let $\mathcal{F}_k$ be the σ -algebra generated by the vectors ξ_{I_j} for $1 \leq j \leq k$, so that $M_k(\sigma)$ (as defined by (1.3)) is measurable with respect to $\mathcal{F}_k$. Recall $\ell = (1 - \gamma)L$, and define

$$\Omega^{(\ell)} \equiv \left\{ \sigma \in \Omega : \text{dist}\left(M_\ell(\sigma), [a, b]\right) \leq L\gamma \cdot \frac{\bar{R}\epsilon^{1/2}}{L^{1/2}} \cdot \nu^{1/2} \right\}.$$

Note that if $\sigma \notin \Omega^{(\ell)}$ then, by the definition (3.5) of $\bar{R}$, we have

$$|M_\ell(\sigma)| \geq L\gamma \cdot \frac{\bar{R}\epsilon^{1/2}}{L^{1/2}} \cdot \nu^{1/2} - \max\{|a|, |b|\} \geq L\gamma \cdot \frac{\bar{R}\epsilon^{1/2}}{2L^{1/2}} \cdot \nu^{1/2},$$

where the last bound holds by using (3.5) together with the lower bound on L from (3.6), along with the assumption $\gamma = \epsilon$ and the fact that $\nu \geq 1$ (see Assumption 1.1). It follows using Markov's inequality and the subgaussian tail bound that

$$\begin{aligned} \text{(I)} &\equiv \mathbb{P}\left(\frac{|\Omega^{(\ell)}|}{|\Omega|} \leq \frac{1}{2}\right) \leq \mathbb{P}\left(\frac{1}{|\Omega|} \left| \left\{ \sigma \in \Omega : |M_\ell(\sigma)| \geq L\gamma \cdot \frac{\bar{R}\epsilon^{1/2}}{2L^{1/2}} \cdot \nu^{1/2} \right\} \right| \geq \frac{1}{2}\right) \\ &\leq 4 \exp\left(-\frac{L\bar{R}^2\gamma^2\epsilon}{8}\right) \stackrel{(3.5)}{\leq} 4 \exp\left(-\frac{L\bar{r}\gamma\epsilon}{8}\right) \leq \frac{4}{\exp(L\epsilon^2)}, \end{aligned}$$

where the last inequality again uses $\gamma = \epsilon$, and the fact that $\bar{r} \geq C'$ (a large absolute constant) by the definition (3.5). Next, for $\ell + 1 \leq k \leq L$ let us define recursively the sets

$$\Omega^{(k)} \equiv \left\{ \sigma \in \Omega^{(k-1)} : \text{dist} \left(M_k(\sigma), [a, b] \right) \leq (L - k) \cdot \frac{\bar{R}\epsilon^{1/2}}{L^{1/2}} \cdot \nu^{1/2} \right\}.$$

For each $\ell \leq k \leq L$ we further define the following bipartition of $\Omega^{(k)}$:

$$\Omega^{(k),+} \equiv \left\{ \sigma \in \Omega^{(k)} : M_k(\sigma) \geq \frac{a+b}{2} \right\}, \quad \Omega^{(k),-} \equiv \Omega^{(k)} \setminus \Omega^{(k),+}.$$

One of these subsets must be at least half the size of $\Omega^{(k)}$; without loss of generality we assume it is $\Omega^{(k),+}$. Assuming $|\Omega^{(k),+}| \geq \exp(K\delta)$, we can extract disjoint subsets $X_1, \dots, X_a \subseteq \Omega^{(k),+}$ with $|X_a| = \exp(K\delta)$ such that the union of the X_a occupies more than half the mass of $\Omega^{(k),+}$. Define

$$\Xi^{(k)} \equiv \left\{ \sigma \in \Omega^{(k),+} : \frac{(\xi_{I_{k+1}}, \sigma_{I_{k+1}})}{N^{1/2}} \in \left[\frac{(\epsilon\nu)^{1/2}\bar{R}}{L^{1/2}}, \chi \right] \right\}.$$

It follows from the assumption that all pairs $\sigma \neq \tau$ in Ω are ϵ -separated on I_{k+1} , so applying Lemma 3.10 gives

$$\mathbb{P} \left(\frac{1}{m} \sum_{a \leq m} \mathbf{1} \left\{ \frac{|X_a \cap \Xi^{(k)}|}{|X_a|} \leq \exp \left(-\frac{C''\nu s^2}{\epsilon} \right) \right\} \leq \frac{1}{2} \right) \leq \frac{10}{\exp(s^2\epsilon)}, \quad (3.7)$$

provided that the parameters satisfy the bounds

$$\bar{R} \leq s \leq s_{\max} \equiv \frac{(L\epsilon)^{1/2}\chi}{2(C'')^{1/2}\nu} \leq s_2 \equiv \frac{1}{C''} \left(\frac{K\epsilon^3}{\nu} \right)^{1/4}. \quad (3.8)$$

Note that in the restriction (3.6), the lower bound on L guarantees that we in fact have $s_{\max} \geq \bar{R}$ above, while the upper bound on L guarantees $s_{\max} \leq s_2$. Note that $X_a \cap \Xi^{(k)} \subseteq \Omega^{(k+1)}$ for all $a \leq m$. It follows that, on the event $|\Omega^{(k)}| \geq 2\exp(K\delta)$, we have

$$\mathbb{P} \left(\frac{|\Omega^{(k+1)}|}{|\Omega^{(k)}|} \leq \frac{1}{4} \exp \left(-\frac{C''\nu s^2}{\epsilon} \right) \middle| \mathcal{F}_k \right) \leq \frac{10}{\exp(s^2\epsilon)}. \quad (3.9)$$

To apply this bound, we define the stopping time $\ell \leq \zeta \leq L$ by

$$\zeta \equiv \min \left\{ k \geq \ell : k = L, \text{ or } \frac{|\Omega^{(k)}|}{2} < \exp(K\delta) \right\}. \quad (3.10)$$

Applying (3.9) with $s = s_{\max}$ and $\ell \leq k < \zeta$ gives, on the event $|\Omega^{(\ell)}| \geq |\Omega|/2$, the bound

$$\begin{aligned} (\text{II}) &\equiv \mathbb{P} \left(\min_{\ell \leq k < \zeta} \frac{|\Omega^{(k+1)}|}{|\Omega^{(k)}|} \leq \frac{1}{4} \exp \left(-\frac{C''\nu(s_{\max})^2}{\epsilon} \right) \middle| \mathcal{F}_\ell \right) \\ &\stackrel{(3.9)}{\leq} \frac{L\gamma \cdot 10}{\exp((s_{\max})^2\epsilon)} = 10\gamma \exp \left(\log L - \frac{L\epsilon\chi^2}{4C''\nu^2} \right) \leq 10\gamma \exp \left(-\frac{L\epsilon\chi^2}{8C''\nu^2} \right), \end{aligned}$$

where the last bound again uses the lower bound on L from (3.6). We then define the truncated random variables

$$V_{k+1} \equiv \min \left\{ \frac{|\Omega^{(k)}|}{|\Omega^{(k+1)}|}, 4 \exp \left(\frac{C''\nu(s_{\max})^2}{\epsilon} \right) \right\}.$$

It follows by simplifying (3.9) that, on the event $|\Omega^{(k)}| \geq 2\exp(K\delta)$, we have for all $s \geq \bar{R}$ that

$$\mathbb{P} \left(V_{k+1} \geq \exp \left(\frac{2C''\nu s^2}{\epsilon} \right) \middle| \mathcal{F}_k \right) \leq \frac{10}{\exp(s^2\epsilon)}.$$

(More precisely, for $\bar{R} \leq s \leq s_{\max}$, the above bound holds due to (3.9). For $s \geq s_{\max}$, the bound holds trivially, because the truncation in the definition of V_{k+1} implies that the left-hand side above is zero.) Making the change of variables $t = \exp(s^2\epsilon/2)$ gives, for all $t \geq t_{\min} = \exp(\bar{R}^2\epsilon/2)$,

$$\mathbb{P} \left((V_{k+1})^{\epsilon^2/(4C''\nu)} \geq t \middle| \mathcal{F}_k \right) \leq \frac{10}{t^2}.$$

Recall from the definition (3.5), and the assumption $\gamma = \epsilon$, that $\bar{R}^2\epsilon = \bar{r} \geq C'$ (a large absolute constant). Integrating the above gives the moment bound

$$\mathbb{E} \left[(V_{k+1})^{\epsilon^2/(4C''\nu)} \middle| \mathcal{F}_k \right] \leq t_{\min} + \int_{t_{\min}}^{\infty} \frac{10}{t^2} dt = t_{\min} + \frac{10}{t_{\min}} \leq \exp(\bar{R}^2\epsilon), \quad (3.11)$$

where the last inequality uses that $\bar{R}^2\epsilon = \bar{r}$ is large by (3.5). Consequently, applying Markov's inequality we obtain, on the event $|\Omega^{(\ell)}| \geq |\Omega|/2$, for any $w \geq 0$,

$$\mathbb{P} \left(\prod_{\ell < k \leq \zeta} V_k \geq \exp \left(\frac{4C''\nu w^2}{\epsilon} \right) \middle| \mathcal{F}_\ell \right) \stackrel{(3.11)}{\leq} \frac{\exp(L\gamma \cdot \bar{R}^2\epsilon)}{\exp(w^2\epsilon)}.$$

Setting $w^2 = 2L\bar{R}^2\gamma$ then gives the bound

$$(III) \equiv \mathbb{P} \left(\prod_{\ell < k \leq \zeta} V_k \geq \exp \left(\frac{8C''\nu L \bar{R}^2\gamma}{\epsilon} \right) \middle| \mathcal{F}_\ell \right) \leq \frac{1}{\exp(L\bar{R}^2\gamma\epsilon)} \stackrel{(3.5)}{=} \frac{1}{\exp(L\bar{r}\epsilon)}.$$

