

Sharp Thresholds for the Overlap Gap Property: Ising p -Spin Glass and Random k -SAT

Eren C. Kızıldağ ✉ 

Department of Statistics, University of Illinois Urbana-Champaign, IL, USA

Abstract

The Ising p -spin glass and random k -SAT are two canonical examples of disordered systems that play a central role in understanding the link between geometric features of optimization landscapes and computational tractability. Both models exhibit hard regimes where all known polynomial-time algorithms fail and possess the multi Overlap Gap Property (m -OGP), an intricate geometrical property that rigorously rules out a broad class of algorithms exhibiting input stability.

We establish that, in both models, the symmetric m -OGP undergoes a sharp phase transition, and we pinpoint its exact threshold. For the Ising p -spin glass, our results hold for all sufficiently large p ; for the random k -SAT, they apply to all k growing mildly with the number of Boolean variables. Notably, our findings yield qualitative insights into the power of OGP-based arguments. A particular consequence for the Ising p -spin glass is that the strength of the m -OGP in establishing algorithmic hardness grows without bound as m increases.

These are the first sharp threshold results for the m -OGP. Our analysis hinges on a judicious application of the second moment method, enhanced by concentration. While a direct second moment calculation fails, we overcome this via a refined approach that leverages an argument of Frieze [33] and exploiting concentration properties of carefully constructed random variables.

2012 ACM Subject Classification Theory of computation → Randomness, geometry and discrete structures; Mathematics of computing → Combinatorial optimization

Keywords and phrases spin glasses, p -spin model, random constraint satisfaction problems, overlap gap property, phase transitions, computational complexity

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2025.48

Category RANDOM

Related Version *Full Version*: <https://arxiv.org/abs/2309.09913> [54]

Acknowledgements I am grateful to David Gamarnik for his guidance and many helpful discussions during the early stages of this work. I would like to thank Brice Huang for carefully proofreading the draft and offering insightful comments, Aukosh Jagannath for numerous valuable discussions regarding spin glasses, Lutz Warnke for a helpful discussion on the bounded differences inequality, and anonymous referees for their valuable feedback.

1 Introduction

In this paper, we focus on two fundamental models in the statistical mechanics of disordered systems: Ising p -spin glass and random k -SAT. Both models have been extensively studied across probability, theoretical computer science, and physics; see [58, 71] and references therein. Sharing profound structural similarities, including rich phase transitions and rugged landscapes, they are central to understanding the interplay between the geometrical properties of optimization landscapes and computational phase transitions in average-case settings. Our particular focus is on their inherent algorithmic barriers.

Both models exhibit regimes where no polynomial-time algorithms are known to succeed. For the Ising p -spin glass, which involves optimizing random polynomials, existing algorithms find strictly suboptimal solutions. For the random k -SAT, a random constraint satisfaction



© Eren C. Kızıldağ;

licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025).

Editors: Alina Ene and Eshan Chattopadhyay; Article No. 48; pp. 48:1–48:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

problem (CSP), known algorithms operate only at constraint densities far below the satisfiability threshold. Due to their average-case nature, classical complexity theory offers little insight into these barriers. An extensive body of work has instead sought unconditional lower bounds, showing that both models exhibit the multi Overlap Gap Property (m -OGP) – an intricate geometrical feature of the optimization landscape that rigorously rules out a broad class of algorithms [41, 17, 37, 15, 46, 36, 34]. These works have focused almost exclusively on establishing such lower bounds via the OGP and exploring their tightness – i.e., whether they match known algorithmic guarantees – often by analyzing increasingly refined versions of the OGP as m grows. For fixed $m \in \mathbb{N}$, a canonical result asserts that there exists a threshold $\gamma(m)$ such that when $\gamma > \gamma(m)$, the corresponding model exhibits m -OGP and certain algorithms necessarily fail; here, γ denotes a natural model parameter. However, a finer, two-sided understanding of the link between OGP and hardness remains elusive. In some models, OGP-based lower bounds match provable average-case hardness [79]. At the same time, certain optimization problems exhibit the OGP, yet still admit efficient algorithms [55]. These examples collectively highlight the need for a more nuanced understanding of OGP-based methods – one that addresses not only how OGP rules out algorithms, but also how its power evolves with m and what its absence implies.

To this end, a natural question arises: does the OGP indeed become more powerful as m increases? This is directly tied to understanding the regime where m -OGP is absent, as elaborated below. Typically, the presence of OGP is established via the first moment method, with the threshold $\gamma(m)$ often improving with m . Consequently, a common rule of thumb in the field is that increasing m leads to stronger algorithmic lower bounds.¹ However, this approach offers no insight into the regime $\gamma < \gamma(m)$. To *truly* establish that the power of m -OGP grows with m , one must precisely locate its threshold – by showing that m -OGP is present for $\gamma > \gamma(m)$ and absent for $\gamma < \gamma(m)$. Without this, the observed improvement in $\gamma(m)$ could be spurious, merely an artifact of the first moment method. Such a sharp analysis would also rigorously justify the prevailing practice of increasing m to tighten lower bounds. This issue is particularly stark in the binary perceptron, where the first moment method fails: in certain regimes, the predicted m -OGP threshold actually worsens with m , and the very presence of OGP remains unknown for $m \geq 4$ [39, Section II.B]. Moreover, the absence of OGP is not merely of theoretical interest: for the Sherrington-Kirkpatrick (SK) model, algorithms are known to succeed under the widely believed assumption that OGP is absent [61]. These examples underscore that understanding both the presence and absence of OGP – especially in the regime $\gamma < \gamma(m)$ – is essential to fully grasp its power and limitations.

In this paper, we pinpoint the precise threshold for the symmetric m -OGP – a widely used variant of the OGP (see below) – for both models, thereby demonstrating a sharp phase transition. These thresholds are particularly transparent and strictly monotonic in m , in the large- p and large- k regimes. Our results yield qualitative insights into the power of OGP-based arguments. Notably, for the p -spin model, we establish that the power of m -OGP in proving algorithmic hardness increases indefinitely with m . In the context of spin glasses and shattering, our results also explain why a certain *soft-OGP* argument by El Alaoui [26] is necessary; see Section 1.3.

For an overview of our results, see Section 1.3. We next formally introduce the two models we study.

¹ A similar monotonicity in m has also been conjectured in the context of spherical spin glasses [14].

1.1 Ising Pure p -Spin Glass Model

For a fixed integer $p \geq 2$, an order- p tensor $\mathbf{J} = (J_{i_1, \dots, i_p} : 1 \leq i_1, \dots, i_p \leq n)$ with i.i.d. standard normal entries, $J_{i_1, \dots, i_p} \sim \mathcal{N}(0, 1)$, and $\boldsymbol{\sigma} \in \mathbb{R}^n$, the p -spin Hamiltonian is:

$$H_{n,p}(\boldsymbol{\sigma}) = n^{-\frac{p+1}{2}} \langle \mathbf{J}, \boldsymbol{\sigma}^{\otimes p} \rangle = n^{-\frac{p+1}{2}} \sum_{1 \leq i_1, \dots, i_p \leq n} J_{i_1, \dots, i_p} \sigma_{i_1} \cdots \sigma_{i_p}. \quad (1)$$

Introduced by Derrida [24], the p -spin model generalizes the SK model ($p = 2$) [70] by incorporating p -body interactions. Our particular focus is on the *Ising* case, where the configuration space is the discrete hypercube, $\boldsymbol{\sigma} \in \Sigma_n := \{-1, 1\}^n$.² A key quantity regarding spin glass models is the limiting *free energy* $F(\beta) := \lim_{n \rightarrow \infty} \ln Z_\beta / n$, where $Z_\beta = \sum_{\boldsymbol{\sigma} \in \Sigma_n} \exp(\beta H_{n,p}(\boldsymbol{\sigma}))$ at *inverse temperature* $\beta > 0$. The limiting free energy for the SK model was predicted in a celebrated work of Parisi [67] using the non-rigorous replica method. Following the work of Guerra [42], Talagrand [75] rigorously verified Parisi's prediction in a breakthrough paper; his proof in fact extends to all even p . For general p , Parisi formula was verified by Panchenko [65], following the work Aizenman, Sims, and Starr [9]. The existence of the free energy limit was shown earlier by Guerra and Toninelli [43] using the so-called interpolation method. Far more is known rigorously regarding the p -spin glass models, including the ultrametricity of the asymptotic Gibbs measure [63, 48, 16], the TAP equations [12] and the TAP free energy [72, 18], a certain phase diagram [78, 10, 49], shattering [14, 29, 36], and algorithmic results (see below). Furthermore, more general spin glass models have also been studied in the literature. A particularly relevant case is the *vector spin models* where the states are tuples of points in $\mathbb{R}^n$ (as opposed to points in $\mathbb{R}^n$) with a prescribed overlap, see [66, 52]. For further pointers to relevant literature, see the textbooks [76, 77, 64].

