

Random Max-CSPs Inherit Algorithmic Hardness from Spin Glasses

Chris Jones* Kunal Marwaha[†] Juspreet Singh Sandhu[‡] Jonathan Shi[§]

March 18, 2023

Abstract

We study random constraint satisfaction problems (CSPs) at large clause density. We relate the structure of near-optimal solutions for *any* Boolean Max-CSP to that for an associated spin glass on the hypercube, using the Guerra-Toninelli interpolation from statistical physics. The *noise stability* polynomial of the CSP's predicate is, up to a constant, the mixture polynomial of the associated spin glass. We show two main consequences:

1. We prove that the maximum fraction of constraints that can be satisfied in a random Max-CSP at large clause density is determined by the ground state energy density of the corresponding spin glass. Since the latter value can be computed with the Parisi formula [Par80, Tal06, AC17], we provide numerical values for some popular CSPs.
2. We prove that a Max-CSP at large clause density possesses generalized versions of the *overlap gap property* if and only if the same holds for the corresponding spin glass. We transfer results from [HS21] to obstruct algorithms with *overlap concentration* on a large class of Max-CSPs. This immediately includes local classical and local quantum algorithms [CLSS22].

*Department of Computer Science, University of Chicago, Chicago, Illinois, USA. Supported by the National Science Foundation under Grant No. CCF-2008920. Email: csj@uchicago.edu.

[†]Department of Computer Science, University of Chicago, Chicago, Illinois, USA. Supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1746045. Email: kmarw@uchicago.edu.

[‡]School of Engineering & Applied Sciences, Harvard University, Cambridge, Massachusetts, USA. Supported by a Simons Investigator Fellowship, NSF grant DMS-2134157, NSF STAQ award PHY-1818914, DARPA grant W911NF2010021, DARPA ONISQ program award HR001120C0068 and DOE grant DE-SC0022199. Email: jus065@g.harvard.edu.

[§]Department of Computer Science, Bocconi University, Milan, Italy. Supported by European Research Council (ERC) award No. 834861 (SO-ReCoDi). Email: jonathan.shi@unibocconi.it.

Contents

1	Introduction	1
1.1	Main results	1
1.2	Related work	4
2	Preliminaries	6
2.1	Random constraint satisfaction problems (CSPs)	6
2.2	Mean-field spin glasses	7
2.3	Solution geometry of optimization problems	8
2.4	Fourier analysis on the hypercube	10
2.5	Concentration inequalities	11
3	Sparse and dense models have the same free energy	12
3.1	Proof of Theorem 3.5	14
4	Optimal value of a random Max-CSP	19
4.1	Numerical calculations	20
5	Overlap gaps in a random Max-CSP	20
5.1	The branching OGP	22
6	Implications for algorithmic hardness	23
7	Discussion	26
A	Concentration Proofs	33
B	Correlation functions of sparse Hamiltonians	36
C	Debiasing	37

1 Introduction

In this work, we formalize a general and deep connection between two intensely studied classes of optimization problems: constraint satisfaction problems (CSPs), studied in computer science, and spin glass models, studied in statistical physics. We demonstrate that as the clause density of the random CSP increases, the geometric properties of the set of nearly-optimal solutions converge to those of a corresponding spin glass model. In these spin glass models, the very same geometric properties imply bounds on the average-case approximability achieved by broad classes of algorithms [AMS21b, GJW20, GJ21]; these bounds are conjectured to be the best possible among all polynomial-time algorithms [Gam21, HS21]. The correspondence we establish here implies that the same lower bounds apply to average-case CSPs.

CSPs are paradigmatic computational tasks. Their study has led to foundational results in computational hardness, approximability, and optimization [Pas13]. In recent years, we have learned more about CSPs through methods inspired by statistical physics, especially when the clauses of the CSP are chosen randomly [DSS16b, DSS16a, DMS17, Sen18, DSS22]. By identifying the solution quality of a variable assignment with the *energy* of a configuration of particles, we can investigate “physical” properties of the CSP, such as phase transitions or solution clustering at different temperatures. Surprisingly, these physical properties can have computational consequences.

We study random CSP instances with Boolean variables, random literal signs, and number of constraints which is a large constant times the number of variables, such that each constraint acts on a constant number of variables. If n is the number of variables, then $m = \alpha n$ is the number of constraints for some constant α . For large enough α , the CSP is unsatisfiable with probability $1 - o_n(1)$. Therefore the goal is to find a variable assignment that maximizes the number of satisfied constraints, and we think of these as *Max-CSPs*.

Given a random Max-CSP, how many constraints can be satisfied? How are the best assignments distributed around the hypercube? Can we find these assignments with efficient algorithms? Statistical physicists use questions like these to investigate the *solution geometry* of a problem. Our main result connects the solution geometry of a Max-CSP (with large enough α) to that of a *spin glass*. As a consequence, much of our mathematical and algorithmic understanding of spin glasses transfers to CSPs at large clause density.

A *spin glass* (more properly a *mixed mean-field spin glass*) is a random system of n particles (variables) specified by a *mixture polynomial* $\xi(s) = \sum_{p \geq 1} c_p^2 s^p$. In this model, the interaction strength between every p -tuple of particles is an independent Gaussian with variance $c_p^2 n^{1-p}$; this can be thought of as a randomly-weighted CSP on the complete p -uniform hypergraph. We show that as the clause density of any random CSP increases ($\alpha \rightarrow \infty$), the solution space starts to resemble that of a spin glass. For example, Max-Cut on random graphs with large constant average degree qualitatively looks like the *Sherrington-Kirkpatrick model*, where $c_2 = 1, \{c_i\}_{i \neq 2} = 0$ [SK75, DMS17]. (Note that our proof applies to Max-2XOR instead of Max-Cut, so we do not recover this result exactly.)

1.1 Main results

Formally, we relate the *free energy density* of a random Max-CSP instance to that of a particular spin glass. The associated spin glass is determined only by the Fourier weights of the CSP. In fact, the mixture polynomial of the spin glass is, up to a constant, the *noise stability polynomial* of the CSP

(see [Section 2](#) for definitions):

Theorem 1.1 (Free energy density). *Generate a random CSP instance \mathcal{I} (with cost function $H_{\mathcal{I}}$) consisting of $\alpha \cdot n$ independent and uniform constraints of a predicate $f : \{\pm 1\}^k \rightarrow \{0, 1\}$ with randomly signed literals. Define the polynomial*

$$\xi(s) = \text{Stab}_s(f) - \widehat{f}(\emptyset)^2 = \sum_{j=1}^k \|f^{=j}\|^2 s^j, \quad (1)$$

where $\text{Stab}_s(f)$ is the noise stability polynomial of f , and $\|f^{=j}\|^2$ is the Fourier weight of f at degree j . Generate a random spin glass instance H^ξ with mixture polynomial ξ .

Let $\beta > 0$, and define $Z_{\mathcal{I}}(\beta) = \sum_{\sigma \in \{\pm 1\}^n} e^{\beta H_{\mathcal{I}}(\sigma)}$ and $Z_{\text{SG}_\xi}(\beta) = \sum_{\sigma \in \{\pm 1\}^n} e^{\beta H^\xi(\sigma)}$ as the respective partition functions. Then

$$\frac{1}{\beta n} \log Z_{\mathcal{I}}(\beta) = \widehat{f}(\emptyset) + \frac{1}{\beta n} \frac{\log Z_{\text{SG}_\xi}(\beta)}{\sqrt{\alpha}} + O\left(\frac{\beta^2}{\alpha^2}\right) + o_n(1). \quad (2)$$

where the second-to-last term (which may depend on n) satisfies $\left|O\left(\frac{\beta^2}{\alpha^2}\right)\right| \leq C \cdot \frac{\beta^2}{\alpha^2}$ whenever $\frac{\beta}{\alpha} \leq \varepsilon_0$ for absolute constants $C, \varepsilon_0 > 0$, and the last term is random and is $o_n(1)$ w.h.p.

Prior work relates the free energy density of specific CSPs such as Max- k XOR [[DMS17](#), [Sen18](#)] and Max- k SAT [[Pan18](#)] to that of a spin glass. [Theorem 1.1](#) generalizes this connection to any random Max-CSP with randomly signed literals.

The asymptotic equivalence of the free energy density implies the equivalence of several properties of the solution geometry for large enough α . We show two specific implications of [Theorem 1.1](#). The first is that the optimal value of a random Max-CSP in the large clause density limit can be found with a spin glass calculation.

Corollary 1.2. (Optimal value equivalence). *Generate a random CSP instance \mathcal{I} consisting of $\alpha \cdot n$ independent and uniform constraints of a predicate $f : \{\pm 1\}^k \rightarrow \{0, 1\}$ with randomly signed literals. Let $v_{\mathcal{I}}$ be the maximum fraction of constraints of \mathcal{I} that can be satisfied. Let ξ be defined as in [Equation \(1\)](#), and $\text{GSED}(\text{SG}_\xi)$ as the (non-random) ground state energy density of the associated spin glass. Then*

$$v_{\mathcal{I}} = \widehat{f}(\emptyset) + \frac{\text{GSED}(\text{SG}_\xi)}{\sqrt{\alpha}} + o\left(\frac{1}{\sqrt{\alpha}}\right). \quad (3)$$

where the last term is at most $\kappa(n, \alpha)$ w.h.p for a function $\kappa(n, \alpha)$ satisfying $\kappa(n, \alpha) \cdot \sqrt{\alpha} \rightarrow 0$ as $n \rightarrow \infty$ then $\alpha \rightarrow \infty$.

Computing the minimum value, or ground state energy, of a spin glass can famously be done using the *Parisi formula*. In [Section 4](#), we use the Parisi formula to compute $\text{GSED}(\text{SG}_\xi)$ for several common CSPs (our code is available online).

