

Frozen 1-RSB Structure of the Symmetric Ising Perceptron^{*}

Will Perkins

University of Illinois at Chicago
Chicago, Illinois, USA
math@willperkins.org

Changji Xu

Harvard University
Cambridge, Massachusetts, USA
cxu@cmsa.fas.harvard.edu

ABSTRACT

We prove, under an assumption on the critical points of a real-valued function, that the symmetric Ising perceptron exhibits the ‘frozen 1-RSB’ structure conjectured by Krauth and Mézard in the physics literature; that is, typical solutions of the model lie in clusters of vanishing entropy density. Moreover, we prove this in a very strong form conjectured by Huang, Wong, and Kabashima: a typical solution of the model is isolated with high probability and the Hamming distance to all other solutions is linear in the dimension. The frozen 1-RSB scenario is part of a recent and intriguing explanation of the performance of learning algorithms by Baldassi, Ingrosso, Lucibello, Saglietti, and Zecchina. We prove this structural result by comparing the symmetric Ising perceptron model to a planted model and proving a comparison result between the two models. Our main technical tool towards this comparison is an inductive argument for the concentration of the logarithm of number of solutions in the model.

CCS CONCEPTS

• **Mathematics of computing** → **Probability and statistics**; • **Theory of computation** → **Randomness, geometry and discrete structures**.

KEYWORDS

Perceptron model, frozen solutions, solution space, Ising perceptron

ACM Reference Format:

Will Perkins and Changji Xu. 2021. Frozen 1-RSB Structure of the Symmetric Ising Perceptron. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC '21)*, June 21–25, 2021, Virtual, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3406325.3451119>

1 INTRODUCTION

The *perceptron model* is a simple model of a neural network storing random patterns. It has been studied in several fields including information theory [19], statistical physics [23, 24, 29, 37], and probability theory [28, 38, 39].

There are several variants of the model, grouped into two main categories: spherical perceptrons in which patterns are N -dimensional

vectors on the unit sphere and Ising perceptrons in which patterns are ± 1 vectors of length N . In each case, we want to understand how many random patterns can be ‘stored’ by a neural network formed by taking random synapses J_{ij} and applying a given activation function. Here we will study the Ising perceptron.

Let $\Sigma_N = \{\pm 1\}^N$ and let $\{X_i\}_{i \geq 1}$ be a sequence of independent N -dimensional standard Gaussian vectors¹. For a real valued function ϕ , a real number κ , and $X \in \mathbb{R}^N$, define

$$H_{\phi, \kappa}(X) = \left\{ \sigma \in \Sigma_N : \phi \left(\langle X, \sigma \rangle / \sqrt{N} \right) \leq \kappa \right\}.$$

The *solution space* of the Ising perceptron with Gaussian disorder, activation function ϕ , threshold κ , and m constraints is the random subset of Σ_N ,

$$S = S_{\phi, \kappa, N, m} = \bigcap_{i=1}^m H_{\phi, \kappa}(X_i).$$

Thus S is a random subset of the Hamming cube Σ_N . We call the vectors X_i *constraint vectors*. The constraints depend on the set of constraint vectors, the activation function ϕ , and the threshold κ . The classic Ising perceptron corresponds to the choice $\phi(x) = x$ where the most studied case is $\kappa = 0$ (e.g. [21, 29]).

1.1 Structure of the Solution Space

We will be concerned with the typical structure of the solution space S as a function of ϕ , κ , and the constraint density $\alpha := m/N$, as $N \rightarrow \infty$.

The most basic structural question is whether S is empty or not. The *capacity* of the perceptron is defined as the random variable

$$M_{\phi, \kappa}(N) = \max\{m : S_{\phi, \kappa, N, m} \neq \emptyset\},$$

and the *critical capacity density*

$$\alpha_{c, \phi}(\kappa) = \inf \left\{ \alpha : \liminf_{N \rightarrow \infty} \mathbb{P}(S_{\phi, \kappa, N, \lfloor \alpha N \rfloor} = \emptyset) = 1 \right\}$$

is the typical constraint density of the capacity.

For densities below the critical capacity density, when the solution space is typically non-empty, we can ask about its structure, and how this structure affects the performance of learning algorithms: algorithms that find some solution in S given the instance defined by $\mathbf{X} = (X_1, \dots, X_m)$.

Basic structural questions include whether solutions appear in connected clusters or are isolated; and what the typical distance is from a solution to the next nearest solution. For these structural

^{*}WP supported in part by NSF DMS-1847451 and CCF-1934915.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

STOC '21, June 21–25, 2021, Virtual, Italy

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8053-9/21/06...\$15.00

<https://doi.org/10.1145/3406325.3451119>

¹In another variant of the model the constraint vectors X_i are given by independent samples from Σ_N (Bernoulli disorder). While there are significant differences between the spherical and Ising perceptrons, the choice of Gaussian or Bernoulli disorder is insignificant for the properties discussed here.

properties we regard Σ_N as the Hamming cube endowed with Hamming distance: for $\sigma, \sigma' \in \Sigma_N$,

$$\text{dist}(\sigma, \sigma') = |\{i : \sigma_i \neq \sigma'_i\}| = \frac{n - \langle \sigma, \sigma' \rangle}{2}.$$

The perspective taken in recent work on Ising perceptrons in both statistical physics and computer science is to view the Ising perceptron as a random constraint satisfaction problem (CSP): Σ_N is the set of possible solutions, and each random vector X_i defines a constraint $\phi(\langle X, \sigma \rangle / \sqrt{N}) \leq \kappa$ on possible solutions $\sigma \in \Sigma_N$. The critical capacity density $\alpha_{c, \phi}$ is then the satisfiability threshold of the model.

Just as in the random k -SAT, random k -NAE-SAT, or random k -XOR-SAT problems, each constraint rules out a constant fraction of all solutions in Σ_N . Where the perceptron differs from these other models is that in the perceptron each constraint involves all of the N coordinates, while in the other models each constraint only involves k , which is held constant as $N \rightarrow \infty$.

Random CSP's have been studied extensively in computer science, statistical physics, probability, and combinatorics since Mitchell, Selman, and Levesque [32] observed empirically that random k -SAT formulae at certain densities proved extremely challenging for widely used SAT solvers. Understanding this phenomenon is a major ongoing challenge that has led to the development of a new field of inquiry at the intersection of computer science, physics, and mathematics.

A key to the current understanding of random CSP's is understanding the typical structure of the solution space at different constraint densities. A beautifully detailed but non-rigorous picture was put forth by Krzakala, Montanari, Ricci-Tersenghi, Semerjian, and Zdeborová [30], based on the cavity method from statistical physics.

In a paper that transformed the computer science perspective on random computational problems, Achlioptas and Coja-Oghlan [2] proved the existence of a ‘clustering threshold’ at which the solution space of certain random CSP's breaks apart into exponentially many clusters separated by linear Hamming distance. They also observed that this clustering threshold coincides asymptotically with the threshold at which known efficient search algorithms for these problems fail.

A further structural property is *freezing*. Given a solution σ , a variable (or coordinate) is free if flipping that coordinate results in another solution σ' . Variables that are not free are frozen. The freezing threshold for a random CSP is the threshold at which typical solutions have a linear number of frozen variables [33, 42].

These two properties – clustering and freezing – are conjectured to be the source of computational hardness in random CSP's. Identifying thresholds for the onset of these and other structural phenomena and rigorously connecting them to the performance of algorithms have been a main focus of the field of random computational problems in the last decade.

For the Ising perceptron, the conjectured structural picture looks strikingly different. Krauth and Mezard [29] conjectured, by means of the replica method, that at all densities below the critical capacity, the solution space of the perceptron is dominated by clusters of vanishing entropy density (that is, of size $e^{o(n)}$), each cluster separated from the others by linear Hamming distance. Wong, Kabashima, and

Huang [27] and Huang and Kabashima [26] refined these conjectures and posited that in fact typical solution in the Ising perceptron are *completely frozen*; that is, all coordinates are frozen and the solutions lie in clusters of size 1, separated from all other solutions by linear Hamming distance. Thus the Ising perceptron exhibits clustering and freezing in the strongest possible form throughout the entire satisfiability regime. Based on the current conjectural understanding of random CSP's (see e.g. [44]), one might venture a guess that finding a solution in the Ising perceptron is computationally hard at all positive densities.