Finding a Near-Ground State in Polynomial-Time. Define the *ground-state energy* by $H^* := \lim_{n \rightarrow \infty} \max_{\boldsymbol{\sigma} \in \Sigma_n} H_{n,p}(\boldsymbol{\sigma})$, which can be recovered as the zero temperature limit of the Parisi formula, $H^* = \lim_{\beta \rightarrow \infty} F(\beta) / \beta$.³ Given H^* , an algorithmic question is to find a near ground-state efficiently. That is, given $\mathbf{J} \in (\mathbb{R}^n)^{\otimes p}$ and $\epsilon > 0$, find in polynomial time a $\boldsymbol{\sigma}_{\text{ALG}} \in \Sigma_n$ such that $H_{n,p}(\boldsymbol{\sigma}_{\text{ALG}}) \geq (1 - \epsilon)H^*$, ideally with high probability (w.h.p.). For the SK model, this problem was solved by Montanari [61], who developed an Approximate Message Passing (AMP) type algorithm. Montanari's algorithm was inspired by an algorithm of Subag [73] for spherical mixed p -spin model; it is contingent on a widely believed conjecture that the SK model does not exhibit the OGP. More recently, [27, 69] extended Montanari's algorithm to Ising mixed p -spin models. For any $\epsilon > 0$, these algorithms also return a $\boldsymbol{\sigma}_{\text{ALG}}$ for which $H_{n,p}(\boldsymbol{\sigma}_{\text{ALG}}) \geq (1 - \epsilon)H^*$ w.h.p., provided the underlying mixed p -spin model does not exhibit the OGP. It is known, however, that the Ising p -spin glass model exhibits the OGP for $p \geq 4$ [17, Theorem 3], and that the OGP serves a rigorous barrier to AMP type algorithms [35]. Specifically, for any even $p \geq 4$, there exists a value $\bar{\mu}$ such that no AMP type algorithm can find a $\boldsymbol{\sigma}_{\text{ALG}}$ with $H_{n,p}(\boldsymbol{\sigma}_{\text{ALG}}) \geq H^* - \bar{\mu}$ [35]. This algorithmic lower bound was later extended to low-degree polynomials [37] and low-depth Boolean circuits [38]. However, classical OGP [17] and the aforementioned lower bounds do not match the best known algorithmic threshold. More recently, Huang and Sellke [46, 47] introduced a very elegant *branching OGP* consisting of an ultrametric tree of solutions and subsequently obtained tight bounds for Lipschitz algorithms.

² The *spherical* case where $\boldsymbol{\sigma}$ lies in the hypersphere of radius $\sqrt{n}$, $\|\boldsymbol{\sigma}\|_2 = \sqrt{n}$, is also studied extensively.

³ See also [11] for a generalized Parisi formula for H^* that does not require $\beta \rightarrow \infty$ limit.

1.2 Random k -SAT Model

Given Boolean variables $x_1, \dots, x_n$, a k -clause $\mathcal{C}$ is defined as a disjunction of k literals $\{y_1, \dots, y_k\}$ chosen from $\{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$, i.e., $\mathcal{C} = y_1 \vee \dots \vee y_k$. A random k -SAT formula Φ is a conjunction of M k -clauses $\mathcal{C}_1, \dots, \mathcal{C}_M$, with each of its kM literals sampled independently and uniformly from $\{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$: $\Phi = \mathcal{C}_1 \wedge \dots \wedge \mathcal{C}_M$.⁴ Throughout, $M = \Theta(n)$ and $n \rightarrow \infty$, where we refer to $\alpha := M/n$ as the *clause density*.

Arguably the most fundamental question about random k -SAT is *satisfiability*: when does a satisfying assignment, i.e., a $\sigma \in \{0, 1\}^n$ with $\Phi(\sigma) = 1$, exist? For fixed $k \in \mathbb{N}$, Franco and Paull [31] showed, using a simple first moment argument, that a random k -SAT formula is w.h.p. unsatisfiable if $\alpha \geq 2^k \ln 2$. This was later refined by Kirousis et al. [51], who established a sharper bound $\alpha \geq 2^k \ln 2 - \frac{1}{2}(\ln 2 + 1) + o_k(1)$, where the $o_k(1) \rightarrow 0$ as $k \rightarrow \infty$. On the algorithmic side, Chvátal and Reed [19] showed that a simple algorithm known as unit clause w.h.p. finds a satisfying assignment if $\alpha < 2^k/k$; see also [60], who showed that the same algorithm succeeds with constant probability. The asymptotic gap between $2^k/k$ and $\Omega_k(2^k)$ was substantially narrowed by Achlioptas and Moore [6]. Using a non-constructive second moment argument, they showed that a random k -SAT formula is w.h.p. satisfiable if $\alpha \leq 2^{k-1} \ln 2 - O_k(1)$. Coja-Oghlan and Panagiotou [23] later refined this bound to $\alpha \leq 2^k \ln 2 - \frac{1}{2}(\ln 2 + 1) + o_k(1)$, nearly matching the lower bound of [51] up to $o_k(1)$ terms. A major breakthrough by Ding, Sly, and Sun [25] confirmed, for large k , that the satisfiability threshold is the value predicted by Mézard, Parisi and Zecchina [59]: for any large k , there exists a threshold $\alpha_*(k)$ such that a random k -SAT formula is satisfiable for $\alpha < \alpha_*(k)$ and unsatisfiable for $\alpha > \alpha_*(k)$, both w.h.p. as $n \rightarrow \infty$. For a detailed discussion on the random k -SAT, see the introduction of [15] and the survey by [3].

Finding a Satisfying Assignment Efficiently. Given the existence of satisfying assignments for $\alpha = O_k(2^k)$, a natural algorithmic question arises: can we find such an assignment efficiently? As mentioned earlier, the simple unit clause algorithm finds a satisfying assignment in polynomial time for $\alpha \leq 2^k/k$. Later, Coja-Oghlan [20] devised a polynomial-time algorithm that succeeds for $\alpha \leq (1 + o_k(1))2^k \ln k/k$, beyond which no efficient algorithm is known. Interestingly, the threshold $(1 + o_k(1))2^k \ln k/k$ also marks the onset of *clustering* in the solution space of random k -SAT, as predicted by Krzakala et al. [53] and rigorously verified by Achlioptas and Coja-Oghlan [4]. Beyond this threshold, the solution space of random k -SAT fragments into exponentially many well-separated sub-dominant clusters, which are $\Omega(n)$ apart, each containing only an exponentially small fraction of solutions. This phenomenon led to the conjecture that $(1 + o_k(1))2^k \ln k/k$ is the algorithmic threshold for the random k -SAT.⁵ Specifically, it is conjectured that no polynomial-time algorithm can find (w.h.p.) a satisfying assignment above this threshold. While this conjecture remains open, lower bounds against particular classes of algorithms have been obtained. Using the symmetric m -OGP, Gamarnik and Sudan [41] showed that *sequential local algorithms* with a moderately growing number of message passing iterations fail to find a satisfying assignment for a variant of the random k -SAT model known as the random Not-All-Equal-

⁴ This model allows clauses to contain repeated literals, or both a variable x_i and its negation $\bar{x}_i$. However, we focus on the regime $k = \Omega(\ln n)$, where such occurrences are rare as $n \rightarrow \infty$. This is standard in the literature [32, 21, 15].

⁵ It is worth noting that clustering alone does not necessarily imply algorithmic hardness. For instance, the binary perceptron model exhibits an extreme form of clustering at all densities, including the “easy” regime where polynomial-time algorithms exist; see [68, 1, 2, 39, 40].

k -SAT (NAE- k -SAT) for $\alpha \geq 2^{k-1} \log^2 k/k$. Hetterich [45] proved that *Survey Propagation Guided Decimation* algorithm fails for $\alpha \geq (1 + o_k(1))2^k \log k/k$ even with an unbounded number of iterations. Coja-Oglan, Haqshenas, and Hetterich [22] showed that *Walksat* stalls for $\alpha \geq C \cdot 2^k \log^2 k/k$, for some absolute constant $C > 0$. More recently, Bresler and Huang [15] demonstrated that *low-degree polynomials* fail to find a satisfying assignment for the random k -SAT when $\alpha \geq \kappa^*(1 + o_k(1))2^k \ln k/k$, where $\kappa^* \approx 4.911$. Their approach leverages a novel asymmetric variant of the m -OGP.