Finally, let us note that on the event

$$E = \left\{ \frac{|\Omega^{(\ell)}|}{|\Omega|} \geq \frac{1}{2}, \min_{\ell \leq k < \zeta} \frac{|\Omega^{(k+1)}|}{|\Omega^{(k)}|} \geq \frac{1}{4} \exp \left(-\frac{C''\nu(s_{\max})^2}{\epsilon} \right), \prod_{\ell < k \leq \zeta} V_k \leq \exp \left(\frac{8C''\nu L \bar{R}^2\gamma}{\epsilon} \right) \right\},$$

we must have $|\Omega^{(k+1)}|/|\Omega^{(k)}| = 1/V_{k+1}$ for all $\ell \leq k < \zeta$. Consequently, recalling $K = N/L$, we have

$$\frac{|\Omega^{(\zeta)}|}{2} = \frac{|\Omega|}{2} \cdot \frac{|\Omega^{(\ell)}|}{|\Omega|} \cdot \prod_{\ell < k \leq \zeta} \frac{1}{V_k} \geq \frac{1}{4} \exp \left(2K\delta - \frac{8C''\nu L \bar{R}^2\gamma}{\epsilon} \right) \geq \exp(K\delta),$$

where the last inequality uses the upper bound on L from (3.6). From the definition (3.10) of ζ , this implies $\zeta = L$. Combining the above bounds gives

$$\begin{aligned} (IV) &\equiv \mathbb{P} \left(\frac{|\Omega^{(L)}|}{|\Omega|} \leq \frac{1}{4} \exp \left(-\frac{8C''\nu L \bar{R}^2\gamma}{\epsilon} \right) \right) \leq 1 - \mathbb{P}(E) \leq (I) + (II) + (III) \\ &\leq \frac{4}{\exp(L\epsilon^2)} + 10\gamma \exp \left(-\frac{L\epsilon\chi^2}{8C''\nu^2} \right) + \frac{1}{\exp(L\bar{r}\epsilon)} \leq \exp \left(-\frac{L\epsilon^2}{2} \right), \end{aligned}$$

where the last bound holds for ϵ small enough (which is guaranteed by taking δ small enough). Finally, recall that $\sigma \in \Omega^{(L)}$ implies $M_L(\sigma) = (\xi, \sigma)/N^{1/2} \in [a, b]$ as desired, so this concludes the proof. $\square$

Proof of Theorem 1.5. Since we can adjust the constants later, we can assume $|S| \geq \exp(9N\delta)$. Recalling the statement of Proposition 3.2, let ϵ, γ be positive constants with $\psi_2(8\epsilon) = \delta$ and $\gamma = \epsilon$. Assume L satisfies (3.6). It follows from Proposition 3.2 that there exists $J_1 \subseteq [L]$ with $|J_1| = L\gamma$, and $\Omega_1 \subseteq S$ with $|\Omega_1| = \exp(N\delta/L)$, such that all pairs $\sigma \neq \tau$ in Ω_1 are ϵ -separated on I_j for all $j \in J_1$. We can apply Proposition 3.2 again on $S \setminus \Omega_1$, and so on, to extract disjoint subsets $\Omega_1, \dots, \Omega_m \subseteq S$, up to the first m such that

$$\sum_{a \leq m} |\Omega_a| \geq \frac{|S|}{2}.$$

Then for each a we will have $|\Omega_a| = \exp(N\delta/L)$, and all pairs $\sigma \neq \tau$ in Ω_a will be ϵ -separated on I_j for all $j \in J_a$, where $J_a \subseteq [L]$ with $|J_a| = L\gamma$. Let

$$\Xi \equiv \left\{ \sigma \in S : \frac{(g, \sigma)}{N^{1/2}} \in [a, b] \right\}, \quad \Xi_a \equiv \Xi \cap \Omega_a.$$

By Markov's inequality and Proposition 3.11,

$$\mathbb{P} \left(\frac{1}{m} \sum_{a \leq m} \mathbf{1} \left\{ \frac{|\Xi_a|}{|\Omega_a|} \leq \frac{1}{4} \exp \left(-\frac{8C''\nu \bar{r} L}{\epsilon} \right) \right\} \geq \frac{1}{2} \right) \leq \exp \left(-\frac{L\epsilon^2}{2} \right).$$

On the complementary event we must have

$$|\Xi| \geq \sum_{a \leq m} |\Xi_a| \geq \frac{m|\Omega_1|}{2} \cdot \frac{1}{4} \exp\left(-\frac{8C''v\bar{r}L}{\epsilon}\right) \geq \frac{|S|}{16} \exp\left(-\frac{8C''v\bar{r}L}{\epsilon}\right),$$

so we have shown that, for any L satisfying (3.6), we have

$$\mathbb{P}\left(\frac{|\Xi|}{|S|} \leq \frac{1}{16} \exp\left(-\frac{8C''v\bar{r}L}{\epsilon}\right)\right) \leq \exp\left(-\frac{L\epsilon^2}{2}\right).$$

The conclusion follows by recalling that we chose γ and ϵ depending on δ . $\square$

Proof of Theorem 1.4. This follows essentially by the [proof of Theorem 1.5](#), the main difference being that in the proof of Proposition 3.11 we can apply Lemma 3.9 in place of Lemma 3.10. Thus in place of (3.7) we will have

$$\mathbb{P}\left(\frac{1}{m} \sum_{a \leq m} \mathbf{1}\left\{\frac{|X_a \cap \Xi^{(k)}|}{|X_a|} \leq \frac{1}{\exp(2C's^2)}\right\} \leq \frac{1}{2}\right) \leq \frac{4}{\exp(s^2\epsilon/C')},$$

where instead of (3.8) we will require (from the conditions of Lemma 3.9)

$$\bar{R} \leq s \leq s_{\max} \equiv \frac{L^{1/2}\chi}{(6C')^{1/2}} \leq s_2 \equiv \frac{(K\delta)^{1/2}}{C'}.$$

Since $KL = N$, this explains why we ultimately require $w \leq N^{1/2}/C_\delta$, in contrast with Theorem 1.5 which requires $w \leq N^{1/3}/C_\delta$. $\square$

4. UNIVERSALITY ON AVERAGE IN PERCEPTRON MODELS

Most of this section is devoted to the proof of the following theorem, which says that the **expectation** of the perceptron free energy (with a suitable truncation) is universal with respect to the disorder. In Section 5 we will combine this with a concentration result (Proposition 5.4) to yield the [proof of Theorem 1.11](#).

Theorem 4.1. *Consider a perceptron model (1.2) where U is $\{0, 1\}$ -valued and piecewise continuous. Let $Z(g)$ be the partition function with standard gaussian disorder, and let $Z(\xi)$ be the partition function with general disorder, assuming the $(\xi^a)_i$ are i.i.d. random variables with zero mean, unit variance, and finite third moment. Then*

$$\frac{1}{N} \left| \mathbb{E} \log_{N\delta} Z(g) - \mathbb{E} \log_{N\delta} Z(\xi) \right| = o_N(1).$$

for any small positive constant δ .

This section is organized as follows:

- In §4.1 we give some multivariate central limit theorem estimates. The main estimate is Corollary 4.4.
- In §4.2 we use Corollary 4.4 to prove a version (Proposition 4.5) of Theorem 4.1 for a “positive-temperature” variant of the model.
- In §4.3 we show how to transfer the result of Proposition 4.5 from positive temperature to zero temperature, leading to the [proof of Theorem 4.1](#).

4.1. Central limit theorem estimates. In this subsection we state and prove a consequence of the multivariate central limit theorem, Corollary 4.4 below, which was already used in the proof of Proposition 3.7. Corollary 4.4 will also be used below in the proof of Proposition 4.5, which is a preliminary universality result. We begin with a central limit theorem for sufficiently smooth functions, which is similar to [Mos10, Thm. 4.1]:

Lemma 4.2. *Let g be a standard gaussian vector in $\mathbb{R}^N$. Let ξ be a random vector in $\mathbb{R}^N$ with i.i.d. coordinates ξ_i , with mean zero, unit variance, and finite third moment. If $F : \mathbb{R}^p \rightarrow \mathbb{R}$ is a bounded function with derivatives up to third order bounded uniformly by a finite constant C_F , then we have*

$$\left| \mathbb{E} \left[F\left(\frac{\Sigma \xi}{N^{1/2}}\right) - F\left(\frac{\Sigma g}{N^{1/2}}\right) \right] \right| \leq \frac{pC_F}{N^{1/2}} \left\{ \mathbb{E}(|\xi_1|^3) + \mathbb{E}(|g_1|^3) \right\}$$

for any matrix $\Sigma \in \{-1, +1\}^{p \times N}$.