However, this theory is at odds with other work in physics on learning algorithms for the perceptron. Braunstein and Zecchina [15] observed empirically that a simple message-passing algorithm is able to find solutions at positive densities in the Ising perceptron (further algorithms followed in [6, 8]). Attempting to reconcile this apparent contradiction between the theory and empirical observations, Baldassi, Ingrosso, Lucibello, Saglietti, and Zecchina [10] conjectured that these successful learning algorithms were in fact finding solutions belonging to rare clusters of positive entropy density. That is, although a $1 - o(1)$ fraction of solutions belong to isolated, frozen clusters, an exponentially small fraction of solutions belong to clusters that are exponentially large; strikingly, the authors observed that learning algorithms find solutions in these rare clusters. Specifically, the solutions that contribute to the dominant portion of the partition function (number of solutions) and determine the equilibrium properties of the model are completely distinct from those solutions that efficient algorithms find. This work was followed by the proposal of several different algorithms to target these subdominant clusters [7, 11].

In [5], a *symmetric Ising perceptron*, with activation function $\phi(x) = |x|$, was studied as a model conjectured to exhibit the same structural and algorithmic properties, but more amenable to mathematical analysis. Baldassi, Della Vecchia, Lucibello, and Zecchina [9] confirmed that on the level of the physics predictions this model has the same qualitative behavior as the Ising perceptron with activation function $\phi(x) = x$.

In summary, the Ising perceptron, and its symmetric variants, are conjectured to exhibit ‘frozen 1-RSB behavior’ at all positive densities below the critical capacity density. The current understanding of the link between clustering, freezing, and the performance of algorithms would suggest that finding a solution in these models is therefore intractable. This, however, is seemingly in contradiction with empirical observations, and one hypothesis suggests that learning algorithms find exponentially rare solutions with atypical structural properties.

Resolving these questions is a pressing problem since the hypothesis about subdominant clusters calls into question the link between the equilibrium properties of these models and algorithmic tractability. In this work, we take a first step in addressing this problem rigorously by establishing the frozen 1-RSB picture for the symmetric Ising perceptron (Theorem 1 below), under an assumption on the critical points of a real-valued function (Assumption 1).

1.2 Previous Results

There are few rigorous results on the Ising perceptron, and most are concerned with bounds on the critical capacity.

For the classic Ising perceptron Krauth and Mézard [29] predicted, using the replica method, that $\alpha_c(\kappa) = \alpha_{\text{KM}}(\kappa)$ for a complicated but explicit function α_{KM} (with $\alpha_{\text{KM}}(0) \approx .83$). Following some previous bounds of Kim and Roche and Talagrand [28, 38], Ding and Sun [21] recently proved that $\alpha_c(\kappa) \geq \alpha_{\text{KM}}(\kappa)$ using a sophisticated form of the second-moment method guided by the Thouless–Anderson–Palmer (TAP) equations [40]. Their result assumes a technical condition on a certain real-valued function, akin to Assumption 1 below.

Much of the technical difficulty of [21] comes from the asymmetry inherent in the activation function $\phi(x) = x$; this necessitates a conditioning argument and the sophisticated second-moment calculation. On the other hand, Aubin, Perkins, and Zdeborová [5] considered two symmetric activation functions: $\phi_r(x) = |x|$ and $\phi_u(x) = -|x|$, which they called the rectangular and ‘u’ activation functions respectively. Studying symmetric constraints has a long history in the random CSP literature: the random k -NAE-SAT model is a symmetric variant of the random k -SAT model. While the qualitative properties of the two models are expected to be very similar, the symmetric model is often more amenable to rigorous analysis, and thus a clearer understanding can be obtained (see e.g., [14, 20, 36] for recent work on the k -NAE-SAT model). Studying symmetric perceptrons allows us to prove stronger and more detailed results than are currently attainable for the classic perceptron, but the phenomena studied are expected to be universal.

Aubin, Perkins, and Zdeborová [5] determine the critical capacity density for the symmetric perceptron with rectangular activation function:

$$\alpha_{r,c}(\kappa) = -\frac{\log 2}{\log p(\kappa)}$$

where $p(\kappa) = \mathbb{P}(|Z| \leq \kappa)$ for a standard Gaussian random variable Z . This result, like that of [21], is contingent on an assumption about a certain real-valued function. This function will also prove useful in our work. Let $H(\beta) = -\beta \log \beta - (1 - \beta) \log(1 - \beta)$ be the Shannon entropy function (all logarithms in this paper are base e) and let

$$q_\kappa(\beta) = \mathbb{P}(|Z_1| \leq \kappa, |Z_2| \leq \kappa) \quad (1)$$

where (Z_1, Z_2) is a jointly Gaussian vector with means 0, variances 1, and covariance $2\beta - 1$.

ASSUMPTION 1 ([5]). *The function*

$$F_\alpha(\beta) = H(\beta) + \alpha \log q_\kappa(\beta),$$

has a single critical point for $\beta \in (0, 1/2)$ whenever $F''_\alpha(1/2) < 0$.

REMARK 1. *We (and the authors of [5]) have plotted $F_\alpha(\beta)$ for many choices of α and κ and in all instances the function has the shape depicted in Figure 1, consistent with Assumption 1. We believe a proof of the assumption might be possible adapting the methods of [4, Proof of Lemma 3].*

Xu [41] proved a general sharp threshold result (an analogue of Friedgut’s sharp threshold result for random graphs and CSP’s [22]), which, combined with [5], gives a sharp threshold for the existence of solutions: for any $\varepsilon > 0$,

$$\begin{aligned} \mathbb{P}(S_{\phi_r, \kappa, N(\alpha_{r,c} + \varepsilon)N} = \emptyset) &\rightarrow 1 \text{ as } N \rightarrow \infty, \\ \mathbb{P}(S_{\phi_r, \kappa, N(\alpha_{r,c} - \varepsilon)N} = \emptyset) &\rightarrow 0 \text{ as } N \rightarrow \infty. \end{aligned}$$

Both statements hold also for the u -function in a range of κ values, and the results of [21, 41] prove that the second statement holds for the classic perceptron. Proving the matching upper bound on the critical capacity for the classic perceptron remains a challenging open problem.

Baldassi, Della Vecchia, Lucibello, and Zecchina [9] used the second-moment method to show the existence of pairs of solutions at arbitrary distances in the symmetric Ising perceptron.

1.3 Main Results

We will study properties of typical solutions in the symmetric Ising perceptron (with the rectangular activation function $\phi(x) = |x|$). We now specialize and simplify the notation from above.

For $X \in \mathbb{R}^N$ and $\kappa > 0$, define

$$H_\kappa(X) := \{\sigma \in \Sigma_N : |\langle X, \sigma \rangle| \leq \kappa \sqrt{N}\}.$$

Let $\{X_i\}_{i \geq 1}$ be a sequence of i.i.d. N -dimensional standard Gaussians, and define the solution space

$$S = S_\alpha(N) = \bigcap_{i=1}^{\lfloor \alpha N \rfloor} H_\kappa(X_i). \quad (2)$$

The critical capacity density, determined in [5], is $\alpha_c = \alpha_c(\kappa) = -\log 2 / \log p(\kappa)$.

The main result of this paper confirms the frozen 1-RSB scenario in the symmetric Ising perceptron: typical solutions are completely frozen with high probability for $\alpha < \alpha_c$. Let

$$\beta_c = \beta_c(\kappa, \alpha) = \beta \in (0, 1/2) : F(\beta) - \alpha \log p(\kappa) = 0. \quad (3)$$

See Figure 1 for a depiction of β_c . We show in Lemma 5 that for $\alpha < \alpha_c$, $\beta_c > 0$ exists and is uniquely defined.




Figure 1: $F_\alpha(\beta) - \alpha p(\kappa)$ plotted for $\alpha = \{1.69, 1.75, 1.81\}$ and $\kappa = 1$. The dots mark β_c for the three values of α .