Significance of Symmetric m -OGP. Our focus is on the symmetric m -OGP, which plays a central role in establishing unconditional algorithmic lower bounds. It yields the sharpest known bounds in central models such as the binary perceptron [39, 56] and discrepancy minimization [40], as well as nearly sharp bounds for NAE- k -SAT, a closely related variant of random k -SAT [41]. In Ising p -spin glasses, symmetric m -OGP is asymptotically tight as p grows: its onset matches the algorithmic threshold of the Random Energy Model, the formal $p \rightarrow \infty$ limit of the Ising p -spin glasses [24, 74, 8, 36]. For fixed p , the sharpest lower bounds follow from a very elegant branching OGP argument [46]; however, this applies only to overlap-concentrated algorithms – a strict subclass of stable algorithms that excludes, for example, low-degree polynomials. In contrast, symmetric m -OGP applies to this broader class, albeit yielding weaker bounds. Moreover, branching OGP yields lower bounds only for even p , whereas symmetric m -OGP applies to both even and odd p . In fact, for odd p , it provides the only known algorithmic lower bounds for Ising p -spin glasses [36]. Further underscoring its relevance, the case $m = 2$, or pairwise OGP, is directly linked to the shattering phenomenon – a notion from spin glass theory that has shaped our understanding of algorithmic phase transitions in random CSPs [7, 4, 5], see below. By contrast, such a connection is less clear for more sophisticated OGP variants. Taken together, these observations underscore the significance of symmetric m -OGP and highlight the need to fully understand its power.

For $m \in \mathbb{N}$, $\gamma > \gamma_{p\text{-spin}}(m) := 1/\sqrt{m}$, and sufficiently large p , Ising p -spin glasses exhibit symmetric m -OGP above the threshold $\gamma\sqrt{2 \ln 2}$ [36, Theorem 2.11].⁶ no tuples of equidistant solutions at a prescribed distance exist beyond this threshold (see Definition 6). Likewise, for $m \in \mathbb{N}$, $\gamma > \gamma_{k\text{-SAT}}(m) := 1/m$, and sufficiently large k , the random k -SAT model exhibits symmetric m -OGP above the constraint density $\gamma 2^k \ln 2$; see [41] or Theorem 13 below. That is, no tuples of nearly equidistant satisfying assignments at a prescribed distance exist beyond the density $\gamma 2^k \ln 2$ (see Definition 12).

1.3 Overview of Our Results

As noted, studying the absence of OGP is essential for a deeper understanding of OGP-based methods and for developing a more complete average-case theory, yet this question has not been rigorously addressed. We initiate this direction by investigating the absence of symmetric m -OGP for Ising p -spin glass and random k -SAT for large p and k – two fundamental average-case models. This is a natural starting point: extending such investigations to other variants of OGP (e.g., asymmetric OGP [15] or branching OGP [46]) appears far more challenging, and symmetric OGP provides a natural foundation for these extensions. Moreover, both Ising p -spin and random k -SAT are more tractable for large p and k . Rigorous results for small, fixed k remain elusive for random k -SAT, and results for fixed p are extremely challenging for spin glasses. This reflects a broad trend in random structures; see Remark 5.

⁶ The value $\sqrt{2 \ln 2}$ corresponds to the ground-state energy of the Random Energy Model [74, 36].

Prior work, focusing exclusively on leveraging symmetric m -OGP to establish algorithmic hardness, overlooked the case $\gamma < \gamma(m)$, where $\gamma_{\text{p-spin}}(m) = 1/\sqrt{m}$ or $\gamma_{\text{k-SAT}}(m) = 1/m$. Our first contribution is to precisely determine the thresholds for both models, providing deeper insights into the power of symmetric m -OGP. All statements below hold w.h.p. as $n \rightarrow \infty$. We begin with the Ising spin glasses.

► **Theorem 1** (Informal, see Theorem 8). *For any $m \in \mathbb{N}$, $\gamma < \gamma_{\text{p-spin}}(m) = 1/\sqrt{m}$ and sufficiently large p , the symmetric m -OGP is absent in the Ising p -spin glass below the value $\gamma\sqrt{2 \ln 2}$.*

Namely, for any $m \in \mathbb{N}$ and sufficiently large p , nearly equidistant m -tuples of solutions exist at all pairwise distances below the value $\sqrt{(2 \ln 2)/m}$. Theorem 1 offers a qualitative insight into the OGP-based arguments. Specifically, in the limit of large p , it reveals that the power of OGP in establishing algorithmic hardness indeed strictly improves with m . We next address the random k -SAT.

► **Theorem 2** (Informal, see Theorem 14). *For any $m \in \mathbb{N}$, $\gamma < \gamma_{\text{k-SAT}}(m) = 1/m$ and $k = \Omega(\ln n)$, the symmetric m -OGP is absent in the random k -SAT below the density $\gamma 2^k \ln 2$.*

Theorem 2 asserts that for any $m \in \mathbb{N}$ and mildly growing k , $k = \Omega(\ln n)$, nearly equidistant m -tuples of satisfying assignments exist at all pairwise distances below the constraint density $(2^k \ln 2)/m$. Recall that the algorithmic threshold for the random k -SAT is asymptotic to $2^k \ln 2/k$ as $k \rightarrow \infty$. Consequently, Theorem 3 implies that achieving matching algorithmic lower bounds via the symmetric m -OGP requires m to be superconstant. Indeed, sharp lower bounds for this model were obtained instead through an asymmetric variant of OGP [15].

Theorem 2 holds under the assumption that k grows mildly in n . For constant k , we establish:

► **Theorem 3** (Informal, see Theorem 18). *For any $m \in \mathbb{N}$, $\gamma < 1/m$ and sufficiently large k , the symmetric m -OGP is absent below the density $\gamma 2^k \ln 2$ for the set of assignments that satisfy a $1 - \exp(-\Theta_k(k))$ fraction of clauses.*

Namely, the m -OGP is absent for the set of assignments violating a small fraction of clauses.

Combined with prior results on the m -OGP, Theorems 1-2 establish a sharp phase transition.

► **Theorem 4** (Informal, see Theorems 11 and 16). *For any $m \in \mathbb{N}$, the symmetric m -OGP undergoes a sharp phase transition in the Ising p -spin glass and the random k -SAT.*

Phase transition points are $\gamma_{\text{p-spin}} = 1/\sqrt{m}$ and $\gamma_{\text{k-SAT}} = 1/m$, both of which are strictly monotonic in m . While the first moment calculations had suggested that the OGP thresholds improve with m – as expected, since m -OGP structure becomes increasingly nested as m increases – a rigorous confirmation had been lacking. Theorem 4 establishes this structural fact for the models we study, thereby justifying the aforementioned common practice of increasing m when applying OGP-based methods.

► **Remark 5.** Notice that Theorems 1-4 hold for large p and k . This mirrors a rich body of work in random structures – e.g., random k -SAT, spin glasses, and random graphs – where rigorous results for fixed parameters are notoriously challenging and foundational progress has largely occurred in the large parameter regime. For random k -SAT, both the satisfiability threshold [25] and algorithmic lower bounds [15] are known only for sufficiently large k . Earlier developments [32, 21, 57] focused on the same regime $k = \Omega(\ln n)$ considered here. For Ising p -spin glasses, rigorous results for large p [74] preceded technically far more

delicate breakthroughs for fixed p [75, 63]. For sparse random graphs, most rigorous results pertain to large average degree d [33, 13, 81]. The large p regime has also proven useful in high-dimensional statistics; for instance, a recent work rigorously analyzes the large average subtensor problem in a challenging regime for p large [44]. Our work, which initiates a new direction into the absence of OGP, also focuses on large parameter regime; extending beyond this, particularly to fixed p or k , is a central open challenge for future work.

Shattering in Spin Glasses. The celebrated work by Kirkpatrick and Thirumalai [50] predicted a shattering phase in the Gibbs measure of the Ising p -spin glass, in which exponentially many well-separated clusters, each with exponentially small mass, collectively contain almost all of the Gibbs mass. Subsequent work [62, 30] refined this conjecture, postulating shattering over a range of inverse temperatures $\beta \in (\beta_d(p), \beta_c(p))$, where $\beta_d(p) = (1 + o_p(1))\sqrt{(2 \ln p)/p}$ as shown in [30, 28] and $\beta_c(p) = (1 - o_p(1))\sqrt{2 \ln 2}$ as established in [74], both asymptotically as $p \rightarrow \infty$. Shattering at zero temperature (i.e., shattering of the solution space) has played a key role since then in the study of algorithmic hardness in random CSPs [7, 4, 5].