The second implication relates to algorithmic hardness.¹ Intuitively, when global minima are located in clusters, some algorithms cannot efficiently find them. A recent body of work (starting

¹For this implication, we need to boost [Theorem 1.1](#) to interpolate the free energy density restricted to any given overlap and correlation structure.

with [GS14]) has pinpointed hardness on the presence of an *overlap gap property* (OGP), showing how this property obstructs noise-stable algorithms from $(1 - \varepsilon)$ -approximating average-case instances. We show that a flexible version of the OGP, called the *branching OGP* [HS21], exists on a spin glass exactly when it exists on the associated Max-CSP at large enough clause density. As a result, the techniques that obstruct algorithms on certain spin glasses also work on the corresponding Max-CSPs.

Corollary 1.3 (OGP equivalence, informal). *Take any Max-CSP. Consider the associated spin glass \mathbf{SG}_ξ , where ξ is defined as in Equation (1). Then \mathbf{SG}_ξ exhibits an OGP at value v if and only if the Max-CSP exhibits an OGP at value $\widehat{f}(\emptyset) + \frac{v}{\sqrt{\alpha}}$ for all sufficiently large α .*

The formal class of algorithms we obstruct is as follows [HS21]. Consider two correlated instances with correlation parameter $t \in [0, 1]$. A deterministic² algorithm is *overlap-concentrated* if for every t , the overlap (a.k.a the Hamming distance) between the two output assignments for the two correlated instances is concentrated inside a narrow interval. This is a notion of noise robustness that many commonly-used algorithms have, including approximately Lipschitz algorithms on spin glasses [HS21] and local classical and local quantum algorithms on random Max-CSPs [CLSS22]. Additionally, survey propagation with a constant number of message-passing rounds likely has this property [BMZ05, Gam21, BH22].

Theorem 1.4 (OGP obstructs algorithms with overlap concentration, informal). *Consider an algorithm \mathcal{A} with overlap concentration. Then \mathcal{A} cannot output arbitrarily good approximate solutions on instances of Max-CSPs which exhibit an OGP.*

Our proof of Theorem 1.4 follows that of Huang and Sellke [HS21], who prove the same result on spin glasses. Only a few changes are needed to transfer their result to Max-CSPs. Note that this result also applies to CSPs with small α if they exhibit an OGP.

It is known that an OGP (specifically, a branching OGP) exists with high probability on spin glasses with even mixture polynomials without quadratic terms [CGPR19, HS21]. Combining this with the results that we have stated, we conclude the following:

Corollary 1.5 (informal). *Consider a random Max-CSP with a predicate f such that the only nonzero Fourier coefficients of f have even degree $j \geq 4$. For almost all instances of the Max-CSP, no algorithm with overlap concentration can output $(1 - \varepsilon)$ -approximately optimal solutions for all $\varepsilon > 0$.*

For example, this unconditionally obstructs overlap concentrated algorithms from approximating random 4XOR instances.

Corollary 1.3 makes progress on [CLSS22, Problem 9.2 (arXiv version)], asking which CSPs exhibit an OGP at large clause density. A full characterization of spin glasses with an OGP, and thereby CSPs in the large α limit, is not complete (although *spherical spin glasses* have been classified [Sub18, Proposition 1]). For example, the Sherrington-Kirkpatrick model is strongly suspected to have no OGP, but this is not fully proven [ACZ20].

We cannot help but mention that these topological properties of the solution space (i.e. the solution geometry) may *precisely characterize* the algorithmic approximability of spin glasses and random

²A randomized algorithm is overlap-concentrated if it is overlap-concentrated for every fixing of the randomness.

CSPs [ACO08]. For spin glasses, the following is known:

Theorem 1.6. *For all spin glasses SG_ξ , there is a value ALG (given by an extended Parisi formula) such that:*

1. [Mon21, AMS21b, Sel21] *Assume that the minimizer of the extended Parisi functional for SG_ξ exists. Then for all $\varepsilon > 0$, there is an efficient algorithm that outputs a solution with value $ALG - \varepsilon$ on almost all instances.*
2. [HS21] *Assume the mixture polynomial ξ is even. For all $\varepsilon > 0$, the spin glass exhibits a branching OGP with value $ALG + \varepsilon$, which therefore obstructs overlap-concentrated algorithms from achieving this value on almost all instances.*

It is likely the case that the same holds for CSPs with sufficiently large α . Specifically, there are explicitly computable constants $ALG \leq OPT$ for any CSP predicate f such that the optimum value of the CSP is about $\widehat{f}(\emptyset) + \frac{OPT}{\sqrt{\alpha}}$ (Corollary 1.2), but the threshold for efficient algorithms appears to be $\widehat{f}(\emptyset) + \frac{ALG}{\sqrt{\alpha}}$. In this paper we prove that the lower bound transfers (Part 2). The upper bound (Part 1) is known for Max-Cut by [AMS21a], and can likely be generalized to an arbitrary CSP (in a similar way that [AMS21b] generalizes [Mon21] for spin glasses). Existence of an OGP may also characterize approximability of random CSPs for small α , but proving this is significantly more difficult.

The paper is organized as follows. We provide formal definitions and additional motivation in Section 2. In Section 3, we give the main interpolation between every Max-CSP and a related spin glass. In Section 4, we show how this implies equivalence of optimal value (Corollary 1.2), and list numerical approximations to optimal values of several common Max-CSPs. In Section 5, we prove that the main interpolation implies the equivalence of OGPs (Corollary 1.3). In Section 6, we prove that an OGP obstructs overlap-concentrated algorithms on Max-CSPs (Theorem 1.4) and conclude Corollary 1.5. We close with a discussion in Section 7. The appendices contain some technical proofs.

1.2 Related work

Spin glasses. Spin glasses as a state of matter have been studied since the early 20th century. They were first considered as metallic alloys with many ground states, in which the magnetic spins of the individual particles in the alloy are frustrated (i.e. many nearby spins are mismatched). In the language of CSPs, spin glasses have exponentially many near-optimal solutions, in which many constraints are unsatisfied. In this work, we reserve the term “spin glass” for *mean-field* spin glass models, where “mean-field” means that all pairwise (or higher arity) interactions between particles are present. The spins (i.e. the possible values for each variable) are Boolean (also called *Ising*) for all spin glasses that we consider, although in other physical and mathematical settings they may be $[q]$ -valued or vector-valued.

The Sherrington-Kirkpatrick model is an early mathematical model of a spin glass [SK75]. It was solved (by deriving an explicit formula for the free energy density) by Parisi [Par80], but the solution relied on non-mathematically rigorous physical arguments. A later series of works [Gue03, Tal06, Pan13] proved that the formula is correct for all mixed spin glass models. Its numerical value was carefully approximated for the Sherrington-Kirkpatrick model [CR02] and

more recently for other mixed spin glasses [AM20, MH22].

Random CSPs through the lens of statistical physics. Statistical physicists have studied random instances of combinatorial optimization problems since at least the 1980s [FA86, MPV87]. A partial “dictionary” converting between the language of computer science and physics is provided in [CLSS22, Table 1].

Several modern works use the Guerra-Toninelli interpolation [GT04] in a similar technical way as our work; the interpolation method is by now a standard tool in spin glass theory. Dembo, Montanari, and Sen [DMS17] applied the interpolation to prove that the size of the Max-Cut and Max-Bisection in a random d -regular or Erdős-Rényi graph is related to the Sherrington-Kirkpatrick model in the same way as Corollary 1.2. The interpolation was later used (with different spin glass models) to determine the optimal value of random Max- k SAT [Pan18], Max- k XOR and Max q -cut [Sen18] instances in the highly unsatisfiable regime. Compared to these works, our Theorem 1.1 generalizes the CSP predicate to arbitrary mixtures of Boolean functions, but they must have random signs on the literals (e.g. Max-2XOR instead of Max-Cut; this is used in the debiasing argument (Appendix C) and to simplify a technical part of the proof (Equation (54)). We are not certain which of our results extend without random signs. We also extend the Guerra-Toninelli interpolation (in Theorem 3.5) to transfer more properties of the solution geometry than just the optimal value; this is exactly what allows us to compare results on algorithmic hardness.

In this paper, we study the *highly unsatisfiable* regime, where the number of clauses of the CSP is αn for some large constant α . When the clause density α is smaller, for example near the satisfiability threshold of the CSP, the exact connection with spin glasses breaks down, and existing results are less unified. Nonetheless, methods inspired by statistical physics continue to give powerful insight into the solution structure of these CSPs [AM02, ART06, PT04, DSS16b, DSS16a, DSS22].

Overlap gaps. When near-optimal solutions are clustered, it becomes impossible for many algorithms to find them. From a geometric perspective, we can’t “move” from one cluster to others without passing through a lower-value assignment. This general phenomenon was named the *overlap gap property* (OGP) and it was shown to obstruct local algorithms [GS14]. Further generalizations of overlaps [CGPR19, HS21] show stronger obstructions on wider classes of algorithms. The OGP and its generalizations have been used to obstruct algorithms from finding near-optimal solutions of various quantities in mixed spin glasses [GJW20, GJ21, HS21, Sel21], CSPs [CGPR19, BH22, CLSS22], sparse random graphs [GS14, RV17], and matrices [GL18, AWZ20]. See [Gam21, Hua22] for a survey of overlaps and solution geometry.

Optimizing spin glasses and random Max-CSPs. Recently, [AMS21b, Sel21] showed that a type of approximate message-passing algorithm finds the ground state energy of spin glasses without overlap gaps. Under the same assumption, a version of this algorithm was also shown to be optimal on Max-Cut for sparse random graphs with constant (but sufficiently large) degree [AMS21a]. Both the Sherrington-Kirkpatrick model and Max-Cut on sparse random graphs are strongly suspected to have no overlap gaps [ACZ20].