THEOREM 1. *Let $\kappa > 0$ and $\alpha < \alpha_c(\kappa)$. Let σ be uniformly sampled from S conditioned on the event $S \neq \emptyset$. Under Assumption 1, for any $\delta \in (0, \beta_c)$,*

$$\{\sigma' \in S : \text{dist}(\sigma, \sigma') \leq (\beta_c - \delta)N\} = \{\sigma\}.$$

with probability $1 - o(1)$ as $N \rightarrow \infty$. In particular, σ is completely frozen with probability $1 - o(1)$.

Note that σ is selected according to two sources of randomness: the randomness of the perceptron instance $\mathbf{X}$ and the random choice of σ from S .

The next result shows that the logarithm of the number of number of solutions in the rectangular Ising perceptron is tightly concentrated below the critical density.

THEOREM 2. *Under Assumption 1, for $\alpha < \alpha_c$*

$$\frac{\log |S|}{N} = \log 2 + \alpha \log p(\kappa) + O_{\mathbb{P}}\left(\frac{\log N}{N}\right) \quad \text{as } N \rightarrow \infty \text{ in probability.}$$

In particular, for $\alpha < \alpha_c$, S is non-empty with probability $1 - o(1)$.

The second statement of the Theorem 2 proves that the symmetric Ising perceptron undergoes a sharp satisfiability phase transition at α_c , answering an open question from [5] (where the complementary statement that for $\alpha > \alpha_c$, $S = \emptyset$ with high probability is proved). This could also be proved by adapting the sharp threshold result of [41] to the symmetric perceptron.

1.4 Overview of the Techniques

We study the properties of a typical solution drawn from S by way of the *planted model*: the experiment of first selecting a uniformly random solution from Σ_N , then choosing a random configuration of constraints consistent with this solution. Planted models have been studied extensively in the random CSP literature and beyond. They are used as a toy model for statistical inference: e.g. the ‘teacher–student model’ [43] or the stochastic block model [1]. They are used to understand the condensation threshold in random CSP’s [13, 16–18]. They are used to understand the structure of the solution space in random CSP’s [2, 3, 33, 34].

Given $N, m \in \mathbb{N}, \kappa > 0$, and following [2], we define two probability distributions on pairs $(\sigma^*, \mathbf{X}) \in \Sigma_N \times (\mathbb{R}^N)^m$ of solutions and configurations of m constraint vectors.

In the **random model** we:

- (1) Sample m i.i.d. N -dimensional standard Gaussian constraint vectors $\mathbf{X} = (X_1, \dots, X_m)$, conditioned on the event that $S(\mathbf{X}) = \bigcap_{i=1}^m H_{\kappa}(X_i) \neq \emptyset$.
- (2) Sample σ^* uniformly at random from S .

We denote the law of the random model with $\mathbb{P}_r, \mathbb{E}_r$ to distinguish the law from both the unconditional perceptron model and the planted model below. The random model is simply the experiment of selecting a uniformly random solution from the symmetric Ising perceptron conditioned on satisfiability.

In the **planted model** we:

- (1) Sample σ^* uniformly at random from Σ_N .
- (2) Sample a configuration of m i.i.d. constraint vectors $\mathbf{X} = (X_1, \dots, X_m)$, with each X_i distributed as a standard N -dimensional Gaussian vector conditioned on the event that $\sigma^* \in H_{\kappa}(X_i)$.

We denote the law of the planted model with $\mathbb{P}_{\text{pl}}, \mathbb{E}_{\text{pl}}$.

The key to using the planted model to understand the original model is to show that at low enough constraint densities, the two distributions on $(\sigma^*, \mathbf{X})$ are close. Proving that the distributions are close, as we do below in Lemma 19, amounts to proving that the number of solutions, $|S|$, is typically not too far from its expectation,

$\mathbb{E}|S|$. The better concentration of $|S|$ one can prove, the more one can deduce about the original model from the planted model. In [2] it is shown (in the case of q -colorings of a random graph) that if $\log |S| = \log \mathbb{E}|S| + o(N)$, then events that occur with probability at most $\exp(-\theta(N))$ in the planted model occur with probability $o(1)$ as $N \rightarrow \infty$ in the random model. This notion of closeness is ‘quiet planting’ [31] and it suffices to prove some structural results on the solution spaces such as clustering [2]. On the other hand, much stronger notions of closeness have been proved: ‘silent planting’ [12] which implies the two distributions are mutually contiguous: any event with probability $o(1)$ in the planted model has probability $o(1)$ in the random model. This has been used to prove stronger structural results [33]. Proving contiguity requires much stronger concentration of $\log |S|$. A very general result on the contiguity of the planted and random model for symmetric CSP’s [16] involves a rigorous implementation of the cavity method and the small subgraph conditioning method. In the setting of the perceptron, neither of these tools exist and so we must prove concentration via another route.

We prove Theorem 1 in three steps.

In Section 2, we prove that the planted solution is isolated and the next nearest solution is at linear Hamming distance with high probability in the planted model (Lemma 6).

In Section 3 we prove Theorem 2, showing that for $\alpha < \alpha_c$, the logarithm of the number of solutions in the random model is concentrated around the logarithm of the expected number of solutions. To the best of our knowledge, our approach to proving concentration in this way is new, and we expect the technique to have further applications.

In Section 4 we transfer our results about the planted model to the random model by showing that events that occur with probability at most $N^{-\omega(1)}$ in the planted model occur with probability $o(1)$ in the standard model (Lemma 19). This relies on the concentration properties of the logarithm of the number of solutions.

1.5 Extensions and Future Work

Both Theorem 1 and Theorem 2 can be extended verbatim to the u -function Ising perceptron studied in [5], for $\kappa \in (0, .817)$ (the same range of κ for which the second-moment method works there).

The main open problem in this area is to resolve the conceptual dilemma described in Section 1.1 and answer the questions raised in [10]. Are there efficient learning algorithms that always find out-of-equilibrium solutions in subdominant clusters? This would raise a serious challenge to the belief that an understanding of the associated equilibrium statistical mechanics model (on the level of the free energy) can explain computational tractability or intractability.

Concretely, now that we have verified the frozen 1-RSB scenario, we can ask for provably efficient learning algorithms for the symmetric Ising perceptron.

QUESTION 3. *Is there a polynomial-time algorithm that, with probability $1 - o(1)$, finds a solution to the symmetric Ising perceptron for some density $\alpha \in (0, \alpha_c)$?*

Alternatively, one could leverage the structure results we have proved here and rule out some class of learning algorithms.

We leave the following additional open problems for future work.

1. Prove that the classic perceptron with activation function $\phi(x) = x$ exhibits the frozen 1-RSB property. It is not at all clear how to extend the method of this paper to this case. As discussed in [16, Sec. 2.4], for asymmetric random CSP's, like the random k -SAT model, the planted model, at least in its straightforward implementation, is not useful to compare to the random model (in particular it is not contiguous with the random model at any positive density). We expect the same with the classic perceptron, and so our strategy of arguing via the planted model will not work.

2. Prove full contiguity between the random and planted models for $\alpha < \alpha_c$. Our comparison result (Lemma 19) suffices for our purposes here, but it is natural to ask for more (as is the case for a large class of symmetric random CSP's [16]).

CONJECTURE 4. *For $\alpha < \alpha_c$, the random and planted models of the symmetric Ising perceptron are mutually contiguous. That is, if $\mathbb{P}_{\text{pl}}(A) = o(1)$ then $\mathbb{P}_{\text{r}}(A) = o(1)$ and vice versa.*

2 THE PLANTED MODEL

Consider the planted model with planted solution σ^* and constraints $\mathbf{X} = (X_1, X_2, \dots, X_m)$. Define $S(\mathbf{X}) = \bigcap_{i=1}^{\lfloor \frac{\alpha N}{2} \rfloor} H_{\kappa}(X_i)$ as in (2). Recall the definition of β_c from (3). We show that β_c exists and is unique for $\alpha \in (0, \alpha_c)$.

