Critical window of the symmetric perceptron

Dylan J. Altschuler*

Abstract

We study the critical window of the symmetric binary perceptron, or equivalently, random combinatorial discrepancy. Consider the problem of finding a ± 1-valued vector σ satisfying $\|A\sigma\|_\infty \leq K$, where A is an $\alpha n \times n$ matrix with iid Gaussian entries. For fixed K, at which constraint densities α is this constraint satisfaction problem (CSP) satisfiable? A sharp threshold was recently established by Perkins and Xu [29], and Abbe, Li, and Sly [2], answering this to first order. Namely, for each K there exists an explicit critical density α_c so that for any fixed $\epsilon > 0$, with high probability the CSP is satisfiable for $\alpha n < (\alpha_c - \epsilon)n$ and unsatisfiable for $\alpha n > (\alpha_c + \epsilon)n$. This corresponds to a bound of $o(n)$ on the size of the critical window.

We sharpen these results significantly, as well as provide exponential tail bounds. Our main result is that, perhaps surprisingly, the critical window is actually at most of order $\log(n)$. More precisely, for a large constant C, with high probability the CSP is satisfiable for $\alpha n < \alpha_c n - C \log(n)$ and unsatisfiable for $\alpha n > \alpha_c n + C$. These results add the the symmetric perceptron to the short list of CSP models for which a critical window is rigorously known, and to the even shorter list for which this window is known to have nearly constant width.

1 Introduction

Discrepancy arises as a fundamental quantity in combinatorics, functional analysis, geometry, and optimization. The **combinatorial discrepancy** of a matrix $A \in \mathbb{R}^{\alpha n \times n}$ is given by the following optimization problem:

$$\operatorname{disc}(A) := \min_{\sigma \in \{-1, +1\}^n} \|A\sigma\|_\infty.$$

In words, discrepancy is the task of finding some $\sigma \in \{\pm 1\}^n$ that has small inner product in absolute value with all rows of A. This task is naturally encoded as a CSP (constraint

---
*Courant Institute, New York University. E-mail: daltschu@andrew.cmu.edu

satisfaction problem). Fix a positive parameter K and look for assignments $\sigma \in \{\pm 1\}^n$ that satisfy each of the αn constraints encoded by the rows of A:

$$\{|\langle A_1, \sigma \rangle| \leq K\} \wedge \cdots \wedge \{|\langle A_{\alpha n}, \sigma \rangle| \leq K\}. \tag{1.1}$$

We say the problem is satisfiable if such an assignment σ exists. For fixed K, if we construct Eq. (1.1) by adding one constraint at a time, how many constraints can we add before the equation becomes unsatisfiable?

Much recent work focuses on understanding, for random A, the parameters K and α for which Eq. (1.1) is satisfiable with positive probability. For random CSP, it is natural to suspect the existence of a sharp threshold [12, 19, 22]. A *sharp threshold* is said to occur if there exists some α_c so that for any $\epsilon > 0$, the probability that Eq. (1.1) is satisfiable tends to one for $\alpha n < (\alpha_c - \epsilon)n$ and to zero for $\alpha n > (\alpha_c + \epsilon)n$, as n grows. We now formally state the model of randomness that will be studied.

Definition 1.1. *Let $\alpha := \alpha(n)$. Say an $\alpha n \times n$ matrix A is (α, n)-Gaussian if it has iid normal entries with mean 0 and variance $1/n$.*

The normalization is so that $\mathrm{disc}(A)$ is on the constant scale. We will only consider $\alpha := \alpha(n)$ tending to a finite, strictly postive limit (see Section 2 for $\alpha \to 0$). The discrepancy of such matrices has gathered significant recent attention in the statistical physics community due to connections with the celebrated *Binary Perceptron*.

The binary perceptron is an idealized model of learning that has enjoyed almost six decades of intense study by the statistical physics community. While there is an extensive body of detailed physics predictions for the perceptron, these predictions have largely resisted rigorous mathematical treatment [16, 21, 27, 33, 39, 10]. Aubin, Perkins, and Zdeborová [6] had the insight of interpreting combinatorial discrepancy as a "symmetric" analogue of the binary perceptron. Happily, discrepancy is more amenable to analysis and is conjectured to capture most of the interesting behaviours of the perceptron. This connection has already lead to several breakthroughs in both problems [1, 2, 20, 28, 29].

In brief, the reduction between the SBP and discrepancy is as follows. We are interested in the set of parameters (K, α) for which Eq. (1.1) is satisfiable. One could equally well fix α and ask for the smallest K, or else fix K and ask for the largest α. The former is adopted in discrepancy and the latter for the SBP. Formally,

Definition 1.2 (Symmetric perceptron). *For a fixed parameter K, an instance of the associated **symmetric binary perceptron** (SBP) problem is given by a sequence of iid Gaussian vectors $\{a_i\}_{i \in \mathbb{N}}$ with*

$$a_i \sim N(0, n^{-1} \mathrm{I}_{n \times n}).$$

*The **storage capacity** of the SBP is the random variable $\boldsymbol{\alpha_*}$ defined as the largest α so that $\alpha n \in \mathbb{Z}$ and there exists $\sigma \in \{\pm 1\}^n$ with*

$$\max_{i \in [\alpha n]} |\langle a_i, \sigma \rangle| \leq K.$$

The original binary perceptron model asks for σ to have positive, rather than small, inner product with rows. A simple but important observation (made formal in Lemma 3.2 below) is the following.

> Let K_c and α_c denote the "typical values" of $\mathrm{disc}(A)$ and $\boldsymbol{\alpha_*}$ respectively. A bound on $|\mathrm{disc}(A) - K_c|$ is equivalent to a bound on $|\boldsymbol{\alpha_*} - \alpha_c|$ of the same order. Thus, fluctuation bounds on the capacity of the SBP or the discrepancy of a Gaussian matrix are equivalent.

As such, we discuss these quantities interchangeably. For matrices with iid Gaussian or Rademacher entries, the existence of a *coarse threshold* was shown by Aubin, Perkins, and Zdeborová [6] via the second moment method. They also noted the second moment method fails to give a sharp threshold and conjectured a sharp threshold should still occur. The failure of the second moment method to yield a sharp threshold—usually a difficult obstacle to overcome—is a key feature of many interesting CSP, including the binary perceptron. (In fact, the second moment method famously fails to even yield a coarse threshold for the binary perceptron). In an impressive series of works, Perkins and Xu [29] and Abbe, Li, and Sly [2] independently established the sharp threshold for the SBP, resolving the conjecture of [6].

When a random CSP admits a sharp threshold, a next natural question is to identify the width of the critical window, sometimes also called the scaling window. Namely, fix some small constant $\delta > 0$; as a function of n, how many constraints must be added for the probability that (1.1) is satisfiable to decrease from $1 - \delta$ to δ? The following proposition summarizes what is known:

Proposition 1.3 (Sharp threshold [29, 2])**.** *For any positive constants K, ϵ, and δ, there exists $\alpha_c := \alpha_c(K)$ (given explicitly in Eq. (3.1)) so that for A drawn from the (α, n)-Gaussian ensemble and n sufficiently large,*

$$\alpha n < \alpha_c n - \epsilon n \implies \mathbb{P}\left[\mathrm{disc}(A) > K\right] < \delta,$$
$$\alpha n > \alpha_c n + \epsilon n \implies \mathbb{P}\left[\mathrm{disc}(A) < K\right] < \delta.$$

Equivalently, the capacity of the symmetric perceptron satisfies

$$\mathbb{P}\left[\alpha_c n - \epsilon n < \boldsymbol{\alpha}_* < \alpha_c n + \epsilon n\right] > 1 - 2\delta$$

No further quantification has been established or even conjectured. Our main contribution is to narrow the current bound on the critical window from $o(n)$ to $\mathcal{O}\left(\log(n)\right)$. Strikingly, the list of CSP for which the critical window is even known to be polynomially smaller than n, let alone constant, is very short. This is discussed further in Section 2 below. We begin with a simplified, qualitative statement of our main result:

Theorem 1.1 (**Main result:** logarithmic critical window)**.** *For any positive constants K and δ, there exists an explicit constant $\alpha_c := \alpha_c(K)$ and some sufficiently large $c > 0$ so that for A drawn from the (α, n)-Gaussian ensemble and n sufficiently large,*

$$\alpha n < \alpha_c n - c\log(n) \implies \mathbb{P}\left[\mathrm{disc}(A) > K\right] < \delta,$$
$$\alpha n > \alpha_c n + c \implies \mathbb{P}\left[\mathrm{disc}(A) < K\right] < \delta.$$

Equivalently, the capacity of the symmetric perceptron satisfies

$$\mathbb{P}\left[\alpha_c n - c < \boldsymbol{\alpha}_* < \alpha_c n + c\log(n)\right] > 1 - 2\delta$$

We actually prove a significant quantitative strengthening of this. Our main technical contribution gives sub-exponential tails on the fluctuations of discrepancy and directly implies Theorem 1.1.

Theorem 1.2 (**Fluctuations**: capacity)**.** *Fix $K > 0$, and let $\boldsymbol{\alpha}_*$ be the storage capacity of the corresponding SBP. There exist positive constants c and $\alpha_c := \alpha_c(K)$ depending only on K so that for any sufficiently large fixed x, the following holds for all sufficiently large n:*

$$\mathbb{P}\left[|\boldsymbol{\alpha}_* - \alpha_c| > \frac{x\log(n)}{n}\right] \le \exp\left\{-cx\log\left(n\right)\right\}.$$

In addition, there exist some other positive constants c and C depending only on K so that for all n sufficiently large,

$$\frac{c}{n} \le \mathrm{Var}(\boldsymbol{\alpha}_*)^{1/2} \le \frac{C}{n} \log(n).$$

Theorem 1.2 implies Theorem 1.1 as an immediate corollary, as well as an analogous bound on discrepancy:

Theorem 1.3 (**Fluctuations**: discrepancy). *Fix $\alpha > 0$ and let A be from the (α, n)-Gaussian ensemble. There exist positive constants c and $K_c := K_c(\alpha)$ depending only on α so that for any sufficiently large fixed x, the following holds for all sufficiently large n:*

$$\mathbb{P}\left[|\mathrm{disc}(A) - K_c| > \frac{x \log(n)}{n}\right] \le \exp\left\{-cx \log(n)\right\}. \tag{1.2}$$

In addition, there exist some other positive constants c and C depending only on α so that for all n sufficiently large,

$$\frac{c}{n} \le \mathrm{Var}(\mathrm{disc}(A))^{1/2} \le \frac{C}{n} \log(n). \tag{1.3}$$

Some remarks:

1. Existing general tools fall quite short of being able to capture when the critical window of a CSP has width polynomially smaller than n—let alone a window of constant width. See e.g. a result of Xu giving a sharp threshold for the (asymmetric) perceptron [41] based on Hatami's theorem [22], which establishes a sharp threshold sequence of width at most $n/(\log(\log(n)))^{1/10}$.

2. It is natural to compare discrepancy to the smallest singular value of a Gaussian matrix, since they solve somewhat similar minimization problems. Theorem 1.2 shows discrepancy is significantly more concentrated, with fluctuations of order n^{-1}, whereas the smallest singular value of a Gaussian matrix follows the Tracy-Widom law with fluctuations of order $n^{-2/3}$ [34].

3. Our lower-bound on the variance shows that the critical window is at least constant. cf. random linear equations over a finite field, which can be written as a CSP with a degenerate window [26, 24]. By "degenerate", we mean that the critical window has size exactly one, as well as a deterministic location; this corresponds to random (regular) equations being solvable w.h.p. exactly until the number of equations αn exceeds the number of variables n.

We conclude by highlighting a key intermediate result. Let us briefly introduce a standard perspective for studying CSP: counting solutions. Fix α and define $Z_K := Z_{K,\alpha}(A)$ to be the set of solutions incurring discrepancy at most K.

$$Z_K := \{x \in \{\pm 1\}^n : \|Ax\|_\infty \le K\}, \quad \mathrm{disc}(A) = \inf\{K : Z_K \ne \emptyset\}.$$

A natural prediction is that $\mathrm{disc}(A)$ should concentrate around $K_c := K_c(\alpha)$, defined as the smallest K for which $\mathbb{E}[|Z_K|] \ge 1$. Showing $|Z_K|$ concentrates around its mean would verify this prediction. Markov's inequality readily implies $\mathrm{disc}(A) > (K_c - \epsilon)$ with exponentially high probability, for any $\epsilon > 0$. The coarse threshold of Aubin, Perkins, and Zdeborová [6], which states $\mathrm{disc}(A) < (K_c + \epsilon)$ with positive probability for any $\epsilon > 0$, corresponds to a variance bound:

$$\forall \epsilon > 0, \quad 1 < c \le \frac{\mathbb{E}\left[|Z_{K_c+\epsilon}|^2\right]}{\mathbb{E}\left[|Z_{K_c+\epsilon}|\right]^2} \le C < \infty. \tag{1.4}$$

While Eq. (1.4) implies that the variance of $|Z_K|$ is small enough to obtain a coarse threshold, it also implies the variance of $|Z_K|$ is too large to obtain a sharp threshold. Rather, it is $\log |Z_K|$ that concentrates well. We will soon check that for any small positive constant ϵ, $\log \mathbb{E}\left[|Z_{K_c+\epsilon}|\right] = \Theta\left(\epsilon n\right)$. The sharp threshold Proposition 1.3 of [29, 2] corresponds to the asymptotic bound that for any positive δ, there exists large M with

$$\lim_{n \to \infty} \mathbb{P}\left[|\log |Z_{K_c+\epsilon}| - \log \mathbb{E}\left[|Z_{K_c+\epsilon}|\right]| > M \log(n)\right] \le \delta. \tag{1.5}$$

Our contribution is to give a non-asymptotic sharpening of Eq. (1.5) in which ϵ and δ have explicit, near-optimal dependence on n. As a necessary intermediate step, we show a significant strengthening of Eq. (1.4). Namely, the second moment method yields a non-trivial bound even *exactly* down to the critical threshold K_c. This is new and perhaps surprising.