There has been some study of a near-term quantum algorithm (the QAOA [FGG14]) optimizing spin glasses, with recently-proven rigorous performance bounds [CvD21, FGGZ22]. In fact, for large enough clause density, the performance of the QAOA is *identical* on a random instance of Max- k XOR and on its corresponding spin glass [BM21, BFM⁺22, BGMZ22].

2 Preliminaries

2.1 Random constraint satisfaction problems (CSPs)

A CSP instance, denoted \mathcal{I} , consists of n variables and a set of constraints, denoted $E(\mathcal{I})$. In a random CSP instance, the constraints $E(\mathcal{I})$ are drawn from a distribution Λ on functions $f : \Sigma^k \rightarrow \mathbb{R}$.

Definition 2.1 (Instance of a random CSP). *Let Λ be a distribution on functions $f : \Sigma^k \rightarrow [-1, 1]$ with alphabet $\Sigma = \{\pm 1\}$. Fix a constant $\alpha > 0$. Then, a random CSP over n variables $\{\sigma_i\}_{i=1}^n$ and $m = \alpha n$ clauses is generated by: for each $i = 1, \dots, m$, draw i_1, \dots, i_k uniformly i.i.d from $[n]$, draw $f \sim \Lambda$, draw k random signs ε_i uniformly i.i.d from $\{\pm 1\}$, and add the constraint e to $E(\mathcal{I})$, describing the clause $f_e(\sigma_e) := f(\varepsilon_1 \sigma_{i_1}, \dots, \varepsilon_k \sigma_{i_k})$.*

Remark 2.2. *The canonical case is to take Λ which is supported on a single predicate $f : \Sigma^k \rightarrow \{0, 1\}$. For example, the OR predicate corresponds to k SAT. Our proofs apply to the more general setting of [Definition 2.1](#), but it does not apply to $|\Sigma| > 2$ or instances without random signs. Note that the constraints are scaled so that the output of f is always in the interval $[-1, +1]$. Predicates of mixed arity may be simulated using k -ary predicates that ignore some input variables.*

We denote this random model as $\text{CSP}_\Lambda(\alpha)$. We study the *highly unsatisfiable* regime; for example, one can think of $\alpha \gg \exp(k)$.

In this work, we use the language of *Hamiltonians*. In other settings this quantity may be called the objective function, the value, or the score of the assignment.

Definition 2.3 (CSP Hamiltonian). *Consider a $\text{CSP}_\Lambda(\alpha)$ instance \mathcal{I} on $[n]$. For any input $\sigma \in \{\pm 1\}^n$ let*

$$H^\alpha(\sigma) = \frac{1}{\alpha} \sum_{e \in E(\mathcal{I})} f_e(\sigma_e). \quad (4)$$

Remark 2.4. *We divide by α so that, regardless of the value of α , the value of the Hamiltonian is in the same interval $[-n, +n]$.*

Definition 2.5 (Optimal value of a CSP instance). *Define the maximum or optimal value of a $\text{CSP}_\Lambda(\alpha)$ instance \mathcal{I} by:*

$$v_{\mathcal{I}} = \frac{1}{n} \max_{\sigma \in \{\pm 1\}^n} H^\alpha(\sigma). \quad (5)$$

When the predicates in the CSP are 0/1-valued, the maximum value of a CSP instance \mathcal{I} is the maximum possible fraction of constraints that can be satisfied.

Remark 2.6. *By a union bound, with high probability all assignments to a random CSP with predicate f satisfy $\widehat{f}(\emptyset) + O(\frac{1}{\sqrt{\alpha}})$ fraction of constraints. The purpose of our work is to precisely study the behavior of the $\frac{1}{\sqrt{\alpha}}$ term.*

Remark 2.7. *The constraints of an instance of $\text{CSP}_\Lambda(\alpha)$ are slightly dependent because of the fixed number of edges $m = \alpha n$. In the proof we will pass to the “Poisson model” in which the number of edges is $\text{Pois}(\alpha n)$, or equivalently, the multiplicity of each constraint is an i.i.d $\text{Pois}\left(\frac{\alpha}{2^{k-1}}\right)$ random variable.*

2.2 Mean-field spin glasses

Let n denote the number of particles or variables in the system. We introduce the mathematical objects describing a spin glass:

Definition 2.8 (Mixture polynomial). A mixture polynomial is $\xi(s) = \sum_{p \geq 1}^k c_p^2 s^p$ for some nonnegative coefficients c_p^2 .

Definition 2.9 (Gaussian disorder). The disorder coefficients $J^{(p)}$ are an order- p tensor where each axis has length n . When $J^{(p)}$ is a Gaussian tensor, we say that there is Gaussian disorder. Specifically, for each $(i_1, \dots, i_p) \in [n]^p$, we have $J^{(p)}[i_1, \dots, i_p] \sim \mathcal{N}(0, 1)$ i.i.d.³

We now have enough to define our spin glasses. The mixture polynomial specifies the random model, and the disorder coefficients determine the instance of the model. For example, $\xi(s) = s^2$ (or $\xi(s) = s^2/2$ in some works) specifies the well-studied *Sherrington-Kirkpatrick spin glass model*.

Definition 2.10 (Finite mixed spin glass). Fix a mixture polynomial ξ . A spin glass is a random Hamiltonian $H^\xi : \{\pm 1\}^n \rightarrow \mathbb{R}$ given by sampling Gaussian disorder $J := (J^{(p)} : p = 1, 2, \dots)$, where the $J^{(p)}$ are independent, then

$$H^\xi(\sigma) = \sum_{p=1}^k \frac{c_p}{n^{(p-1)/2}} \langle J^{(p)}, \sigma^{\otimes p} \rangle = \sum_{p=1}^k \frac{c_p}{n^{(p-1)/2}} \sum_{(i_1, \dots, i_p) \in [n]^p} J^{(p)}[i_1, \dots, i_p] \sigma_{i_1} \dots \sigma_{i_p}. \quad (6)$$

We say that H^ξ is an instance of the random model \mathbf{SG}_ξ .

\mathbf{SG}_ξ is also known as the *mixed p -spin model*⁴ (in contrast to the pure p -spin model, in which the mixture polynomial is s^p , i.e. all interactions have the same size p).

An early motivator of spin glass theory was the conjecture (and eventual proof) of the following representation for the ground state energy of the model:

Theorem 2.11 (Parisi formula [Par80, Gue03, Tal06, Pan14, AC17]). Fix any spin glass model \mathbf{SG}_ξ . Sample a sequence of independent instances $\{H_{n=1}^\xi, H_{n=2}^\xi, \dots\}$ from \mathbf{SG}_ξ for increasing values of n . Then the limit

$$\text{GSED}(\mathbf{SG}_\xi) := \lim_{n \rightarrow \infty} \max_{\sigma \in \{\pm 1\}^n} \frac{1}{n} H_n^\xi(\sigma), \quad (7)$$

almost surely exists. Furthermore,

$$\text{GSED}(\mathbf{SG}_\xi) \stackrel{\text{a.s.}}{=} \min_{\zeta \in \mathcal{U}} \mathcal{P}_\infty^\xi(\zeta), \quad (8)$$

where \mathcal{P}_∞^ξ is the Parisi functional at zero temperature with mixture polynomial ξ . The Parisi functional at zero temperature and the set \mathcal{U} are defined in [Definition 4.2](#).

³We follow the definition in [Pan13, Chapter 2]. Some definitions of Gaussian disorder use a *symmetric* Gaussian tensor, which is equivalent up to scaling of the c_k and $o_n(1)$ change in the free energy density. This is because the models only differ on tensor entries $[i_1, \dots, i_p]$ such that i_1, \dots, i_p are not all distinct, which only make up $o_n(1)$ fraction of all entries of the tensor.

⁴The $p = 1$ term is a Gaussian external field. Some spin glass models are instead defined with a *fixed* external field.

We state one other fact about spin glasses. One way to characterize a spin glass is as a Gaussian process on the state space $\{\pm 1\}^n$ with covariance structure given by the mixture polynomial ξ . The covariance of the Gaussians $H^\xi(\sigma_1), H^\xi(\sigma_2)$ is a function of the normalized inner product, or *overlap*, of σ_1 and σ_2 .

Fact 2.12 (Gaussian characterization of a spin glass). *Let $\{H^\xi(\sigma)\}_{\sigma \in \{\pm 1\}^n}$ be a spin glass instance of the random model SG_ξ . Then these variables are jointly Gaussian, with mean zero and covariance*

$$\mathbb{E} \left[H^\xi(\sigma_1) H^\xi(\sigma_2) \right] = n \cdot \xi \left(\frac{\langle \sigma_1, \sigma_2 \rangle}{n} \right), \quad (9)$$

where $\langle \cdot, \cdot \rangle$ is the standard inner product.

2.3 Solution geometry of optimization problems

Solution geometry refers to the distribution of the optimal and near-optimal solutions on the Boolean hypercube. Given a Hamiltonian $H : \{\pm 1\}^n \rightarrow \mathbb{R}$, what do we mean by “near-optimal solutions”? We consider two notions. The first, used when considering computational tasks, is the set of σ such that $H(\sigma) \geq v$ for some value v . The second, a “smoother” notion from statistical physics, is that near-optimal solutions are samples from a low-temperature Gibbs distribution.