Proof. Write $\sigma^{(\ell)} \in \{-1, +1\}^N$ for the ℓ -th row of Σ . Let $\Sigma_{j,\circ}$ denote the $p \times N$ matrix that results from zeroing the j -th row of Σ . We interpolate between g and ξ by defining $\xi^j \equiv (\xi_1, \dots, \xi_j, g_{j+1}, \dots, g_N)$. Let $\xi^{j,\circ}$ denote the vector that results from zeroing the j -th entry of ξ^j . We then define the p -dimensional vectors

$$\Delta_j \equiv \frac{\Sigma \xi^j}{N^{1/2}} = \left(\sum_{i=1}^j \frac{\xi_i(\sigma^{(\ell)})_i}{N^{1/2}} + \sum_{i=j+1}^N \frac{g_i(\sigma^{(\ell)})_i}{N^{1/2}} \right)_{\ell \leq p},$$

$$\Delta_{j,\circ} \equiv \frac{\Sigma_{j,\circ} \xi^{j,\circ}}{N^{1/2}} = \left(\sum_{i=1}^{j-1} \frac{\xi_i(\sigma^{(\ell)})_i}{N^{1/2}} + \sum_{i=j+1}^N \frac{g_i(\sigma^{(\ell)})_i}{N^{1/2}} \right)_{\ell \leq p}.$$

Then the quantity of interest can be bounded as

$$\left| \mathbb{E}[F(\Delta_N) - F(\Delta_0)] \right| \leq \sum_{j \leq N} \left| \mathbb{E}[F(\Delta_j) - F(\Delta_{j-1})] \right|,$$

so it suffices to bound $|\mathbb{E}[F(\Delta_j) - F(\Delta_{j-1})]|$ for each $1 \leq j \leq N$. To this end we let

$$\Delta_{j,\ell} \equiv \Delta_{j,\circ} + \frac{1}{N^{1/2}} \left(\xi_j(\sigma^{(1)})_j, \dots, \xi_j(\sigma^{(\ell)})_j, g_j(\sigma^{(\ell+1)})_j, \dots, g_j(\sigma^{(p)})_j \right),$$

$$\Delta_{j,\ell,\circ} \equiv \Delta_{j,\circ} + \frac{1}{N^{1/2}} \left(\xi_j(\sigma^{(1)})_j, \dots, \xi_j(\sigma^{(\ell-1)})_j, 0, g_j(\sigma^{(\ell+1)})_j, \dots, g_j(\sigma^{(p)})_j \right),$$

and note that $F(\Delta_{j-1}) = F(\Delta_{j,0})$ while $F(\Delta_j) = F(\Delta_{j,p})$. It follows that

$$\left| \mathbb{E}[F(\Delta_j) - F(\Delta_{j-1})] \right| \leq \sum_{\ell=1}^p \left| \mathbb{E}[F(\Delta_{j,\ell}) - F(\Delta_{j,\ell-1})] \right|.$$

We then apply the assumption on F to Taylor expand

$$F(\Delta_{j,\ell}) = F(\Delta_{j,\ell,\circ}) + \partial_\ell F(\Delta_{j,\ell,\circ}) \frac{\xi_j(\sigma^{(\ell)})_j}{N^{1/2}} + \frac{(\partial_\ell)^2 F(\Delta_{j,\ell,\circ})}{2} \frac{(\xi_j)^2}{N} + \frac{R^\xi}{N^{3/2}},$$

$$F(\Delta_{j,\ell-1}) = F(\Delta_{j,\ell,\circ}) + \partial_\ell F(\Delta_{j,\ell,\circ}) \frac{g_j(\sigma^{(\ell)})_j}{N^{1/2}} + \frac{(\partial_\ell)^2 F(\Delta_{j,\ell,\circ})}{2} \frac{(g_j)^2}{N} + \frac{R^g}{N^{3/2}},$$

where $\mathbb{E}|R^\xi| \leq C_F \mathbb{E}(|\xi_j|^3)$ and $\mathbb{E}|R^g| \leq C_F \mathbb{E}(|g_j|^3)$. Matching the first and second moments of ξ_j and g_j gives

$$\left| \mathbb{E}[F(\Delta_{j,\ell}) - F(\Delta_{j,\ell-1})] \right| \leq \frac{C_F}{N^{3/2}} \left\{ \mathbb{E}(|\xi_j|^3 + |g_j|^3) \right\}.$$

Summing over $1 \leq \ell \leq p$ and $1 \leq j \leq N$ gives the desired overall bound on $|\mathbb{E}[F(\Delta_N) - F(\Delta_0)]|$. $\square$

Lemma 4.3. *Let g be a standard gaussian vector in $\mathbb{R}^N$. Let ξ be a random vector in $\mathbb{R}^N$ with i.i.d. coordinates ξ_i , with mean zero, unit variance, and finite third moment. If $F : \mathbb{R}^p \rightarrow \mathbb{R}$ is a bounded continuous function, then*

$$\lim_{N \rightarrow \infty} \max \left\{ \left| \mathbb{E} \left[F \left(\frac{\Sigma \xi}{N^{1/2}} \right) - F \left(\frac{\Sigma g}{N^{1/2}} \right) \right] \right| : \Sigma \in \{-1, +1\}^{p \times N} \right\} = 0.$$

Proof. For $\eta > 0$ define the smoothed function $F_\eta(x) \equiv \mathbb{E}F(x + \eta z)$ where z is a standard gaussian vector in $\mathbb{R}^p$. For any $\eta > 0$ the function F_η satisfies the conditions of Lemma 4.2, so we have

$$(I) \equiv \max \left\{ \left| \mathbb{E} \left[F_\eta \left(\frac{\Sigma \xi}{N^{1/2}} \right) - F_\eta \left(\frac{\Sigma g}{N^{1/2}} \right) \right] \right| : \Sigma \in \{-1, +1\}^N \right\} \leq \frac{p C_{F,\eta}}{N^{1/2}} \left\{ \mathbb{E}(|\xi_j|^3 + |g_j|^3) \right\}. \quad (4.1)$$

Since F is bounded and continuous, the function F_η converges locally uniformly to F as $\eta \downarrow 0$, so we have

$$(II) \equiv \max \left\{ \left| \mathbb{E} \left[F_\eta \left(\frac{\Sigma \xi}{N^{1/2}} \right) - F \left(\frac{\Sigma \xi}{N^{1/2}} \right) \right] \right| : \frac{\|\Sigma \xi\|_\infty}{N^{1/2}} \leq R \right\} \\ \leq \sup \left\{ |F_\eta(x) - F(x)| : \|x\|_\infty \leq R \right\} \xrightarrow{\eta \downarrow 0} 0$$

for any finite R . For the complementary event $\|\Sigma\xi\|_\infty > R$, we can use Chebychev's inequality to bound

$$\begin{aligned} \text{(III)} &\equiv \max \left\{ \mathbb{E} \left[\left| F_\eta \left(\frac{\Sigma\xi}{N^{1/2}} \right) - F \left(\frac{\Sigma\xi}{N^{1/2}} \right) \right| ; \frac{\|\Sigma\xi\|_\infty}{N^{1/2}} > R \right] : \Sigma \in \{-1, +1\}^N \right\} \\ &\leq 2\|F\|_\infty \mathbb{P} \left(\frac{\|\Sigma\xi\|_\infty}{N^{1/2}} > R \right) \leq \frac{2p\|F\|_\infty}{R^2}, \end{aligned}$$

which can be made arbitrarily small by choosing $R = R(p)$ large. The claim follows by combining the estimates for (I), (II), and (III). $\square$

Corollary 4.4. *Let g be a standard gaussian vector in $\mathbb{R}^N$. Let ξ be a random vector in $\mathbb{R}^N$ with i.i.d. coordinates ξ_i , with mean zero, unit variance, and finite third moment. If $f : \mathbb{R} \rightarrow \mathbb{R}$ is a bounded piecewise continuous function, and we define $F : \mathbb{R}^p \rightarrow \mathbb{R}$ by $F(x_1, \dots, x_p) = f(x_1) \cdots f(x_p)$, then*

$$\lim_{N \rightarrow \infty} \max \left\{ \left| \mathbb{E} \left[F \left(\frac{\Sigma\xi}{N^{1/2}} \right) - F \left(\frac{\Sigma g}{N^{1/2}} \right) \right] \right| : \Sigma \in \{-1, +1\}^{p \times N} \right\} = 0.$$

Proof. As in the proof of Lemma 4.3, consider the smoothed function

$$F_\eta(x) = \mathbb{E} F(x + \eta z) = \mathbb{E} \left(\prod_{\ell=1}^p f(x_\ell + \eta z_\ell) \right) = \prod_{\ell=1}^p f_\eta(x_\ell).$$

We then have the approximation result of Lemma 4.2, which gives (as in (4.1))

$$\max \left\{ \left| \mathbb{E} \left[F_\eta \left(\frac{\Sigma\xi}{N^{1/2}} \right) - F_\eta \left(\frac{\Sigma g}{N^{1/2}} \right) \right] \right| : \Sigma \in \{-1, +1\}^N \right\} \leq \frac{pC_{F,\eta}}{N^{1/2}} \left\{ \mathbb{E} (|\xi_j|^3 + |g_j|^3) \right\}. \quad (4.2)$$

We next bound

$$\mathbb{E} \left[\left| F_\eta \left(\frac{\Sigma\xi}{N^{1/2}} \right) - F \left(\frac{\Sigma\xi}{N^{1/2}} \right) \right| \right] \leq (\|f\|_\infty)^{p-1} \sum_{\ell=1}^p \mathbb{E} \left[\left| f_\eta \left(\frac{(\xi, \sigma^{(\ell)})}{N^{1/2}} \right) - f \left(\frac{(\xi, \sigma^{(\ell)})}{N^{1/2}} \right) \right| \right], \quad (4.3)$$

and likewise with g in place of ξ . Recall the assumption that f is piecewise continuous, so its set of discontinuities D_f is finite. Therefore, given any $\theta > 0$, we can construct a continuous function $\Upsilon : \mathbb{R} \rightarrow [0, 1]$ such that $\Upsilon = 1$ on an open set $O_f \supseteq D_f$ with $\mathbb{R} \setminus O_f$ compact, but $\mathbb{E} \Upsilon(g_1) \leq \theta$ if g_1 is a standard gaussian random variable. Applying Lemma 4.3 gives, for any fixed choice of O_f and Υ ,

$$\Delta_N \equiv \max \left\{ \left| \mathbb{E} \left[\Upsilon \left(\frac{(\xi, \sigma)}{N^{1/2}} \right) - \Upsilon \left(\frac{(g, \sigma)}{N^{1/2}} \right) \right] \right| : \sigma \in \{-1, +1\}^N \right\} \xrightarrow{N \rightarrow \infty} 0.$$