Despite major advances in spin glass theory, shattering phenomena have only recently begun to be rigorously understood. [36] established shattering in the Ising p -spin glass model for $\beta \in (\sqrt{\ln 2}, \sqrt{2 \ln 2})$ and all sufficiently large p . Their argument, based on the pairwise OGP (2-OGP), falls short of establishing shattering when $\beta_d(p) < \beta < \sqrt{\ln 2}$. In light of Theorems 1 and 8, the 2-OGP is absent when $\beta < \sqrt{\ln 2}$, which accounts for the fact that the techniques of [36] apply only to the range $\sqrt{\ln 2} < \beta < \sqrt{2 \ln 2}$. Very recently, El Alaoui [26] introduced a “soft” version of the 2-OGP, which led to a near-optimal shattering result extending down to the conjectured $\beta_d(p)$ threshold. Our results confirm that the bounds in [36] are tight. This highlights the necessity of the soft-OGP approach developed in [26] to obtain optimal shattering results.

Notation. For any event E , $\mathbb{1}\{E\}$ denotes its indicator. For $r \in \mathbb{R}$, $\exp(r) := e^r$. For $n \in \mathbb{N}$, $\Sigma_n := \{-1, 1\}^n$. For any $\sigma, \sigma' \in \Sigma_n$ or $\sigma, \sigma' \in \{0, 1\}^n$, $d_H(\sigma, \sigma') := \sum_{1 \leq i \leq n} \mathbb{1}\{\sigma(i) \neq \sigma'(i)\}$. For $\sigma, \sigma' \in \Sigma_n$, $\langle \sigma, \sigma' \rangle := \sum_{1 \leq i \leq n} \sigma(i)\sigma'(i) = n - 2d_H(\sigma, \sigma')$. We employ standard asymptotic notation: $\Theta(\cdot)$, $O(\cdot)$, $o(\cdot)$, $\Omega(\cdot)$, and $\omega(\cdot)$; the underlying asymptotics is often w.r.t. $n \rightarrow \infty$. Asymptotics other than $n \rightarrow \infty$ are reflected with a subscript, e.g. $\Omega_k(\cdot)$.

2 Sharp Phase Transition for the m -OGP: Ising p -Spin Glass Model

In this section, we establish a sharp phase transition for the multi Overlap Gap Property (m -OGP) in the Ising p -spin model. We first formalize the set of m -tuples under consideration.

► **Definition 6.** Let $m \in \mathbb{N}$, $0 < \gamma < 1$, and $0 < \eta < \xi < 1$. Denote by $\mathcal{S}_{p\text{-spin}}(\gamma, m, \xi, \eta)$ the set of all m -tuples $\sigma^{(t)} \in \Sigma_n := \{-1, 1\}^n$, $1 \leq t \leq m$, that satisfy the following:

- **γ -Optimality:** For any $1 \leq t \leq m$, $H(\sigma^{(t)}) \geq \gamma\sqrt{2 \ln 2}$.
- **Overlap Constraint:** For any $1 \leq t < \ell \leq m$, $n^{-1} \langle \sigma^{(t)}, \sigma^{(\ell)} \rangle \in [\xi - \eta, \xi]$.

Definition 6 regards m -tuples of near-optimal solutions whose (pairwise) overlaps are constrained; parameter γ quantifies the near-optimality and ξ, η control the region of overlaps. The model exhibits symmetric m -OGP at threshold $\gamma\sqrt{2 \ln 2}$ if $\mathcal{S}_{p\text{-spin}}$ is w.h.p. empty for suitable m, ξ and η .

Gamarnik, Jagannath, and Kızıldağ establish the following result.

► **Proposition 7** ([36, Theorem 2.11]). *For any $m \in \mathbb{N}$ and any $\gamma > 1/\sqrt{m}$, there exists $0 < \eta < \xi < 1$ and a $P^* \in \mathbb{N}$ such that the following holds. Fix any $p \geq P^*$. Then, as $n \rightarrow \infty$*

$$\mathbb{P}[\mathcal{S}_{p\text{-spin}}(\gamma, m, \xi, \eta) \neq \emptyset] \leq e^{-\Theta(n)}.$$

That is, for any $m \in \mathbb{N}$ and $\gamma > 1/\sqrt{m}$, Ising pure p -spin model exhibits the symmetric m -OGP above $\gamma\sqrt{2\ln 2}$ for all large enough p .⁷

Proposition 7 focuses on $\gamma > 1/\sqrt{m}$. Our first main result addresses the regime $\gamma < 1/\sqrt{m}$.

► **Theorem 8.** *For any $m \in \mathbb{N}$, $\gamma < 1/\sqrt{m}$, and any $0 < \eta < \xi < 1$, there exists a $P^* \in \mathbb{N}$ such that the following holds. Fix any $p \geq P^*$. Then, as $n \rightarrow \infty$*

$$\mathbb{P}[\mathcal{S}_{p\text{-spin}}(\gamma, m, \xi, \eta) \neq \emptyset] \geq 1 - e^{-\Theta(n)}.$$

The proof of Theorem 8 relies on multiple techniques, which we overview in Section 2.1. The complete proof can be found in the full version [54, Section 5.2].

Taken together, Proposition 7 and Theorem 8 collectively imply that for sufficiently large p , the value $\sqrt{(2\ln 2)/m}$ is tight: for any $m \in \mathbb{N}$, and sufficiently large p , the onset of m -OGP is $\sqrt{(2\ln 2)/m}$. We make this more precise in Theorem 11, establishing a sharp phase transition.

We emphasize that Theorem 8 regards a single Hamiltonian, rather than a correlated ensemble. Crucially, it holds for *any* ξ and η , strengthening the ensuing results: the phase transition for the m -OGP is independent of ξ, η or the ensemble, making it inherently definition-robust.

Sharp Phase Transition. We define the notion of *admissible* values of γ .

► **Definition 9.** *Fix an $m \in \mathbb{N}$. A value $\gamma > 0$ is called m -admissible if for any $0 < \eta < \xi < 1$ there exists a $P^*(m, \gamma, \eta, \xi) \in \mathbb{N}$ such that for every fixed $p \geq P^*(m, \gamma, \eta, \xi)$,*

$$\lim_{n \rightarrow \infty} \mathbb{P}[\mathcal{S}_{p\text{-spin}}(\gamma, m, \xi, \eta) \neq \emptyset] = 1.$$

Similarly, γ is called m -inadmissible if there exists $0 < \eta < \xi < 1$, η sufficiently small, and a $\hat{P} \in \mathbb{N}$ such that for every fixed $p \geq \hat{P}$,

$$\lim_{n \rightarrow \infty} \mathbb{P}[\mathcal{S}_{p\text{-spin}}(\gamma, m, \xi, \eta) \neq \emptyset] = 0.$$

The sharp phase transition we investigate is formalized as follows.

► **Definition 10.** *Fix $m \in \mathbb{N}$. The m -OGP exhibits a sharp phase transition at the threshold $\gamma(m)$ if*

$$\sup\{\gamma > 0 : \gamma \text{ is } m\text{-admissible}\} = \gamma(m) = \inf\{\gamma > 0 : \gamma \text{ is } m\text{-inadmissible}\}.$$

As noted above, the phase transition is definition-robust, i.e., it remains independent of ξ, η or the ensemble. We establish that:

► **Theorem 11.** *For any $m \in \mathbb{N}$, the m -OGP for the Ising p -spin glass model exhibits a sharp phase transition in the sense of Definition 10 at $\gamma_{p\text{-spin}}(m) = 1/\sqrt{m}$.*

Theorem 11 follows directly from Proposition 7 and Theorem 8; the proof is omitted. We highlight that the phase transition point $\gamma_{p\text{-spin}}(m)$ is strictly monotonic in m .

⁷ We emphasize that [36, Theorem 2.11] is, in fact, stronger. It establishes the *ensemble* m -OGP, concerning tuples of near-optima w.r.t. correlated Hamiltonians, which is necessary for proving algorithmic hardness.