Definition 2.13 (Gibbs distribution). *Given a Hamiltonian $H : \{\pm 1\}^n \rightarrow \mathbb{R}$ and an inverse temperature parameter $\beta \geq 0$, the Gibbs distribution is the distribution on $\{\pm 1\}^n$ with probability proportional to $e^{\beta H(\cdot)}$.*

Definition 2.14 (Partition function). *The partition function of H is the normalizing constant of the Gibbs distribution, i.e. the exponentially-weighted sum*

$$Z_H(\beta) = \sum_{\sigma \in \{\pm 1\}^n} e^{\beta H(\sigma)}. \quad (10)$$

The most important geometric notion is the *overlap* of two assignments, which is exactly the normalized inner product and is linearly proportional to the Hamming distance between the assignments:

Definition 2.15 (Overlap). *The overlap between two assignments (configurations) $\sigma_1, \sigma_2 \in \{\pm 1\}^n$ is the normalized inner product,*

$$R(\sigma_1, \sigma_2) = \frac{\langle \sigma_1, \sigma_2 \rangle}{n}. \quad (11)$$

Note that the overlap is always in $[-1, 1]$. We study the distribution of $R(\sigma_1, \sigma_2)$ when σ_1, σ_2 are independently sampled near-optimal solutions. Roughly speaking, a gap in the support of this distributions (i.e. an *overlap gap*) implies obstructions for some types of algorithms.

We also study pairwise overlaps between more than two samples. An example is the distribution of the 3-tuple of overlaps in a “triangle” of independent near-optimal solutions $\sigma_1, \sigma_2, \sigma_3$. Topological gaps in this distribution will also obstruct algorithms. The generalization of overlap used in this paper is the *I-overlap*; given a set σ of $|\sigma| = \ell$ assignments, each in $\{\pm 1\}^n$, we define an overlap value for every subset of σ .

Definition 2.16 (*I*-overlap). Consider any $\sigma := (\sigma_1, \dots, \sigma_\ell) \in (\{\pm 1\}^n)^\ell$ and any $I \subseteq [\ell]$. Define the *I*-overlap of σ as

$$R_I(\sigma) = \frac{1}{n} \sum_{j=1}^n \prod_{i \in I} (\sigma_i)_j. \quad (12)$$

The *I*-overlap recovers the definition of overlap when *I* has two elements.

We define an *overlap vector* to associate a set of ℓ assignments with all of its possible *I*-overlaps:

Definition 2.17 (Overlap vector). Consider any $\sigma := (\sigma_1, \dots, \sigma_\ell) \in (\{\pm 1\}^n)^\ell$. Then its overlap vector $Q(\sigma) \in [-1, 1]^{2^{[\ell]}}$ lists all possible *I*-overlaps; that is, $Q(\sigma)_I = R_I(\sigma)$ for every $I \in 2^{[\ell]}$.

Definition 2.18 (Overlap polytope). Consider the set of all possible overlap vectors $Q(\sigma)$ for any positive $n \in \mathbb{N}$ and sets of vectors $\sigma \in (\{\pm 1\}^n)^\ell$. Then $\mathcal{R}^{(\ell)} \subseteq [-1, 1]^{2^{[\ell]}}$ is the closure of this set. Formally,

$$\mathcal{R}^{(\ell)} = \overline{\{Q(\sigma) : n \in \mathbb{N}, \sigma \in (\{\pm 1\}^n)^\ell\}}. \quad (13)$$

Remark 2.19. Since $R_\emptyset(\sigma) = 0$ for all σ , we can ignore this coordinate, and then $\mathcal{R}^{(\ell)}$ is a non-degenerate convex polytope in $\mathbb{R}^{2^\ell - 1}$. For example, $\mathcal{R}^{(2)}$ is a regular tetrahedron in \mathbb{R}^3 .

Definition 2.20 (Preimage of overlap vectors). For any open subset $S \subseteq \mathcal{R}^{(\ell)}$ (in the Euclidean subset topology of $\mathcal{R}^{(\ell)}$), let $U_n^{(\ell)}(S) \subseteq (\{\pm 1\}^n)^\ell$ be the preimage of S in dimension n ; i.e. the set of $\sigma \in (\{\pm 1\}^n)^\ell$ such that $Q(\sigma) \in S$. We drop the superscript when ℓ is in context.

We are interested in the distribution on $\mathcal{R}^{(\ell)}$ of the overlap vector $Q(\sigma_1, \dots, \sigma_\ell)$ when $\sigma_1, \dots, \sigma_\ell$ are independently sampled near-optimal solutions.⁵ Our main result shows that the support of this distribution for a Max-CSP in the $n \rightarrow \infty$ then $\alpha \rightarrow \infty$ limit equals that of a corresponding spin glass. Hence, *overlap gaps* (formally defined in Section 5) transfer between the two models.

The proof uses the concept of *free energy*, defined as the logarithm of the partition function.

Definition 2.21 (Free energy). The free energy of H is $\log Z_H(\beta)$.

The free energy at low temperature ($\beta \rightarrow \infty$) matches the optimal value of the Hamiltonian. This is a commonly-used fact (for example, [Pan13, Equation 1.7]); we include the short proof.

Fact 2.22. The maximum value of a Hamiltonian $H : \Omega \rightarrow \mathbb{R}$ on a finite domain Ω is related to its free energy by

$$\max_{\sigma \in \Omega} H(\sigma) \leq \frac{1}{\beta} \log Z_H(\beta) \leq \max_{\sigma \in \Omega} H(\sigma) + \frac{\log |\Omega|}{\beta}. \quad (14)$$

Proof. We have

$$\frac{1}{\beta} \log \sum_{\sigma \in \Omega} \exp(\beta H(\sigma)) \geq \max_{\sigma \in \Omega} \frac{1}{\beta} \log \exp(\beta H(\sigma)) = \max_{\sigma \in \Omega} H(\sigma), \quad (15)$$

⁵It is likely that the distribution converges in some sense as $n \rightarrow \infty$, and then converges again as $\beta \rightarrow \infty$, but our proof does not show this nor use this.

and

$$\frac{1}{\beta} \log \sum_{\sigma \in \Omega} \exp(\beta H(\sigma)) \leq \max_{\sigma \in \Omega} \frac{1}{\beta} \log |\Omega| \exp(\beta H(\sigma)) = \max_{\sigma \in \Omega} H(\sigma) + \frac{\log |\Omega|}{\beta}. \quad (16)$$

□

2.3.1 Extended remark on multi-overlaps

The I -overlap for $|I| > 2$ is a type of *multi-overlap*, i.e. a higher-order version of the standard overlap. In fact, the structures we will use to obstruct algorithms only constrain the pairwise overlaps with $|I| = 2$. The pairwise overlaps characterize functions that are rotationally-invariant. Multi-overlaps may be useful when studying CSPs at small α ; see the concluding remarks in [Section 7](#).

The I -overlap is general enough to capture any property of the solution geometry which is *permutation-invariant* under the action of S_n permuting the bits of $\{\pm 1\}^n$. Formally, we have the following:

Definition 2.23 (Permutation-invariant function). *Let a permutation $\pi \in S_n$ act on $\sigma \in \{\pm 1\}^n$ as $\sigma_i^\pi = \sigma_{\pi(i)}$ and on a function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$ (extended coordinate-wise to $f : (\{\pm 1\}^n)^\ell \rightarrow \mathbb{R}$) as $f^\pi(\sigma) = f(\sigma^\pi)$. Then f is permutation-invariant if it is fixed by all π .*

Fact 2.24. *Any permutation-invariant function $f : (\{\pm 1\}^n)^\ell \rightarrow \mathbb{R}$ can be expressed as a function on $\mathcal{R}^{(\ell)}$,*

$$f(\sigma_1, \dots, \sigma_\ell) = f(Q(\sigma_1, \dots, \sigma_\ell)). \quad (17)$$

If $f(\sigma_1, \dots, \sigma_\ell)$ is additionally invariant under the action of S_ℓ which permutes the inputs then f is determined by the overlaps $R_{\{1\}}, R_{\{1,2\}}, R_{\{1,2,3\}}, \dots, R_{\{1,2,3,\dots,\ell\}}$. The corresponding overlap polytope is $[-1, +1]^\ell$, which is significantly smaller and simpler. These overlaps are a more common definition of multi-overlap than [Definition 2.16](#) [\[BP22\]](#).

2.4 Fourier analysis on the hypercube

Fourier analysis is commonly used to study Boolean functions [\[O'D14\]](#). For example, the Fourier basis provides a convenient way to understand the action of linear operators on a Boolean function.

We consider the space of Boolean functions $f : \{\pm 1\}^k \rightarrow \mathbb{R}$ with the expectation inner product, over the uniform distribution on $\{\pm 1\}^n$. These functions have a canonical decomposition.

Definition 2.25 (Fourier spectrum of a Boolean function). *Every function $f : \{\pm 1\}^k \rightarrow \mathbb{R}$ permits a unique decomposition as a linear combination of parity functions. Specifically,*

$$f(\sigma) = \sum_{S \subseteq [k]} \widehat{f}(S) \prod_{i \in S} \sigma_i, \quad (18)$$

where $\widehat{f}(S)$ are called the Fourier coefficients.

One can verify that the monomials $\{\prod_{i \in S} \sigma_i\}_{S \subseteq [k]}$ form an orthonormal basis, often called the *Fourier basis*.

Recall that the average value of f is exactly the Fourier coefficient of the empty set:

Fact 2.26. $\mathbb{E}_{\sigma \in \{\pm 1\}^n} [f(\sigma)] = \widehat{f}(\emptyset)$.

We also consider the *noise stability* of a Boolean function, which describes how resistant the function is to independent noise on its input bits.

Definition 2.27 (ρ -correlated). Fix $\sigma \in \{\pm 1\}^k$ and $\rho \in [-1, +1]$. Then a random sample τ of the distribution $N_\rho(\sigma)$ chooses each spin τ_i independently as

$$\tau_i = \begin{cases} \sigma_i & \text{with probability } \frac{1+\rho}{2}, \\ -\sigma_i & \text{with probability } \frac{1-\rho}{2}. \end{cases} \quad (19)$$

Since $\mathbb{E}_{\tau \sim N_\rho(\sigma)} [\tau_i \sigma_i] = \rho$, we say that τ is ρ -correlated with σ .