For any fixed $\sigma \in \{-1, +1\}^N$, the scalar product $(g, \sigma)/N^{1/2}$ is a standard gaussian random variable. It follows that

$$\begin{aligned} \text{(I)} &\equiv \max \left\{ \mathbb{E} \left[\left| f_\eta \left(\frac{(\xi, \sigma^{(\ell)})}{N^{1/2}} \right) - f \left(\frac{(\xi, \sigma^{(\ell)})}{N^{1/2}} \right) \right| ; \frac{(\xi, \sigma^{(\ell)})}{N^{1/2}} \in O_f \right] : \sigma^{(\ell)} \in \{-1, +1\}^N \right\} \\ &\leq 2\|f\|_\infty \mathbb{E} \left[\Upsilon \left(\frac{(\xi, \sigma)}{N^{1/2}} \right) \right] \leq 2\|f\|_\infty \left\{ \Delta_N + \mathbb{E} \left[\Upsilon \left(\frac{(g, \sigma)}{N^{1/2}} \right) \right] \right\} \leq 2\|f\|_\infty \left\{ \Delta_N + \theta \right\} \end{aligned}$$

We have f_η converging to f uniformly on the compact set $\mathbb{R} \setminus O_f$, so

$$\text{(II)} \equiv \max \left\{ \mathbb{E} \left[\left| f_\eta \left(\frac{(\xi, \sigma^{(\ell)})}{N^{1/2}} \right) - f \left(\frac{(\xi, \sigma^{(\ell)})}{N^{1/2}} \right) \right| ; \frac{(\xi, \sigma^{(\ell)})}{N^{1/2}} \in \mathbb{R} \setminus O_f \right] : \sigma^{(\ell)} \in \{-1, +1\}^N \right\} \xrightarrow{\eta \downarrow 0} 0.$$

Of course, the estimates for (I) and (II) also hold with g in place of ξ . The claim follows by substituting those estimates into (4.3), and combining with (4.2). $\square$

4.2. Universality at positive temperature. In this subsection we establish universality for a perceptron model at “positive temperature,” more precisely, of the form (1.2) where $\log U : \mathbb{R} \rightarrow (-\infty, 0]$ is uniformly bounded. Note that the Azuma–Hoeffding martingale inequality implies in this case that $N^{-1} \log Z(U; \xi)$ is exponentially well concentrated around its mean, although we will not ultimately use this fact. The following result shows that the mean value is universal with respect to the disorder model:

Proposition 4.5. *For the “positive-temperature” perceptron model (1.2) where $\log U : \mathbb{R} \rightarrow (-\infty, 0]$ is uniformly bounded and piecewise continuous, let $Z(U; g)$ be the partition function with standard gaussian disorder, and let $Z(U; \xi)$ be the partition function with general disorder, assuming the $(\xi^a)_i$ are i.i.d. random variables with zero mean, unit variance, and finite third moment. Then*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \left| \mathbb{E} \log Z(U; g) - \mathbb{E} \log Z(U; \xi) \right| = 0.$$

Proof. It follows from Corollary 4.4 that, uniformly over $\sigma^{(1)}, \dots, \sigma^{(\ell)} \in \{-1, +1\}^N$, we have

$$\Delta_j(\sigma^{(1)}, \dots, \sigma^{(\ell)}) \equiv \mathbb{E} \left[\prod_{k=1}^{\ell} U \left(\frac{(\xi^j, \sigma^{(k)})}{N^{1/2}} \right) - \prod_{k=1}^{\ell} U \left(\frac{(g^j, \sigma^{(k)})}{N^{1/2}} \right) \right] \xrightarrow{N \rightarrow \infty} 0. \quad (4.4)$$

We interpolate between $Z(U; g)$ and $Z(U; \xi)$ by defining

$$\begin{aligned} Z_j &\equiv \sum_{\sigma \in \{-1, +1\}^N} \prod_{i=1}^j U \left(\frac{(\xi^i, \sigma)}{N^{1/2}} \right) \prod_{i=j+1}^M U \left(\frac{(g^i, \sigma)}{N^{1/2}} \right), \\ Z_{j,\circ} &\equiv \sum_{\sigma \in \{-1, +1\}^N} \prod_{i=1}^{j-1} U \left(\frac{(\xi^i, \sigma)}{N^{1/2}} \right) \prod_{i=j+1}^M U \left(\frac{(g^i, \sigma)}{N^{1/2}} \right). \end{aligned}$$

Then $Z_0 = Z(U; g)$ and $Z_M = Z(U; \xi)$, so

$$\log Z(U; \xi) - \log Z(U; g) = \sum_{j \leq M} \left(\log Z_j - \log Z_{j-1} \right) \equiv \sum_{j \leq M} Y_j.$$

Writing $\langle \cdot \rangle_{j,\circ}$ for expectation with respect to the Gibbs measure $\mu_{j,\circ}$ corresponding to $Z_{j,\circ}$, we have

$$Y_j = \log \frac{Z_j}{Z_{j,\circ}} - \log \frac{Z_{j-1}}{Z_{j,\circ}} = \log \left\langle U \left(\frac{(\xi^j, \sigma)}{N^{1/2}} \right) \right\rangle_{j,\circ} - \log \left\langle U \left(\frac{(g^j, \sigma)}{N^{1/2}} \right) \right\rangle_{j,\circ}$$

Given any $A < \infty$ and $\eta > 0$ we can choose $p = p(A, \eta) < \infty$ such that

$$\sup \left\{ \left| \log(1+x) - \sum_{\ell \leq p} \frac{(-1)^{\ell+1}}{\ell} x^\ell \right| : \frac{1}{\exp(A)} \leq 1+x \leq 0 \right\} \leq \eta$$

It follows that $|\mathbb{E} Y_\ell - \bar{y}_\ell| \leq \eta$ where

$$\bar{y}_j \equiv \sum_{\ell \leq p} \frac{(-1)^{\ell+1}}{\ell} \mathbb{E} \left[\left(\left\langle U \left(\frac{(\xi^j, \sigma)}{N^{1/2}} \right) \right\rangle_{j,\circ} - 1 \right)^\ell - \left(\left\langle U \left(\frac{(g^j, \sigma)}{N^{1/2}} \right) \right\rangle_{j,\circ} - 1 \right)^\ell \right].$$

Let $c_{p,\ell}$ denote the coefficients such that

$$\sum_{\ell \leq p} \frac{(-1)^{\ell+1}}{\ell} (x-1)^\ell = \sum_{\ell \leq p} c_{p,\ell} x^\ell.$$

Recalling that $\mu_{j,\circ}$ is the Gibbs measure corresponding to $Z_{j,\circ}$, we can rewrite

$$\begin{aligned} \bar{y}_j &= \sum_{\ell \leq p} c_{p,\ell} \mathbb{E} \left[\left(\left\langle U \left(\frac{(\xi^j, \sigma)}{N^{1/2}} \right) \right\rangle_{j,\circ} \right)^\ell - \left(\left\langle U \left(\frac{(g^j, \sigma)}{N^{1/2}} \right) \right\rangle_{j,\circ} \right)^\ell \right] \\ &= \sum_{\ell \leq p} c_{p,\ell} \sum_{\sigma^{(1)}, \dots, \sigma^{(\ell)}} \Delta_j(\sigma^{(1)}, \dots, \sigma^{(\ell)}) \prod_{k=1}^{\ell} \mu_{j,\circ}(\sigma^{(k)}), \end{aligned}$$

which by (4.4) converges to zero. The claim follows by recalling that $|\mathbb{E}Y_\ell - \bar{y}_\ell| \leq \eta$, and $\eta > 0$ was arbitrary. $\square$

4.3. Positive temperature to zero temperature. We first review a simple inequality (4.5) which was previously proved by [Tal11, Lem. 8.3.7]. If $x \geq y \geq 0$, then for any $\Gamma < 0$ we have

$$\log_\Gamma y - \log_\Gamma x = \begin{cases} \log y - \log x & \text{if } x \geq y \geq \exp(\Gamma), \\ \Gamma - \log x & \text{if } x \geq \exp(\Gamma) \geq y, \\ 0 & \text{if } \exp(\Gamma) \geq x \geq y. \end{cases}$$

If we further assume $1 \geq x \geq y \geq 0$, then the above quantities are all in $[\Gamma, 0]$. It follows that

$$0 \geq \log_\Gamma y - \log_\Gamma x \geq \max \left\{ \log \frac{y}{x}, \Gamma \right\} = \log_\Gamma \frac{y}{x}. \quad (4.5)$$

We will use this inequality in the proof of the next result.

Proposition 4.6. *Consider a perceptron model (1.2) where U is $\{0, 1\}$ -valued and piecewise continuous. Let $Z(\xi)$ denote the partition function with general disorder, assuming the $(\xi^a)_i$ are i.i.d. random variables with zero mean, unit variance, and finite third moment. Let $u_A(x) \equiv \max\{\log U(x), -A\}$. Then for any positive constant δ we have*

$$\frac{1}{N} \mathbb{E} \left| \log_{N\delta} Z(u_A; \xi) - \log_{N\delta} Z(\xi) \right| \leq o_N(1) + o_A(1)$$

in the limit $N \rightarrow \infty$ followed by $A \rightarrow \infty$.