LEMMA 5. *Under Assumption 1, for $\alpha \in (0, \alpha_c)$, there exists a unique $\beta \in (0, 1/2)$ so that*

$$F_{\alpha}(\beta) - \alpha \log p(\kappa) = 0.$$

Also, for any $\delta \in (0, \beta_c/2)$,

$$\sup_{\delta < \beta < \beta_c - \delta} F_{\alpha}(\beta) - \alpha \log p(\kappa) < 0. \quad (4)$$

PROOF. Let $G(\beta) = F_{\alpha}(\beta) - \alpha \log p(\kappa)$. Then since $q_{\kappa}(0) = p(\kappa)$ and $H(0) = 0$, we have $G(0) = 0$. As observed in [5], $G'(0) = -\infty$ and so for some $\varepsilon > 0$, $G(x) < 0$ for $x \in (0, \varepsilon)$. Moreover, since $\alpha < \alpha_c$, $G(1/2) > 0$, and so by continuity there exists $\beta \in (0, 1/2)$ with $G(\beta) = 0$. By Assumption 1 this β is unique. In addition, Assumption 1 implies that $G(\beta)$ is first strictly decreasing and then strictly increasing on $\beta \in (0, 1/2)$. This gives (4). $\square$

The following result says the planted solution is completely frozen with high probability in the planted model.

LEMMA 6. *Under Assumption 1, for any $\alpha \in (0, \alpha_c)$ and any $\delta \in (0, \beta_c)$, there exists a constant $c_{\delta} > 0$ such that*

$$\mathbb{P}_{\text{pl}}\left(\{\sigma \in S : \text{dist}(\sigma, \sigma^*) \leq (\beta_c - \delta)N\} \neq \{\sigma^*\}\right) \leq \exp\left\{-c_{\delta}\sqrt{N}\right\}.$$

Before proving Lemma 6 we introduce some notation. Let $q(m) = \mathbb{P}(\sigma, \sigma' \in H_{\kappa}(X))$, where $\sigma, \sigma' \in \Sigma_N$ are two arbitrary vectors with $\langle \sigma, \sigma' \rangle = m$ and X is a standard N -dimensional Gaussian vector. Then

$$q(m) = q_{\kappa}\left(\frac{1}{2} + \frac{m}{2N}\right), \quad (5)$$

where q_{κ} is defined in (1).

Lemma 6 will follow from the following results.

LEMMA 7. *There exists a constant $\epsilon > 0$ sufficiently small such that for all $m \leq \epsilon N$,*

$$\log \mathbb{E}_{\text{pl}}[|\{\sigma \in S : \langle \sigma, \sigma^* \rangle = N - m\}|] \leq -c\sqrt{mN}.$$

PROOF. Let σ_0 be a vector in Σ_N . Note that

$$\begin{aligned} \mathbb{E}_{\text{pl}}[|\{\sigma \in S : \langle \sigma, \sigma^* \rangle = N - m\}|] &= \sum_{\sigma : \langle \sigma, \sigma_0 \rangle = N - m} \mathbb{P}(\sigma \in S \mid \sigma_0 \in S) \\ &= \binom{N}{\frac{m}{2}} \left(\frac{q(m)}{p(\kappa)} \right)^{\alpha N}. \end{aligned} \quad (6)$$

We claim that uniformly over all $|m| \leq \epsilon N$ with sufficiently small ϵ ,

$$q(N - m) \leq p(\kappa) - c\sqrt{m/N} \quad (7)$$

Provided with (7), we have

$$(6) \leq \exp\left\{m \log(2N/m)/2 - c\sqrt{mN}\right\} \leq \exp\left\{-c\sqrt{mN}\right\}.$$

Hence Lemma 7 follows. Now it remains to prove (7). To this end, let $\mathbf{1}$ be the all 1's vector of length N and $\mathbf{1}_m$ be an N -dimensional vector with the first $(N - m)$ coordinates +1 and the remaining m coordinates -1. We write

$$\begin{aligned} p(\kappa) - q(N - m) &= \mathbb{P}(|\langle \mathbf{1}, X \rangle| \leq \kappa\sqrt{N}, |\langle \mathbf{1}_m, X \rangle| > \kappa\sqrt{N}) \\ &\geq \mathbb{P}\left(\sum_{i=1}^{N-\frac{m}{2}} X_i \in (\kappa\sqrt{N}, \kappa\sqrt{N} + \sqrt{m}), \sum_{i=N-\frac{m}{2}+1}^N X_i \in (-3\sqrt{m}, -2\sqrt{m})\right) \\ &= \mathbb{P}\left(Z_1 \in \left(\kappa\sqrt{\frac{N}{N-\frac{m}{2}}}, \kappa\sqrt{\frac{N}{N-\frac{m}{2}}} + \sqrt{\frac{m}{N-\frac{m}{2}}}\right), Z_2 \in (-3\sqrt{2}, -2\sqrt{2})\right) \\ &\geq c\sqrt{m/N} \end{aligned}$$

where Z_1, Z_2 are two independent standard Gaussian random variables. This proves (7). $\square$

LEMMA 8. *For any $\epsilon > 0$, uniformly in $|m| \leq N - \epsilon N$,*

$$\begin{aligned} \lim_{N \rightarrow \infty} \left| \frac{1}{N} \log \mathbb{E}_{\text{pl}}[|\{\sigma \in S : \langle \sigma, \sigma^* \rangle = m\}|] - \left[F\left(\frac{1}{2} + \frac{m}{2N}\right) - \alpha \log p(\kappa) \right] \right| \\ = 0. \end{aligned} \quad (8)$$

PROOF. Note that

$$\mathbb{E}_{\text{pl}}[|\{\sigma \in S : \langle \sigma, \sigma^* \rangle = m\}|] = \binom{N}{\frac{N+m}{2}} \cdot \left(\frac{q(m)}{p(\kappa)} \right)^{\alpha N} \quad (9)$$

and that for any $\epsilon > 0$, uniformly in $\epsilon N \leq k \leq N - \epsilon N$,

$$\frac{\frac{1}{N} \log \binom{N}{k}}{H(k/N)} \rightarrow 1 \quad \text{as } N \rightarrow \infty.$$

Combining these with (5) yields (8). $\square$

Proof of Lemma 6. Let $\delta \in (0, \beta_c)$. It suffices to show that

$$\mathbb{E}_{\text{pl}}[|\{\sigma \in S : 0 < \text{dist}(\sigma, \sigma^*) \leq (\beta_c - \delta)N\}|] \leq \exp\left\{-c_{\delta}\sqrt{N}\right\}.$$

This bound follows from Lemmas 7 and 8 and (4). $\square$

3 CONCENTRATION OF THE NUMBER OF SOLUTIONS

Fixing N , consider the symmetric perceptron as a discrete-time stochastic process with one constraint vector added at each time step. The solution space at time $t \in \mathbb{N}$ is defined as

$$S_t := \bigcap_{i=1}^t H_{\kappa}(X_i)$$

which is the intersection of t random rectangles.

The following strengthening of Theorem 2 is the main result of this section.

THEOREM 9. *Under Assumption 1, for every $\epsilon > 0$ there exists $M = M(\epsilon)$ such that for any $\alpha < \alpha_c$,*

$$\limsup_{N \rightarrow \infty} \sup_{0 \leq t \leq \alpha N} \mathbb{P} \left(\left| \log \left(\frac{|S_t|}{\mathbb{E}[|S_t|]} \right) \right| \geq M \log N \right) \leq \epsilon. \quad (10)$$

Theorem 2 follows immediately since $\frac{1}{N} \log \mathbb{E}[|S_t|] = \log 2 + \frac{t}{n} \log p(\kappa)$.