Definition 2.28 (Noise stability of a Boolean function around a point). For $\rho \in [-1, +1]$, define the noise stability of the Boolean function $f : \{\pm 1\}^k \rightarrow \mathbb{R}$ around point $\sigma \in \{\pm 1\}^n$ as

$$\text{Stab}_\rho[f](\sigma) = \mathbb{E}_{\tau \sim N_\rho(\sigma)} \left[f(\sigma) f(\tau) \right]. \quad (20)$$

Definition 2.29 (Noise stability of a Boolean function). The noise stability of a Boolean function $f : \{\pm 1\}^k \rightarrow \mathbb{R}$ is defined as

$$\text{Stab}_\rho[f] = \mathbb{E}_{\sigma \sim \{\pm 1\}^k} \left[\mathbb{E}_{\tau \sim N_\rho(\sigma)} [f(\sigma) f(\tau)] \right] = \mathbb{E}_{\sigma \sim \{\pm 1\}^k} [\text{Stab}_\rho[f](\sigma)]. \quad (21)$$

Fact 2.30 ([O'D14, Theorem 2.49]).

$$\text{Stab}_\rho[f] = \sum_{S \subseteq [k]} \rho^{|S|} \widehat{f}(S)^2 = \sum_{j=0}^k \rho^j \|f^{=j}\|^2, \quad (22)$$

where $\|f^{=j}\|^2 = \sum_{T \subseteq [k], |T|=j} \widehat{f}(T)^2$ is the Fourier weight of the j th Fourier level of f .

2.5 Concentration inequalities

In the probabilistic combinatorics literature, “with high probability” means with probability $1 - o_n(1)$, where $f = o_n(g)$ denotes $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0$.

Remark 2.31. All of the quantities that we consider are exponentially concentrated. As a consequence, for all $\varepsilon > 0$, the stated theorems hold with ε slack with probability at least $1 - \exp(-Cn)$, for some $C > 0$ depending on ε . For ease of exposition, we have opted to simply state the results with high probability.

Lemma 2.32 (Concentration under spin glass Gibbs distribution, [Pan13, Theorem 1.2]). Let $\beta > 0$ and $f : \{\pm 1\}^n \rightarrow \mathbb{R}_{\geq 0}$ be arbitrary. Let $\{H(\sigma)\}_{\sigma \in \{\pm 1\}^n}$ be a Gaussian process such that $\mathbb{E} H(\sigma)^2 \leq s^2$ for all $\sigma \in \{\pm 1\}^n$. Let $X = \log \sum_{\sigma \in \{\pm 1\}^n} \exp(\beta H(\sigma)) f(\sigma)$. Then for all $x \geq 0$,

$$\Pr \left[|X - \mathbb{E} X| \geq x \right] \leq 2 \exp \left(\frac{-x^2}{4s^2\beta^2} \right). \quad (23)$$

For a spin glass H^ξ , we will think of X as having mean $C \cdot n + o(n)$, which formally holds if f satisfies $\exp(-C' \cdot n) \leq f(\sigma) \leq \exp(C' \cdot n)$ for all σ . By [Fact 2.12](#), $\mathbb{E} H^\xi(\sigma)^2 = n \cdot \xi(1) = O(n)$, so X typically fluctuates by only $O(\sqrt{n})$. An important special case of this concentration inequality is concentration of the maximum value of H . The proof is in [Appendix A](#).

Corollary 2.33. *Let $\{H(\sigma)\}_{\sigma \in \{\pm 1\}^n}$ be a Gaussian process such that $\mathbb{E} H(\sigma)^2 \leq s^2$ for all $\sigma \in \{\pm 1\}^n$. Then for all $x \geq 0$,*

$$\Pr \left[\left| \max_{\sigma \in \{\pm 1\}^n} H(\sigma) - \mathbb{E} \max_{\sigma \in \{\pm 1\}^n} H(\sigma) \right| \geq x \right] \leq 2 \exp \left(\frac{-x^2}{4s^2} \right). \quad (24)$$

We use a similar concentration inequality for the CSP setting. The proof is found in [Appendix A](#).

Lemma 2.34 (Concentration under CSP Gibbs distribution). *Let $\beta > 0$ and $f : \{\pm 1\}^n \rightarrow \mathbb{R}_{\geq 0}$ be arbitrary. Let $\{H(\sigma)\}_{\sigma \in \{\pm 1\}^n}$ be a sample from $\text{CSP}_\Lambda(\alpha)$. Let $X = \log \sum_{\sigma \in \{\pm 1\}^n} \exp(\beta H(\sigma)) f(\sigma)$. Then for all $x \geq 0$,*

$$\Pr \left[|X - \mathbb{E} X| \geq x \right] \leq 2 \exp \left(\frac{-x^2 \alpha}{4(n + \frac{x}{\beta}) \beta^2} \right). \quad (25)$$

Exactly analogously to [Corollary 2.33](#), we conclude concentration of the max.

Corollary 2.35. *Let $\{H(\sigma)\}_{\sigma \in \{\pm 1\}^n}$ be a sample from $\text{CSP}_\Lambda(\alpha)$. Then for all $x \geq 0$,*

$$\Pr \left[\left| \max_{\sigma \in \{\pm 1\}^n} H(\sigma) - \mathbb{E} \max_{\sigma \in \{\pm 1\}^n} H(\sigma) \right| \geq x \right] \leq 2 \exp \left(\frac{-x^2 \alpha}{4(n + x)} \right). \quad (26)$$

3 Sparse and dense models have the same free energy

In this section we prove [Theorem 1.1](#). We first define a *linearly coupled model* for both $\text{CSP}_\Lambda(\alpha)$ and SG_ξ . We connect the free energy of the coupled models via the Guerra-Toninelli interpolation as in several prior works [[Pan05](#), [CGPR19](#), [JS20](#), [CLSS22](#)]; our proof is especially inspired by [[Pan18](#)].

Definition 3.1 ((A, b) -coupled models). *Let $A \in \{0, 1\}^{\ell \times \ell'}$ be a 0-1 matrix and $b \in \mathbb{R}_{\geq 0}^{\ell'}$ a nonnegative vector satisfying $(Ab)_i = 1$ for all $i \in [\ell]$.*

An (A, b) -coupled model of SG_ξ is a collection of random Hamiltonians $\mathcal{G}_1, \dots, \mathcal{G}_\ell$ related by

$$\begin{pmatrix} \mathcal{G}_1 \\ \vdots \\ \mathcal{G}_\ell \end{pmatrix} = A \begin{pmatrix} \mathcal{G}'_1 \\ \vdots \\ \mathcal{G}'_{\ell'} \end{pmatrix}, \quad (27)$$

where each \mathcal{G}'_i is the Hamiltonian for an independent instance of $\text{SG}_{s \rightarrow \sqrt{b_i} \xi(s)}$. The grand Hamiltonian is $\mathcal{G} : (\{\pm 1\}^n)^\ell \rightarrow \mathbb{R}$ defined by $\mathcal{G}(\mathbf{x}) = \sum_{i=1}^\ell \mathcal{G}_i(\mathbf{x}_i)$.

An (A, b) -coupled model of $\text{CSP}_\Lambda(\alpha)$ is a collection of random Hamiltonians $\mathcal{H}_1, \dots, \mathcal{H}_\ell$ related by

$$\begin{pmatrix} \mathcal{H}_1 \\ \vdots \\ \mathcal{H}_\ell \end{pmatrix} = A \begin{pmatrix} \mathcal{H}'_1 \\ \vdots \\ \mathcal{H}'_{\ell'} \end{pmatrix}, \quad (28)$$

where each \mathcal{H}_i is the Hamiltonian for an independent instance of $\text{CSP}_\Lambda(b_i, \alpha)$. The grand Hamiltonian is $\mathcal{H} : (\{\pm 1\}^n)^\ell \rightarrow \mathbb{R}$ defined by $\mathcal{H}(\mathbf{x}) = \sum_{i=1}^\ell \mathcal{H}_i(\mathbf{x}_i)$.

In this section we use \mathcal{G} for a grand Hamiltonian of a spin glass and \mathcal{H} for that of a CSP. Expectations are over the randomness of the Hamiltonians.

Notice that A is $\{0, 1\}$ -valued. One can think of this matrix as choosing which of the ℓ independent “hidden” instances are connected to each of the ℓ “observed” instances. The vector b gives the variances of the hidden instances, and the scaling is set so that the variance of each observed instance is the same as a single instance.

Remark 3.2. *This model describes any linear correlation structure on ℓ Hamiltonians $\mathcal{H}_1, \dots, \mathcal{H}_\ell$, in the sense that it is equivalent to the following alternative model. For the CSP model, there are 2^ℓ parameters $(\eta_S)_{S \subseteq [\ell]}$ such that η_S specifies the fraction of constraints that are simultaneously shared between all of the Hamiltonians in S and the remaining constraints are sampled independently.*

We study the free energy of our models, restricted to sets of assignments with particular overlap vectors.

Definition 3.3 (Free energy of states with given overlap). *Let $S \subseteq \mathcal{R}^{(\ell)}$ be an open set. Let $\mathcal{H} : (\{\pm 1\}^n)^\ell \rightarrow \mathbb{R}$ be a grand Hamiltonian. Define $Z_{\mathcal{H}, S}(\beta)$ as the partition function of \mathcal{H} when the configuration tuples are restricted to those having overlap vector in S ; that is,*

$$Z_{\mathcal{H}, S}(\beta) = \sum_{\mathbf{x} \in U_n^{(\ell)}(S)} e^{\beta \mathcal{H}(\mathbf{x})}. \quad (29)$$

Define the free energy density of \mathcal{H} as

$$\phi_{\mathcal{H}, S}(\beta) = \frac{1}{\beta \ell n} \log Z_{\mathcal{H}, S}(\beta). \quad (30)$$

Remark 3.4. *We ignore the edge case $U_n^{(\ell)}(S) = \emptyset$.*

Now we can state our main interpolation.