Theorem 1.4 (Coarse threshold, up to criticality). *Fix any $\alpha > 0$ and let $K_c := K_c(\alpha)$. There exists $C := C(\alpha) > 0$ so that for all n sufficiently large,*

$$\frac{\mathbb{E}\left[|Z_{K_c,\alpha}|^2\right]}{\mathbb{E}\left[|Z_{K_c,\alpha}|\right]^2} \ge C > 0. \tag{1.6}$$

Consequently,

$$\mathbb{P}\left[\mathrm{disc}(A) \le K_c\right] > C. \tag{1.7}$$

Since the second moment method can be carried out all the way to criticality, it is natural to suspect similar improvements are possible in Eq. (1.5). Indeed, this is the content of Theorem 1.3. We choose to refine the techniques of [29] because they naturally lead to strong tail bounds, albeit at the cost of an extra logarithm for the width of the critical window. It seems likely that the techniques of [2] can be used to show $\mathrm{disc}(A) = K_c + \Theta\left(1/n\right)$, resolving this extra log. However, the techniques of [2] are asymptotic and do not easily lead to tail bounds.

Simultaneously providing strong tail bounds and capturing the exact size of fluctuations is left as an open question; we conjecture that our main result can be improved both by removing the logarithmic mismatch in Eq. (1.3), and also by giving subgaussian rather than subexponential tails in Eq. (1.2).

Conjecture 1.4. $|\mathrm{disc}(A) - K_c|$ *is $1/n^2$-subgaussian.*

Author's note: after the appearance of this work, Sah and Sawhney [35] showed the lower tails of discrepancy are actually sub-exponential. For applications in discrepancy theory, analysis of the upper tails is essential and remains open.

1.1 Notation

Throughout this work, we use c and C to denote positive universal constants that may take different values on different lines. Any important constants will be distinguished with subscripts. For functions (namely p, q, F, and Z, all defined in the following section) that depend on the parameters n, K, and α, we frequently suppress some of these dependencies when there is no ambiguity, particularly when one of these parameters is held fixed.

2 Previous works

Critical windows of random CSP The previous sharp threshold results [2] and [29] on the symmetric perceptron correspond to a bound of $o(n)$ on the width of the critical window. In the literature of random CSP, it is rare to find optimal, or even non-trivial, improvements over $o(n)$ on the critical window. General tools from Fourier Analysis

such as the sharp threshold sequence theorems of Friedgut and Bourgain [19] and Hatami [22], are the best known bounds for most CSP (the perceptron [41], k-SAT [15], and k-coloring [3], etc.). Even when applicable, the Friedgut-Bourgain sharp threshold theorem gives only the qualitative upper-bound $o(n)$, and Hatami's theorem gives a slightly improved $n/\log(\log(n))$. It is possible that many of these standard CSP should have critical windows *polynomially* smaller than n.

We are only aware of two other CSP for which such an improvement has been shown: "constrained" XOR-SAT [30], and 2-SAT [9]. Both enjoy very special structure: XOR-SAT is essentially a "matrix invertibility" problem, and its critical window is at most constant due to dimension-counting arguments. The 2-SAT model has critical window of size $\Theta\left(n^{2/3}\right)$, and the analysis employs from a series of clever reductions to a well-studied percolation problem.

Let us also clarify the difference between our result and sharp threshold results for "regular" models, e.g. random regular NAE-SAT. A random regular constraint satisfaction problem on n variables is a collection of clauses each involving k variables, with each variable in d clauses. The number of clauses must then be

$$nd/k.$$

Ding, Sly, and Sun showed for random regular NAE-SAT [14] that the largest value of d for which there are solutions concentrates around an explicit constant, with at most constant fluctuations. Since a constant change in d results in a change of order n to the number of clauses, this kind of sharp threshold for d seems closer to a qualitative sharp threshold rather than the type of critical window bound investigated here.

Finally, it is quite special that the symmetric perceptron enjoys such a small critical window: there are many sources of fluctuations that could preclude a CSP from having a constant-size critical window. For example, Wilson showed the critical window of the infamous K-SAT problem is at least $\Omega\left(\sqrt{n}\right)$. His argument proceeds by noting that the number of trivially satisfiable clauses is of order n with CLT-like fluctuations of order $\Theta\left(\sqrt{n}\right)$. Another example is the critical window of the binary perceptron, which is conjectured by Ding and Sun [16] to be of order $\Theta\left(\sqrt{n}\right)$. They posit that such large fluctuations should occur due to fluctuations in the barycenter of the solution set.

Coarse threshold   The connection between discrepancy and the perceptron model was introduced by Aubin, Perkins and Zdeborová in 2019 [6]. They established a *coarse threshold* via the a second moment method. More precisely, they showed—conditional on a numerical conjecture (see Section 3)—for any $\epsilon > 0$, that $\mathrm{disc}(A) > K_c - \epsilon$ with high probability and $\mathrm{disc}(A) < K_c - \epsilon$ with positive probability. Their work also establishes the negative result that the second moment method does not suffice to prove a sharp threshold, but they nonetheless conjectured a sharp threshold (i.e. the upper-bound should hold with high, rather than positive, probability).

This conjecture was proven by Abbe, Li, and Sly [2], and Perkins and Xu [29] independently and simultaneously.

Sharp threshold (cycle counting)   Abbe, Li, and Sly used a "dense graph" analogue of the so-called small subgraph conditioning technique of Robinson and Wormald [32]. They established both the sharp threshold as well as an explicit asymptotic description of the number of solutions. The small subgraph conditioning method is a celebrated technique to correct the failure of the second moment method for counting objects in sparse random graphs. The idea is to explain the extra variance by explicitly quantifying how the number of cycles in a graph affects this count.

Interestingly, the techniques of [2] have a strong analogue in the literature of the SK (Sherrington-Kirkpatrick) model. A seminal paper of Aizenman, Lebowitz, and Ruelle [4]—written in 1987, actually predating the appearance of the small subgraph conditioning technique—established the asymptotic log-normal limit for the fluctuations in the free energy of the SK model at high temperatures, also using cycle counts in a dense graph, with many of the same objects appearing. The connection between the techniques of [4] and [32] were further investigated and made explicit in a recent work of Banerjee [7] that establishes a log-normal limit for a perturbation of the SK model.

Sharp threshold (cavity method)   Perkins and Xu independently and simultaneously resolved the sharp threshold conjecture for the SBP with different techniques. They adapted and refined an argument of Talagrand [38] originally used for the binary perceptron. The technique is a "rigorous cavity method", in which martingale techniques are used to show concentration of the log-partition function as rows of the matrix are added. We discuss this method further in Section 5.

Sun and Nakajima [28] recently extended the method of Talagrand to give a sharp threshold *sequence* for wide class of random matrices, yielding comparable results to [2, 29] in a broader setting, notably including the binary perceptron. (A sharp threshold sequence is a weaker notion of a sharp threshold in which the capacity concentrates around some $\alpha_c(n)$ for each n, but $(\alpha_c(n))_{n=1}^\infty$ may not converge to some value that is independent of n.)

Algorithms   The foundational result of combinatorial discrepancy is Spencer's "Six standard deviations suffice" [37], which asserts for *any* square matrix A with entries in $[-n^{-1/2}, n^{1/2}]$, deterministically $\mathrm{disc}(A) \leq 6$. Finding an efficient algorithm that achieves Spencer's bound took almost 25 years. While a number of such algorithms were discovered in the last decade, and while many of these easily generalize beyond bounded entries to our setting of iid random entries (e.g. [11]), they are all far from achieving the optimal constant K_c. There are intimate connections between average-case complexity and the geometry of the "solution space" for optimization problems that heavily motivated the recent work on sharp thresholds for the perceptron. We refer the interested reader to [6, 2, 1, 29, 20] for further discussion.

Wide matrices   The discrepancy of $\alpha n \times n$ matrices with $\alpha \to 0$ has applications in a surprising range of combinatorial optimization problems. The case of matrices with iid Bernoulli(p) entries is of particular interest due to practical reasons as well as connections with the long-standing Beck-Fiala conjecture. The second moment method readily gives a sharp threshold for Gaussian matrices [40]. However for Bernoulli matrices, there are fundamental obstacles to analyzing the second moment if $p := p(n)$ vanishes quickly. A sharp characterization of discrepancy was first established for large values of p [23, 31, 8, 17, 18, 25], and eventually for all p [5].

3   Preliminaries

3.1   First moment

Let us adopt the language of random Constraint Satisfaction Problems (CSP's). For an (α, n)-Gaussian matrix A, define the sublevel set Z by:

$$Z_{K,\alpha} := \{x \in \{\pm 1\}^n : \|Ax\|_\infty \leq K\},$$

so that

$$\mathrm{disc}(A) = \inf\{K : Z_{K,\alpha} \neq \emptyset\}.$$

Then $Z_{K,\alpha}$ is the set of solutions to the "discrepancy instance" with parameters K and α given by the matrix A. The matrix A encodes a set of constraints: for each row a_i, a solution x needs to satisfy $|\langle x, a_i\rangle| \leq K$. It is natural to suspect that for fixed α, the probability of $Z_{K,\alpha}$ being empty undergoes a rapid transition from zero to one as K increases.

A standard first approach to understanding the value of an optimization problem is the first moment method. If $\mathbb{E}[|Z_{K,\alpha}|]$ is vanishing, then by Markov's inequality there are no solutions with high probability. This provides a lower bound on $\mathrm{disc}(A)$. For an (α, n)-Gaussian matrix A, it is an elementary computation that for W a standard normal,

$$\frac{1}{n}\log\mathbb{E}[|Z_{K,\alpha}|] = \log(2) + \alpha\log p_K,$$

$$p_K := \mathbb{P}[|W| \leq K] = \frac{1}{\sqrt{2\pi}}\int_{-K}^{K} e^{-w^2/2}dw.$$

For fixed K, we define $\alpha_c(K)$ the "critical value" of α so that the expected number of solutions is one:

$$\alpha_c(K) := -\frac{\log(2)}{\log p_K}. \tag{3.1}$$

Conversely, for fixed α, we define $K_c(\alpha)$ as the unique solution to

$$\alpha_c(K_c) = \alpha.$$

We end this section by showing that for α even slightly greater than α_c—or equivalently, K slightly smaller than K_c—the expected size of $Z_{K,\alpha}$ is exponentially smaller than one. Applying Markov's inequality yields an upper-bound on the capacity of the symmetric perceptron, completing the elementary half of Theorem 1.2.

Proposition 3.1. *Consider an arbitrary fixed $\alpha > 0$ and let $K_c := K_c(\alpha)$. There exist positive constants ϵ, c, and C depending only on α so that for all $K > 0$ with $|K - K_c| < \epsilon$, we have:*

$$c \leq \frac{\frac{1}{n}\log\mathbb{E}[|Z_{K,\alpha}|]}{K - K_c} \leq C. \tag{3.2}$$

Proof. Let $\alpha_c := \alpha_c(K)$. Recall that $\mathbb{E}[|Z_{K,\alpha_c}|] = 1$ by construction of α_c, so that

$$\mathbb{E}[|Z_{K,\alpha}|] = (p_K)^{n(\alpha-\alpha_c)}, \quad \text{i.e.} \quad \frac{1}{n}\log\mathbb{E}[|Z_{K,\alpha}|] = (\log p_K)(\alpha - \alpha_c). \tag{3.3}$$

For fixed $\alpha > 0$, K_c is bounded away from 0. For ϵ small enough, K is bounded from 0 also. It follows that p_K and p_{K_c} are bounded from zero as well, so that $\log p_K =: c_K$ and the first two derivatives in K of $\log p_K$ are bounded in absolute value by some constants depending only on α. (We have eliminated dependence on ϵ by recall that ϵ is a constant depending only on α). Thus, we have by Taylor's expansion:

$$\alpha - \alpha_c = \log(2)\left(\frac{1}{\log p_K} - \frac{1}{\log p_{K_c(\alpha)}}\right) = c(K_c - K) + \mathcal{O}(|K_c - K|^2). \tag{3.4}$$

Combining with Eq. (3.3) yields the result. $\square$

We will repeatedly use Eq. (3.4) to justify the idea that studying α close to α_c and K close to K_c are equivalent. It will be convenient extract a formal statement for later use.