Proof. We now interpolate between $Z(\xi)$ and $Z(u_A; \xi)$ by defining

$$\begin{aligned} Z_j &\equiv \sum_{\sigma \in \{-1, +1\}^N} \exp \left\{ \sum_{i=1}^j u_A \left(\frac{(\xi^i, \sigma)}{N^{1/2}} \right) \right\} \cdot \prod_{i=j+1}^M U \left(\frac{(g^i, \sigma)}{N^{1/2}} \right), \\ Z_{j,\circ} &\equiv \sum_{\sigma \in \{-1, +1\}^N} \exp \left\{ \sum_{i=1}^{j-1} u_A \left(\frac{(\xi^i, \sigma)}{N^{1/2}} \right) \right\} \cdot \prod_{i=j+1}^M U \left(\frac{(g^i, \sigma)}{N^{1/2}} \right). \end{aligned}$$

Then $Z_0 = Z(\xi) \leq Z_1 \leq \dots \leq Z_M = Z(u_A; \xi)$, so

$$\frac{1}{N} \left| \log_{N\delta} Z(u_A; \xi) - \log_{N\delta} Z(\xi) \right| = \frac{1}{N} \sum_{j \leq M} \left(\log_{N\delta} Z_j - \log_{N\delta} Z_{j-1} \right) \equiv \frac{1}{N} \sum_{j \leq M} Y_j.$$

Let $E_{j,\circ}$ denote the event that $Z_{j,\circ} \geq \exp(N\delta)$. Note that $Z_{j,\circ} \geq \max\{Z_j, Z_{j-1}\}$, so $Y_j = 0$ on the complement of $E_{j,\circ}$. It follows that $Y_j = \dot{Y}_j - \ddot{Y}_j$ where

$$\begin{aligned} \dot{Y}_j &\equiv \mathbf{1}\{E_{j,\circ}\} \left(\log_{N\delta} Z_j - \log_{N\delta} Z_{j,\circ} \right) = \mathbf{1}\{E_{j,\circ}\} \log \max \left\{ \frac{Z_j}{Z_{j,\circ}}, \frac{\exp(N\delta)}{Z_{j,\circ}} \right\}, \\ \ddot{Y}_j &\equiv \mathbf{1}\{E_{j,\circ}\} \left(\log_{N\delta} Z_{j-1} - \log_{N\delta} Z_{j,\circ} \right) = \mathbf{1}\{E_{j,\circ}\} \log \max \left\{ \frac{Z_{j-1}}{Z_{j,\circ}}, \frac{\exp(N\delta)}{Z_{j,\circ}} \right\}. \end{aligned}$$

Writing $\langle \cdot \rangle_{j,\circ}$ for expectation with respect to the Gibbs measure corresponding to $Z_{j,\circ}$, we have

$$\begin{aligned} \ddot{x}_j &\equiv \frac{Z_{j-1}}{Z_{j,\circ}} = \left\langle U \left(\frac{(\xi^j, \sigma)}{N^{1/2}} \right) \right\rangle_{j,\circ}, \\ \dot{x}_j &\equiv \frac{Z_j}{Z_{j,\circ}} = \left\langle \exp u_A \left(\frac{(\xi^j, \sigma)}{N^{1/2}} \right) \right\rangle_{j,\circ} = e^{-A} + (1 - e^{-A}) \ddot{x}_j, \end{aligned}$$

and clearly $0 \leq \ddot{x}_j \leq \dot{x}_j \leq 1$. Moreover we can bound

$$1 \leq \frac{\dot{x}_j}{\ddot{x}_j} = \frac{e^{-A} + (1 - e^{-A}) \ddot{x}_j}{\ddot{x}_j} = 1 + \frac{e^{-A}(1 - \ddot{x}_j)}{\ddot{x}_j} \leq 1 + \frac{1}{e^A \ddot{x}_j}. \quad (4.6)$$

Combining with [Tal11, Lem. 8.3.7] (i.e., the bound (4.5) above) gives

$$\begin{aligned} 0 \geq -Y_j = \ddot{y}_j - \dot{y}_j &\stackrel{(4.5)}{\geq} \mathbf{1}\{E_{j,\circ}\} \log \max \left\{ \frac{\ddot{x}_j}{\dot{x}_j}, \frac{\exp(N\delta)}{Z_{j,\circ}} \right\} \\ &\stackrel{(4.6)}{\geq} \mathbf{1}\{E_{j,\circ}\} \max \left\{ -\log \left(1 + \frac{1}{e^A \ddot{x}_j} \right), -N \log 2 \right\}. \end{aligned} \quad (4.7)$$

Note $\log(1+x) \leq x$ for all $x \geq 0$, and $\log(1+x) \leq \log(2x)$ for $x \geq 1$. Writing $w_{\max} \equiv N^{1/3}/C_\delta$, we hereafter assume $A \geq C_\delta$, and decompose

$$\begin{aligned} \text{(I)} &\equiv \mathbb{E} \left[Y_j; \ddot{x}_j \leq \frac{1}{\exp(w_{\max})} \right] \stackrel{(4.7)}{\leq} N \log 2 \cdot \mathbb{P} \left(\ddot{x}_j \leq \frac{1}{\exp(w_{\max})}; E_{j,\circ} \right), \\ \text{(II)} &\equiv \mathbb{E} \left[Y_j; \frac{1}{\exp(w_{\max})} \leq \ddot{x}_j \leq \frac{1}{e^A} \right] \stackrel{(4.7)}{\leq} \mathbb{E} \left[\mathbf{1}\{E_{j,\circ}\} \log \left(\frac{2}{e^A \ddot{x}_j} \right); \frac{1}{\exp(w_{\max})} \leq \ddot{x}_j \leq \frac{1}{e^A} \right], \\ \text{(III)} &\equiv \mathbb{E} \left[Y_j; \frac{1}{e^A} \leq \ddot{x}_j \leq \frac{1}{\exp(C_\delta)} \right] \stackrel{(4.7)}{\leq} \mathbb{E} \left[\frac{\mathbf{1}\{E_{j,\circ}\}}{e^A \ddot{x}_j}; \frac{1}{e^A} \leq \ddot{x}_j \leq \frac{1}{\exp(C_\delta)} \right]. \end{aligned}$$

If $U \equiv 0$ there is nothing to prove, so we may assume $U(x) \geq \mathbf{1}\{x \in [a, b]\}$ for some $-\infty < a < b < \infty$. We can then apply Theorem 1.5 to bound the above quantities:

$$\begin{aligned} \text{(I)} &\leq N \log 2 \cdot \exp \left(-\frac{w_{\max}}{C_\delta} \right), \\ \text{(III)} &\leq \frac{1}{e^A} \int_{\exp(C_\delta)}^{\exp(A)} \frac{du}{u^{1/C_\delta}} \leq \frac{1}{e^A} \frac{(e^A)^{1-1/C_\delta}}{1-1/C_\delta} \leq \frac{2}{e^{A/C_\delta}}. \end{aligned}$$

Lastly, making a change of variables gives

$$\mathbb{P} \left(\log \left(\frac{2}{e^A \ddot{x}_j} \right) \geq t; E_{j,\circ} \right) = \mathbb{P} \left(\frac{1}{\ddot{x}_j} \geq \frac{e^{A+t}}{2}; E_{j,\circ} \right) \leq \left(\frac{2}{e^{A+t}} \right)^{1/C_\delta},$$

and integrating this tail bound over $t \geq 0$ gives

$$\text{(II)} \leq \int_0^\infty \left(\frac{2}{e^{A+t}} \right)^{1/C_\delta} dt \leq \frac{2C_\delta}{e^{A/C_\delta}}.$$

Combining the bounds for (I), (II), and (III) gives $0 \leq \mathbb{E}Y_j \leq o_A(1) + o_N(1)$, and the claim follows. $\square$

Proof of Theorem 4.1. Follows by combining Proposition 4.5 with Proposition 4.6. $\square$

The concentration of the free energy is addressed in the next section; see the [proof of Theorem 1.7](#).

5. CONCENTRATION, SHARP THRESHOLD SEQUENCE, AND UNIVERSALITY

In this section we prove Theorems 1.7–1.11. Recall the abstract model (1.4). Since we mainly consider how the system behaves as M varies, we will mostly drop N from the notation, e.g. we will abbreviate $Z_M \equiv Z_{M,N}$. Let $\mathcal{F}_a$ be the σ -field generated by the random functions $\Theta_1, \dots, \Theta_a$. The following is a stronger version of Assumption 1.6, which accommodates the perceptron models considered in this paper:

Assumption 5.1. For the model (1.4), let Z_{M+1} be the partition function that results from adding one more factor to Z_M . Suppose for all $\delta > 0$ small enough that on the event $Z_M \geq \exp(N\delta)$ we have

$$\mathbb{P} \left(\frac{Z_{M+1}}{Z_M} \leq \frac{1}{\exp(w)} \middle| \mathcal{F}_M \right) \leq \exp \left(-\frac{w}{C_\delta} \right)$$

for all $C_\delta \leq w \leq w_{\max}$, where C_δ is a finite constant that depends only on the model and on δ , and $w_{\max}$ can depend on the model as well as on δ and N .

In the case of the perceptron model, Assumption 5.1 holds with the following parameters:

(a) For the half-space perceptron (1.1) with gaussian disorder, Theorem 1.2 gives $w_{\max} = N/C_\delta$.

(b) For the half-space perceptron (1.1) with general disorder, Theorem 1.3 gives $w_{\max} = N^{1/2}/C_\delta$.

(c) For the U -perceptron (1.2) with gaussian disorder, Theorem 1.4 gives $w_{\max} = N^{1/2}/C_\delta$.

(d) For the U -perceptron (1.2) with general disorder, Theorem 1.5 gives $w_{\max} = N^{1/3}/C_\delta$.

(For both (c) and (d) above, we assume U is $\{0, 1\}$ -valued with $U(x) \geq \mathbf{1}\{x \in [a, b]\}$ for some $-\infty < a < b < \infty$.)

We also will show that the weaker Assumption 1.6 suffices for our main claims. This section is organized as follows:

- In §5.1 we show that Assumption 1.6 or 5.1 implies that the partition function of the model (1.4) is unlikely to decrease very sharply after the addition of a small linear number of constraints.
- In §5.2 we show that Assumption 1.6 or 5.1 implies concentration of the truncated free energy of the model (1.4), leading to the proof of Theorem 1.7.
- In §5.3 we combine the results obtained thus far to give the proof of Theorem 1.9 (on the sharp threshold sequence) and finally the proof of Theorem 1.11 (on universality).

5.1. Slow decrease of partition function with new constraints. In this subsection we prove Proposition 5.2, which says that under Assumption 5.1, the partition function of the model (1.4) is unlikely to decrease very sharply after adding a small linear number of constraints. We also prove Proposition 5.3 which gives a similar but weaker estimate under Assumption 1.6.