Theorem 9 says that the cardinality of the solution space will only deviate from its expectation slightly after adding αN random constraints. To prove this theorem, we will look at the change in this deviation at each time when a new constraint is added. Write

$$Q_t := \log \left(\frac{|S_t|}{\mathbb{E}[|S_t|]} \right) = \sum_{i=1}^t \log \left(\frac{|S_i|/|S_{i-1}|}{\mathbb{E}[|S_i|]/\mathbb{E}[|S_{i-1}|]} \right).$$

Note that $\mathbb{E}[|S_i|]/\mathbb{E}[|S_{i-1}|] = p(\kappa)$. Let

$$Y_t := \frac{1}{p(\kappa)} \left(\frac{|S_t|}{|S_{t-1}|} - p(\kappa) \right),$$

so that

$$Q_t = \sum_{i=1}^t \log(1 + Y_i). \quad (11)$$

Since $0 \leq |S_t| \leq |S_{t-1}|$, we have that $-1 \leq Y_t \leq (1 - p(\kappa))/p(\kappa)$; however, we expect that the Y_t 's are very close to zero with high probability. Hence by a Taylor expansion, as $N \rightarrow \infty$,

$$Q_t = \sum_{i=1}^t Y_i - \frac{Y_i^2}{2} + o(1).$$

In fact, we will prove in Lemma 11 that Y_t is roughly of order $N^{-1/2}$ provided that S_{t-1} is “regular” (see Definition 10). Thus if the S_t 's are all regular, we expect the second term $\sum_{i=1}^{\alpha N} Y_i^2$ to be of order $O(1)$. In addition, notice that $(\sum_{i=0}^t Y_i)_{t \geq 0}$ is a martingale with respect to the filtration

$$\mathcal{F}_t = \sigma(S_1, S_2, \dots, S_t), \quad t \geq 1.$$

Hence if the S_t 's are all regular, we also expect the first term $\sum_{i=1}^{\alpha N} Y_i$ to be of order $O(1)$, and hence $Q_t = O(1)$.

DEFINITION 10. *For each $t \geq 0$, we let $(\sigma_i^{(t)})_{i \geq 1}$ be independent uniform random samples in S_t and denote $\mathbb{P}_t(\cdot) := \mathbb{P}(\cdot \mid \mathcal{F}_t)$, $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot \mid \mathcal{F}_t]$. We say S_t is regular if*

$$\mathbb{P}_t \left(|\langle \sigma_1^{(t)}, \sigma_2^{(t)} \rangle| \leq C_2 \sqrt{N} \sqrt{|Q_t| + \log N} \right) \geq 1 - N^{-10}.$$

Roughly speaking, S_t is regular if two random samples from S_t are almost orthogonal with high probability. Define the stopping time

$$\tau_S := \inf \{t \geq 0 : S_t \text{ is not regular}\}.$$

The following lemma says that for regular S_t , Y_{t+1} is roughly of order $N^{-1/2}$.

LEMMA 11. *There exists a constant $C > 0$ such that for all $t \geq 0$,*

$$\mathbf{1}_{\tau_S > t} \mathbb{P}_t \left(|Y_{t+1}| \geq C \sqrt{\frac{|Q_t| + \log N}{N}} \right) \leq C \exp(-x).$$

By (11), Lemma 11 provides an upper bound on Y_{t+1} , which can be used to control the increment of $|Q_t|$. This will be one of the key ingredients in proving Lemma 12, which gives an upper bound on $|Q_t|$ for time t before a stopping time defined below.

With $C_3 > 0$ a constant to be determined later in Lemma 18, we define

$$\tau_Y := \inf \left\{ t \geq 1 : |Y_t| \geq C_3 \sqrt{\frac{|Q_{t-1}| + \log N}{N}} \log N \right\}, \quad (12)$$

$$\tau_Q := \inf \{t \geq 0 : |Q_t| \geq (\log N)^2\},$$

and

$$\tau = \tau_S \wedge \tau_Y \wedge \tau_Q. \quad (13)$$

LEMMA 12. *There exists a constant $C > 0$ such that for all $t \geq 1$,*

$$\mathbb{E}[|Q_{t \wedge \tau-1}|] \leq \exp(Ct/N) \log N. \quad (14)$$

Lemma 13 says that the stopping time τ will occur later than time αN with high probability. Theorem 9 will be a direct consequence of Lemmas 12 and 13.

LEMMA 13. *Under Assumption 1, for any $\alpha < \alpha_c$,*

$$\lim_{N \rightarrow \infty} \mathbb{P}(\tau < \alpha N) = 0.$$

Proof of Theorem 9. By Lemmas 12 and 13, the probability bound (10) follows from Markov's inequality. $\square$

3.1 Proof of Lemma 11

In this section, we consider an arbitrary subset $A \subset \Sigma_N$ and define the probability measure $\mathbb{P}_A$ under which $(\sigma_i)_{i \geq 1}$ are independent uniformly random vectors in A .

Lemma 11 follows from the following lemma.

LEMMA 14. *There exists constant $C, c > 0$ depending only on κ such that for all N sufficiently large, all $A \subset \Sigma_N$, and $\varphi_N \in \mathbb{R}$ satisfying*

$$\delta_A := \mathbb{P}_A(|\langle \sigma_1, \sigma_2 \rangle| > \varphi_N \sqrt{N}) \leq \varphi_N^{-2} N^{-2}, \quad (15)$$

we have

$$\mathbb{P} \left(\left| \frac{|H_{\kappa}(X) \cap A|}{|A|} - p(\kappa) \right| > \frac{\varphi_N x}{\sqrt{N}} \right) \leq C \exp(-cx),$$

where X is a standard N -dimensional Gaussian vector.

The rest of this section is devoted to prove Lemma 14.

LEMMA 15. *For all N sufficiently large and every $B \subset A \subset \Sigma_N$, if we denote*

$$\delta_A = \mathbb{P}_A(|\langle \sigma_1, \sigma_2 \rangle| > \varphi_N \sqrt{N}) \quad \text{and} \quad q = \mathbb{P}_A(\sigma_1 \notin B),$$

then for all $1 \leq k \leq \frac{q^2}{4\delta_A}$

$$\mathbb{P}_A(E_k(B)) \geq \left(q - \frac{2k\delta_A}{q}\right)^k,$$

where

$$E_k(B) := \{\sigma_i \notin B, \forall 1 \leq i \leq k\} \cap \{|\langle \sigma_i, \sigma_j \rangle| \leq \varphi_N \sqrt{N} \forall 1 \leq i, j \leq k, i \neq j\}. \quad (16)$$

PROOF. We will prove the following by induction

$$\mathbb{P}_A(E_{t+1} | E_t) \geq q - \frac{2t\delta_A}{q} \quad \text{for } 1 \leq t \leq \frac{q^2}{4\delta_A}. \quad (17)$$

First, since $\mathbb{P}_A(E_1) = \mathbb{P}_A(\sigma_1 \notin B) = 1 - q$, we see that (17) holds for $t = 1$.

Next, we suppose (17) is true for $t = k$ with $k \leq \frac{q^2}{4\delta_A} - 1$. Since σ_{k+1} is independent of E_k , we have

$$\begin{aligned} \mathbb{P}_A(E_{k+1} | E_k) &= \mathbb{P}_A(\sigma_{k+1} \notin B | E_k) - \sum_{i=1}^k \mathbb{P}_A(|\langle \sigma_i, \sigma_{k+1} \rangle| > \varphi_N \sqrt{N} | E_k) \\ &= q - k \mathbb{P}_A(|\langle \sigma_k, \sigma_{k+1} \rangle| > \varphi_N \sqrt{N} | E_k). \end{aligned}$$

Also, since (σ_{k+1}, σ_k) is independent of E_{k-1} and $E_{k-1} \supset E_k$,

$$\begin{aligned} \mathbb{P}_A(|\langle \sigma_k, \sigma_{k+1} \rangle| > \varphi_N \sqrt{N} | E_k) &\leq \frac{\mathbb{P}_A(\{|\langle \sigma_k, \sigma_{k+1} \rangle| > \varphi_N \sqrt{N}\} \cap E_{k-1})}{\mathbb{P}_A(E_k)} \\ &= \mathbb{P}_A(|\langle \sigma_1, \sigma_{k+1} \rangle| > \varphi_N \sqrt{N}) \cdot \frac{\mathbb{P}(E_{k-1})}{\mathbb{P}(E_k)}. \end{aligned}$$

By the induction hypothesis

$$\frac{\mathbb{P}(E_{k-1})}{\mathbb{P}(E_k)} = \frac{1}{\mathbb{P}(E_k | E_{k-1})} \leq \frac{q}{2},$$

where we used the assumption $k \leq \frac{q^2}{4\delta_A} - 1$. Combining with the previous two inequalities, we get

$$\mathbb{P}_A(E_{k+1} | E_k) \geq q - \frac{2k\delta_A}{q}.$$

Hence (17) holds. $\square$

LEMMA 16. For any $\kappa > 0$ and $m \in \mathbb{N}$, there exists constants $c_\kappa, C > 0$ such that if we let $(X_i)_{i \geq 1}$ be a Gaussian vector with

$$\mathbb{E}[X_i] = 0, \quad \mathbb{E}[X_i^2] = 1, \quad |\mathbb{E}[X_i X_j]| \leq c_\kappa m^{-1},$$

then

$$\begin{aligned} \mathbb{P}\left(\min_{1 \leq i \leq m} |X_i| > \kappa\right) &\leq C \mathbb{P}(|X_1| > \kappa)^m, \\ \mathbb{P}\left(\max_{1 \leq i \leq m} |X_i| \leq \kappa\right) &\leq C \mathbb{P}(|X_1| \leq \kappa)^m. \end{aligned}$$

We omit the proof of Lemma 16 in this extended abstract, and proceed to prove Lemma 14.