Lemma 3.2 (Equivalence of perturbations to α and K)**.** *Let $\alpha := \alpha(n)$ and $K := K(n)$ be two sequences strictly bounded away from zero. There exist positive constants ϵ, c, and C so that if*

$$\limsup_{n\to\infty}\max(|K - K_c(\alpha)|, |\alpha - \alpha_c(K)|) < \epsilon,$$

then for all n sufficiently large:

$$c \leq \frac{\alpha_c(K) - \alpha}{K - K_c(\alpha)} \leq C.$$

The following are immediate consequences of applying Markov's inequality to Proposition 3.1 and then invoking Eq. (3.3).

Lemma 3.3 (First moment method). *Fix $K > 0$ and let $\alpha_c := \alpha_c(K)$. For any $y > 0$,*

$$\mathbb{P}\left[\boldsymbol{\alpha}_* n \geq \alpha_c n + y\right] \leq (p_K)^y. \tag{3.5}$$

Alternatively, fix $\alpha > 0$ and let $K_c := K_c(\alpha)$. There exist positive constants c and ϵ so that for all $y \in [0, \epsilon]$,

$$\mathbb{P}\left[(\mathrm{disc}(A) - K_c)_- > y\right] < e^{-cny}. \tag{3.6}$$

Integrating the tail bound Eq. (3.6) easily yields a one-sided variance bound:

Corollary 3.4.

$$\mathbb{E}\left[(\mathrm{disc}(A) - K_c)_-^2\right]^{1/2} = \mathcal{O}\left(n^{-1}\right).$$

Proof of Corollary 3.4. Integrating the tail bound Eq. (3.6) for $y \in [0, \epsilon]$ and noting that trivially $(\mathrm{disc}(A) - K_c)_- \leq K_c$ deterministically,

$$\mathbb{E}\left[(\mathrm{disc}(A) - K_c)_-^2\right] \leq \int_0^\epsilon 2y e^{-cny} dy + (K_c)^2 e^{-cn\epsilon} \leq \mathcal{O}\left(n^{-2}\right). \qquad \square$$

In summary, the first moment method applied to $|Z_{K,\alpha}|$ shows $\mathrm{disc}(A)$ cannot be much smaller than K_c. Far more difficult is showing $\mathrm{disc}(A)$ is not much larger than K_c.

3.2 Second moment

We turn to the second moment method for an upper bound on $\mathrm{disc}(A)$. Fix α and let $Z_K := Z_{K,\alpha}$. If the variance of $|Z_K|$ is small, then $\mathrm{disc}(A)$ will concentrate around the first value K for which $\mathbb{E}\left[|Z_K|\right] \geq 1$.

Aubin, Perkins, and Zdeborová [6] established (originally contingent on a numerical hypothesis, later removed by [2])

Theorem 3.1 (Coarse threshold; [6]). *Fix any $\alpha > 0$ and let $K_c := K_c(\alpha)$, $Z_K := Z_{K,\alpha}$. There exist positive constants $1 < C_1 < C_2 < \infty$ so that for any $\epsilon > 0$, for all n sufficiently large,*

$$C_1 \leq \frac{\mathbb{E}\left[|Z_{K_c+\epsilon}|^2\right]}{\mathbb{E}\left[|Z_{K_c+\epsilon}|\right]^2} \leq C_2.$$

An upper bound on $\mathrm{disc}(A)$ immediately follows. By the Paley-Zygmund inequality, for any constants $\alpha > 0$ and $\epsilon > 0$,

$$\liminf_{n \to \infty} \mathbb{P}\left[\mathrm{disc}(A) < K_c + \epsilon\right] > 0.$$

However, since $C_1 > 1$, Theorem 3.1 also establishes that the variance of $|Z_K|$ is too large to directly show that $\mathrm{disc}(A) < K_c + \epsilon$ with high probability. Of course, $|Z_K|$ having large variance certainly does not imply that $\mathrm{disc}(A)$ has large variance. A sharp threshold for $\mathrm{disc}(A)$ only requires that the random variable $\mathbb{1}\left(|Z_K| > 0\right)$ is well-concentrated.

Moving beyond the failure of the second moment method to establish a sharp threshold is often quite difficult, and indeed the results of Abbe, Li, and Sly [2], and Perkins and Xu [29] required significant new ideas. Our work builds on [29], which in turn builds on an idea of Talagrand [38]. Let us begin by formally stating the result of [29] that we wish to refine:

Theorem 3.2 ([29], Theorem 9)**.** *Fix $K > 0$ and $\epsilon > 0$. Let $Z_\alpha := Z_{K,\alpha}$ and $\alpha_c := \alpha_c(K)$. For $\alpha \leq \alpha_c - \epsilon$ and for every $\delta > 0$, there exists $M = M(\delta, \epsilon, K)$ so that*

$$\limsup_{n \to \infty} \mathbb{P}\left[\left|\log\left(\frac{|Z_\alpha|}{\mathbb{E}[|Z_\alpha|]}\right)\right| \geq M \log n\right] \leq \delta. \qquad (3.7)$$

Similarly to Theorem 3.1, the original statement of Theorem 3.2 was originally contingent on a numerical conjecture (Eq. (3.10) below) that can be removed by the results of [2]. Recall

$$\mathbb{E}[|Z_\alpha|] = (p_K)^{n(\alpha_c - \alpha)},$$

so that $\mathbb{E}[|Z_\alpha|] = \exp\{\Omega(\epsilon n)\}$ is exponentially large. In contrast, Theorem 3.2 yields that the fluctuations of $\log|Z_\alpha|$ are at most polynomial, establishing that $|Z_\alpha|$ is non-empty with high probability. The shortcoming we seek to improve is that Theorem 3.2 only holds for $|\alpha - \alpha_c|$ of constant order. In order to capture the critical window of the SBP, we need to show the bound of Theorem 3.2 actually continues to hold even as $|\alpha - \alpha_c|$ approaches order $1/n$. We begin with the important first step of showing that this improvement holds with positive probability.

Theorem 3.3 (Second moment method)**.** *Fix positive $K > 0$. There exist positive constants $1 < C_1 < C_2 < \infty$, so that for n sufficiently large and any $\alpha := \alpha(n) \leq \alpha_c(K)$,*

$$C_1 < \frac{\mathbb{E}[|Z_{K,\alpha}|^2]}{\mathbb{E}[|Z_{K,\alpha}|]^2} < C_2. \qquad (3.8)$$

Consequently, for A an (α, n)-Gaussian matrix (see Definition 1.1),

$$\mathbb{P}[\mathrm{disc}(A) \leq K_c(\alpha)] > C_2^{-1}. \qquad (3.9)$$

Eq. (3.8) will be used as an important step in our strengthening of Theorem 3.2, and Eq. (3.9) will be the main ingredient for our lower bound on the variance of $\mathrm{disc}(A)$. We prove Theorem 3.3 in the next section.

In order to compute the second moment of $|Z_{K,\alpha}|$, we study a quantity we will call the *second moment profile function* (sometimes called in physics the "annealed free energy of pairs"). The second moment profile function is a measure of how much pairs of corners (σ, τ) from the discrete hypercube contribute to the second moment $\mathbb{E}\left[|Z_{K,\alpha}|^2\right]$, as a function of the normalized Hamming distance $n^{-1}\sum_{i=1}^n \mathbb{1}(\sigma_i \neq \tau_i)$.

Definition 3.5 (Pair probability)**.** *For X and Y independent standard normals, let*

$$q_K(\beta) := \mathbb{P}\left[\left|\sqrt{\beta}X + \sqrt{1-\beta}Y\right| \leq K, \ \left|\sqrt{\beta}X - \sqrt{1-\beta}Y\right| \leq K\right].$$

Then $q_K(\beta)^{\alpha n}$ is the probability that σ and τ with $\langle \sigma, \tau \rangle = (2\beta - 1)n$ are both in $Z_{K,\alpha}$.

Definition 3.6 (Second moment profile function)**.** *Let H be binary entropy; define the second moment profile function $F : [0, 1] \to \mathbb{R}$ by*

$$F(\beta) := F_{K,\alpha}(\beta) = H(\beta) + \alpha \log q_K(\beta).$$

Let us collect some trivial facts about F for intuition. First, Gaussian random variables which are orthogonal are also independent, so $q(1/2) = p^2$ for any K and α. Second, the functions $q(\beta)$ and $H(\beta)$ are symmetric around $\beta = 1/2$, so F is as well. Finally, by construction of α_c, for any K and n,

$$F_{K,\alpha_c(K),n}\left(\frac{1}{2}\right) = F_{K,\alpha_c(K)}(0) = -\log(2).$$

Regarding the shape of F, the following was conjectured and assumed in [6] and [29] (as well as an earlier version of [2]) in order for second moment method computations to be tractable:
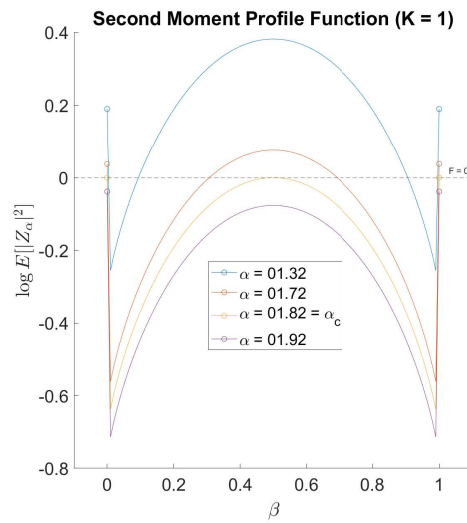
Figure 1: The shape of the second moment profile function $F(\beta)$ for $K = 1$ and various α. (An identical plot has previously appeared in Figure 1a of [6]). Note that $F(\beta)$ is symmetric around $\beta = 1/2$; $F(1/2) = 2F(0) = 2F(1)$; and most importantly, F has a unique maximum at $\beta = 1/2$ for $\alpha \leq \alpha_c$.

Conjecture 3.7 (Shape of F; conjecture). *For all positive K and α with $F''_{K,\alpha}(1/2) < 0$, it further holds that $F := F_{K,\alpha}$ has a single critical point in $(0, 1/2)$. In particular, for any $0 \leq a \leq b \leq 1/2$,*

$$\max_{a \leq \beta \leq b} F(\beta) \in \max\{F(a),\ F(b)\}. \tag{3.10}$$

Compelling numerical evidence was supplied in [6]. Later, in [2], a slightly weaker form of this conjecture was recovered, for $K > 0$ and $\alpha < \alpha_c$. The main difference is that for some K, it is only easy to rigorously establish that F is decreasing near 0 and increasing near $1/2$. In between, it is far easier to simply show that F is much less than $F(1/2)$—rather than controlling the number of critical points—which is more than enough for second moment method applications. We now collect the facts, rigorously established in [2], that we will need:

Lemma 3.8 (Shape of F [2]). *Recall $F := F_{K,\alpha}$.*

1. *For any α and K*

$$F''\left(\frac{1}{2}\right) = 4\left(-1 + \frac{2}{\pi}\frac{\alpha K^2 e^{-K^2}}{p^2}\right). \tag{3.11}$$

2. *Fix any $K > 0$. There exists $\epsilon := \epsilon(K) > 0$ sufficiently small so that for any $\alpha \leq \alpha_c(K)$,*

$$F''\left(\frac{1}{2}\right) < -\epsilon. \tag{3.12}$$

3. *There exists $b := b(K) > 0$ such that for all $\alpha \leq \alpha_c$, $F_\alpha(\beta)$ is decreasing for $\beta \in [0, b]$.*

4. *Let $K > 0$ and $F_{\alpha_c} := F_{K,\alpha_c(K)}$. There exists $\epsilon := \epsilon(K) > 0$ such that, for any $x, y \in [0, 1/2]$,*

$$\max_{\beta \in [x,y]} F_{\alpha_c}(\beta) \leq \max\{F_{\alpha_c}(x),\ F_{\alpha_c}(y),\ F_{\alpha_c}(1/2) - \epsilon\}. \tag{3.13}$$

The proof of Lemma 3.8 is deferred to the appendix, since it is directly reproduced from [2] with only trivial modifications. Upon first read, one could also simply assume Eq. (3.10).

3.3 Sharp threshold

We seek to boost the "positive probability" guarantee of Theorem 3.3 into a "high probability" guarantee. Our goal is: if $\log\left(\mathbb{E}\left[|Z_{K,\alpha}|\right]\right) \gg \log n$, then $|Z_r| \gg 1$ with high probability. Making this precise and also giving a quantitative tail bound, we will show:

Theorem 3.4 (Upper tail). *Fix $K > 0$. There exists a constant $c > 0$ so that if*

$$\alpha := \alpha(n) = \alpha_c(K) - xn^{-1}\log(n)$$

for a sufficiently large constant $x > 0$, then for all n sufficiently large:

$$\mathbb{P}\left[|Z_{K,\alpha}| = 0\right] < n^{-cx}. \tag{3.14}$$

Equivalently, for any $\alpha > 0$ there exists some constant $c > 0$ so that for any fixed x sufficiently large, for all n sufficiently large:

$$\mathbb{P}\left[\operatorname{disc}(A) > K_c + \frac{x\log(n)}{n}\right] < n^{-cx}. \tag{3.15}$$

It is fairly straightforward to control the variance of $\operatorname{disc}(A)$ by integrating this tail bound. The only (small) complication is that our tail bound does not allow for x growing as a function of n.