Proposition 5.2. *For the model (1.4), under Assumption 5.1, for all $\delta > 0$ there exists $\rho_\delta > 0$ such that for all $0 \leq \rho \leq \rho_\delta$ we have*

$$\mathbb{P}\left(Z_{M+N\rho} < \exp(N\delta) \mid Z_M \geq \exp(2N\delta)\right) \leq N\rho \exp\left(-\frac{w_{\max}}{C_\delta}\right) + \exp\left(-\frac{N\delta}{4C_\delta}\right).$$

In particular, if $w_{\max}$ grows faster than $\log N$, then this is $o_N(1)$.

Proof. For $0 \leq \ell \leq N\rho$ let us abbreviate $X_\ell = Z_{M+\ell}$ and $\mathcal{G}_\ell \equiv \mathcal{F}_{M+\ell}$. It follows by Assumption 5.1 that on the event $X_\ell \geq \exp(N\delta)$, we have the bound

$$\mathbb{P}\left(\frac{X_{\ell+1}}{X_\ell} \leq \frac{1}{\exp(w)} \mid \mathcal{G}_\ell\right) \leq \exp\left(-\frac{w}{C_\delta}\right), \quad (5.1)$$

provided $C_\delta \leq w \leq w_{\max}$. To apply this bound, we define a stopping time $0 \leq \zeta \leq N\rho$ by

$$\zeta \equiv \min\left\{\ell \geq 0 : \ell = N\rho \text{ or } X_\ell < \exp(N\delta)\right\}. \quad (5.2)$$

Applying (5.1) with $w = w_{\max}$ gives, on the event $X_0 \geq \exp(2N\delta)$,

$$(I) \equiv \mathbb{P}\left(\min_{\ell < \zeta} \frac{X_{\ell+1}}{X_\ell} \leq \frac{1}{\exp(w_{\max})} \mid \mathcal{G}_0\right) \leq N\rho \exp\left(-\frac{w_{\max}}{C_\delta}\right). \quad (5.3)$$

Next define the truncated random variables

$$V_{\ell+1} \equiv \min\left\{\frac{X_\ell}{X_{\ell+1}}, \exp(w_{\max})\right\} = \min\left\{\frac{Z_{M+\ell}}{Z_{M+\ell+1}}, \exp(w_{\max})\right\}.$$

On the event $X_\ell \geq \exp(N\delta)$, (5.1) implies the tail bound

$$\mathbb{P}\left((V_{\ell+1})^{1/(2C_\delta)} \geq y \mid \mathcal{G}_\ell\right) \leq \frac{1}{y^2},$$

for all $y \geq \exp(C_\delta/2)$. Integrating over y gives the moment bound

$$\mathbb{E}\left[(V_{\ell+1})^{1/(2C_\delta)} \mid \mathcal{G}_\ell\right] \leq \exp\left(\frac{C_\delta}{2}\right) + \int_{\exp(C_\delta/2)}^\infty \frac{dy}{y^2} \leq \exp(C_\delta). \quad (5.4)$$

Applying Markov's inequality gives, on the event $Z_M = X_0 \geq \exp(2N\delta)$,

$$(II) \equiv \mathbb{P}\left(\prod_{\ell \leq \zeta} V_\ell \geq \exp(N\delta) \mid \mathcal{G}_0\right) \leq \frac{\exp(N\rho C_\delta)}{\exp(N\delta/(2C_\delta))} \leq \exp\left(-\frac{N\delta}{4C_\delta}\right),$$

where the last bound holds by taking $\rho \leq \rho_\delta$ small enough. Thus, on the event $Z_M = X_0 \geq \exp(2N\delta)$, we have

$$\mathbb{P}\left(X_\zeta < \exp(N\delta) \mid \mathcal{G}_0\right) \leq \text{(I)} + \text{(II)} \leq N\rho \exp\left(-\frac{w_{\max}}{C_\delta}\right) + \exp\left(-\frac{N\delta}{4C_\delta}\right).$$

On the complementary event $X_\zeta \geq \exp(N\delta)$, it follows from the definition (5.2) that we must have $\zeta = N\rho$, which concludes the proof. $\square$

Proposition 5.3. *For the model (1.4), under Assumption 1.6, for all $\delta > 0$ there exists $\rho_\delta > 0$ such that for all $0 \leq \rho \leq \rho_\delta$ we have*

$$\mathbb{P}\left(Z_{M+N\rho} < \exp(N\delta) \mid Z_M \geq \exp(2N\delta)\right) \leq N\rho f_\delta(w_{\max}) + \frac{4\rho C_{\delta,2}}{N\delta^2}.$$

In particular, if $f_\delta(w_{\max})$ tends to zero superpolynomially in N , then this is $o_N(1)$.

Proof. Following the first few steps of the proof of Proposition 5.2, with $X_\ell \equiv Z_{M+\ell}$, $\mathcal{G}_\ell \equiv \mathcal{F}_{M+\ell}$, and the same stopping time ζ as defined by (5.2), yields that on the event $X_0 \geq \exp(2N\delta)$ we have

$$\text{(I)} \equiv \mathbb{P}\left(\min_{\ell < \zeta} \frac{X_{\ell+1}}{X_\ell} \leq \frac{1}{\exp(w_{\max})} \mid \mathcal{G}_0\right) \leq N\rho f_\delta(w_{\max}).$$

(instead of (5.3)). Next define the truncated random variables

$$V_{\ell+1} \equiv \min\left\{\frac{X_\ell}{X_{\ell+1}}, \exp(w_{\max})\right\},$$

and note that $V_{\ell+1} \geq 1$. Assumption 1.6 implies that on the event $X_\ell \geq \exp(N\delta)$ we have

$$\begin{aligned} \mathbb{E}\left(\log V_{\ell+1} \mid \mathcal{G}_\ell\right) &\leq C_\delta + \int_{C_\delta}^{\infty} f_\delta(w) dw \leq C_{\delta,1}. \\ \mathbb{E}\left((\log V_{\ell+1})^2 \mid \mathcal{G}_\ell\right) &\leq (C_\delta)^2 + \int_{C_\delta}^{\infty} 2w f_\delta(w) dw \leq C_{\delta,2}. \end{aligned}$$

Now consider the Doob martingale decomposition

$$\sum_{k=1}^{\ell} \log V_k = \sum_{k=1}^{\ell} \left(\log V_k - \mathbb{E}(\log V_k \mid \mathcal{G}_{k-1})\right) + \sum_{k=1}^{\ell} \mathbb{E}(\log V_k \mid \mathcal{G}_{k-1}) \equiv M_\ell + A_\ell$$

It follows from the preceding bounds that $0 \leq A_\zeta \leq N\rho C_{\delta,1}$ and $\mathbb{E}[(M_\zeta)^2] \leq N\rho C_{\delta,2}$. We can take ρ small enough to guarantee that $\rho C_{\delta,1} < \delta/2$. It follows using Chebychev's inequality that, on the event $X_0 \geq \exp(2N\delta)$,

$$\text{(II)} \equiv \mathbb{P}\left(\sum_{k \leq \zeta} \log V_k \geq N\delta \mid \mathcal{G}_0\right) \leq \mathbb{P}\left(M_\zeta \geq N\delta - N\rho C_{\delta,1} \geq \frac{N\delta}{2} \mid \mathcal{G}_0\right) \leq \frac{4\rho C_{\delta,2}}{N\delta^2}.$$

Combining the above bounds for (I) and (II) gives, on the event $X_0 \geq \exp(2N\delta)$,

$$\mathbb{P}\left(X_\zeta < \exp(N\delta) \mid \mathcal{G}_0\right) \leq \text{(I)} + \text{(II)} \leq N\rho f_\delta(w_{\max}) + \frac{4\rho C_{\delta,2}}{N\delta^2}.$$

On the complementary event $X_\zeta \geq \exp(N\delta)$, it follows from the definition (5.2) that we must have $\zeta = N\rho$, and this concludes the proof. $\square$

5.2. Concentration of truncated free energy. The first result below, Proposition 5.4, is fairly similar to the result of [Tal11, Propn. 9.2.6], and makes use of Assumption 5.1. The next result, Proposition 5.5, is a similar but weaker estimate that uses only Assumption 1.6. We conclude the subsection with the proof of Theorem 1.7.

Proposition 5.4. *For the model (1.4), under Assumption 5.1, for all $\delta > 0$ we have*

$$\mathbb{P}\left(\left|\log_{N\delta} Z - \mathbb{E} \log_{N\delta} Z\right| \geq 2Nt\right) \leq \frac{M \log 2}{t} \exp\left(-\frac{w_{\max}}{C_\delta}\right) + \exp\left(-\min\left\{\frac{Nt^2}{4(A_\delta)^2}, \frac{Nt}{2A_\delta}\right\}\right),$$

where A_δ is a large constant which depends only on C_δ (where C_δ is the constant appearing in Assumption 5.1). In particular, if $w_{\max}$ grows faster than $\log N$, then the above is $o_N(1)$ provided $t \gg 1/N^{1/2}$.