Theorem 3.5 (Interpolation of random Max-CSPs and spin glasses; generalized version of [Theorem 1.1](#)). *Choose a positive integer ℓ . Consider any set $S \subseteq \mathcal{R}^{(\ell)}$. Let \mathcal{G} and \mathcal{H} be grand Hamiltonians for (A, b) -coupled models of SG_ξ and $\text{CSP}_\Lambda(\alpha)$, respectively, where ξ is related to Λ as in [Equation \(1\)](#). Then:*

$$\mathbb{E} \phi_{\mathcal{H}, S}(\beta) = \mathbb{E}_{f \sim \Lambda} [f] + \frac{1}{\sqrt{\alpha}} \mathbb{E} \phi_{\mathcal{G}, S}(\beta) + O\left(\frac{\beta^2}{\alpha^2}\right) + o_n(1). \quad (31)$$

where the second-to-last term (which may depend on n) satisfies $\left|O\left(\frac{\beta^2}{\alpha^2}\right)\right| \leq C \cdot \frac{\beta^2}{\alpha^2}$ whenever $\frac{\beta}{\alpha} \leq \varepsilon_0$ for constants $C, \varepsilon_0 > 0$ depending on ℓ .

Furthermore, by the concentration arguments in [Section 2.5](#), we may drop the expectations over \mathcal{G}, \mathcal{H} with the qualitative change that the last term is random and is $o_n(1)$ with high probability.

In order to use the Guerra-Toninelli interpolation, we modify the CSP model so that the number of constraints is $m \sim \text{Pois}(\alpha n)$ rather than $m = \alpha n$ fixed. Because $\text{Pois}(\alpha n)$ concentrates around its mean as $n \rightarrow \infty$, the free energy density is asymptotically the same. The (straightforward) proof of the following lemma is in [Appendix A](#). Using this lemma contributes the $o_n(1)$ term to [Equation \(31\)](#).

Lemma 3.6. *Let $\phi_{\mathcal{H},S}^{(\text{Pois})}$ and $\phi_{\mathcal{H},S}^{(\text{exact})}$ denote the free energy density of a $\text{CSP}_\Lambda(\alpha)$ instance with $m \sim \text{Pois}(\alpha n)$ and $m = \alpha n$ clauses, respectively. Then*

$$\left| \mathbb{E} \phi_{\mathcal{H},S}^{(\text{Pois})} - \mathbb{E} \phi_{\mathcal{H},S}^{(\text{exact})} \right| \leq \frac{1}{\sqrt{\alpha n}} = o_n(1). \quad (32)$$

3.1 Proof of [Theorem 3.5](#)

For notational convenience, we assume in this proof that Λ is supported on a single predicate f . The general case is handled by converting $\widehat{f}(\emptyset)$ back to $\mathbb{E}_{f \sim \Lambda}[f]$.

Fix positive integers ℓ, ℓ' . We define an interpolated Hamiltonian

$$\mathcal{K}(t, \mathbf{x}) = \mathcal{H}^{\alpha(1-t)n}(\mathbf{x}) - (1-t)\ell n \widehat{f}(\emptyset) + \sqrt{\frac{t}{\alpha}} \mathcal{G}(\mathbf{x}), \quad (33)$$

where $\mathcal{H}^{\alpha(1-t)n}$ is the grand Hamiltonian of an (A, b) -coupled model of $\text{CSP}_\Lambda(\alpha(1-t))$. The parameter t controls the interpolation from the Max-CSP (when $t = 0$) to the spin glass (when $t = 1$).

We let $\beta > 0$ be a parameter independent of n . We will later choose β such that $1 \ll \beta \ll \alpha$ as $\alpha \rightarrow \infty$.

Fix an open subset $S \subseteq \mathcal{R}^{(\ell)}$. We write the average free energy density of \mathcal{K} , at inverse temperature β , among states that produce overlap vectors in S :

$$\phi_\beta(t) := \mathbb{E}_{\mathcal{H}, \mathcal{G}} \phi_{\mathcal{K}(t, \cdot), S}(\beta) \quad (34)$$

$$= \frac{1}{\ell n} \frac{1}{\beta} \mathbb{E}_{\mathcal{H}, \mathcal{G}} \log \sum_{\mathbf{x} \in U_n(S)} \exp(\beta \mathcal{K}(t, \mathbf{x})). \quad (35)$$

In this proof, we upper-bound the derivative $\frac{d}{dt} \phi_\beta(t)$, and thereby the difference between $\phi_\beta(0)$ and $\phi_\beta(1)$, showing that the free energy densities of \mathcal{G} and \mathcal{H} are close.

3.1.1 Taking the derivative of $\phi_\beta(t)$

Let's calculate $\frac{d}{dt} \phi_\beta(t)$. First, we generalize \mathcal{K} and ϕ_β so the t -dependence of \mathcal{G} and \mathcal{H} are controlled by independent parameters:

$$\widetilde{\mathcal{K}}(\rho, \gamma, \mathbf{x}) := \mathcal{H}^\rho(\mathbf{x}) - \frac{\rho \ell}{\alpha} \widehat{f}(\emptyset) + \frac{\gamma}{\sqrt{\alpha}} \mathcal{G}(\mathbf{x}), \quad (36)$$

$$\widetilde{\phi}_\beta(\rho, \gamma) := \frac{1}{\beta \ell n} \mathbb{E}_{\mathcal{H}^\rho, \mathcal{G}} \log \sum_{\mathbf{x} \in U_n(S)} \exp(\beta \widetilde{\mathcal{K}}(\rho, \gamma, \mathbf{x})). \quad (37)$$

When $\rho = (1-t)\alpha n$ and $\gamma = \sqrt{t}$, we recover the original expressions:

$$\mathcal{K}(t, \mathbf{x}) = \tilde{\mathcal{K}}((1-t)\alpha n, \sqrt{t}, \mathbf{x}) \quad (38)$$

$$\phi_\beta(t) = \tilde{\phi}_\beta((1-t)\alpha n, \sqrt{t}) \quad (39)$$

We then take a derivative using the chain rule:

$$\frac{d}{dt}\phi_\beta(t) = \left(\frac{\partial}{\partial \rho}\tilde{\phi}_\beta(\rho, \gamma)\right)\frac{d\rho}{dt} + \left(\frac{\partial}{\partial \gamma}\tilde{\phi}_\beta(\rho, \gamma)\right)\frac{d\gamma}{dt} \quad (40)$$

$$= -\alpha n \left(\frac{\partial}{\partial \rho}\tilde{\phi}_\beta(\rho, \gamma)\right) + \frac{1}{2\sqrt{t}} \left(\frac{\partial}{\partial \gamma}\tilde{\phi}_\beta(\rho, \gamma)\right) \quad (41)$$

We also introduce a Gibbs expectation operator to use when computing the partial derivatives. Given a function $p(\mathbf{x})$, the average of p with respect to the Gibbs distribution of $\tilde{\mathcal{K}}(\rho, \gamma, \mathbf{x})$ is

$$\langle p \rangle_{\mathbf{x}} := \frac{\sum_{\mathbf{x} \in U_n(S)} p(\mathbf{x}) \cdot \exp(\beta \tilde{\mathcal{K}}(\rho, \gamma, \mathbf{x}))}{\sum_{\mathbf{x} \in U_n(S)} \exp(\beta \tilde{\mathcal{K}}(\rho, \gamma, \mathbf{x}))}. \quad (42)$$

3.1.2 The Poisson derivative

To calculate $\frac{\partial}{\partial \rho}\tilde{\phi}_\beta(\rho, \gamma)$, we introduce ℓ' new variables $\rho_1, \dots, \rho_{\ell'}$ to parameterize the independent Poisson instances used to construct \mathcal{H}^ρ . We introduce intermediate functions

$$\tilde{\mathcal{K}}(\rho_1, \dots, \rho_{\ell'}, \gamma, \mathbf{x}) = \sum_{i \in [\ell']} \left[\left(A \begin{pmatrix} \mathcal{H}_1^{\rho_1} \\ \vdots \\ \mathcal{H}_{\ell'}^{\rho_{\ell'}} \end{pmatrix} \right)_i(\mathbf{x}_i) - \left(A \begin{pmatrix} \rho_1 \\ \vdots \\ \rho_{\ell'} \end{pmatrix} \right)_i \frac{\tilde{f}(\emptyset)}{\alpha} + \frac{\gamma}{\sqrt{\alpha}} \mathcal{G}_i(\mathbf{x}_i) \right], \quad (43)$$

$$\tilde{\phi}_\beta(\rho_1, \dots, \rho_{\ell'}, \gamma) = \frac{1}{\beta \ell n} \mathbb{E}_{\mathcal{H}_1^{\rho_1}, \dots, \mathcal{H}_{\ell'}^{\rho_{\ell'}}} \mathbb{E}_{\mathcal{G}} \left[\log \sum_{\mathbf{x} \in U_n^{(\ell)}(S)} \exp(\beta \tilde{\mathcal{K}}(\rho_1, \dots, \rho_{\ell'}, \gamma, \mathbf{x})) \right]. \quad (44)$$

When $\rho_{i'} = b_{i'}\rho$ for all $i' \in [\ell']$, we recover the original functions $\tilde{\mathcal{K}}$ and $\tilde{\phi}_\beta$. In this case, through another application of the chain rule,

$$\frac{\partial}{\partial \rho}\tilde{\phi}_\beta(\rho, \gamma) = \sum_{i' \in [\ell']} \left(\frac{\partial}{\partial \rho_{i'}}\tilde{\phi}_\beta(\rho_1, \dots, \rho_{\ell'}, \gamma) \right) \frac{\partial \rho_{i'}}{\partial \rho} = \sum_{i' \in [\ell']} b_{i'} \frac{\partial}{\partial \rho_{i'}}\tilde{\phi}_\beta(\rho_1, \dots, \rho_{\ell'}, \gamma). \quad (45)$$