Corollary 3.9. *Fix $\alpha > 0$ and let n be sufficiently large. We have the one-sided variance bound*

$$\mathbb{E}\left[(\operatorname{disc}(A) - K_c)_+^2\right]^{1/2} = \mathcal{O}\left(\frac{\log(n)}{n}\right). \tag{3.16}$$

Proof of Corollary 3.9. A crude bound suffices. We can always upper-bound $\operatorname{disc}(A)$ by $\|A\sigma\|_\infty$ for some arbitrary σ. Recall that $A\sigma$ is a centered Gaussian vector with an $\alpha n \times \alpha n$ identity covariance matrix. It is easy to check by a direct computation the very sub-optimal bound: $\mathbb{E}\left[\|A\sigma\|_\infty^4\right] < n^4$. Define

$$U := (\operatorname{disc}(A) - K_c)_+.$$

Then for sufficiently large $x > 0$, we have by Cauchy-Schwarz:

$$\mathbb{E}\left[U^2\right]^{1/2} = \mathbb{E}\left[U^2 \mathbb{1}(U \leq xn^{-1}\log n)\right]^{1/2} + \mathbb{E}\left[U^2 \mathbb{1}(xn^{-1}\log n \leq U)\right]^{1/2}$$

$$\leq \frac{x\log(n)}{n} + \mathbb{E}\left[\|A\sigma\|_\infty^4\right]^{1/4} \mathbb{P}\left[\operatorname{disc}(A) > K_c + xn^{-1}\log n\right]^{1/4}$$

$$\leq \frac{x\log(n)}{n} + n\left(n^{-cx/4}\right)$$

$$= \mathcal{O}\left(\frac{\log(n)}{n}\right). \qquad \square$$

4 Second moment method

Here we establish that solutions exist up to criticality with positive probability.

Proof of Theorem 3.3. Fix $K > 0$ and some sequence $\alpha := \alpha(n) \leq \alpha_c$. Our goal is to upper bound the ratio

$$\frac{\mathbb{E}\left[|Z_{K,\alpha}|^2\right]}{\mathbb{E}\left[|Z_{K,\alpha}|\right]^2} = 2^{-n} \sum_{s=0}^{n} \binom{n}{s} \left(\frac{q\left(\frac{s}{n}\right)}{p^2}\right)^{\alpha n}. \tag{4.1}$$

We claim that Eq. (4.1) is monotone increasing in α, so that if we establish the claim for α set to equal α_c, then the claim also follows for all $\alpha \leq \alpha_c$. Indeed, for any β, it easy to

check that $p^2 \leq q(\beta)$ by a direct computation. Alternatively, if the reader is familiar with such tools, the Gaussian Correlation Inequality (e.g. Theorem 1 of [36]) also immediately yields $p^2 \leq q(\beta)$. Indeed, $q(\beta)$ is the probability that the Gaussian vector A_i is in the intersection of two origin-symmetric convex slabs:

$$\{a : |\langle a, \sigma \rangle| \leq K\} \cap \{a : |\langle a, \tau \rangle| \leq K\},$$

where σ and τ are corners of the discrete cube with Hamming distance βn.

So, it suffices to upper bound Eq. (4.1) for $\alpha = \alpha_c$. Note that for this choice of α, the denominator of the left of Eq. (4.1) is simply one. For convenience, we adopt the shorthand $F := F_{K, \alpha_c(K), n}$.

Let $m := \alpha_c n$ be the number of rows in the matrix A. Consider $\beta_s := \frac{1}{2} + \frac{s}{2n}$. We would like to Taylor expand $q(\beta_s)$; first note that $q(0) = p^2$. Next,

$$q'(\beta) = \frac{1}{2\pi} \partial_\beta \int_{x=-K}^{K} \int_{y=\frac{-K-(2\beta-1)x}{2\sqrt{\beta(1-\beta)}}}^{\frac{K-(2\beta-1)x}{2\sqrt{\beta(1-\beta)}}} \exp\left\{ -\frac{x^2+y^2}{2} \right\} dy dx$$

$$= \frac{1}{\pi\sqrt{\beta(1-\beta)}} e^{-\frac{K^2}{2(1-\beta)}} \left(e^{\frac{(2\beta-1)K^2}{2\beta(1-\beta)}} - 1 \right).$$

Note that $q'(1/2) = 0$. Differentiating again and then evaluating at $\beta = 1/2$ yields:

$$q''(1/2) = \frac{8e^{-K^2}K^2}{\pi} = 4p^2 \left(1 - \mu_2\right)^2,$$

where we have adopted the following notation of [2]:

$$\mu_2 := p^{-1} \int_{-K}^{K} x^2 e^{-x^2/2} \frac{dx}{\sqrt{2\pi}} = -p^{-1} \sqrt{\frac{2}{\pi}} K e^{-K^2/2} + 1. \tag{4.2}$$

Thus, restricted to $|\beta_s| < \epsilon$ for some small constant $\epsilon > 0$ depending only on K, Taylor's expansion yields

$$q(\beta_s) = p^2 \left(1 + \frac{(1-\mu_2)^2(s^2/n)}{2n} + \mathcal{O}\left(\frac{s^4}{n^4}\right) \right). \tag{4.3}$$

We are ready to bound Eq. (4.1). Split the sum into three annular regions and treat them separately. (For the first region, highly similar computations appear in e.g. Lemma 3.1 in [2]; Lemma 7 of [6]; or Lemma 3 of [5].)

1. $(1/2 - \delta \leq \beta \leq 1/2)$ Recall $\beta_s := 1/2 + s/(2n)$. Let $\delta > 0$ be some arbitrarily small constant that we fix later. For δ sufficiently small, we apply Stirling's formula and Eq. (4.3). Denote for shorthand $B := \sqrt{\alpha}(1 - \mu_2)/2$. Counting over even-valued s,

$$I_1 := \sum_{|s| \leq 2\delta n} 2^{-n} \binom{n}{\frac{n+s}{2}} \left(\frac{q(\beta_s)}{p^2} \right)^m$$

$$= \sqrt{\frac{2}{n\pi}} \sum_{|s| \leq 2\delta n} \exp\left\{ -\frac{s^2}{2n}\left(1 + \mathcal{O}(\delta)\right) \right\} \exp\left\{ \frac{\alpha}{2}(1-\mu_2)^2 \left(\frac{s^2}{n}\right) \left(1 + \mathcal{O}(\delta)\right) \right\}$$

$$= \sqrt{\frac{2}{n\pi}} \sum_{|s| \leq 2\delta n} \exp\left\{ -\frac{s^2}{2n}\left(1 - 4B^2\right)\left(1 + \mathcal{O}(\delta)\right) \right\} e^{-2B^2}.$$

Since δ can be taken arbitrarily small, it suffices to show $|B|$ is strictly less than $1/2$, uniformly for sufficiently large n. This follows immediately by combining Eq. (3.11) with Eq. (4.2). Indeed, by Eq. (3.12), there is some $\epsilon := \epsilon(K) > 0$ so that

$$1 - 4B^2 = -\frac{1}{4}F''\left(\frac{1}{2}\right) > \epsilon > 0.$$

Thus $|B|$ is strictly bounded from $1/2$, yielding

$$I_1 = (1 + \mathcal{O}(\delta)) e^{-2B^2} \sqrt{\frac{1}{1-4B^2}}. \tag{4.4}$$

(The factor of 2 from $\sqrt{2/(n\pi)}$ was cancelled by the fact that we were only summing over even s). Taking δ sufficiently small, we obtain for some implicit constant depending only on K that $I_1 = \mathcal{O}(1)$.

2. ($\delta \leq \beta \leq 1/2 - \delta$) Let I_2 be

$$I_2 := \sum_{s:\ 2\delta n \leq |s| < (1-2\delta)n} 2^{-n} \binom{n}{\frac{n+s}{2}} \left(\frac{q(\beta_s)}{p^2}\right)^m.$$

Applying Stirling's formula,

$$I_2 < n(2^{-n}p^{-2m}) \max_{\beta \in [\delta, \frac{1}{2}-\delta]} \exp\{nF(\beta) + \mathcal{O}(\log n)\}$$
$$= \max_{\beta \in [\delta, \frac{1}{2}-\delta]} \exp\{n(F(\beta) - F(1/2)) + \mathcal{O}(\log n)\}.$$

It suffices to show the exponent is strictly negative and of order n on this interval. Applying Eq. (3.13), we have for some $\epsilon_0 := \epsilon_0(K, \delta) > 0$,

$$\max_{\beta \in [\delta, 1/2-\delta]} F(\beta) - F(1/2) \leq \max\{F(\delta) - F(1/2),\ F(1/2-\delta) - F(1/2),\ -\epsilon\} \leq -\epsilon_0.$$

Thus for some implicit constant depending only on K and δ, $I_2 = \exp\{-|\Theta(n)|\}$.

3. ($0 \leq \beta < \delta$) Define I_3, the contribution of the remaining terms, by

$$I_3 := \sum_{|s| > (1-2\delta)n} 2^{-n} \binom{n}{\frac{n+s}{2}} \left(\frac{q(\beta_s)}{p^2}\right)^m.$$

By Stirling's formula—and recalling we set α to be $\alpha_c(K)$, where α_c is given in Eq. (3.1)—we also have the identity

$$I_3 = \sum_{|s| > (1-2\delta)n} p^{-m} \exp\{nF(\beta_s) + \mathcal{O}(\log n)\}. \tag{4.5}$$

Since we are able to take δ an arbitrarily small constant in all the previous parts of this proof, let δ be sufficiently small so that by Lemma 3.8, F is decreasing on $[0, \delta]$. The idea is to say $F(\beta)$ is already extremely negative even for $\beta = 1/n$. Intuitively, such a dramatic decay is possible due to the "frozen" nature of typical solutions or, equivalently, the fact that $F'(\beta) \to -\infty$ as $\beta \to 0$ from above. Note the minor subtlety that F being decreasing does not directly imply that the summand of I_3 is decreasing, due to the logarithmic error terms in Stirling's formula. Conveniently, $F(1/n)$ is negative enough to outweigh these lower-order corrections.

First computing the summand for $\beta = 0$, since $q(0) = p$, we have:

$$2^{-n} \binom{n}{0} \left(\frac{q(0)}{p^2}\right)^m = 2^{-n}p^{-m} = 1.$$

Next, in order to compute the summand for $\beta = 1/n$, we claim: (cf. Lemma 6 of [29] for a related result on the "planted" model)

$$\frac{q(1/n)}{q(0)} = \frac{q(1/n)}{p} \leq 1 - \Omega\left(\frac{1}{\sqrt{n}}\right). \tag{4.6}$$

Indeed, let $\beta := 1/n$ and define $\Delta := c_0 n^{-1/2}$ for some new positive constant c_0 that we fix to be sufficiently small later. We have by definition of q:

$$q(\beta) = \frac{1}{2\pi} \int_{-K}^{K} e^{-y^2/2} \int_{\frac{-K+(1-2\beta)y}{2\sqrt{\beta(1-\beta)}}}^{\frac{K+(1-2\beta)y}{2\sqrt{\beta(1-\beta)}}} \exp\left\{-x^2/2\right\} dx dy$$

$$\leq \frac{1}{\sqrt{2\pi}} \int_{-K}^{K-\Delta} e^{-y^2/2} dy + \frac{1}{2\pi} \int_{K-\Delta}^{K} e^{-y^2/2} \int_{\frac{-K+(1-2\beta)y}{2\sqrt{\beta(1-\beta)}}}^{\frac{K+(1-2\beta)y}{2\sqrt{\beta(1-\beta)}}} \exp\left\{-x^2/2\right\} dx dy$$

$$\leq p_K - \Delta \left(\frac{e^{-K^2/2}}{\sqrt{2\pi}}\right) + \frac{1}{2\pi} \int_{K-\Delta}^{K} e^{-y^2/2} \int_{\frac{-K+(1-2\beta)(K-\Delta)}{2\sqrt{\beta(1-\beta)}}}^{\infty} \exp\left\{-x^2/2\right\} dx dy.$$

Examining the bounds of integration,

$$\frac{-K + (1-2\beta)(K-\Delta)}{2\sqrt{\beta(1-\beta)}} = \frac{-2K\beta - \Delta}{2\sqrt{1/n}} \left(1 + \mathcal{O}\left(\frac{1}{n}\right)\right) = -\frac{\Delta\sqrt{n}}{2} \left(1 + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right)\right).$$

Thus, we obtain in total:

$$q\left(\frac{1}{n}\right)$$

$$\leq p_K - \Delta \left(\frac{e^{-K^2/2}}{\sqrt{2\pi}}\right) + \frac{1 + \mathcal{O}\left(n^{-1}\right)}{\sqrt{2\pi}} \int_{K-\Delta}^{K} e^{-y^2/2} \left(\frac{1}{2} + \int_{-\frac{c_0}{2}}^{0} \exp\left\{-x^2/2\right\} dx\right) dy$$

$$= p_K + \Delta \left(-\frac{e^{-K^2/2}}{\sqrt{2\pi}} + \frac{e^{-K^2/2}}{2\sqrt{2\pi}} + \Theta\left(c_0 e^{-K^2/2}\right)\right) + \mathcal{O}\left(\frac{1}{n}\right)$$

$$\leq p_K - \Theta\left(\frac{1}{\sqrt{n}}\right).$$

The last line follows by taking c_0 sufficiently small. We have thus established the desired bound Eq. (4.6). Raising this inequality to the m'th power yields

$$q(1/n)^m \leq p^m \exp\left\{-\Omega\left(\sqrt{n}\right)\right\},$$

which is far smaller than even what is needed. Indeed, returning to Eq. (4.5) and recalling that F is strictly decreasing, as well as noting crudely that $H(1/n) < n^{-.9}$, we obtain:

$$I_3 = \sum_{s=0}^{\delta n} p^{-m} \exp\left\{nF(\beta_s) + \mathcal{O}\left(\log n\right)\right\}$$

$$\leq 2 \left(1 + \delta n \left(\frac{q(1/n)}{p}\right)^m \exp\left\{n^{.1} + \mathcal{O}\left(\log n\right)\right\}\right)$$

$$\leq 2 \left(1 + \exp\left\{-\Omega\left(\sqrt{n}\right)\right\}\right).$$

Thus, $I_3 = 2 + \mathrm{o}(1)$.