Proof. Recall that $\mathcal{F}_k$ denotes the σ -field generated by the random functions $\Theta_1, \dots, \Theta_k$. Consider the martingale decomposition

$$\log_{N\delta} Z - \mathbb{E}(\log_{N\delta} Z) = \sum_{k \leq M} \left\{ \mathbb{E}(\log_{N\delta} Z \mid \mathcal{F}_k) - \mathbb{E}(\log_{N\delta} Z \mid \mathcal{F}_{k-1}) \right\} \equiv \sum_{k \leq M} Y_k,$$

Let $Z_{k,\circ}$ be the partition function without the k -th factor:

$$Z_{k,\circ} \equiv \sum_{\sigma \in \{-1, +1\}^N} \prod_{\ell \leq M, \ell \neq k} \Theta_\ell(\sigma).$$

Since we assumed the Θ functions are $\{0, 1\}$ -valued, we have $Z \leq Z_{k,\circ}$, so if $Z_{k,\circ} \leq \exp(N\delta)$ then we must have $\log_{N\delta} Z = \log_{N\delta} Z_{k,\circ}$. It follows that

$$L_k \equiv \log_{N\delta} Z - \log_{N\delta} Z_{k,\circ} = \left(\log_{N\delta} Z - \log Z_{k,\circ} \right) \mathbf{1}_{\{Z_{k,\circ} \geq \exp(N\delta)\}}.$$

Let $\mathbb{E}_k$ denote expectation over Θ_k only, and let $\mathbb{E}^k$ denote expectation over all the $(\Theta_\ell)_{\ell \geq k}$. We can then rewrite

$$Y_k = \mathbb{E} \left(\log_{N\delta} Z - \log_{N\delta} Z_{k,\circ} \mid \mathcal{F}_k \right) - \mathbb{E} \left(\log_{N\delta} Z - \log_{N\delta} Z_{k,\circ} \mid \mathcal{F}_{k-1} \right) = \mathbb{E}^{k-1}(L_k - \mathbb{E}_k L_k).$$

For comparison, let us also define $y_k \equiv \mathbb{E}^{k-1}(\ell_k - \mathbb{E}_k \ell_k)$ where

$$\ell_k \equiv \left(\log_{N\delta} Z - \log Z_{k,\circ} \right) \mathbf{1}_{\{Z_{k,\circ} \geq \exp(N\delta); \frac{Z}{Z_{k,\circ}} \geq \frac{1}{\exp(w_{\max})}\}}$$

Since $-N \log 2 \leq \log_{N\delta} Z - \log_{N\delta} Z_{k,\circ} \leq 0$, we can use Assumption 5.1 to bound

$$\begin{aligned} \text{(I)} &\equiv \mathbb{P} \left(\sum_{k \leq M} |Y_k - y_k| \geq Nt \right) \leq \frac{1}{Nt} \mathbb{E} \left(\sum_{k \leq M} |Y_k - y_k| \right) \\ &\leq \frac{MN \log 2}{Nt} \mathbb{P} \left(Z_{k,\circ} \geq \exp(N\delta); \frac{Z}{Z_{k,\circ}} \geq \frac{1}{\exp(w_{\max})} \right) \leq \frac{MN \log 2}{Nt} \exp \left(-\frac{w_{\max}}{C_\delta} \right). \end{aligned} \quad (5.5)$$

We will bound, for small enough θ , the exponential moment

$$\begin{aligned} \mathbb{E} \left(\exp(\theta |y_k|) \mid \mathcal{F}_{k-1} \right) &= \mathbb{E}^{k-1} \mathbb{E}_k \exp \left\{ \theta \left| \mathbb{E}^{k-1}(\ell_k - \mathbb{E}_k \ell_k) \right| \right\} \leq \mathbb{E}^{k-1} \mathbb{E}_k \exp \left(\theta |\ell_k - \mathbb{E}_k \ell_k| \right) \\ &\leq \mathbb{E}^{k-1} \left[\left(\mathbb{E}_k \exp(\theta |\ell_k|) \right) \cdot \exp(\theta |\mathbb{E}_k \ell_k|) \right] \leq \mathbb{E}^{k-1} \left[\left(\mathbb{E}_k \exp(\theta |\ell_k|) \right)^2 \right]. \end{aligned}$$

We now proceed to bound $\mathbb{E}_k \exp(\theta |\ell_k|)$: for $\theta = 1/(2C_\delta) \geq 0$, Assumption 5.1 gives

$$\mathbb{E}_k \exp(\theta |\ell_k|) \leq \mathbb{E}_k \left[\left(\frac{Z_{k,\circ}}{Z} \right)^{1/(2C_\delta)}; Z_{k,\circ} \geq \exp(N\delta), \frac{Z}{Z_{k,\circ}} \leq \frac{1}{\exp(w_{\max})} \right] \leq \exp(C_\delta),$$

by the same calculation as in (5.4). It follows using Jensen's inequality that for A_δ a large enough constant (depending only on C_δ) we have

$$\mathbb{E}_k \exp \left(\frac{|\ell_k|}{A_\delta} \right) \leq \left(\mathbb{E}_k \exp \left(\frac{|\ell_k|}{2C_\delta} \right) \right)^{2C_\delta/A_\delta} \leq \exp \left(\frac{2(C_\delta)^2}{A_\delta} \right) \leq 2,$$

where the last inequality holds by choosing A_δ large enough (depending only on C_δ). It follows by the martingale Bernstein inequality (see [Tal11, Thm. A.6.1]) that

$$\text{(II)} \equiv \mathbb{P} \left(\left| \sum_{k \leq M} y_k \right| \geq Nt \right) \leq \exp \left(-\min \left\{ \frac{Nt^2}{4(A_\delta)^2}, \frac{Nt}{2A_\delta} \right\} \right). \quad (5.6)$$

The claim follows by combining (5.5) and (5.6). $\square$

Proposition 5.5. *For the model (1.4), under Assumption 1.6, for all $\delta > 0$ we have*

$$\mathbb{P}\left(\left|\log_{N\delta} Z - \mathbb{E} \log_{N\delta} Z\right| \geq 2Nt\right) \leq \frac{M \log 2}{t} f_\delta(w_{\max}) + \frac{MC_{\delta,2}}{(Nt)^2}.$$

In particular, if $f_\delta(w_{\max})$ tends to zero superpolynomially in N , then this is $o_N(1)$ provided $t \gg 1/N^{1/2}$.

Proof. Following the first few steps of the proof of Proposition 5.4 gives (instead of (5.5))

$$(I) \equiv \mathbb{P}\left(\sum_{a \leq M} |Y_a - y_a| \geq Nt\right) \leq \frac{1}{Nt} \mathbb{E}\left(\sum_{a \leq M} |Y_a - y_a|\right) \leq \frac{MN \log 2}{Nt} f_\delta(w_{\max}).$$

Recall that $\mathbb{E}_k$ denotes expectation over Θ_a only, while $\mathbb{E}^k$ denotes expectation over all $(\Theta_\ell)_{\ell \geq k}$. Then

$$\mathbb{E}_a[(y_a)^2] \leq \mathbb{E}_a[(\ell_a)^2] \leq \mathbb{E}_a\left[\left(\log \frac{Z}{Z_{a,\circ}}\right)^2; \frac{Z}{Z_{a,\circ}} \geq \frac{1}{\exp(w_{\max})}\right] \leq C_{\delta,2},$$

by Assumption 1.6. It follows using Chebychev's inequality that

$$(II) \equiv \mathbb{P}\left(\left|\sum_{a \leq M} y_a\right| \geq Nt\right) \leq \frac{MC_{\delta,2}}{(Nt)^2}.$$

The claim follows by combining the bounds on (I) and (II). $\square$

Proof of Theorem 1.7. Follows from Proposition 5.5. In the perceptron model, a sharper concentration result can be obtained using Proposition 5.4, where Assumption 5.1 is satisfied by Theorems 1.2–1.5. $\square$

5.3. Sharp threshold sequence and universality. In this subsection we conclude with the proofs of Theorems 1.9 and 1.11. We summarize the results of the preceding subsections with the following assumptions:

Assumption 5.6. For all $\delta > 0$ there exists $\rho_\delta > 0$ such that for all $0 \leq \rho \leq \rho_\delta$ we have, for all $M \geq 0$,

$$\mathbb{P}\left(Z_M \geq \exp(2N\delta), Z_{M+N\rho} < \exp(N\delta)\right) \leq o_N(1).$$

(Under Assumptions 5.1 or 1.6, this estimate is implied by Propositions 5.2 and 5.3.)

Assumption 5.7. For all $\delta > 0$ and any constant $t > 0$,

$$\mathbb{P}\left(\left|\log_{N\delta} Z - \mathbb{E} \log_{N\delta} Z\right| \geq Nt\right) \leq o_N(1).$$

(Under Assumptions 5.1 or 1.6, this estimate is implied by Propositions 5.4 and 5.5.)

Proposition 5.8. *For the model (1.4), suppose $\sup\{\mathbb{E}\Theta_a(x) : x \in \{-1, +1\}^N\} \leq \exp(-c)$ for a positive constant c . Assumption 5.7 implies a sharp threshold sequence: that is to say, there is a sequence α_N such that $\mathbb{P}(Z_{N\alpha} > 0)$ transitions from $1 - o_N(1)$ to $o_N(1)$ in an $o_N(1)$ window around α_N . Assumption 5.6 further implies $\alpha_N \asymp 1$ in the limit $N \rightarrow \infty$.*

Proof. Note that $Z_0 = 2^N$. The bound on $\mathbb{E}\Theta_a$ implies $\mathbb{E}Z_M \leq 2^N \exp(-Mc)$, so $\mathbb{P}(Z_{N\alpha} > 0) = o_N(1)$ as soon as $\alpha > (\log 2)/c$. Thus $\mathbb{P}(Z_{N\alpha} > 0)$ transitions from $1 - o_N(1)$ to $o_N(1)$ as α increases from 0 to $2(\log 2)/c$. Now suppose for the sake of contradiction that there is not a sharp threshold sequence: this means that there exists arbitrarily large N such that we have

$$\epsilon \leq \inf\left\{\mathbb{P}(Z_{N\alpha} > 0) : \alpha \in [\alpha_1, \alpha_2]\right\} \leq \sup\left\{\mathbb{P}(Z_{N\alpha} > 0) : \alpha \in [\alpha_1, \alpha_2]\right\} \leq 1 - \epsilon, \quad (5.7)$$

where the α_i can depend on N , but $\alpha_2 - \alpha_1$ stays bounded away from zero as $N \rightarrow \infty$. Then

$$\begin{aligned} \mathbb{P}\left(Z_{N\alpha_1} \leq \exp\left(\frac{N(\alpha_2 - \alpha_1)c}{2}\right)\right) &\leq \mathbb{P}(Z_{N\alpha_2} = 0) + \mathbb{P}\left(Z_{N\alpha_1} \leq \exp\left(\frac{N(\alpha_2 - \alpha_1)c}{2}\right), Z_{N\alpha_2} > 0\right) \\ &\leq \mathbb{P}(Z_{N\alpha_2} = 0) + \exp\left(-\frac{N(\alpha_2 - \alpha_1)c}{2}\right) \leq 1 - \epsilon + o_N(1). \end{aligned}$$