We use the explicit derivative of a Poisson variable:

Fact 3.7 (Derivative of a Poisson variable).

$$\frac{\partial}{\partial \lambda} \mathbb{E}_{X \sim \text{Pois}(\lambda)} [f(X)] = \mathbb{E}_{X \sim \text{Pois}(\lambda)} [f(X+1) - f(X)] \quad (46)$$

Because of [Fact 3.7](#),

$$\frac{\partial}{\partial \rho_{i'}} \tilde{\phi}_\beta(\rho_1, \dots, \rho_{\ell'}, \gamma) = \tilde{\phi}_\beta(\rho_1, \dots, \rho_{i'} + 1, \dots, \rho_{\ell'}, \gamma) - \tilde{\phi}_\beta(\rho_1, \dots, \rho_{i'}, \dots, \rho_{\ell'}, \gamma) \quad (47)$$

$$= \frac{1}{\beta \ell n} \mathbb{E}_{\mathcal{H}'_1{}^{\rho_1}, \dots, \mathcal{H}'_{\ell'}{}^{\rho_{\ell'}}} \mathbb{E}_{\mathcal{G}} \log \left\langle \exp \left(\beta \tilde{\mathcal{K}}(\rho_1, \dots, \rho_{i'} + 1, \dots, \rho_{\ell'}, \gamma, \mathbf{x}) - \beta \tilde{\mathcal{K}}(\rho_1, \dots, \rho_{i'}, \dots, \rho_{\ell'}, \gamma, \mathbf{x}) \right) \right\rangle_{\mathbf{x}} \quad (48)$$

$$= \frac{1}{\beta \ell n} \mathbb{E}_{\mathcal{H}'_1{}^{\rho_1}, \dots, \mathcal{H}'_{\ell'}{}^{\rho_{\ell'}}} \log \left\langle \exp \left(\beta \sum_{j \in [\ell]} A_{ji'} (\mathcal{H}'_i{}^{\rho_{i'}+1} - \mathcal{H}'_i{}^{\rho_{i'}})(\mathbf{x}_j) - \widehat{f}(\emptyset) \right) \right\rangle_{\mathbf{x}}. \quad (49)$$

The difference of $\mathcal{H}'_i{}^{\rho_{i'}+1}$ and $\mathcal{H}'_i{}^{\rho_{i'}}$ is a single extra clause (normalized by α), applied to k random indices of the input with random signs. Let f^* be the extra clause. Then

$$= \frac{1}{\beta \ell n} \mathbb{E}_{f^*} \log \left\langle \exp \left(\frac{\beta}{\alpha} \sum_{j \in [\ell]} A_{ji'} (f^*(\mathbf{x}_j) - \widehat{f}(\emptyset)) \right) \right\rangle_{\mathbf{x}}. \quad (50)$$

When $\frac{\beta}{\alpha}$ is small, this term can be Taylor-expanded as $\log z = \log(1 - (1 - z)) = -\sum_{p \geq 1} \frac{(1-z)^p}{p}$:

$$= \frac{-1}{\beta \ell n} \mathbb{E}_{f^*} \sum_{p \geq 1} \frac{1}{p} \left(1 - \left\langle \exp \left(\frac{\beta}{\alpha} \sum_{j \in [\ell]} A_{ji'} (f^*(\mathbf{x}_j) - \widehat{f}(\emptyset)) \right) \right\rangle_{\mathbf{x}} \right)^p. \quad (51)$$

We introduce additional ‘‘replicas’’ $\mathbf{x}_{(s)}$ of \mathbf{x} . Precisely, each $\mathbf{x}_{(s)}$ is an i.i.d. copy of \mathbf{x} . For any function w and any set of replicas y_0, y_1, \dots , we have the identity $\langle w(y_0) \rangle^p = \langle \prod_{s \in [p]} w(y_s) \rangle$. As a result,

$$= \frac{-1}{\beta \ell n} \mathbb{E}_{f^*} \left\langle \sum_{p \geq 1} \frac{1}{p} \prod_{s=1}^p \left(1 - \exp \left(\frac{\beta}{\alpha} \sum_{j \in [\ell]} A_{ji'} (f^*(\mathbf{x}_{(s)j}) - \widehat{f}(\emptyset)) \right) \right) \right\rangle_{\mathbf{x}_{(1)}, \mathbf{x}_{(2)}, \dots}. \quad (52)$$

We Taylor-expand this expression in $\frac{\beta}{\alpha}$. The first line comes from the $p = 1$ term and the second line comes from the $p = 2$ term:

$$= \frac{-1}{\beta \ell n} \mathbb{E}_{f^*} \left\langle -\frac{\beta}{\alpha} \sum_{j \in [\ell]} A_{ji'} (f^*(\mathbf{x}_j) - \widehat{f}(\emptyset)) - \frac{\beta^2}{2\alpha^2} \left(\sum_{j \in [\ell]} A_{ji'} (f^*(\mathbf{x}_j) - \widehat{f}(\emptyset)) \right)^2 + \frac{\beta^2}{2\alpha^2} \left(\sum_{j \in [\ell]} A_{ji'} (f^*(\mathbf{x}_j) - \widehat{f}(\emptyset)) \right) \left(\sum_{j \in [\ell]} A_{ji'} (f^*(\mathbf{y}_j) - \widehat{f}(\emptyset)) \right) + O\left(\frac{\beta^3}{\alpha^3}\right) \right\rangle_{\mathbf{x}, \mathbf{y}}. \quad (53)$$

Notice that because the clause applies random signs to its input,

$$\mathbb{E}_{f^*} [\langle f^*(\mathbf{x}_j) \rangle_{\mathbf{x}_j}] = \widehat{f}(\emptyset). \quad (54)$$

As a result, the only terms that remain are

$$= \frac{-\beta}{2\alpha^2 \ell n} \mathbb{E}_{f^*} \left\langle \sum_{i, j \in [\ell]} A_{ii'} A_{jj'} (-f^*(\mathbf{x}_i) f^*(\mathbf{x}_j) + f^*(\mathbf{x}_i) f^*(\mathbf{y}_j)) \right\rangle_{\mathbf{x}, \mathbf{y}} + \frac{1}{\ell n} O\left(\frac{\beta^2}{\alpha^3}\right). \quad (55)$$

From here, we rewrite the correlation of f^* as a function of the noise stability of f . Specifically, let $u_a = \varepsilon_a(\mathbf{x}_i)_{d_a}$ and $v_a = \varepsilon_a(\mathbf{y}_j)_{d_a}$ for some uniformly chosen $\varepsilon \in_{\mathbb{R}} \{\pm 1\}^k$ and $d \in_{\mathbb{R}} [\ell]^k$. Then u, v are marginally uniform points in the hypercube that are $R(\mathbf{x}_i, \mathbf{y}_j)$ -correlated.

Write $f^*(\sigma) = f(\varepsilon_1 \sigma_{d_1}, \dots, \varepsilon_k \sigma_{d_k})$. Then

$$\langle f^*(\mathbf{x}_i) f^*(\mathbf{y}_j) \rangle_{\mathbf{x}_i, \mathbf{y}_j} = \langle f(u) f(v) \rangle_{\mathbf{x}_i, \mathbf{y}_j} = \langle \text{Stab}_{R(\mathbf{x}_i, \mathbf{y}_j)}[f] \rangle_{\mathbf{x}_i, \mathbf{y}_j} \quad (56)$$

where $\text{Stab}_{R(\mathbf{x}_i, \mathbf{y}_j)}[f]$ is defined as in [Definition 2.29](#). Using this, we get

$$= \frac{-\beta}{2\alpha^2 \ell n} \sum_{i, j \in [\ell]} A_{ii'} A_{jj'} \langle -\text{Stab}_{R(\mathbf{x}_i, \mathbf{x}_j)}[f] + \text{Stab}_{R(\mathbf{x}_i, \mathbf{y}_j)}[f] \rangle_{\mathbf{x}, \mathbf{y}} + \frac{1}{\ell n} O\left(\frac{\beta^2}{\alpha^3}\right) \quad (57)$$

$$= \frac{\beta}{2\alpha^2 \ell n} \sum_{i, j \in [\ell]} A_{ii'} A_{jj'} \langle \xi(R(\mathbf{x}_i, \mathbf{x}_j)) - \xi(R(\mathbf{x}_i, \mathbf{y}_j)) \rangle_{\mathbf{x}, \mathbf{y}} + \frac{1}{\ell n} O\left(\frac{\beta^2}{\alpha^3}\right). \quad (58)$$

We can now write the ρ -derivative of $\tilde{\phi}_\beta$:

$$\frac{\partial}{\partial \rho} \tilde{\phi}_\beta = \frac{\beta}{2\alpha^2 \ell n} \sum_{i' \in [\ell']} \sum_{i, j \in [\ell]} b_{i'} A_{ii'} A_{jj'} \langle \xi(R(\mathbf{x}_i, \mathbf{x}_j)) - \xi(R(\mathbf{x}_i, \mathbf{y}_j)) \rangle_{\mathbf{x}, \mathbf{y}} + \frac{1}{\ell n} O\left(\frac{\beta^2}{\alpha^3}\right) \quad (59)$$

3.1.3 The Gaussian derivative

We compute the derivative of $\tilde{\phi}_\beta(\rho, \gamma)$ with respect to γ . Since the variance of each \mathcal{G}_i is independent of γ , the derivative pulls into the expectation operator:

$$\frac{\partial}{\partial \gamma} \tilde{\phi}_\beta(\rho, \gamma) = \frac{1}{\beta \ell n} \mathbb{E}_{\mathcal{H}^\rho, \mathcal{G}} \frac{\partial}{\partial \gamma} \log \left[\sum_{\mathbf{x} \in \mathcal{U}_n(S)} \exp(\beta \tilde{\mathcal{K}}(\rho, \gamma, \mathbf{x})) \right] \quad (60)$$