This completes our casework. Summing I_1, I_2, and I_3, we have shown that the second moment ratio (4.1) is $\mathcal{O}(1)$. By our previous considerations on the monotonicity of Eq. (4.1) with respect to α, Eq. (3.8) is then established for all $\alpha := \alpha(n) \leq \alpha_c$. By the Paley-Zygmund inequality, Eq. (3.9) follows immediately and the lemma is complete. $\square$

5 Sharp threshold

5.1 Sketch of Theorem 3.4

By Lemma 3.2 ("equivalence of perturbations to α_c and K_c"), Eq. (3.14) implies Eq. (3.15). In order to establish Eq. (3.14), we need to refine Theorem 3.2 of Perkins and Xu to hold even for α allowed to vary with n. In particular, it should hold for $\alpha \leq \alpha_0 := \alpha_c - xn^{-1}\log(n)$. We also would like quantitative tail-bounds. Let us briefly survey the proof.

We follow an approach that Talagrand developed for the binary perceptron [38]. Talagrand's approach was carefully adapted to the setting of the symmetric perceptron by Perkins and Xu [29]. The outline is to build a constraint satisfaction problem by adding one clause at a time, and then show the logarithm of the number of solutions concentrates via martingale arguments.

Reusing the notation of [29] for clarity of comparison, define a time-indexed process $(S_t)_t$ of the solutions to the first t rows:

$$S_t := \{\sigma \in \{\pm 1\}^n : |(A\sigma)_j| \leq K, \ \forall \ 1 \leq j \leq t\}, \quad \mathbb{E}\left[|S_t|\right] = 2^n(p_K)^t. \tag{5.1}$$

In terms of our previous notation, S_t is equal to $Z_{K,\alpha}$ with $\alpha = t/n$. Denote the deviation of $\log|S_t|$ from its mean by Q_t, which we write as a telescoping sum:

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

Rewriting Q_t in terms of a martingale that we call Y_t,

$$Q_t = \sum_{i=1}^{t}\log(1 + Y_i), \quad Y_t := \frac{1}{p}\left(\frac{|S_t|}{|S_{t-1}|} - p\right).$$

Then, as long as Y_i is small for each $i \leq \alpha_0 n$, we can use the Taylor Expansion

$$Q_t = \sum_{i=1}^{t} Y_i - \frac{Y_i^2}{2} + \mathcal{O}\left(Y_i^3\right).$$

The Y_i are centered and should be roughly $Y_i \asymp n^{-1/2}$, so that both the sum of the Y_i and the sum of the Y_i^2 are constant order, and the sum of the Y_i^3 is vanishing. We then expect $|Q_t| = \mathcal{O}\left(\log(n)\right)$ with exponential tails, which would yield the theorem. Indeed, $\log\mathbb{E}\left[|S_{\alpha_0 n}|\right]$ is at least $\Omega\left(x\log(n)\right)$, so if we take x sufficiently large, then $|Q_t| \ll \log\mathbb{E}\left[|S_t|\right]$ for all $t \leq \alpha_0 n$ with very high probability. This would imply that the fluctuations of $|S_t|$ are much less than its mean, yielding the result. More precisely, we would have with high probability:

$$\begin{aligned}
\log|S_{\alpha_0 n}| &= \log\mathbb{E}\left[|S_{\alpha_0 n}|\right] + \log\left(\frac{|S_{\alpha_0 n}|}{\mathbb{E}\left[|S_{\alpha_0 n}|\right]}\right) \\
&= \log\mathbb{E}\left[|S_{\alpha_0 n}|\right] + Q_{\alpha_0 n} \\
&\geq cx\log n - C\log n \\
&\gg 0.
\end{aligned}$$

The last line is achieved by taking x sufficiently large, completing the theorem.

The analysis is by induction. The key idea for this induction is a geometric notion of regularity for S_t that asserts solutions are sufficiently "well-spread". We begin by assuming for induction that $S_{t'}$ is *regular* for all $t' \leq t$ (precise definition given shortly).

Then we follow with three estimates (cf. Lemmas 11, 12, and 13 of [29] respectively). First, if S_t is regular, then Y_{t+1} enjoys good tail bounds. Second, if $Y_{t'}$ has good tails for all $t' \leq t + 1$, then martingale concentration yields exponential tail bounds for Q_{t+1}. Third, if Q_{t+1} is small, then S_{t+1} is also regular with good probability.

We emphasize that this structure was innovated in [38] and refined in [29]. We will directly use the first estimate (tail bound on Y_t) from [29]. For the second and third estimates, we also significantly borrow from the structure of the proofs in [29], albeit with exponential improvements. This concludes our sketch.

5.2 Proof of Theorem 3.4 (cf. Theorem 9 of [29])

Let S_t, Q_t, and Y_t be as defined in the sketch. Fix $K > 0$. All other constants in what follows will depend on K; we treat K as a universal constant and suppress this dependence in our notation. Fix x sufficiently large. Define $\alpha_0 := \alpha_0(n)$ by $\alpha_0 n := \alpha_c n - x \log(n)$. We emphasize that no constants in what follows will depend implicitly on x. Recall our convention that c and C denote generic constants that may vary between lines; important constants will have a subscript.

In order to use the martingale structure of Y_t, define the filtration $\mathcal{F}_t$ generated by revealing the first t rows of the matrix A. Define the conditional measure $\mathbb{P}_t[\cdot]$ by

$$\mathbb{P}_t[\cdot] := \mathbb{P}[\cdot \mid \mathcal{F}_t]. \tag{5.2}$$

With this notation, we are ready to formally state our main definitions and lemmas.

Definition 5.1 (Regular). *We say $S_t \subset \{\pm 1\}^n$ is **regular** if $S_t \neq \emptyset$ and, for $\sigma_t^{(1)}$, $\sigma_t^{(2)}$ drawn uniformly and independently (with replacement) from S_t,*

$$\mathbb{P}_t\left[\left|\left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle\right| > \sqrt{C_r n(x \log(n) + |Q_t|)}\right] \leq n^{-c_r x},$$

where C_r and c_r are some positive universal constants.

The parameters C_r and c_r are explicitly chosen below in Eq. (5.16) and Eq. (5.21) respectively. Define two stopping times: τ_S the first time t that S_t is not regular and τ_Q the first time that $|Q_t|$ is large. Also define the first time τ that either of these "bad" events occurs. More precisely, for $C_q \in (0, 1)$ some small universal constant (explicitly chosen in Eq. (5.14) below),

$$\tau_S := \min\{t : S_t \text{ is not regular}\},$$
$$\tau_Q := \min\{t : |Q_t| > C_q x \log(n)\},$$
$$\tau := \tau_S \wedge \tau_Q.$$

The first estimate of the induction is that if S_t is regular then Y_{t+1} is sub-exponential. Up to some trivial rewriting (we use τ_Q to suppress $|Q|$), the following is exactly Lemma 11 of [29].

Lemma 5.2 (Y_{t+1} is small; [29], Lemma 11). *There exist positive constants c and C so that for any t, for all sufficiently large n and y,*

$$\mathbb{1}(\tau > t) \, \mathbb{P}_t\left[|Y_{t+1}| > Cy\sqrt{\frac{(1 + C_q x) \log n}{n}}\right] \leq \exp\{-cy\}. \tag{5.3}$$

Next, we would like to say that Q_t is small for $t \leq \tau_S$. Only a first-moment bound on $|Q_t|$ is given in [29]; the refinement we require is a bound on the moment-generating function. This will both allow us take t much closer to α_c as well as yield strong tail bounds on $|S_t|$.

Lemma 5.3 (Q_{t+1} is small). *There exists a constant $c > 0$ so that for all sufficiently large n and any $t \leq \alpha_0 n$,*

$$\mathbb{P}\left[|Q_{t \wedge \tau}| > C_q x \log n\right] \leq n^{-cx}. \tag{5.4}$$

Assuming S_t is regular, these lemmas yield that $|Y_{t+1}|$ and $|Q_{t+1}|$ are small. The induction is complete if $|Q_{t+1}|$ being small implies S_{t+1} is regular.

Lemma 5.4 (S_{t+1} is regular; τ is large). *There exists a constant $c > 0$ so that for all sufficiently large n,*

$$\mathbb{P}\left[\tau_S \leq \alpha_0 n\right] < n^{-cx}.$$

We prove these lemmas in the next subsection. Let us first formally check that they imply Theorem 3.4, our desired bound on the upper tail of $\mathrm{disc}(A)$.

Proof of Theorem 3.4. For convenience, define the constant $c_K := -\log(p_K) > 0$. Since K is fixed, c_K may be treated as a universal constant. By Eq. (3.3),

$$\log \mathbb{E}\left[|S_t|\right] = c_K(\alpha_c n - t) \geq c_K x \log n, \quad \forall\, t \leq \alpha_0 n. \tag{5.5}$$

We will eventually choose C_q explicitly (in Eq. (5.14)) so that $C_q \leq c_K/2$. Thus, on the event $\{\tau_Q > t\}$, we have $|S_t| \gg 0$ almost surely. Indeed, as sketched in Section 5.1,

$$\log |S_t| = Q_t + \log \mathbb{E}\left[|S_t|\right] \geq -C_q x \log n + c_K x \log n \geq \frac{c_K}{2} x \log n,$$

where the equality is by definition of Q_t and the first inequality holds almost surely on the event $\{\tau_Q > t\}$.

The theorem will follow if the event $\{\tau_Q > \alpha_0 n\}$ holds with very high probability. By Lemma 5.4, with probability at least $1 - n^{-cx}$, we have $\{\tau_S > \alpha_0 n\}$, i.e. the solution space is regular simultaneously for all $t \leq \alpha_0 n$. Then, by Lemma 5.3,

$$\mathbb{P}\left[\tau_Q \leq \alpha_0 n\right] \leq \mathbb{P}\left[\tau_S \leq \alpha_0 n\right] + \sum_{t=1}^{\alpha_0 n} \mathbb{P}\left[\tau_S > \alpha_0 n, \tau_Q = t\right]$$

$$\leq n^{-cx} + \sum_{t=1}^{\alpha_0 n} \mathbb{P}\left[|Q_{t \wedge \tau}| > C_q x \log n\right]$$

$$\leq n^{-cx}.$$

In summary, we have shown:

$$\mathbb{P}\left[|Z_{K,\,\alpha_c n - x \log n}| = 0\right] \leq \mathbb{P}\left[|S_{\alpha_0 n}| \leq e^{\frac{c_K}{2} x \log n}\right] \leq \mathbb{P}\left[\tau_Q \leq \alpha_0 n\right] \leq n^{-cx}. \qquad \square$$

5.3 Proofs of lemmas

Proof of Lemma 5.3. Fix $t \leq \alpha_0 n$ and define the event E that all the Y_i are bounded for $i \leq t \wedge \tau$:

$$E := \bigcap_{i \leq t-1} \left\{\mathbb{1}_{\tau > i}|Y_{i+1}| < n^{-.4}\right\}.$$

By Lemma 5.2 and union bound, the event E holds with probability at least $1 - n^{-C_E}$ for any (arbitrarily large) constant C_E. Let $\lambda \in (0, 1)$, where λ does not vary with n and will be taken sufficiently small later; our goal is to show

$$\mathbb{E}\left[\exp\left\{\lambda|Q_{t \wedge \tau}|\right\} \mathbb{1}_E\right] \leq n^{\lambda^2 Cx}. \tag{5.6}$$