2.1 Proof Overview for Theorem 8

Given $m \in \mathbb{N}$ and $0 < \eta < \xi < 1$, denote by $\mathcal{F}(m, \xi, \eta)$ the set of all $(\sigma^{(1)}, \dots, \sigma^{(m)})$, satisfying the overlap constraints, $n^{-1} \langle \sigma^{(i)}, \sigma^{(j)} \rangle \in [\xi - \eta, \xi]$ for all $i \neq j$. To begin with, we must show that $\mathcal{F}(m, \xi, \eta)$ is non-empty; otherwise the set $\mathcal{S}_{\text{p-spin}}$ in Definition 6 would be trivially empty. We establish this using *probabilistic method*. Specifically, we generate $\sigma^{(1)}, \dots, \sigma^{(m)}$ by independently sampling each coordinate $\sigma^{(i)}(k) \in \{-1, 1\}$ with an appropriate mean. Applying large deviations inequalities, we show that the resulting tuple satisfies the overlap constraints with positive probability, thereby ensuring $\mathcal{F}(m, \xi, \eta) \neq \emptyset$.

The next step is applying the second moment method. Suppose that M is a non-negative integer-valued random variable with finite variance. Then, $\mathbb{P}[M \geq 1] \geq \mathbb{E}[M]^2 / \mathbb{E}[M^2]$ by the Paley-Zygmund Inequality. We fix $\gamma < 1/\sqrt{m}$ and apply this to $M := |\mathcal{S}_{\text{p-spin}}(\gamma, m, \xi, \eta)|$. To estimate $\mathbb{E}[M^2]$, we analyze a summation over *all pairs* of m -tuples in $\mathcal{F}(m, \xi, \eta)$ satisfying Definition 6. Building on an overcounting argument, we show that for large p , these pairs behave approximately independently unless their corresponding coordinates exhibit extreme correlation. At the same time, enforcing a high degree of correlation incurs a steep entropic cost. This allows us to decompose $\mathbb{E}[M^2]$ into two parts, each of which can be controlled separately. This leads to

$$\mathbb{P}[|\mathcal{S}_{\text{p-spin}}(\gamma, m, \xi, \eta)| \geq 1] \geq \exp(-no_p(1)) \quad (2)$$

where $o_p(1) \rightarrow 0$ as $p \rightarrow \infty$. Thus, the second moment method alone falls short of establishing the desired result that $\mathcal{S}_{\text{p-spin}} \neq \emptyset$ w.p. $1 - o(1)$. While this issue could be resolved for large p if $|\mathcal{S}_{\text{p-spin}}|$ were concentrated around its mean, such a concentration property is unclear. Instead, we proceed by defining a suitable proxy random variable that is closely related to $|\mathcal{S}_{\text{p-spin}}|$ and more tractable:

$$T_{m, \xi, \eta} := \max_{(\sigma^{(1)}, \dots, \sigma^{(m)}) \in \mathcal{F}(m, \xi, \eta)} \min_{1 \leq j \leq m} H(\sigma^{(j)}).$$

We prove that $T_{m, \xi, \eta}$, viewed as a function of the tensor $\mathbf{J}$ with i.i.d. $\mathcal{N}(0, 1)$ entries (cf. (1)), satisfies a Lipschitz property. Consequently, standard concentration results for Lipschitz functions of normal random variables yield that for all $\epsilon \geq 0$:

$$\mathbb{P}[|T_{m, \xi, \eta} - \mathbb{E}[T_{m, \xi, \eta}]| \geq \epsilon] \leq 2 \exp(-n\epsilon^2/2). \quad (3)$$

Second Moment Bound for Large p . We now leverage an elegant argument of Frieze [33] to combine the second moment bound with concentration. The key observation is that:

$$|\mathcal{S}_{\text{p-spin}}(\gamma, m, \xi, \eta)| \geq 1 \iff T_{m, \xi, \eta} \geq \gamma\sqrt{2 \ln 2}, \quad (4)$$

for any $\gamma > 0$. Fixing $\gamma' \in (\gamma, 1/\sqrt{m})$, we obtain that for all sufficiently large p and n ,

$$\mathbb{P}[T_{m, \xi, \eta} \geq \gamma'\sqrt{2 \ln 2}] \geq \exp(-no_p(1)) \geq 2 \exp(-n\epsilon^2/2) \geq \mathbb{P}[T_{m, \xi, \eta} \geq \mathbb{E}[T_{m, \xi, \eta}] + \epsilon]$$

using (2), (3), and (4). Consequently, $\mathbb{E}[T_{m, \xi, \eta}] \geq \gamma'\sqrt{2 \ln 2} - \epsilon$. Applying (3) once more, we obtain w.h.p. $T_{m, \xi, \eta} \geq \mathbb{E}[T_{m, \xi, \eta}] - \epsilon \geq \gamma'\sqrt{2 \ln 2} - 2\epsilon$, which at least $\gamma\sqrt{2 \ln 2}$ provided $\epsilon > 0$ is sufficiently small. This bound, together with (4), establishes Theorem 8.

3 Sharp Phase Transition for the m -OGP: Random k -SAT Model

We now turn to the random k -SAT. The set of m -tuples we investigate is as follows.

- **Definition 12.** Let $k \in \mathbb{N}$, $\gamma \in (0, 1)$, $m \in \mathbb{N}$, and $0 < \eta < \xi < 1$. Denote by $\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta)$ the set of all m -tuples $\sigma^{(t)} \in \{0, 1\}^n$, $1 \leq t \leq m$, that satisfy the following.
- **Satisfiability:** For any $1 \leq t \leq m$, $\Phi(\sigma^{(t)}) = 1$, where $\Phi(\cdot)$ is a conjunction of $n\alpha_\gamma$ independent k -clauses and $\alpha_\gamma = \gamma 2^k \ln 2$.
 - **Pairwise Distance:** For any $1 \leq t < \ell \leq m$, $n^{-1}d_H(\sigma^{(t)}, \sigma^{(\ell)}) \in [\xi - \eta, \xi]$.

Definition 12 regards m -tuples of satisfying solutions whose pairwise distances are constrained. The parameter α_γ is the clause density: the formula Φ consists of $n\alpha_\gamma$ independent k -clauses. Parameters ξ and η collectively control the region of overlaps. Similar to above, the model exhibits m -OGP above α_γ if $\mathcal{S}_{k\text{-SAT}}$ is empty (w.h.p.) for suitable m, ξ and η .

Gamarnik and Sudan [41] established the symmetric m -OGP for a variant of random k -SAT known as NAE- k -SAT. We first extend their result to random k -SAT model.

- **Theorem 13.** For any $m \in \mathbb{N}$ and any $\gamma > 1/m$, there exists $0 < \eta < \xi \leq \frac{1}{2}$ and $K^* := K^*(\gamma, m, \xi, \eta) \in \mathbb{N}$ such that the following holds. Fix any $k \geq K^*$. Then, as $n \rightarrow \infty$

$$\mathbb{P}[\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta) \neq \emptyset] \leq e^{-\Theta(n)}.$$

Theorem 13 is shown via the first moment method, see [54, Section 5.4] for the proof. For any $m \in \mathbb{N}$ and $\gamma > 1/m$, the random k -SAT model exhibits m -OGP at clause density $\gamma 2^k \ln 2$.

Our next result addresses the regime $\gamma < 1/m$.

- **Theorem 14.** For any $m \in \mathbb{N}$, $\gamma < 1/m$ and $0 < \eta < \xi \leq \frac{1}{2}$, there exists a constant $C(\gamma, m, \xi, \eta) > 0$ such that the following holds. Fix any $k \geq C(\gamma, m, \xi, \eta) \ln n$. Then as $n \rightarrow \infty$,

$$\mathbb{P}[\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta) \neq \emptyset] \geq 1 - e^{-\Theta(n)}.$$

See [54, Section 5.5] for the proof, which is based on the second moment method.

Theorem 14 yields that for any $m \in \mathbb{N}$, $\alpha < (2^k \ln 2)/m$ and k growing mildly in n , the m -OGP is absent: nearly equidistant m -tuples of satisfying solutions exist at all pairwise distances. Taken together, Theorems 13 and 14 collectively yield that for $k = \Omega(\ln n)$, the onset of m -OGP occurs at density $(2^k \ln 2)/m$. In Theorem 16 below, we combine Theorems 13 and 14 to establish a sharp phase transition.

We remark on the condition $k = \Omega(\log n)$. Several earlier developments in the random k -SAT literature, particularly those concerning the satisfiability threshold, were also established in the same regime $k = \Omega(\ln n)$, see [32, 21, 57]. In Theorem 18 below, we establish the absence of the m -OGP for all sufficiently large $k = O(1)$, in the set of nearly satisfying assignments that violate only a small fraction of clauses.