Proof of Lemma 14. We will prove that

$$\mathbb{P}\left(\frac{|H_\kappa(X) \cap A|}{|A|} < p(\kappa) - \frac{\varphi_N x}{\sqrt{N}}\right) \leq C \exp(-cx). \quad (18)$$

The other direction can be proved similarly.

Recall $E_k(\cdot)$ as in (16) and denote

$$A_k^\perp := \{(\sigma_i)_{i=1}^k \in A^k : |\langle \sigma_i, \sigma_j \rangle| \leq \varphi_N \sqrt{N} \quad \forall 1 \leq i, j \leq k, i \neq j\}.$$

Then

$$\mathbb{E}[\mathbb{P}_A(E_k(H_\kappa(X)))] = \sum_{(\sigma_i)_{i=1}^k \in A_k^\perp} |A|^{-k} \mathbb{P}(\sigma_i \notin H_\kappa(X), \forall 1 \leq i \leq k).$$

By Lemma 16, we have that for $k = \lfloor c_\kappa \varphi_N^{-1} \sqrt{N} \rfloor$ and all $(\sigma_i)_{i=1}^k \in A_k^\perp$,

$$\mathbb{P}(\sigma_i \notin H_\kappa(X), \forall 1 \leq i \leq k) \leq C \mathbb{P}(|Z| > \kappa)^k.$$

Hence

$$\mathbb{E}[\mathbb{P}_A(E_k(H_\kappa(X)))] \leq C \mathbb{P}(|Z| > \kappa)^k. \quad (19)$$

On the other hand, we set

$$\Delta = \frac{\varphi_N x}{\sqrt{N}}.$$

Then (15) implies $\Delta \geq \frac{4k\delta_A}{q}$. Therefore, it follows from Lemma 15 that on the event $\{\mathbb{P}_A(\sigma_1 \notin H_\kappa(X) | H_\kappa(X)) \geq \mathbb{P}(|Z| > \kappa) + \Delta\}$, we have

$$\begin{aligned} \mathbb{P}_A(E_k(H_\kappa(X)) | H_\kappa(X)) &\geq \left(\mathbb{P}_A(\sigma_1 \notin H_\kappa(X) | H_\kappa(X)) - \Delta/2\right)^k \\ &\geq \mathbb{P}(|Z| > \kappa)^k e^{k\Delta/2}. \end{aligned}$$

Combined with (19), this yields (18). $\square$

3.2 Proof of Lemma 12

By the definitions (12) and (13), we see that for all $0 \leq i \leq \tau - 1$,

$$|Y_i| \leq C \sqrt{\frac{|Q_{t-1}| + \log N}{N}} \leq C \sqrt{\frac{(\log N)^2 + \log N}{N}}.$$

Hence it follows from $|\log(1+x) - x| \leq x^2$ for all $x \geq -3/5$ that

$$\left| \sum_{i=1}^{t \wedge \tau - 1} \log(1+Y_i) \right| \leq \left| \sum_{i=1}^{t \wedge \tau - 1} Y_i \right| + \sum_{i=1}^{t \wedge \tau - 1} Y_i^2 \leq \left| \sum_{i=1}^{t \wedge \tau} Y_i \right| + \sum_{i=1}^{t \wedge \tau} Y_i^2 + |Y_{t \wedge \tau}|.$$

Since $\tau \leq \tau_S$, Lemma 11 implies

$$\mathbb{E}[Y_{t \wedge \tau}^2 \mathbf{1}_{\tau \geq i}] = \mathbb{E}[\mathbb{E}_{i-1}[Y_i^2] \mathbf{1}_{\tau \geq i-1}] \leq C' \frac{\mathbb{E}[|Q_{i-1}| \mathbf{1}_{\tau \geq i-1}] + \log N}{N}.$$

Therefore, since $(\sum_{i=1}^{t \wedge \tau} Y_i)_{t \geq 1}$ is a martingale and $t \wedge \tau \leq \tau \leq \tau_S$, we have

$$\begin{aligned} \mathbb{E}\left|\sum_{i=1}^{t \wedge \tau - 1} \log(1+Y_i)\right| &\leq \sqrt{\mathbb{E}\sum_{i=1}^{t \wedge \tau} Y_i^2} + \mathbb{E}\sum_{i=1}^{t \wedge \tau} Y_i^2 + \mathbb{E}|Y_{t \wedge \tau}| \\ &\leq 2\mathbb{E}\sum_{i=1}^t Y_i^2 \mathbf{1}_{\tau \geq i} + 1 + \mathbb{E}|Y_{t \wedge \tau}| \\ &\leq 2C' \cdot \frac{t \log N + \sum_{i=1}^t \mathbb{E}[|Q_{i \wedge \tau - 1}|]}{N} + \frac{2}{p(\kappa)}, \end{aligned}$$

where C' is a constant. Hence

$$\mathbb{E}[|Q_{t \wedge \tau - 1}|] \leq 3C' \left(\frac{t \log N + \sum_{i=1}^{t-1} \mathbb{E}[|Q_{i \wedge \tau - 1}|]}{N} + \frac{1}{p(\kappa)} \right)$$

We claim that this yields (14). To prove this, we define $b_1 := \mathbb{E}[|Q_{1 \wedge \tau-1}|]$ and for $t \geq 2$ define

$$b_t := 3C' \left(\frac{t \log N + \sum_{i=1}^{t-1} b_i}{N} + \frac{1}{p(\kappa)} \right). \quad (20)$$

Then a straightforward induction argument shows that

$$b_t \geq \mathbb{E}[|Q_{t \wedge \tau-1}|] \quad \text{for } t \geq 1. \quad (21)$$

On the other hand, (20) implies

$$\begin{aligned} b_t &= 3C' \left(\frac{\log N + b_{t-1}}{N} + \frac{(t-1) \log N + \sum_{i=1}^{t-2} b_i}{N} + \frac{1}{p(\kappa)} \right) \\ &= 3C' \left(\frac{\log N + b_{t-1}}{N} + \frac{b_{t-1}}{3C'} \right). \end{aligned}$$

This implies

$$b_t + \log N = \left(1 + \frac{3C'}{N}\right)(b_{t-1} + \log N) = \left(1 + \frac{3C'}{N}\right)^{t-1} (b_1 + \log N). \quad (22)$$

Note that $S_0 = \{-1, 1\}^N$, $Q_0 = 0$ and $\mathbb{P}(\tau \geq 1) = 1$. We get $b_1 = 0$. Hence Lemma (14) follows from (21) and (22). We complete the proof of Lemma 12.