Let us simplify Eq. (5.6). Since e^{Q_t} is a martingale for all t by construction and $\lambda \in (0, 1)$, the optional stopping theorem and Jensen's inequality yield

$$\mathbb{E}\left[\exp\left\{\lambda Q_{t \wedge \tau}\right\}\right] \leq 1.$$

Thus, to establish Eq. (5.6), it suffices to show:

$$\mathbb{E}\left[\exp\left\{-\lambda Q_{t\wedge\tau}\right\}\mathbb{1}_E\right] \leq n^{\lambda^2 Cx}. \tag{5.7}$$

By a Taylor expansion, we have on the event E that

$$\exp\left\{-\lambda\log(1+Y_i)\right\} \leq 1 - C\lambda Y_i + C\lambda^2 Y_i^2, \quad \forall\, i \leq t\wedge\tau. \tag{5.8}$$

Additionally, by integrating the tail bound of Lemma 5.2, we obtain for some new positive constant $C > 0$:

$$\mathbb{1}_{\tau>i}\mathbb{E}_i\left[Y_{i+1}^2\right] \leq Cn^{-1}(1+C_q x)\log n, \quad \forall\, i. \tag{5.9}$$

Recall that $(Y_i)_i$ is a martingale under the filtration generated by revealing the rows of A. Combining Eq. (5.8) and Eq. (5.9) yields

$$
\begin{aligned}
&\mathbb{E}\left[\exp\left\{-\lambda Q_{t\wedge\tau}\right\}\mathbb{1}_E\right] \\
&= \mathbb{E}\left[\exp\left\{-\lambda Q_{(t-1)\wedge\tau} - \mathbb{1}_{\tau>t-1}\lambda\log(1+Y_t)\right\}\mathbb{1}_E\right] \\
&= \mathbb{E}\left[\exp\left\{-\lambda Q_{(t-1)\wedge\tau}\right\}\mathbb{E}_{t-1}\left[\exp\left\{-\mathbb{1}_{\tau>t-1}\lambda\log(1+Y_t)\right\}\right]\mathbb{1}_E\right] \\
&\leq \mathbb{E}\left[\exp\left\{-\lambda Q_{(t-1)\wedge\tau}\right\}\left(1 - C\mathbb{1}_{\tau>t-1}\lambda\mathbb{E}_{t-1}[Y_t] + C\mathbb{1}_{\tau>t-1}\lambda^2\mathbb{E}_{t-1}\left[Y_t^2\right]\right)\mathbb{1}_E\right] \\
&\leq \mathbb{E}\left[\exp\left\{-\lambda Q_{(t-1)\wedge\tau}\right\}\mathbb{1}_E\left(1 + \lambda^2\frac{C(1+C_q x)\log n)}{n}\right)\right].
\end{aligned}
$$

We have used that the linear term of Y in Eq. (5.8) vanishes in expectation. Since $Q_0 = 0$ by definition, we obtain the desired result (5.7) by an easy induction:

$$\mathbb{E}\left[\exp\left\{-\lambda Q_{t\wedge\tau}\right\}\mathbb{1}_E\right] \leq \left(1 + \lambda^2\frac{C(1+C_q x)\log n}{n}\right)^t \leq n^{\lambda^2 C(1+C_q x)} \leq n^{\lambda^2 Cx},$$

where the last inequality follows for some new constant C, since C_q is an explicit universal constant independent of x, and x is assumed to be sufficiently large. Finally, fix $\lambda = \min\left\{C_q/(2C),\ 1/2\right\}$, so that $\lambda \in (0,1)$ as promised. Then

$$\mathbb{P}\left[|Q_{t\wedge\tau}| > C_q x\log(n)\right] \leq n^{\lambda^2 Cx}n^{-\lambda C_q x} + \mathbb{P}\left[E^c\right] \leq \min\left\{n^{-C_q^2 x/(2C)},\ n^{-C_q x/2}\right\} + n^{-C_E}$$
$$=: n^{-cx}.$$

The last line follows by taking C_E arbitrarily large. $\qquad\square$

Proof of Lemma 5.4. By union bound and the tower property,

$$
\begin{aligned}
\mathbb{P}\left[\tau \leq \alpha_0 n\right] &= \sum_{t=1}^{\alpha_0 n}\mathbb{P}\left[\tau = t\right] \\
&\leq \sum_{t=1}^{\alpha_0 n}\mathbb{P}\left[\tau_S = t,\ \tau_Q > t\right] + \mathbb{P}\left[\tau_Q = t,\ \tau > t-1\right] \\
&\leq n^{-cx} + \sum_{t=1}^{\alpha_0 n}\mathbb{P}\left[\tau_S = t,\ \tau_Q > t\right].
\end{aligned}
$$

The last line follows by Lemma 5.3. It then suffices to show each summand is bounded above by n^{-cx}. Formally, our goal is:

$$\mathbb{P}\left[\tau_S = t,\ \tau_Q > t\right] \leq n^{-cx}, \quad \forall t \leq \alpha_0 n. \tag{5.10}$$

By Lemma 3.8, there exists some constant $c_f > 0$ so that for all n sufficiently large,

$$\max_{\beta\in[0,1/2-\lambda n^{-1/2}]} F_{K,\,t/n,\,n}(\beta) \leq -c_f\min\left\{\lambda^2,\ \mathbb{E}\left[|Z_{K,\,t/n}|\right]\right\}. \tag{5.11}$$

As a convenient short-hand, we write

$$R_t(\lambda) := c_f(\lambda^2 \wedge L_t), \quad L_t := \log\left(\mathbb{E}\left[|Z_{K,t/n}|\right]\right).$$

Then, for two vectors $\sigma_t^{(1)}$ and $\sigma_t^{(2)}$ drawn uniformly and independently from S_t, it can easily be extracted from the proof of Theorem 3.3 that:

$$\mathbb{E}\left[\left|\left\{(\sigma_t^{(1)}, \sigma_t^{(2)}) \in S_t : \left|\left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle\right| \geq \lambda\sqrt{n}\right\}\right|\right] \leq \exp\left\{-R_t(\lambda)\right\} \mathbb{E}\left[|S_t|^2\right].$$

In particular, since we showed in Theorem 3.3 that there exists a positive constant C so that $\mathbb{E}\left[|S_t|^2\right] \leq C\mathbb{E}\left[|S_t|\right]^2$ uniformly for all t,

$$\mathbb{E}\left[\mathbb{P}_t\left[\left|\left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle\right| \geq \lambda\sqrt{n}\right] \frac{|S_t|^2}{\mathbb{E}\left[|S_t|\right]^2}\right] \leq C\exp\left\{-R_t(\lambda)\right\}. \tag{5.12}$$

We would now like to show an exponential tail bound on the overlap of a pair of solutions. Let us give this a short-hand name for convenience:

$$\mathcal{O}^t := \frac{1}{\sqrt{n}}\left|\left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle\right|.$$

Recall that by definition $\exp\{Q_t\} = |S_t|/\mathbb{E}\left[|S_t|\right]$. Then,

$$\mathbb{E}\left[\exp\left\{\frac{R_t(\mathcal{O}^t)}{2} + 2Q_t\right\}\right]$$

$$= \mathbb{E}\left[\exp\left\{\frac{1}{2}\left(\frac{c_f\left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle^2}{n} \wedge L_t\right) + 2Q_t\right\}\right]$$

$$= \mathbb{E}\left[\int_0^\infty \mathbb{P}_t\left[\exp\left\{\frac{1}{2}\left(\frac{c_f\left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle^2}{n} \wedge L_t\right)\right\} \geq y\right] dy \frac{|S_t|^2}{\mathbb{E}\left[|S_t|\right]^2}\right]$$

$$= \mathbb{E}\left[\int_0^{\exp\{L_t/2\}} \mathbb{P}_t\left[\exp\left\{\frac{c_f}{2n}\left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle^2\right\} \geq y\right] dy \frac{|S_t|^2}{\mathbb{E}\left[|S_t|\right]^2}\right].$$

From the second to the third line, we have conditioned on S_t. Applying the change of variables $z := \sqrt{2\log(y)/c_f}$ and then Eq. (5.12),

$$\mathbb{E}\left[\exp\left\{\frac{R_t(\mathcal{O}^t)}{2} + 2Q_t\right\}\right] = \int_0^{\sqrt{L_t/c_f}} c_f z e^{c_f z^2/2} \mathbb{E}\left[\mathbb{P}_t\left[\left|\left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle\right| \geq z\sqrt{n}\right] \frac{|S_t|^2}{\mathbb{E}\left[|S_t|\right]^2}\right] dz$$

$$\leq C\int_0^{\sqrt{L_t/c}} c_f z e^{-c_f z^2/2} dz$$

$$\leq C. \tag{5.13}$$

We would like to apply Markov's inequality using Eq. (5.13). Recall Eq. (5.5), namely that there exists some universal constant $c_K > 0$ with

$$L_t \geq c_K x \log(n), \quad \forall t \leq \alpha_0 n.$$

We have not yet fixed C_q up to this point; here we will need $C_q \ll C$, e.g. $C_q \leq C/10$. For the proof of Theorem 3.4, we also would like $C_q < c_K$. Thus, set:

$$C_q := \min\left\{C/10, \ c_K/2\right\}. \tag{5.14}$$

In order to establish Eq. (5.10), note that if $\{\tau_Q > t\}$, we have

$$|Q_t| < .1L_t \tag{5.15}$$

We also have yet to fix C_r, one of the two parameters in the definition of regular. Set

$$C_r := 2C/c_f \vee 1. \tag{5.16}$$

(We only take the maximum with one for notation convenience later). On the event $\{\tau_Q > t\}$, this choice yields

$$R_t \left(\sqrt{C_r x \log(n) + |Q_t|} \right) + 2Q_t > \min \{(c_f C_r - 2C_q)x \log(n), \ .9c_f L_t\}$$
$$\geq .9 R_t \left(\sqrt{C_r x \log(n)} \right). \tag{5.17}$$

Since $R_t(\lambda)$ is weakly increasing in λ, we have:

$$\mathbb{P}\left[\left| \left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle \right| > \sqrt{n(C_r x \log(n) + |Q_t|)}, \ \tau_Q > t \right]$$
$$\leq \mathbb{P}\left[\exp\left\{ R_t(\mathcal{O}^t) + 2Q_t \right\} > \exp\left\{ R_t\left(\sqrt{C_r x \log(n) + |Q_t|} \right) + 2Q_t \right\}, \ \tau_Q > t \right]. \tag{5.18}$$

By tower property, and then applying Markov's inequality to Eq. (5.18) via Eq. (5.13) and Eq. (5.17),

$$\mathbb{E}\left[\mathbb{1}_{\tau_Q > t} \mathbb{P}_t \left[\left| \left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle \right| > \sqrt{n(C_r x \log(n) + |Q_t|)} \right] \right]$$
$$= \mathbb{P}\left[\left| \left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle \right| > \sqrt{n(C_r x \log(n) + |Q_t|)}, \ \tau_Q > t \right]$$
$$\leq C e^{-.9 R_t \left(\sqrt{C_r x \log(n) + |Q_t|} \right)}.$$

By a final application of Markov's inequality,

$$\mathbb{P}\left[\mathbb{P}_t \left[\left| \left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle \right| > \sqrt{n(C_r x \log(n) + |Q_t|)} \right] > e^{-R_t \left(\sqrt{C_r x \log(n)} \right)/2}, \ \tau_Q > t \right] \tag{5.19}$$
$$\leq \mathcal{O}\left(e^{-\frac{2}{5} R_t \left(\sqrt{C_r x \log(n)} \right)} \right)$$
$$\leq C\left(n^{-cx} + e^{-cL_t} \right). \tag{5.20}$$

In conclusion, since $\mathbb{E}\left[|Z_{k,t}|\right] \geq Cx \log(n)$ for all $t \leq \alpha_0 n$, we may simply upper bound e^{-cL_t} by n^{-cx} and also set the constant c_r from the definition of S_t being regular to

$$c_r := c_f C_r / 2. \tag{5.21}$$

With this notation, Eq. (5.20) implies our goal Eq. (5.10): since $C_r \geq 1$,

$$\mathbb{P}\left[\tau_S = t, \ \tau_Q > t \right]$$
$$:= \mathbb{P}\left[\mathbb{P}_t \left[\left| \left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle \right| > \sqrt{C_r n(x \log(n) + |Q_t|)} \right] > n^{-c_r x}, \ \tau_Q > t \right]$$
$$\leq \mathbb{P}\left[\mathbb{P}_t \left[\left| \left\langle \sigma_t^{(1)}, \sigma_t^{(2)} \right\rangle \right| > \sqrt{n(C_r x \log(n) + |Q_t|)} \right] > e^{-R_t \left(\sqrt{C_q x \log(n)} \right)/2}, \ \tau_Q > t \right]$$
$$\leq n^{-cx}.$$

Taking x sufficiently large immediately yields the lemma. $\qquad \square$

6 Lower bound on fluctuations

Lemma 6.1. *Fix positive α and let $K_c := K_c(\alpha)$, $Z_K := Z_{K,\alpha}$. There exists a positive constant C such that for K with $0 < K \leq 2K_c$, n sufficiently large, and positive $\epsilon > 0$ sufficiently small:*

$$\mathbb{P}\left[K - \frac{\epsilon}{n} \leq \mathrm{disc}(A) \leq K\right] < C\epsilon\, \mathbb{E}\left[|Z_K|\right]. \tag{6.1}$$

As an easily corollary, we obtain a lower bound on the variance of $\mathrm{disc}(A)$, completing our main result Theorem 1.2. In fact, we obtain a lower bound on fluctuations, which is strictly stronger than a bound on the variance.