Our final remark concerns the technical condition $\xi \leq \frac{1}{2}$ in Theorem 14. For Theorem 14 to be non-trivial, there must exist m -tuples $\sigma^{(1)}, \dots, \sigma^{(m)} \in \{0, 1\}^n$ such that for all sufficiently small $\eta > 0$, their Hamming distances satisfy $d_H(\sigma^{(i)}, \sigma^{(j)}) \in [\xi - \eta, \xi]$ for $1 \leq i < j \leq m$. (In fact, our proof extends to any ξ for which such tuples exist.) An application of the probabilistic method confirms the existence of such tuples for $\xi \leq \frac{1}{2}$; however, the case $\xi > \frac{1}{2}$ remains unclear. Thus, we restrict our attention to $\xi \leq \frac{1}{2}$, leaving $\xi > \frac{1}{2}$ for future work.

A Sharp Phase Transition. Similar to the Ising p -spin glass model, we define the notion of *admissible* values of γ by modifying Definition 9.

► **Definition 15.** Fix an $m \in \mathbb{N}$. A value $\gamma > 0$ is called *m -admissible* if for any $0 < \eta < \xi < 1$ there exists a $C(\gamma, m, \xi, \eta) \in \mathbb{N}$ such that for every $k \geq C(\gamma, m, \xi, \eta) \ln n$,

$$\lim_{n \rightarrow \infty} \mathbb{P}[\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta) \neq \emptyset] = 1.$$

Similarly, γ is called *m -inadmissible* if there exists $0 < \eta < \xi < 1$, η sufficiently small, and a $C(\gamma, m, \xi, \eta) \in \mathbb{N}$ such that for every $k \geq C(\gamma, m, \xi, \eta) \ln n$,

$$\lim_{n \rightarrow \infty} \mathbb{P}[\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta) \neq \emptyset] = 0.$$

Equipped with Definition 15, we now analyze the sharp phase transition of the m -OGP in the sense of Definition 10 and establish the following result.

► **Theorem 16.** For any $m \in \mathbb{N}$, the m -OGP for the random k -SAT model exhibits a sharp phase transition in the sense of Definition 10 at $\gamma_{k\text{-SAT}} = 1/m$.

Theorem 16 follows directly from Theorems 13 and 14 [54, Section 5.3]. Notably, the phase transition point $\gamma_{k\text{-SAT}}(m)$ is, once again, strictly monotonic in m .

3.1 Case of Constant k

The m -OGP result for the random k -SAT (Theorem 13) holds for all sufficiently large k , while Theorem 14 along with the sharp phase transition result (Theorem 16) apply to $k = \Omega(\ln n)$. A fundamental question arises: can we extend the analysis to the challenging regime where k is constant?

In this section, we address this question by analyzing a variant of the k -SAT model that focuses on nearly satisfying assignments – assignments that violate only a small fraction of clauses. This setting is closely related to the maximum satisfiability (MAX-SAT) model. For this model, we establish the absence of the m -OGP for sufficiently large constant k .

An Objective Function: Number of Violated Constraints. Given a random k -SAT formula Φ which is a conjunction of M independent k -clauses $\mathcal{C}_1, \dots, \mathcal{C}_M$, and a truth assignment $\sigma \in \{0, 1\}^n$, define $\mathcal{L}(\sigma)$ as the number of clauses violated by σ : $\mathcal{L}(\sigma) = \sum_{1 \leq i \leq M} \mathbb{1}\{\mathcal{C}_i(\sigma) = 0\}$. In particular, $\Phi(\sigma) = 1$ iff $\mathcal{L}(\sigma) = 0$. With this, we modify Definition 12 accordingly.

► **Definition 17.** Let $k \in \mathbb{N}$, $\gamma \in (0, 1)$, $m \in \mathbb{N}$, $0 < \eta < \xi < 1$, and $\kappa \in [0, 1]$. Denote by $\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta, \kappa)$ the set of all m -tuples $\sigma^{(t)} \in \{0, 1\}^n$, $1 \leq t \leq m$, that satisfy the following.

- **Near-Optimality:** For any $1 \leq i \leq m$, $M^{-1}\mathcal{L}(\sigma^{(i)}) \leq \kappa$, where $M = n\alpha_\gamma$, $\alpha_\gamma = \gamma 2^k \ln 2$.
- **Overlap Constraint:** For any $1 \leq t < \ell \leq m$, $n^{-1}d_H(\sigma^{(t)}, \sigma^{(\ell)}) \in [\xi - \eta, \xi]$.

That is, $\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta, \kappa)$ denotes the set of all nearly equidistant m -tuples of truth assignments such that for any $(\sigma^{(1)}, \dots, \sigma^{(m)}) \in \mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta, \kappa)$ any $1 \leq i \leq m$, fraction of clauses violated by $\sigma^{(i)}$ is at most κ . In particular, when $\kappa = 0$, we recover the original definition: $\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta) = \mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta, 0)$.

We are now ready to state our final main result.

► **Theorem 18.** For any $m \in \mathbb{N}$, $\gamma < 1/m$, and $0 < \eta < \xi < 1$, there exists constants $C > 0$, $K^* \in \mathbb{N}$ and $\zeta \in (0, 1)$ such that the following holds. Fix any $k \geq K^*$. Then as $n \rightarrow \infty$,

$$\mathbb{P}[\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta, C\zeta^{k/2}) \neq \emptyset] \geq 1 - e^{-\Theta(n)}.$$

48:12 Sharp Thresholds for the Overlap Gap Property

See below for the proof sketch and [54, Section 5.6] for the full proof.

Theorem 18 asserts that for any $m \in \mathbb{N}$, $\gamma < 1/m$ and $0 < \eta < \xi < 1$, there exists m -tuples $(\sigma^{(1)}, \dots, \sigma^{(m)})$ such that the normalized Hamming distances obey $n^{-1}d_H(\sigma^{(i)}, \sigma^{(j)}) \in [\xi - \eta, \xi]$, and each $\sigma^{(i)}$ satisfies at least $(1 - O(\zeta^{k/2}))M$ many clauses. In other words, there exists m -tuples of nearly equidistant and nearly satisfying assignments below the density $(2^k \ln 2)/m$.

Proof Sketch for Theorem 18. Recall from above that $\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta) := \mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta, 0)$. An inspection of the second moment calculation for Theorem 14 reveals:

$$\mathbb{P}[\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta) \neq \emptyset] \geq \exp(-2\alpha_\gamma n m^2 \bar{\zeta}^k), \quad (5)$$

for some $\bar{\zeta} \in (0, 1)$, where $\alpha_\gamma = \gamma 2^k \ln 2$. Next, we define the proxy random variable

$$Z_{\xi, \eta} := \max_{\sigma^{(1)}, \dots, \sigma^{(m)} \in \{0,1\}^n} \min_{1 \leq i \leq m} \widehat{\mathcal{L}}(\sigma^{(i)}), \quad (6)$$

$n^{-1}d_H(\sigma^{(i)}, \sigma^{(j)}) \in [\xi - \eta, \xi]$

where $\widehat{\mathcal{L}}(\sigma) = M - \mathcal{L}(\sigma)$ is the number of *satisfied* clauses. Viewing $Z_{\xi, \eta}$ as a function $\mathcal{C}_1, \dots, \mathcal{C}_M$, we obtain $|Z(\mathcal{C}_1, \dots, \mathcal{C}_M) - Z(\mathcal{C}_1, \dots, \mathcal{C}_{i-1}, \mathcal{C}'_i, \mathcal{C}_{i+1}, \dots, \mathcal{C}_M)| \leq 1$ for all i , where $\mathcal{C}'_i$ is an independent copy of $\mathcal{C}_i$. So, $Z_{\xi, \eta}$ concentrates due to McDiarmid's inequality [80]:

$$\mathbb{P}[|Z_{\xi, \eta} - \mathbb{E}[Z_{\xi, \eta}]| \geq t_0] \leq 2 \exp(-t_0^2/M), \quad (7)$$

for any $t_0 \geq 0$. Notice that $Z_{\xi, \eta} = M$ iff $|\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta)| \geq 1$. Using (5) and (7), we have

$$\mathbb{P}[Z_{\xi, \eta} = M] \geq \exp(-2\alpha_\gamma n m^2 \bar{\zeta}^k) \geq \exp(-(t^*)^2/M) \geq \mathbb{P}[Z_{\xi, \eta} \geq \mathbb{E}[Z_{\xi, \eta}] + t^*],$$

where $t^* = \sqrt{3n\alpha_\gamma m \bar{\zeta}^{k/2}}$. Consequently, $\mathbb{E}[Z_{\xi, \eta}] \geq M - t^* = n\alpha_\gamma - t^*$. Applying (7) with $t_0 = \sqrt{\alpha_\gamma}n$, we obtain that w.h.p., $Z_{\xi, \eta} \geq \mathbb{E}[Z_{\xi, \eta}] - \sqrt{\alpha_\gamma}n \geq n\alpha_\gamma(1 - C\zeta^{k/2})$, where $C > 0$ and $\zeta \in (0, 1)$ are suitable constants. Finally, we complete the proof by recalling:

$$Z_{\gamma, \eta} \geq n\alpha_\gamma(1 - C\zeta^{k/2}) \iff |\mathcal{S}_{k\text{-SAT}}(\gamma, m, \xi, \eta, C\zeta^{k/2})| \geq 1.$$

4 Future Directions & Further Background on OGP

An immediate direction is to complement Theorem 18 with an m -OGP result and to establish a sharp threshold for “almost solutions” to random k -SAT, analogous to Theorems 11 and 16. Extensions to other spin glass models, such as mixed p -spin or multi-species models, are also intriguing directions for future work (thanks to an anonymous referee for suggesting this).