3.3 Proof of Lemma 13

LEMMA 17. *Under Assumption 1, there exists a constant C_2 such that for every $1 \leq t \leq \alpha N$,*

$$\mathbb{P} \left(\mathbb{P}_t \left(|\langle \sigma_1^{(t)}, \sigma_2^{(t)} \rangle| \geq C_2 \sqrt{N} \sqrt{Q_t + \log N} \right) \geq N^{-10} \right) \leq N^{-10}.$$

As a result,

$$\mathbb{P}(\tau_S \leq \alpha N) \leq N^{-10}. \quad (23)$$

PROOF. It follows from (9), Assumption 1, and $F''(\frac{1}{2}) < 0$ (see [5]) that there exists a constant $c > 0$ such that for any $\lambda > 0$ and $t \geq 0$,

$$\mathbb{E}[|\{(\sigma_1, \sigma_2) \in S_t : |\langle \sigma_1, \sigma_2 \rangle| \geq \lambda \sqrt{N}\}|] \leq \exp(-c\lambda^2) \mathbb{E}[|S_t|^2].$$

Under Assumption 1, we have $\mathbb{E}[|S_t|^2] \leq C \mathbb{E}[|S_t|]^2$ (see [5]). Hence we have that

$$\mathbb{E} \left[\mathbb{P}_t \left(|\langle \sigma_1^{(t)}, \sigma_2^{(t)} \rangle| \geq \lambda \sqrt{N} \right) \frac{|S_t|^2}{\mathbb{E}[|S_t|]^2} \right] \leq C \exp(-c\lambda^2),$$

Therefore,

$$\begin{aligned} &\mathbb{E} \left[\exp \left(\frac{c}{2} \frac{\langle \sigma_1^{(t)}, \sigma_2^{(t)} \rangle^2}{N} + 2Q_t \right) \right] \\ &= \mathbb{E} \left[\int_0^\infty c\lambda e^{\frac{c\lambda^2}{2}} \mathbb{P}_t \left(\frac{\langle \sigma_1^{(t)}, \sigma_2^{(t)} \rangle^2}{N} \geq \lambda \right) d\lambda \cdot \frac{|S_t|^2}{\mathbb{E}[|S_t|]^2} \right] \\ &\leq C \int_0^\infty c\lambda e^{\frac{-c\lambda^2}{2}} d\lambda \\ &\leq C. \end{aligned}$$

Hence, there exists a constant C such that

$$\mathbb{P}(|\langle \sigma_1^{(t)}, \sigma_2^{(t)} \rangle| \geq C \sqrt{N} \sqrt{Q_t + \log N}) \leq N^{-100},$$

which yields the desired result. $\square$

LEMMA 18. *There exists a constant $C_3 > 0$ such that*

$$\begin{aligned} \mathbb{P}(\tau_Y \leq \tau_S, \tau_Y \leq \alpha N) &\leq N^{-8} \\ \mathbb{P}(|Q_{\tau_Q-1}| \leq (\log N)^2/2, \tau_Q \leq \tau_S, \tau_Q \leq \alpha N) &\leq N^{-8} \end{aligned} \quad (24)$$

PROOF. (11) follows from Lemma 11 and a union bound. Lemma 11 gives

$$\begin{aligned} \mathbb{P}(|Q_{\tau_Q-1}| \leq (\log N)^2/2, \tau_Q \leq \tau_S, \tau_Q = t+1) \\ &= \mathbb{E} \left[\mathbb{P}_t(|Q_t + \log(1 + Y_{t+1})| \geq (\log N)^2) \mathbf{1}_{|Q_t| \leq (\log N)^2/2, \tau_S \geq t+1} \right] \\ &\leq \mathbb{E} \left[\mathbb{P}_t(Y_t \leq -1/2) \mathbf{1}_{|Q_t| \leq (\log N)^2/2, \tau_S \geq t+1} \right] \\ &\leq N^{-10}. \end{aligned}$$

This yields (24). $\square$

Proof of Lemma 13. Write

$$\begin{aligned} \mathbb{P}(\tau \geq t) &\leq \mathbb{P}(\tau_S \leq \alpha N) + \mathbb{P}(\tau_Y \leq \tau_S, \tau_Y \leq \alpha N) \\ &\quad + \mathbb{P}(\tau_Q \leq t, \tau_Q = \tau, |Q_{\tau_Q-1}| \leq (\log N)^2/2) \\ &\quad + \mathbb{P}(\tau_Q \leq t, \tau_Q = \tau, |Q_{\tau_Q-1}| \geq (\log N)^2/2). \end{aligned}$$

Lemma 18 and (23) give upper bounds on the first three probability on the right hand. In addition, Lemma 12 and the Markov inequality yield

$$\begin{aligned} \mathbb{P}(\tau_Q \leq t, \tau_Q = \tau, |Q_{\tau_Q-1}| \geq (\log N)^2/2) \\ &\leq ((\log N)^2/2)^{-1} \mathbb{E}[|Q_{t \wedge \tau-1}|; \tau_Q \leq t, \tau_Q = \tau, |Q_{\tau_Q-1}| \geq (\log N)^2/2] \\ &\leq ((\log N)^2/2)^{-1} \mathbb{E}[|Q_{t \wedge \tau-1}|] \\ &\leq (\log N)^{-1/2}. \end{aligned}$$

Combining these bounds yields $\mathbb{P}(\tau \geq t) \leq 2(\log N)^{-1/2}$. $\square$

4 FROM THE PLANTED MODEL TO THE RANDOM MODEL

Theorem 1 will follow from Lemma 6 and the following lemma on the planted model.

LEMMA 19. *Suppose $A \subseteq \{(\sigma, U) : \sigma \in U \subseteq \Sigma_N\}$, and suppose $\alpha < \alpha_c$. If $\mathbb{P}_{\text{pl}}((\sigma^*, S(\mathbf{X})) \in A) \leq N^{-m_N}$ for some $m_N \rightarrow \infty$, then*

$$\mathbb{P}_{\text{r}}((\sigma^*, S(\mathbf{X})) \in A) \rightarrow 0 \quad \text{as } N \rightarrow \infty.$$

We start by characterizing the planted distribution.

LEMMA 20. *For any $N, m \geq 1$, with $A \subseteq \{(\sigma, U) : \sigma \in U \subseteq \Sigma_N\}$,*

$$\mathbb{P}_{\text{pl}}((\sigma^*, S(\mathbf{X})) \in A) = \mathbb{E}_{\text{r}} \left[\mathbf{1}_{(\sigma_N, S(\mathbf{X})) \in A} \cdot \frac{|S(\mathbf{X})|}{\mathbb{E}[|S(\mathbf{X})|]} \cdot \mathbb{P}(S(\mathbf{X}) \neq \emptyset) \right].$$

We omit the Proof of Lemma 20 in this extended abstract.

Proof of Lemma 19. Theorem 2 implies the event

$$G = \left\{ \frac{|S|}{\mathbb{E}[|S|]} \geq \exp\{-m_N \log N\} \right\}$$

has probability $1 - o(1)$. Therefore, as $N \rightarrow \infty$,

$$\begin{aligned} \mathbb{P}_r((\sigma^*, S) \in A) &\leq \mathbb{P}_r\left(\{(\sigma^*, S) \in A\} \cap G\right) + \mathbb{P}(G)/\mathbb{P}(S \neq \emptyset) \\ &\leq \exp\{m_N \log N\} \mathbb{E} \left[\mathbf{1}_{(\sigma^*, S) \in A} \cdot \frac{|S|}{|\mathbb{E}[S]|} \right] + o(1) \\ &= \exp\{m_N \log N\} \mathbb{P}_{\text{pl}}((\sigma^*, S) \in A)/\mathbb{P}(S \neq \emptyset) + o(1) \\ &= o(1). \end{aligned}$$

where we used Lemma 20 and Theorem 2. This yields Lemma 19. $\square$

Finally we prove Theorem 1.

PROOF OF THEOREM 1. Let

$$A = \{(\tau, U) : \{\sigma \in U : \langle \sigma, \tau \rangle \geq (d_c + \delta)N\} \neq \{\tau\}\}.$$

Then Lemma 6 implies that

$$\mathbb{P}_{\text{pl}}((\sigma^*, S(X)) \in A) \leq \exp\{-c\sqrt{N}\}$$

Combined with Lemma 19, this yields

$$\mathbb{P}_r((\sigma^*, S(X)) \in A) = o(1),$$

and thus Theorem 1. $\square$

ACKNOWLEDGEMENTS

We thank Benjamin Aubin and Lenka Zdeborová for inspiring discussions and introducing us to this problem.