Definition 6.2 (Fluctuations). *We say that a sequence of random variables X_n has fluctuations at least of order σ_n if there exist constants $\delta > 0$ and $C > 0$ such that, for n sufficiently large and any $|a - b| > C\sigma_n$,*

$$\mathbb{P}\left[a < X_n < b\right] < 1 - \delta.$$

Corollary 6.3. *The fluctuations of $\mathrm{disc}(A)$ are at least order $1/n$.*

Proof of Corollary 6.3. By Theorem 3.3, we know that for some $\delta > 0$

$$\mathbb{P}\left[\mathrm{disc}(A) \leq K_c(\alpha)\right] > \delta > 0.$$

So, for any a and b possibly functions of n, with $K_c \leq a \leq b$, we have uniformly

$$\mathbb{P}\left[\mathrm{disc}(A) \in [a, b]\right] < 1 - \delta.$$

All that remains is the case that $a := a(n)$ and $b := b(n)$ with $a < b$ and $a \leq K_c$. Assume $|a - b| \leq \epsilon/n$ for some sufficiently small constant $\epsilon > 0$ we choose shortly. By definition of K_c as well as the expansion Eq. (3.2), we have for some $C > 0$,

$$\mathbb{E}\left[|Z_b|\right] \leq Ce^{Cn(b - K_c)_+} \leq C\exp\left\{C\epsilon\right\}.$$

By Eq. (6.1), for ϵ sufficiently small,

$$\mathbb{P}\left[\mathrm{disc}(A) \in [a, b]\right] < C\epsilon\exp\left\{C\epsilon\right\} \leq \frac{1}{2}. \qquad \square$$

Proof of Lemma 6.1. Let Ξ be the set of solutions with discrepancy inside an ϵ/n window around K, namely

$$\Xi := \Xi_{K,\epsilon} := \left(Z_{K - \epsilon/n}\right)^c \cap Z_K.$$

We will show Ξ is empty with positive probability by Markov's inequality. Let $T_i(x)$ denote the event that row i is *tight*, namely

$$T_i := \{|(Ax)_i| \in I_{K,\epsilon}\}, \quad I_{K,\epsilon} := [K - \epsilon/n,\ K].$$

Correspondingly, define the number of tight rows T by

$$T(x) := \sum_{i=1}^{\alpha n} \mathbb{1}(T_i(x)).$$

Let x be an arbitrary vector in $\{\pm 1\}^n$. Conditioning on T and using the independence of the rows of A, we have for some sufficiently large $C_0 := C_0(\alpha) > 0$ that

$$\begin{aligned}
\mathbb{E}\left[|\Xi|\right] &= 2^n \sum_{t=0}^{\alpha n} \mathbb{P}\left[x \in \Xi | T(x) = t\right] \mathbb{P}\left[T(x) = t\right] \\
&\leq 2^n \sum_{t=1}^{\alpha n} \binom{\alpha n}{t} \mathbb{P}\left[(Ax)_1 \in I_{K,\epsilon}\right]^t \mathbb{P}\left[|(Ax)_1| < K\right]^{\alpha n - t} \\
&\leq 2^n \sum_{t=1}^{\alpha n} \binom{\alpha n}{t} \left(\frac{C_0\epsilon}{n} p_K\right)^t p_K^{\alpha n - t}.
\end{aligned}$$

where, as usual, $p_K = \mathbb{P}\left[|Z| \leq K\right]$ for Z a standard Gaussian. Note that we are able to let C_0 only depend on α because K is bounded above by a constant, namely $2K_c$. Thus, by the binomial theorem, for some other sufficiently large $C := C(\alpha) > 0$,

$$\mathbb{E}\left[|\Xi|\right] \leq (2^n p_K^{\alpha n}) \sum_{t=1}^{\alpha n} \binom{\alpha n}{t} \left(\frac{C_0 \epsilon}{n}\right)^t 1^{\alpha n - t} = \mathbb{E}\left[|Z_K|\right] \left(\left(1 + \frac{C_0 \epsilon}{n}\right)^{\alpha n} - 1\right) \leq \mathbb{E}\left[|Z_K|\right] C\epsilon.$$

Note that $|\Xi|$ is a non-negative integer, so the events $\{|\Xi| \leq 1/2\}$ and $\{|\Xi| = 0\}$ are almost surely equal. Thus, Markov's inequality applied to $|\Xi|$ immediately yields the theorem. $\square$

Supplementary Material

Proof of Lemma 3.8. All claims follow directly from a combination of lemmas in [2] and [6]. The goal of this proof is simply to organize their contributions. We refer the reader to these two papers for details.

Claims 1 and 2 The first claim is an easy computation. The second claim is proven in [6] (page 7). Indeed, they show that $F''(1/2) < 0$ for any $K > 0$ if $\alpha \leq \alpha_c$. Fixing K, we thus have by monotonicity in α of Eq. (3.11), a uniform upper bound on $F''(1/2)$.

Claim 3 Simply compute $F'(0)$ (see e.g. Eq. 12 of [6]) and note

$$\lim_{\beta \to 0^+} F'(\beta) = -\infty.$$

Since F is differentiable on $(0, 1/2]$ and continuous at 0, there is then some positive radius around 0 for which F is decreasing.

Claim 4 It will suffice to organize the results developed in [2] for the proof of their Lemma 3.4. They treat separately the three cases of: $K > 4$, $K \in [.1, 4]$, and $K \in [0, .1)$ respectively in their Lemmas 4.6; 4.10 and 4.11; and 4.7. We summarize these results. For $K > 4$, they show that F' transitions from negative to positive as β from 0 to $1/2$, with no other sign changes. (We communicate their proof of this below). It then follows for $K > 4$ and $0 \leq a \leq b \leq 1/2$ that

$$\max_{\beta \in [a,b]} F_{\alpha_c}(\beta) \leq \max\{F_{\alpha_c}(a), F_{\alpha_c}(b)\}.$$

In the remaining two cases, while conjecturally the same picture holds, the actual picture that is manageable to prove is slightly different. It is established in [2] only that there are some constants $0 < b_1 \leq b_2 < 1/2$ (depending only on K) so that F is decreasing on $(0, b_1)$, increasing on $(b_2, 1/2)$, and $F_{\alpha_c}(\beta) < F_{\alpha_c}(1/2) - \epsilon$ for some $\epsilon := \epsilon(K) > 0$ in between. This will yield Eq. (3.13), which is enough for our purposes. We will reproduce most of the proof for $K > 4$ to make this more concrete, since this case has the cleanest analysis. Then, we will remark on how b_1 and b_2 are chosen for the remaining cases and try to highlight which parts are easy and difficult for the three cases. We leave many details of this summary—especially the actual implementation of the grid search—for the reader to find in [2]. We also emphasize again that the following arguments and ideas we attempt to outline are not our own, but rather completely from [2].

Fix K and $\alpha = \alpha_c(K)$, and denote for short-hand $F := F_{K, \alpha_c, n}$. By an easy computation,

$$F'(\beta) = -\log(\beta) + \log(1 - \beta) + \alpha_c \frac{-\exp\left\{-\frac{K^2}{2(1-\beta)}\right\} + \exp\left\{-\frac{K^2}{2\beta}\right\}}{\pi q_k(\beta)\sqrt{\beta(1-\beta)}}. \tag{7.1}$$

Rearranging this equation, define another quantity $\mathcal{L}$ that we will aim to bound:

$$\mathcal{L}(\beta) := \frac{\alpha_c}{\pi} \frac{\exp\left\{-\frac{K^2}{2(1-\beta)}\right\} - \exp\left\{-\frac{K^2}{2\beta}\right\}}{\sqrt{\beta(1-\beta)}(\log(1-\beta) - \log(\beta))}. \tag{7.2}$$

Note that $F'(\beta) > 0$ is equivalent to $\mathcal{L}(\beta) < q(\beta)$.

Case 1: $4 < K$ Consider $\beta \in [.2, .5]$. It is easy to check that

$$q_K(\beta) \geq q_4(.2) > .9,$$

so our goal is to show $\mathcal{L}(\beta) < .9$ for all $\beta \in [.2, .5]$. We have

$$\frac{\log(1-\beta) - \log(\beta)}{1/2 - \beta} \geq 4, \tag{7.3}$$

and

$$\frac{d^2}{d\beta^2} \exp\left\{-\frac{K^2}{2\beta}\right\} > 0. \tag{7.4}$$

Combined with Taylor's expansion, Eq. (7.4) yields

$$\frac{\exp\left\{-\frac{K^2}{2(1-\beta)}\right\} - \exp\left\{-\frac{K^2}{2\beta}\right\}}{-(\beta - 1/2)} \leq \frac{K^2 \exp\left\{-\frac{K^2}{2(1-\beta)}\right\}}{(1-\beta)^2}. \tag{7.5}$$

By Eq. (7.3) and Eq. (7.5), we simplify the desired inequality:

$$\mathcal{L}(\beta) \leq \frac{\alpha_c}{\pi} \frac{\exp\left\{-\frac{K^2}{2(1-\beta)}\right\} - \exp\left\{-\frac{K^2}{2\beta}\right\}}{\sqrt{\beta(1-\beta)}(1/2 - \beta)}$$

$$\leq \frac{\alpha_c}{\pi} \frac{K^2 \exp\left\{-\frac{K^2}{2(1-\beta)}\right\}}{\sqrt{\beta(1-\beta)}(1-\beta)^2}.$$

Next, using standard asymptotics for the error function,

$$1 - p_K \geq \exp\left\{-K^2/2\right\} \sqrt{\frac{2}{\pi}} \left(\frac{1}{K} - \frac{1}{K^3}\right),$$

so we have

$$\alpha_c := -\frac{\log(2)}{\log(p)} \leq \frac{\log 2}{1 - p_K} \leq \frac{16}{15} \log(2) \sqrt{\frac{\pi}{2}} \exp\left\{K^2/2\right\}. \tag{7.6}$$

These inequalities suffice for the sub-case that $x \in [.35, .5]$. In total:

$$\mathcal{L}(\beta) \leq \frac{1}{4} \log(2) \sqrt{\frac{\pi}{2}} \exp\left\{K^2/2\right\} \frac{16}{15} \frac{K^2 \exp\left\{-\frac{K^2}{2(1-\beta)}\right\}}{\sqrt{\beta(1-\beta)}(1-\beta)^2} \frac{1}{\sqrt{.2 \cdot .8}}$$

$$\leq \exp\left\{K^2 \left(\frac{1}{2} - \frac{1}{2(1-\beta)}\right)\right\} \frac{K^3}{(1-x)^2} .19$$

$$< .9.$$

In the sub-case that $x \in [.2, .35]$, replace Eq. (7.5) with the trivial inequality

$$\frac{\exp\left\{-\frac{K^2}{2(1-\beta)}\right\} - \exp\left\{-\frac{K^2}{2\beta}\right\}}{-(\beta - 1/2)} \leq \frac{\exp\left\{-\frac{K^2}{2(1-\beta)}\right\}}{-(\beta - 1/2)}.$$

Then an identical computation easily yields $\mathcal{L}(\beta) < .9$. So, for all $\beta \in [.2, .5]$, we have $F'(\beta) > 0$.

Suppose that $\mathcal{L}'(\beta) < q'(\beta)$ for $\beta \le .2$. Then L and q can cross at most once. Note $F' < 0$ for some ball of positive radius around $\beta = 0$ by claim three, so $L > q$ for all β sufficiently small. We have also already shown that $L < q$ for $\beta > .2$. Thus, $\mathcal{L}$ and q cross exactly once, and we have the desired picture: F is decreasing and then increasing, with no other changes. In order to establish $\mathcal{L}'(\beta) < q'(\beta)$ for $\beta \le .2$, one can differentiate the definition of L, and then apply standard tail-bounds on the error function.