Random k -SAT for Constant k . Extending Theorem 14 to constant k remains an intriguing question we leave for future work. A promising starting point is the NAE- k -SAT model, a variant of k -SAT where each clause must contain at least one true and one false literal. The inherent symmetry of NAE- k -SAT makes it particularly amenable to rigorous analysis. While the proof of satisfiability conjecture for random k -SAT was a long and arduous journey, culminating in the breakthrough tour de force by Ding, Sly, and Sun [25], analogous results for the NAE- k -SAT were established much earlier, using fairly standard second moment calculations by Achlioptas and Moore [6]. This suggests that sharper results may be attainable for the NAE- k -SAT, potentially avoiding $k = \Omega(\log n)$ growth.

OGP and Algorithms. The interplay between the OGP and algorithms remains a key frontier, with far-reaching implications for average-case computational complexity. While the OGP formally implies the failure of large classes of algorithms, certain tractable optimization problems – solvable by linear programming – still exhibit the OGP [55]. What makes these cases exceptional? What is the broadest class of algorithms that OGP can provably rule out?

Likewise, what are the implications of the absence of the OGP? For instance, the algorithm by Montanari [61], finding a near-ground state for the SK model, crucially relies on the widely-believed conjecture that the model does not exhibit OGP. Our results indicate that the absence of OGP at a certain level m does not necessarily imply algorithmic tractability. Could there exist models without OGP that still pose insurmountable computational barriers? Finally, our findings reveal an intriguing scaling phenomenon: for certain models, the power of m -OGP in establishing algorithmic hardness strengthens indefinitely as m increases. When does this trend persist and when does it saturate?

References

- 1 Emmanuel Abbe, Shuangping Li, and Allan Sly. Binary perceptron: efficient algorithms can find solutions in a rare well-connected cluster. *arXiv preprint*, 2021. [arXiv:2111.03084](#).
- 2 Emmanuel Abbe, Shuangping Li, and Allan Sly. Proof of the contiguity conjecture and lognormal limit for the symmetric perceptron. *arXiv preprint*, 2021. [arXiv:2102.13069](#).
- 3 Dimitris Achlioptas. Random satisfiability. *Handbook of Satisfiability*, 185:245–270, 2009. doi:10.3233/978-1-58603-929-5-245.
- 4 Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. In *2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 793–802. IEEE, 2008. doi:10.1109/FOCS.2008.11.
- 5 Dimitris Achlioptas, Amin Coja-Oghlan, and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. *Random Structures & Algorithms*, 38(3):251–268, 2011. doi:10.1002/RS.20323.
- 6 Dimitris Achlioptas and Cristopher Moore. The asymptotic order of the random k-SAT threshold. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 779–788. IEEE, 2002. doi:10.1109/SFCS.2002.1182003.
- 7 Dimitris Achlioptas and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 130–139, 2006. doi:10.1145/1132516.1132537.
- 8 Louigi Addario-Berry and Pascal Maillard. The algorithmic hardness threshold for continuous random energy models. *Mathematical Statistics and Learning*, 2(1):77–101, 2020.
- 9 Michael Aizenman, Robert Sims, and Shannon L Starr. Extended variational principle for the Sherrington-Kirkpatrick spin-glass model. *Physical Review B*, 68(21):214403, 2003.
- 10 Antonio Auffinger and Wei-Kuo Chen. On properties of Parisi measures. *Probability Theory and Related Fields*, 161(3-4):817–850, 2015.
- 11 Antonio Auffinger and Wei-Kuo Chen. Parisi formula for the ground state energy in the mixed p -spin model. *The Annals of Probability*, 45(6B):4617–4631, 2017.
- 12 Antonio Auffinger and Aukosh Jagannath. Thouless–Anderson–Palmer equations for generic p -spin glasses. *The Annals of Probability*, 47(4):2230–2256, 2019. doi:10.1214/18-AOP1307.
- 13 Mohsen Bayati, David Gamarnik, and Prasad Tetali. Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 105–114, 2010. doi:10.1145/1806689.1806706.
- 14 Gérard Ben Arous and Aukosh Jagannath. Shattering versus metastability in spin glasses. *arXiv preprint*, 2021. [arXiv:2104.08299](#).

- 15 Guy Bresler and Brice Huang. The algorithmic phase transition of random k -SAT for low degree polynomials. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 298–309. IEEE, 2022.
- 16 Sourav Chatterjee and Leila Sloman. Average Gromov hyperbolicity and the Parisi ansatz. *Advances in Mathematics*, 376:107417, 2021.
- 17 Wei-Kuo Chen, David Gamarnik, Dmitry Panchenko, and Mustazee Rahman. Suboptimality of local algorithms for a class of max-cut problems. *The Annals of Probability*, 47(3):1587–1618, 2019.
- 18 Wei-Kuo Chen, Dmitry Panchenko, and Eliran Subag. Generalized TAP Free Energy. *Communications on Pure and Applied Mathematics*, 76(7):1329–1415, 2023.
- 19 V. Chvatal and B. Reed. Mick gets some (the odds are on his side) (satisfiability). In *Proceedings., 33rd Annual Symposium on Foundations of Computer Science*, pages 620–627, 1992.
- 20 Amin Coja-Oghlan. A better algorithm for random k -SAT. *SIAM Journal on Computing*, 39(7):2823–2864, 2010. doi:10.1137/09076516X.
- 21 Amin Coja-Oghlan and Alan Frieze. Random k -SAT: The Limiting Probability for Satisfiability for Moderately Growing k . *The Electronic Journal of Combinatorics*, 15(1):N2, 2008. URL: http://www.combinatorics.org/Volume_15/Abstracts/v15i1n2.html.
- 22 Amin Coja-Oghlan, Amir Haqshenas, and Samuel Hetterich. Walksat stalls well below satisfiability. *SIAM Journal on Discrete Mathematics*, 31(2):1160–1173, 2017. doi:10.1137/16M1084158.
- 23 Amin Coja-Oghlan and Konstantinos Panagiotou. The asymptotic k -SAT threshold. *Advances in Mathematics*, 288:985–1068, 2016.
- 24 Bernard Derrida. Random-energy model: Limit of a family of disordered models. *Physical Review Letters*, 45(2):79, 1980.
- 25 Jian Ding, Allan Sly, and Nike Sun. Proof of the satisfiability conjecture for large k . In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 59–68, 2015. doi:10.1145/2746539.2746619.
- 26 Ahmed El Alaoui. Near-optimal shattering in the ising pure p -spin and rarity of solutions returned by stable algorithms. *arXiv preprint*, 2024. arXiv:2412.03511.
- 27 Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Optimization of mean-field spin glasses. *The Annals of Probability*, 49(6):2922–2960, 2021.
- 28 Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Sampling from mean-field gibbs measures via diffusion processes. *arXiv preprint*, 2023. arXiv:2310.08912.
- 29 Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Shattering in pure spherical spin glasses. *arXiv preprint*, 2023. arXiv:2307.04659.
- 30 Ulisse Ferrari, Luca Leuzzi, Giorgio Parisi, and Tommaso Rizzo. Two-step relaxation next to dynamic arrest in mean-field glasses: Spherical and ising p -spin model. *Physical Review B—Condensed Matter and Materials Physics*, 86(1):014204, 2012.
- 31 John Franco and Marvin Paull. Probabilistic analysis of the Davis Putnam procedure for solving the satisfiability problem. *Discrete Applied Mathematics*, 5(1):77–87, 1983. doi:10.1016/0166-218X(83)90017-3.
- 32 Alan Frieze and Nicholas C Wormald. Random k -SAT: A tight threshold for moderately growing k . *Combinatorica*, 25(3):297–306, 2005. doi:10.1007/S00493-005-0017-3.
- 33 Alan M Frieze. On the independence number of random graphs. *Discrete Mathematics*, 81(2):171–175, 1990. doi:10.1016/0012-365X(90)90149-C.
- 34 David Gamarnik. The overlap gap property: A topological barrier to optimizing over random structures. *Proceedings of the National Academy of Sciences*, 118(41), 2021.
- 35 David Gamarnik and Aukosh Jagannath. The overlap gap property and approximate message passing algorithms for p -spin models. *The Annals of Probability*, 49(1):180–205, 2021.
- 36 David Gamarnik, Aukosh Jagannath, and Eren C Kızıldağ. Shattering in the Ising Pure p -Spin Model. *arXiv preprint*, 2023. arXiv:2307.07461.