In the above, the second-to-last step is by the assumption on $\mathbb{E}\Theta$, and the last step is by (5.7). Rearranging the above inequality gives

$$\epsilon - o_N(1) \leq \mathbb{P}\left(Z_{N\alpha_1} > \exp\left(\frac{N(\alpha_2 - \alpha_1)c}{2}\right)\right) \leq \mathbb{P}(Z_{N\alpha_1} > 0) \leq 1 - \epsilon.$$

It follows that for any $\delta < (\alpha_2 - \alpha_1)c/4$, the quantity $\log_{N\delta} Z_{N\alpha_1}$ is not well-concentrated: with probability at least $\epsilon - o_N(1)$ it equals $N\delta$, but with probability at least $\epsilon - o_N(1)$ it exceeds $2N\delta$. This contradicts Assumption 5.7. This implies that $\mathbb{P}(Z_{N\alpha} > 0)$ transitions from $1 - o_N(1)$ to $o_N(1)$ in an $o_N(1)$ window around a sharp threshold sequence α_N . Finally, Assumption 5.6 implies that α_N stays bounded away from zero in the limit $N \rightarrow \infty$. $\square$

Proof of Theorem 1.9. Follows from Proposition 5.8, where Assumptions 5.6 and 5.7 are satisfied by Propositions 5.3 and 5.5. $\square$

Proof of Theorem 1.11. For the U -perceptron (1.2), let $Z(\xi)$ be the partition function with general (subgaussian) disorder. It follows from the proof of Theorem 1.9 that there is a sharp threshold sequence $\alpha_{N,\xi}$ (potentially depending on the disorder), and the free energy $\log Z_{N,N\alpha}(\xi)$ is exponentially large for $\alpha < \alpha_{N,\xi}$. Suppose the threshold sequence depends nontrivially on the disorder, meaning that on a subsequence $N \rightarrow \infty$ we have $\alpha_{N,\xi} - \alpha_{N,g} \geq \epsilon$ or $\alpha_{N,g} - \alpha_{N,\xi} \geq \epsilon$. Without loss suppose $\alpha_{N,\xi} - \alpha_{N,g} \geq \epsilon$. Then, for $\alpha_{N,g} < \alpha < \alpha_{N,\xi}$, $Z_{N,N\alpha}(g)$ will be exponentially large while $Z_{N,N\alpha}(\xi)$ is zero with high probability. By the concentration result Theorem 1.7, this will yield a contradiction to Theorem 4.1. It follows $Z(g)$ and $Z(\xi)$ share the same sharp threshold sequence α_N . Moreover, since the partition function must be exponentially large for $\alpha < \alpha_N$, it follows from Theorem 1.7 and Theorem 4.1 that $N^{-1} \log Z(\xi)$ and $N^{-1} \log Z(g)$ are $o_N(1)$ -close with high probability for all $\alpha < \alpha_N$. $\square$

REFERENCES

- [ALS21a] E. Abbe, S. Li, and A. Sly. Binary perceptron: efficient algorithms can find solutions in a rare well-connected cluster. arXiv:2111.03084, 2021.
- [ALS21b] E. Abbe, S. Li, and A. Sly. Proof of the contiguity conjecture and lognormal limit for the symmetric perceptron. arXiv:2102.13069, 2021.
- [ALS21c] R. Alweiss, Y. P. Liu, and M. Sawhney. Discrepancy minimization via a self-balancing walk. In *Proc. 53rd STOC*, pages 14–20. ACM, New York, 2021.
- [APZ19] B. Aubin, W. Perkins, and L. Zdeborová. Storage capacity in symmetric binary perceptrons. *J. Phys. A*, 52(29):294003, 32, 2019.
- [AS20] A. E. Alaoui and M. Sellke. Algorithmic pure states for the negative spherical perceptron. arXiv:2010.15811, 2020.
- [Ban10] N. Bansal. Constructive algorithms for discrepancy minimization. In *Proc. 51st FOCS*, pages 3–10. IEEE Computer Soc., Los Alamitos, CA, 2010.
- [BBC⁺16] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina. Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes. *Proc. Nat. Acad. Sci. U.S.A.*, 113(48):E7655–E7662, 2016.
- [BLM13] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013. A nonasymptotic theory of independence, With a foreword by Michel Ledoux.
- [BNSX21] E. Bolthausen, S. Nakajima, N. Sun, and C. Xu. Gardner formula for Ising perceptron models at small densities. arXiv:2111.02855, 2021.
- [Cov65] T. M. Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Trans. Electron.*, 3:326–334, 1965.
- [DS18] J. Ding and N. Sun. Capacity lower bound for the Ising perceptron. arXiv:1809.07742, 2018.
- [FP16] S. Franz and G. Parisi. The simplest model of jamming. *J. Phys. A*, 49(14):145001, 2016.
- [FPS⁺17] S. Franz, G. Parisi, M. Sevelev, P. Urbani, and F. Zamponi. Universality of the SAT-UNSAT (jamming) threshold in non-convex continuous constraint satisfaction problems. *SciPost Physics*, 2(3):019, 2017.
- [Gar88] E. Gardner. The space of interactions in neural network models. *J. Phys. A*, 21(1):257, 1988.
- [GD88] E. Gardner and B. Derrida. Optimal storage properties of neural network models. *J. Phys. A*, 21(1):271, 1988.
- [GKPX22] D. Gamarnik, E. C. Kızıldağ, W. Perkins, and C. Xu. Algorithms and barriers in the symmetric binary perceptron model. arXiv:2203.15667, 2022.
- [Hat12] H. Hatami. A structure theorem for Boolean functions with small total influences. *Ann. of Math. (2)*, 176(1):509–533, 2012.
- [KM89] W. Krauth and M. Mézard. Storage capacity of memory networks with binary couplings. *J. Physique*, 50(20):3057–3066, 1989.
- [KR98] J. H. Kim and J. R. Roche. Covering cubes by random half cubes, with applications to binary neural networks. *J. Comput. System Sci.*, 56(2):223–252, 1998. Eighth Annual Workshop on Computational Learning Theory (COLT) (Santa Cruz, CA, 1995).
- [Lat97] R. Latała. Sudakov minoration principle and supremum of some processes. *Geom. Funct. Anal.*, 7(5):936–953, 1997.

- [Led01] M. Ledoux. *The concentration of measure phenomenon*, volume 89 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2001.
- [LM15] S. Lovett and R. Meka. Constructive discrepancy minimization by walking on the edges. *SIAM J. Comput.*, 44(5):1573–1582, 2015.
- [Mé89] M. Mézard. The space of interactions in neural networks: Gardner’s computation with the cavity method. *J. Phys. A*, 22(12):2181, 1989.
- [Mos10] E. Mossel. Gaussian bounds for noise correlation of functions. *Geom. Funct. Anal.*, 19(6):1713–1756, 2010.
- [MRSY19] A. Montanari, F. Ruan, Y. Sohn, and J. Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. arXiv:1911.01544, 2019.
- [MZZ21] A. Montanari, Y. Zhong, and K. Zhou. Tractability from overparametrization: The example of the negative perceptron. arXiv:2110.15824, 2021.
- [PX21] W. Perkins and C. Xu. Frozen 1-RSB structure of the symmetric Ising perceptron. In *Proc. 53rd STOC*, pages 1579–1588, 2021.
- [Spe85] J. Spencer. Six standard deviations suffice. *Trans. Amer. Math. Soc.*, 289(2):679–706, 1985.
- [ST03] M. Shcherbina and B. Tirozzi. Rigorous solution of the Gardner problem. *Comm. Math. Phys.*, 234(3):383–422, 2003.
- [Sto13] M. Stojnic. Another look at the Gardner problem. arXiv:1306.3979, 2013.
- [Tal87] M. Talagrand. Regularity of Gaussian processes. *Acta Math.*, 159(1-2):99–149, 1987.
- [Tal93] M. Talagrand. Regularity of infinitely divisible processes. *Ann. Probab.*, 21(1):362–432, 1993.
- [Tal94] M. Talagrand. The supremum of some canonical processes. *Amer. J. Math.*, 116(2):283–325, 1994.
- [Tal99a] M. Talagrand. Intersecting random half cubes. *Random Structures Algorithms*, 15(3-4):436–449, 1999. Statistical physics methods in discrete probability, combinatorics, and theoretical computer science (Princeton, NJ, 1997).
- [Tal99b] M. Talagrand. Self-averaging and the space of interactions in neural networks. *Random Structures Algorithms*, 14(3):199–213, 1999.
- [Tal00] M. Talagrand. Intersecting random half-spaces: toward the Gardner-Derrida formula. *Ann. Probab.*, 28(2):725–758, 2000.
- [Tal05] M. Talagrand. *The generic chaining*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2005. Upper and lower bounds of stochastic processes.
- [Tal11] M. Talagrand. *Mean field models for spin glasses. Volume II*, volume 55 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics*. Springer, Heidelberg, 2011. Advanced replica-symmetry and low temperature.
- [TMR20] P. Turner, R. Meka, and P. Rigollet. Balancing gaussian vectors in high dimension. In *Proc. 33rd COLT*, pages 3455–3486. PMLR, 2020.
- [Ver18] R. Vershynin. *High-dimensional probability*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2018. An introduction with applications in data science, With a foreword by Sara van de Geer.
- [Wen62] J. G. Wendel. A problem in geometric probability. *Math. Scand.*, 11:109–111, 1962.
- [Xu21] C. Xu. Sharp threshold for the Ising perceptron model. *Ann. Probab.*, 49(5):2399–2415, 2021.
- [Xu22] C. Xu. Personal communication, 2022.