Case 2: $.1 \le K \le 4$ The strategy is almost the same. We know that F is decreasing in some neighborhood of 0, so a natural first step is to, again, show that F is increasing in some neighborhood of $1/2$. In the previous case, we were able to establish F is increasing in $[.2, .5]$. Here, [2] were only able to show this for $\beta \in [.3, .5]$. The proof is quite similar, except that, roughly speaking, while bounds derived from first-order Taylor expansion sufficed for $K > 4$, second-order expansions are needed for $K \in [.1, .4]$. Next, we again try and make Claim 3 (namely that F is decreasing near 0) quantitative. By looking very close to 0, namely $\beta \in [0, .005]$, we can actually prove a stronger result with easier computations. Since trivially $p \ge q$, it of course suffices to show $\mathcal{L} > p_K$ for β near 0. The strategy is to then show that $\mathcal{L}(.005) > p_K$ and $\mathcal{L}' < 0$ for $\beta \in [0, .005]$. These are simple computations. Finally, in the remaining region $\beta \in [.005, .3]$, we would like to check that $F(\beta) - F(1/2) < -\epsilon$ for some $\epsilon := \epsilon(K) > 0$. This is done in two steps: first, it is relatively straight-forward to check that the derivative of $F_K(\beta) - F_K(1/2)$ with respect to β and the derivative with respect to K are both bounded in absolute value by e.g. 6. Second, do a computerized grid search over the region

$$\{(\beta, K) : \ \beta \in [.005, .3], \ K \in [.1, 4]\}.$$

Since we have a bound on how fast $F_K(\beta) - F_K(1/2)$ can change, a step-size can be picked appropriately to obtain a provably tolerable error and conclude. In summary, taking b_1 and b_2 as $.005$ and $.3$ respectively, the picture we have established is that F is decreasing on $[0, b_1]$; increasing on $[b_2, 1/2]$; and strictly negative between.

Case 3: $0 \le K \le .1$ Here, K being small will actually make the analysis quite tractable. Indeed, we can for example get decent control over p and q, which will be used repeatedly. Using a zero-order and first-order Taylor expansion of $\exp\left\{-x^2/2\right\}$ for $x \in [-K, K]$ for the upper and lower bound respectively, we have trivially

$$\frac{.992K}{\sqrt{2\pi}} \le p_K \le \frac{2K}{\sqrt{2\pi}}. \tag{7.7}$$

One can also bound $q_K(\beta)$ in a similarly crude way:

$$\frac{K^2}{\pi\sqrt{\beta(1-\beta)}} \exp\left\{-\frac{K^2}{2\beta}\right\} \le q_K(\beta) \le \frac{K^2}{\pi\sqrt{\beta(1-\beta)}}. \tag{7.8}$$

Note that since K is small, the gap between the upper and lower bounds here is very small if β is not too small as well. With these bounds, we are now ready to prove the desired picture. First, we claim for $\beta \in [.27, .5]$ that $\mathcal{L}(\beta) \le q(\beta)$. The main difference from the first case ($K > 4$) is that now

$$\frac{d^2}{d\beta^2} \exp\left\{-\frac{K^2}{2\beta}\right\} < 0, \tag{7.9}$$

so we replace Eq. (7.5) by

$$\frac{\exp\left\{-\frac{K^2}{2(1-\beta)}\right\} - \exp\left\{-\frac{K^2}{2\beta}\right\}}{-(\beta - 1/2)} \leq \frac{K^2 \exp\left\{-\frac{K^2}{2\beta}\right\}}{\beta^2}.$$

Next, for $\beta \in (0, K^2/2)$, we need to verify that $\mathcal{L}(\beta) > q(\beta)$. Indeed, since $\sqrt{\beta}\log(\beta)$ is decreasing for $\beta \leq .005$ (and here we are only considering $\beta \leq K^2/2 \leq .005$),

$$\begin{aligned}
-\log(p)\mathcal{L}(\beta) &:= \frac{\log 2}{\pi} \frac{\exp\left\{-\frac{K^2}{2(1-\beta)}\right\} - \exp\left\{-\frac{K^2}{2\beta}\right\}}{\sqrt{\beta(1-\beta)}(\log(1-\beta) - \log(\beta))} \\
&\geq \frac{\log 2}{\pi} \frac{1/2}{\sqrt{\beta}(-\log(\beta))} \\
&\geq .29,
\end{aligned}$$

where the last line follows from monotonicity of the penultimate formula in β, and plugging in $\beta = .005$. From here, it remains to check that $-\log(p_K)q(\beta) > .29$, so that F is decreasing on $\beta \in [0, K^2/2]$. This is readily verified by applying Eq. (7.8) to bound q; and then crudely bounding $-\log p_K \geq -\log p_{.1}$ using Eq. (7.7).

Carrying out the Taylor expansions used in Eq. (7.8) and Eq. (7.7) to another term, the same argument shows that F is actually decreasing on $[0, K/12]$ and then increasing on $[.04, 1/2]$. Finally, for the remaining region of $[K/12, .04]$, we can show directly that $F(\beta) < F(1/2) - \epsilon$ for some $\epsilon := \epsilon(K) > 0$. Indeed, recalling that $F_{\alpha_c(K)}(1/2) = \log(2) + 2\alpha_c(K)\log(p)$, the desired claim follows quickly from the inequality

$$\log\left(\frac{q(\beta)}{p^2}\right) < \log\left(\frac{1}{.99^2 2\sqrt{\beta}}\right) < \log\left(\frac{2}{\sqrt{K}}\right).$$

In summary, for b_0 and b_1 equal to $K/12$ and $.04$ respectively, we have again established the picture that F is decreasing on $[0, b_1]$; increasing on $[b_2, 1/2]$; and strictly negative between. This concludes the analysis of the final case $K \leq .01$, and thus the outline of claim four. $\qquad\square$

References

[1] Emmanuel Abbe, Shuangping Li, and Allan Sly, *Binary perceptron: efficient algorithms can find solutions in a rare well-connected cluster*, in: STOC'21—Proceedings of the 2021 ACM Symposium on Theory of Computing, 2021. MR4490046

[2] Emmanuel Abbe, Shuangping Li, and Allan Sly, *Proof of the contiguity conjecture and lognormal limit for the symmetric perceptron*, in: FOCS'21—IEEE 62nd Annual Symposium on Foundations of Computer Science, 2021. MR4399694

[3] Dimitris Achlioptas and Ehud Friedgut, *A sharp threshold for k-colorability*, Random Struct. Algorithms (1999), no. 1, 63–70. MR1662274

[4] Michael Aizenman, Joel L. Lebowitz, and David Ruelle, *Some rigorous results on the Sherrington-Kirkpatrick spin glass model*, Comm. Math. Phys. **112** (1987), no. 1, 3–20. MR904135

[5] Dylan J. Altschuler and Jonathan Niles-Weed, *The discrepancy of random rectangular matrices*, Random Struct. Algorithms (2021). MR4429837

[6] Benjamin Aubin, Will Perkins, and Lenka Zdeborová, *Storage capacity in symmetric binary perceptrons*, J. Phys. A **52** (2019), no. 29, 294003, 32. MR3983947

[7] Debapratim Banerjee, *Contiguity and non-reconstruction results for planted partition models: the dense case*, Electron. J. Probab. **23** (2018), Paper No. 18, 28. MR3771755

[8] Nikhil Bansal and Raghu Meka, *On the discrepancy of random low degree set systems*, Random Struct. Algorithms **57** (2020), no. 3, 695–705. MR4144080

[9] Béla Bollobás, Christian Borgs, Jennifer T. Chayes, Jeong Han Kim, and David B. Wilson, *The scaling window of the 2-SAT transition*, Random Struct. Algorithms **18** (2001), no. 3, 201–256. MR1824274

[10] Erwin Bolthausen, Shuta Nakajima, Nike Sun, and Changji Xu, *Gardner formula for Ising perceptron models at small densities*, in: Conference on Learning Theory, 2022, pp. 1787–1911.

[11] Karthekeyan Chandrasekaran and Santosh S. Vempala, *Integer feasibility of random polytopes*, in: ITCS'14—Proceedings of the 2014 Conference on Innovations in Theoretical Computer Science, ACM, New York, 2014, pp. 449–458. MR3359497

[12] Jian Ding, Allan Sly, and Nike Sun, *Proof of the satisfiability conjecture for large k [extended abstract]*, in: STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing, ACM, New York, 2015, pp. 59–68. MR3388183

[13] Jian Ding, Allan Sly, and Nike Sun, *Maximum independent sets on random regular graphs*, Acta Math. **217** (2016), no. 2, 263–340. MR3689942

[14] Jian Ding, Allan Sly, and Nike Sun, *Satisfiability threshold for random regular NAE-SAT*, Comm. Math. Phys. **341** (2016), no. 2, 435–489. MR3440193

[15] Jian Ding, Allan Sly, and Nike Sun, *Proof of the satisfiability conjecture for large k*, Ann. of Math. (2) **196** (2022), no. 1, 1–388. MR4429261

[16] Jian Ding and Nike Sun, *Capacity lower bound for the Ising perceptron*, in: STOC'19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, ACM, New York, 2019, pp. 816–827. MR4003386

[17] Esther Ezra and Shachar Lovett, *On the Beck-Fiala conjecture for random set systems*, Random Struct. Algorithms **54** (2019), no. 4, 665–675. MR3957362

[18] Cole Franks and Michael Saks, *On the discrepancy of random matrices with many columns*, Random Struct. Algorithms **57** (2020), no. 1, 64–96. MR4120593

[19] Ehud Friedgut, *Sharp thresholds of graph properties, and the k-sat problem*, J. Amer. Math. Soc. **12** (1999), no. 4, 1017–1054. With an appendix by Jean Bourgain. MR1678031

[20] David Gamarnik, Eren C Kızıldağ, Will Perkins, and Changji Xu, *Algorithms and barriers in the symmetric binary perceptron model*, in: FOCS'22—IEEE 62nd Annual Symposium on Foundations of Computer Science, 2022.

[21] E. Gardner and B. Derrida, *Optimal storage properties of neural network models*, J. Phys. A **21** (1988), no. 1, 271–284. MR939731

[22] Hamed Hatami, *A structure theorem for Boolean functions with small total influences*, Ann. of Math. (2) **176** (2012), no. 1, 509–533. MR2925389

[23] Rebecca Hoberg and Thomas Rothvoss, *A Fourier-analytic approach for the discrepancy of random set systems*, in: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2019. MR3909625

[24] Jiaoyang Huang, *Invertibility of adjacency matrices for random d-regular graphs*, Duke Math. J. **170** (2021), no. 18, 3977–4032. MR4348232

[25] Calum MacRury, Tomáš Masařík, Leilani Pai, and Xavier Pérez-Giménez, *The phase transition of discrepancy in random hypergraphs*, SIAM Journal on Discrete Mathematics **37** (2023), no. 3, 1818–1841. MR4628929

[26] András Mészáros, *The distribution of sandpile groups of random regular graphs*, Trans. Amer. Math. Soc. **373** (2020), no. 9, 6529–6594. MR4155185

[27] Marc Mézard, *The space of interactions in neural networks: Gardner's computation with the cavity method*, J. Phys. A **22** (1989), no. 12, 2181–2190. Special issue in memory of Elizabeth Gardner, 1957–1988. MR1004920

[28] Shuta Nakajima and Nike Sun, *Sharp threshold sequence and universality for ising perceptron models*, in: Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms, 2023. MR4538160

[29] Will Perkins and Changji Xu, *Frozen 1-rsb structure of the symmetric ising perceptron*, in: STOC'21—Proceedings of the Annual ACM SIGACT Symposium on Theory of Computing, ACM, New York, 2021. MR4398942

[30] Boris Pittel and Gregory B. Sorkin, *The satisfiability threshold for k-XORSAT*, Combin. Probab. Comput. **25** (2016), no. 2, 236–268. MR3455676

[31] Aditya Potukuchi, *Discrepancy in random hypergraph models*, ArXiv preprint arXiv:1811.01491 (2018).

[32] R. W. Robinson and N. C. Wormald, *Almost all regular graphs are Hamiltonian*, Random Struct. Algorithms **5** (1994), no. 2, 363–374. MR1262985

[33] Frank Rosenblatt, *The perceptron: a probabilistic model for information storage and organization in the brain*, Psychological Review **65** (1958), no. 6, 386. MR0122606

[34] Mark Rudelson and Roman Vershynin, *Non-asymptotic theory of random matrices: extreme singular values*, in: Proceedings of the International Congress of Mathematicians. Volume III, Hindustan Book Agency, New Delhi, 2010, pp. 1576–1602. MR2827856

[35] Ashwin Sah and Mehtaab Sawhney, *Distribution of the threshold for the symmetric perceptron*, in: FOCS'23—IEEE 64nd Annual Symposium on Foundations of Computer Science, 2023.

[36] Zbyněk Šidák, *On multivariate normal probabilities of rectangles: Their dependence on correlations*, Ann. Math. Statist. (1968). MR230403

[37] Joel Spencer, *Six standard deviations suffice*, Trans. Amer. Math. Soc. **289** (1985), no. 2, 679–706. MR784009

[38] Michel Talagrand, *Intersecting random half cubes*, Random Structures Algorithms **15** (1999), no. 3-4, 436–449, Statistical physics methods in discrete probability, combinatorics, and theoretical computer science (Princeton, NJ, 1997). MR1716771

[39] Michel Talagrand, *Intersecting random half-spaces: toward the Gardner-Derrida formula*, Ann. Probab. **28** (2000), no. 2, 725–758. MR1782273

[40] Paxton Turner, Raghu Meka, and Philippe Rigollet, *Balancing Gaussian vectors in high dimension*, in: Conference on Learning Theory, 2020, pp. 3455–3486.

[41] Changji Xu, *Sharp threshold for the Ising perceptron model*, Ann. Probab. **49** (2021), no. 5, 2399–2415. MR4317708