- 37 David Gamarnik, Aukosh Jagannath, and Alexander S Wein. Low-degree hardness of random optimization problems. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 131–140. IEEE, 2020. doi:10.1109/FOCS46700.2020.00021.
- 38 David Gamarnik, Aukosh Jagannath, and Alexander S Wein. Circuit lower bounds for the p-spin optimization problem. *arXiv preprint*, 2021. arXiv:2109.01342.
- 39 David Gamarnik, Eren C Kızıldağ, Will Perkins, and Changji Xu. Algorithms and barriers in the symmetric binary perceptron model. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 576–587. IEEE, 2022.
- 40 David Gamarnik, Eren C Kızıldağ, Will Perkins, and Changji Xu. Geometric barriers for stable and online algorithms for discrepancy minimization. *arXiv preprint*, 2023. arXiv:2302.06485.
- 41 David Gamarnik and Madhu Sudan. Performance of sequential local algorithms for the random NAE-k-SAT problem. *SIAM Journal on Computing*, 46(2):590–619, 2017. doi:10.1137/140989728.
- 42 Francesco Guerra. Broken replica symmetry bounds in the mean field spin glass model. *Communications in mathematical physics*, 233:1–12, 2003.
- 43 Francesco Guerra and Fabio Lucio Toninelli. The Thermodynamic Limit in Mean Field Spin Glass Models. *Communications in Mathematical Physics*, 230:71–79, 2002.
- 44 Abhishek Hegade KR and Eren C Kızıldağ. Large average subtensor problem: Ground-state, algorithms, and algorithmic barriers. *arXiv preprint*, 2025. arXiv:2506.17118.
- 45 Samuel Hetterich. Analysing survey propagation guided decimation on random formulas. *arXiv preprint*, 2016. arXiv:1602.08519.
- 46 Brice Huang and Mark Sellke. Tight Lipschitz hardness for optimizing mean field spin glasses. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 312–322. IEEE, 2022. doi:10.1109/FOCS54457.2022.00037.
- 47 Brice Huang and Mark Sellke. Algorithmic threshold for multi-species spherical spin glasses. *arXiv preprint*, 2023. doi:10.48550/arXiv.2303.12172.
- 48 Aukosh Jagannath. Approximate ultrametricity for random measures and applications to spin glasses. *Communications on Pure and Applied Mathematics*, 70(4):611–664, 2017.
- 49 Aukosh Jagannath and Ian Tobasco. Some properties of the phase diagram for mixed p-spin glasses. *Probability Theory and Related Fields*, 167:615–672, 2017.
- 50 Theodore R Kirkpatrick and Devarajan Thirumalai. p-spin-interaction spin-glass models: Connections with the structural glass problem. *Physical Review B*, 36(10):5388, 1987.
- 51 Lefteris M Kirousis, Evangelos Kranakis, Danny Krizanc, and Yannis C Stamatiou. Approximating the unsatisfiability threshold of random formulas. *Random Structures & Algorithms*, 12(3):253–269, 1998. doi:10.1002/(SICI)1098-2418(199805)12:3%3C253::AID-RSA3%3E3.0.CO;2-U.
- 52 Justin Ko. Free energy of multiple systems of spherical spin glasses with constrained overlaps. *Electronic Journal of Probability*, 25:1–34, 2020. doi:10.1214/20-EJP431.
- 53 Florent Krzakala, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, 104(25):10318–10323, 2007. doi:10.1073/PNAS.0703685104.
- 54 Eren C. Kızıldağ. Sharp Thresholds for the Overlap Gap Property: Ising p -Spin Glass and Random k -SAT, 2025. arXiv:2309.09913.
- 55 Shuangping Li and Tselil Schramm. Some easy optimization problems have the overlap-gap property. *arXiv preprint*, 2024. doi:10.48550/arXiv.2411.01836.
- 56 Shuangping Li, Tselil Schramm, and Kangjie Zhou. Discrepancy algorithms for the binary perceptron. *arXiv preprint*, 2024. doi:10.48550/arXiv.2408.00796.
- 57 Jun Liu, Zongsheng Gao, and Ke Xu. A Note on Random k -SAT for Moderately Growing k . *The Electronic Journal of Combinatorics*, pages P24–P24, 2012.
- 58 Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.

- 59 Marc Mézard, Giorgio Parisi, and Riccardo Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297(5582):812–815, 2002.
- 60 Chao Ming-Te and John Franco. Probabilistic analysis of a generalization of the unit-clause literal selection heuristics for the k satisfiability problem. *Information Sciences*, 51(3):289–314, 1990. doi:10.1016/0020-0255(90)90030-E.
- 61 Andrea Montanari. Optimization of the Sherrington-Kirkpatrick Hamiltonian. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1417–1433, 2019. doi:10.1109/FOCS.2019.00087.
- 62 Andrea Montanari and Federico Ricci-Tersenghi. On the nature of the low-temperature phase in discontinuous mean-field spin glasses. *The European Physical Journal B-Condensed Matter and Complex Systems*, 33:339–346, 2003.
- 63 Dmitry Panchenko. The Parisi Ultrametricity Conjecture. *Annals of Mathematics*, pages 383–393, 2013.
- 64 Dmitry Panchenko. *The Sherrington-Kirkpatrick model*. Springer Science & Business Media, 2013.
- 65 Dmitry Panchenko. The Parisi formula for mixed p -spin models. *The Annals of Probability*, 42(3):946–958, 2014.
- 66 Dmitry Panchenko. Free energy in the mixed p -spin models with vector spins. *The Annals of Probability*, 46(2):865–896, 2018. doi:10.1214/17-AOP1194.
- 67 Giorgio Parisi. Infinite number of order parameters for spin-glasses. *Physical Review Letters*, 43(23):1754, 1979.
- 68 Will Perkins and Changji Xu. Frozen 1-RSB structure of the symmetric Ising perceptron. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1579–1588, 2021. doi:10.1145/3406325.3451119.
- 69 Mark Sellke. Optimizing mean field spin glasses with external field. *arXiv preprint*, 2021. arXiv:2105.03506.
- 70 David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical review letters*, 35(26):1792, 1975.
- 71 Daniel L Stein and Charles M Newman. *Spin glasses and complexity*, volume 4. Princeton University Press, 2013.
- 72 Eliran Subag. Free energy landscapes in spherical spin glasses. *arXiv preprint*, 2018. arXiv:1804.10576.
- 73 Eliran Subag. Following the ground states of full-rsb spherical spin glasses. *Communications on Pure and Applied Mathematics*, 74(5):1021–1044, 2021.
- 74 Michel Talagrand. Rigorous low-temperature results for the mean field p -spins interaction model. *Probability theory and related fields*, 117:303–360, 2000.
- 75 Michel Talagrand. The Parisi formula. *Annals of Mathematics*, pages 221–263, 2006.
- 76 Michel Talagrand. *Mean field models for spin glasses: Volume I: Basic examples*, volume 54. Springer Science & Business Media, 2010.
- 77 Michel Talagrand. *Mean Field Models for Spin Glasses: Advanced replica-symmetry and low temperature*. Springer, 2011.
- 78 Fabio Lucio Toninelli. About the Almeida-Thouless transition line in the Sherrington-Kirkpatrick mean-field spin glass model. *Europhysics letters*, 60(5):764, 2002.
- 79 Neekon Vafa and Vinod Vaikuntanathan. Symmetric perceptrons, number partitioning and lattices. *arXiv preprint*, 2025. arXiv:2501.16517.
- 80 Lutz Warnke. On the method of typical bounded differences. *Combinatorics, Probability and Computing*, 25(2):269–299, 2016. doi:10.1017/S0963548315000103.
- 81 Alexander S Wein. Optimal low-degree hardness of maximum independent set. *Mathematical Statistics and Learning*, 4(3):221–251, 2022.