REFERENCES

- [1] Emmanuel Abbe. 2017. Community detection and stochastic block models: recent developments. *The Journal of Machine Learning Research* 18, 1 (2017), 6446–6531.
- [2] Dimitris Achlioptas and Amin Coja-Oghlan. 2008. Algorithmic barriers from phase transitions. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 793–802.
- [3] Dimitris Achlioptas, Amin Coja-Oghlan, and Federico Ricci-Tersenghi. 2011. On the solution-space geometry of random constraint satisfaction problems. *Random Structures & Algorithms* 38, 3 (2011), 251–268.
- [4] Dimitris Achlioptas and Cristopher Moore. 2002. The asymptotic order of the random k-SAT threshold. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science*, 2002. *Proceedings*. IEEE, 779–788.
- [5] Benjamin Aubin, Will Perkins, and Lenka Zdeborová. 2019. Storage capacity in symmetric binary perceptrons. *Journal of Physics A: Mathematical and Theoretical* (2019).
- [6] Carlo Baldassi. 2009. Generalization learning in a perceptron with binary synapses. *Journal of Statistical Physics* 136, 5 (2009), 902–916.
- [7] Carlo Baldassi, Christian Borgs, Jennifer T Chayes, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. 2016. Unreasonable effectiveness of learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes. *Proceedings of the National Academy of Sciences* 113, 48 (2016), E7655–E7662.
- [8] Carlo Baldassi, Alfredo Braunstein, Nicolas Brunel, and Riccardo Zecchina. 2007. Efficient supervised learning in networks with binary synapses. *Proceedings of the National Academy of Sciences* 104, 26 (2007), 11079–11084.
- [9] Carlo Baldassi, Riccardo Della Vecchia, Carlo Lucibello, and Riccardo Zecchina. 2020. Clustering of solutions in the symmetric binary perceptron. *Journal of Statistical Mechanics: Theory and Experiment* 2020, 7 (2020), 073303.
- [10] Carlo Baldassi, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. 2015. Subdominant dense clusters allow for simple learning and high computational performance in neural networks with discrete synapses. *Physical review letters* 115, 12 (2015), 128101.
- [11] Carlo Baldassi, Alessandro Ingrosso, Carlo Lucibello, Luca Saglietti, and Riccardo Zecchina. 2016. Local entropy as a measure for sampling solutions in constraint satisfaction problems. *Journal of Statistical Mechanics: Theory and Experiment* 2016, 2 (2016), 023301.
- [12] Victor Bapst, Amin Coja-Oghlan, and Charilaos Efthymiou. 2017. Planting colourings silently. *Combinatorics, probability and computing* 26, 3 (2017), 338–366.
- [13] Victor Bapst, Amin Coja-Oghlan, Samuel Hetterich, Felicia Raßmann, and Dan Vilenchik. 2016. The condensation phase transition in random graph coloring. *Communications in Mathematical Physics* 341, 2 (2016), 543–606.
- [14] Zsolt Bartha, Nike Sun, and Yumeng Zhang. 2019. Breaking of 1RSB in random regular MAX-NAE-SAT. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 1405–1416.
- [15] Alfredo Braunstein and Riccardo Zecchina. 2006. Learning by message passing in networks of discrete synapses. *Physical review letters* 96, 3 (2006), 030201.
- [16] Amin Coja-Oghlan, Charilaos Efthymiou, Nor Jaafari, Mi Hyun Kang, and Tobias Kapetanopoulos. 2018. Charting the replica symmetric phase. *Communications in Mathematical Physics* 359, 2 (2018), 603–698.
- [17] Amin Coja-Oghlan, Florent Krzakala, Will Perkins, and Lenka Zdeborová. 2018. Information-theoretic thresholds from the cavity method. *Advances in Mathematics* 333 (2018), 694–795.
- [18] Amin Coja-Oghlan and Lenka Zdeborová. 2012. The condensation transition in random hypergraph 2-coloring. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 241–250.
- [19] Thomas M Cover. 1965. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Transactions on Electronic Computers* 3 (1965), 326–334.
- [20] Jian Ding, Allan Sly, and Nike Sun. 2014. Satisfiability threshold for random regular NAE-SAT. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. 814–822.
- [21] Jian Ding and Nike Sun. 2019. Capacity lower bound for the Ising perceptron. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 816–827.
- [22] Ehud Friedgut. 1999. Sharp thresholds of graph properties, and the k-sat problem. *Journal of the American mathematical Society* 12, 4 (1999), 1017–1054.
- [23] Elizabeth Gardner. 1987. Maximum storage capacity in neural networks. *EPL (Europhysics Letters)* 4, 4 (1987), 481.
- [24] E Gardner and B Derrida. 1988. Optimal storage properties of neural network models. *Journal of Physics A: Mathematical and general* 21, 1 (1988), 271.
- [25] D. L. Hanson and F. T. Wright. 1971. A bound on tail probabilities for quadratic forms in independent random variables. *Ann. Math. Statist.* 42 (1971), 1079–1083. <https://doi.org/10.1214/aoms/1177693335>
- [26] Haiping Huang and Yoshiyuki Kabashima. 2014. Origin of the computational hardness for learning with binary synapses. *Physical Review E* 90, 5 (2014), 052813.
- [27] Haiping Huang, KY Michael Wong, and Yoshiyuki Kabashima. 2013. Entropy landscape of solutions in the binary perceptron problem. *Journal of Physics A: Mathematical and Theoretical* 46, 37 (2013), 375002.
- [28] Jeong Han Kim and James R Roche. 1998. Covering cubes by random half cubes, with applications to binary neural networks. *J. Comput. System Sci.* 56, 2 (1998), 223–252.
- [29] Werner Krauth and Marc Mézard. 1989. Storage capacity of memory networks with binary couplings. *Journal de Physique* 50, 20 (1989), 3057–3066.
- [30] Florent Krzakala, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. 2007. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences* 104, 25 (2007), 10318–10323.
- [31] Florent Krzakala and Lenka Zdeborová. 2009. Hiding quiet solutions in random constraint satisfaction problems. *Physical review letters* 102, 23 (2009), 238701.
- [32] David Mitchell, Bart Selman, and Hector Levesque. 1992. Hard and easy distributions of SAT problems. In *AAAI*, Vol. 92. 459–465.
- [33] Michael Molloy. 2018. The freezing threshold for k-colourings of a random graph. *Journal of the ACM (JACM)* 65, 2 (2018), 1–62.
- [34] Andrea Montanari, Ricardo Restrepo, and Prasad Tetali. 2011. Reconstruction and clustering in random constraint satisfaction problems. *SIAM Journal on Discrete Mathematics* 25, 2 (2011), 771–808.
- [35] Alexander Markowitsch Ostrowski. 1938. Sur l’approximation du déterminant de Fredholm par les déterminants des systèmes d’équations linéaires. *Ark. Math. Stockholm* 26A (1938), 1–15.
- [36] Allan Sly, Nike Sun, and Yumeng Zhang. 2016. The number of solutions for random regular NAE-SAT. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 724–731.
- [37] Haim Sompolinsky, Naftali Tishby, and H Sebastian Seung. 1990. Learning from examples in large neural networks. *Physical Review Letters* 65, 13 (1990), 1683.
- [38] Michel Talagrand. 1999. Intersecting random half cubes. *Random Structures & Algorithms* 15, 3–4 (1999), 436–449.
- [39] Michel Talagrand. 2010. *Mean field models for spin glasses: Volume I: Basic examples*. Vol. 54. Springer Science & Business Media.
- [40] David J Thouless, Philip W Anderson, and Robert G Palmer. 1977. Solution of ‘Solvable model of a spin glass’. *Philos. Mag.* 35, 3 (1977), 593–601.
- [41] Changji Xu. 2019. Sharp threshold for the Ising perceptron model. *arXiv preprint arXiv:1905.05978* (2019).
- [42] Lenka Zdeborová and Florent Krzakala. 2007. Phase transitions in the coloring of random graphs. *Physical Review E* 76, 3 (2007), 031131.

- [43] Lenka Zdeborová and Florent Krzakala. 2016. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics* 65, 5 (2016), 453–552.
- [44] Lenka Zdeborová and Marc Mézard. 2008. Constraint satisfaction problems with isolated solutions are hard. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 12 (2008), P12004.