By the chain rule, the γ -derivative of the partition function is proportional to the Gibbs average of the γ -derivative of the Hamiltonian $\tilde{\mathcal{K}}$:

$$= \frac{1}{\ell n} \mathbb{E}_{\mathcal{H}^\rho, \mathcal{G}} \left[\left\langle \frac{\partial}{\partial \gamma} (\tilde{\mathcal{K}}(\rho, \gamma, \mathbf{x})) \right\rangle_{\mathbf{x}} \right] \quad (61)$$

Notice that the expectation is a correlation between a Gaussian process and a Gibbs measure, so we can use the following formula:

Lemma 3.8 (Stein's lemma with Gibbs average [[Pan13](#), Lemma 1.1]). *Consider two jointly Gaussian processes $\{Y(\sigma)\}_\sigma$ and $\{Z(\sigma)\}_\sigma$. For any w , let $\langle \cdot \rangle_{\sigma_1, \dots, \sigma_w}$ be the w -product Gibbs measure with respect to the process $\{Z(\sigma)\}_\sigma$. Then*

$$\mathbb{E}_{Y, Z} [\langle Y(\sigma) \rangle_\sigma] = \mathbb{E}_Z \left[\left\langle \mathbb{E}_Y [Y(\sigma) Z(\sigma)] \right\rangle_\sigma - \left\langle \mathbb{E}_Y [Y(\sigma) Z(\sigma')] \right\rangle_{\sigma, \sigma'} \right]. \quad (62)$$

Consider the following Gaussian processes:

$$S_{\mathcal{H}^\rho, \mathcal{G}} : (\mathbf{x}) \rightarrow \frac{\partial}{\partial \gamma} (\widetilde{\mathcal{K}}(\rho, \gamma, \mathbf{x})) \quad (63)$$

$$T_{\mathcal{H}^\rho, \mathcal{G}} : (\mathbf{x}) \rightarrow \beta \widetilde{\mathcal{K}}(\rho, \gamma, \mathbf{x}) \quad (64)$$

Applying [Lemma 3.8](#) then yields

$$= \frac{1}{\ell n} \mathbb{E}_{\mathcal{H}^\rho, \mathcal{G}} \left[\left\langle \mathbb{E}_{\mathcal{H}^\rho, \mathcal{G}} [S_{\mathcal{H}^\rho, \mathcal{G}}(\mathbf{x}) T_{\mathcal{H}^\rho, \mathcal{G}}(\mathbf{x})] \right\rangle_{\mathbf{x}} - \left\langle \mathbb{E}_{\mathcal{H}^\rho, \mathcal{G}} [S_{\mathcal{H}^\rho, \mathcal{G}}(\mathbf{x}) T_{\mathcal{H}^\rho, \mathcal{G}}(\mathbf{y})] \right\rangle_{\mathbf{x}, \mathbf{y}} \right]. \quad (65)$$

The derivative only acts on the Gaussian components of the Hamiltonian $\widetilde{\mathcal{K}}$:

$$= \frac{1}{\sqrt{\alpha \ell n}} \mathbb{E}_{\mathcal{H}^\rho, \mathcal{G}} \left[\left\langle \mathbb{E}_{\mathcal{G}} [\mathcal{G}(\mathbf{x}) T_{\mathcal{H}^\rho, \mathcal{G}}(\mathbf{x})] \right\rangle_{\mathbf{x}} - \left\langle \mathbb{E}_{\mathcal{G}} [\mathcal{G}(\mathbf{x}) T_{\mathcal{H}^\rho, \mathcal{G}}(\mathbf{y})] \right\rangle_{\mathbf{x}, \mathbf{y}} \right] \quad (66)$$

$$= \frac{\beta}{\sqrt{\alpha \ell n}} \mathbb{E}_{\mathcal{H}^\rho, \mathcal{G}} \left[\left\langle \mathbb{E}_{\mathcal{G}} \left[\mathcal{G}(\mathbf{x}) \left(\mathcal{H}^\rho(\mathbf{x}) - \frac{\rho \ell}{\alpha} \widehat{f}(\emptyset) + \frac{\gamma}{\sqrt{\alpha}} \mathcal{G}(\mathbf{x}) \right) \right] \right\rangle_{\mathbf{x}} - \left\langle \mathbb{E}_{\mathcal{G}} \left[\mathcal{G}(\mathbf{x}) \left(\mathcal{H}^\rho(\mathbf{y}) - \frac{\rho \ell}{\alpha} \widehat{f}(\emptyset) + \frac{\gamma}{\sqrt{\alpha}} \mathcal{G}(\mathbf{y}) \right) \right] \right\rangle_{\mathbf{x}, \mathbf{y}} \right] \quad (67)$$

The constant terms (proportional to $\widehat{f}(\emptyset)$) cancel. Furthermore, \mathcal{G} is centered and is independent of \mathcal{H} : For each independent instance \mathcal{G}'_i , $\mathbb{E}[\mathcal{G}'_i \mathcal{H}^\rho] = 0$, and by linearity of expectation, $\mathbb{E}[\mathcal{G} \mathcal{H}^\rho] = 0$. All that remains is

$$= \frac{\beta \gamma}{\alpha \ell n} \mathbb{E}_{\mathcal{G}} \left[\left\langle \mathbb{E}_{\mathcal{G}} [\mathcal{G}(\mathbf{x}) \mathcal{G}(\mathbf{x})] \right\rangle_{\mathbf{x}} - \left\langle \mathbb{E}_{\mathcal{G}} [\mathcal{G}(\mathbf{x}) \mathcal{G}(\mathbf{y})] \right\rangle_{\mathbf{x}, \mathbf{y}} \right]. \quad (68)$$

Using the definition of an (A, b) -coupled model, this is

$$= \frac{\beta \gamma}{\alpha \ell n} \mathbb{E}_{\mathcal{G}} \left[\left\langle \sum_{i, j \in [\ell]} \sum_{i', j' \in [\ell']} A_{ii'} A_{jj'} \sqrt{b_{i'} b_{j'}} \mathbb{E}_{\mathcal{G}} [\mathcal{G}'_{i'}(\mathbf{x}_i) \mathcal{G}'_{j'}(\mathbf{x}_j)] \right\rangle_{\mathbf{x}} - \left\langle \sum_{i, j \in [\ell]} \sum_{i', j' \in [\ell']} A_{ii'} A_{jj'} \sqrt{b_{i'} b_{j'}} \mathbb{E}_{\mathcal{G}} [\mathcal{G}'_{i'}(\mathbf{x}_i) \mathcal{G}'_{j'}(\mathbf{y}_j)] \right\rangle_{\mathbf{x}, \mathbf{y}} \right]. \quad (69)$$

By [Fact 2.12](#), we can write the covariance as a function of ξ :

$$\mathbb{E}_{\mathcal{G}} [\mathcal{G}'_{i'}(\mathbf{x}_i) \mathcal{G}'_{j'}(\mathbf{y}_j)] = \begin{cases} n \xi(R(\mathbf{x}_i, \mathbf{y}_j)) & i' = j' \\ 0 & \text{otherwise} \end{cases} \quad (70)$$

Thus, we have

$$= \frac{\beta \gamma}{\alpha \ell} \mathbb{E}_{\mathcal{G}} \left[\left\langle \sum_{i, j \in [\ell]} \sum_{i' \in [\ell']} A_{ii'} A_{jj'} b_{i'} (\xi(R(\mathbf{x}_i, \mathbf{x}_j)) - \xi(R(\mathbf{x}_i, \mathbf{y}_j))) \right\rangle_{\mathbf{x}, \mathbf{y}} \right]. \quad (71)$$

3.1.4 Putting it all together

Now we can calculate the total derivative of $\phi_\beta(t)$:

$$\frac{d}{dt}\phi_\beta(t) = -\alpha n \left(\frac{\partial}{\partial \rho} \tilde{\phi}_\beta(\rho, \gamma) \right) + \frac{1}{2\sqrt{t}} \left(\frac{\partial}{\partial \gamma} \tilde{\phi}_\beta(\rho, \gamma) \right) \quad (72)$$

$$= \left(\frac{-\beta}{2\alpha\ell} + \frac{\beta\gamma}{2\alpha\ell\sqrt{t}} \right) \sum_{i' \in [\ell']} \sum_{i, j \in [\ell]} b_{i' A_{ii'} A_{jj'}} \langle \xi(R(\mathbf{x}_i, \mathbf{x}_j)) - \xi(R(\mathbf{x}_i, \mathbf{y}_j)) \rangle_{\mathbf{x}, \mathbf{y}} - \frac{1}{\ell} O\left(\frac{\beta^2}{\alpha^2}\right) \quad (73)$$

$$= O\left(\frac{\beta^2}{\alpha^2}\right), \quad (74)$$

since $\gamma = \sqrt{t}$ and ℓ is a constant. So

$$|\phi_\beta(1) - \phi_\beta(0)| \leq \max_{t \in [0, 1]} \frac{d}{dt} \phi_\beta(t) = O\left(\frac{\beta^2}{\alpha^2}\right). \quad (75)$$

This proves [Theorem 3.5](#).

4 Optimal value of a random Max-CSP

As a corollary of [Theorem 3.5](#), we prove that in the large α limit, the optimal value of a coupled Max-CSP among solutions with given overlap structure is determined by that of a spin glass.

Corollary 4.1. *The following holds for all ℓ, ℓ' and (A, b) -coupled models \mathcal{G} of SG_ξ and \mathcal{H} of $\text{CSP}_\Lambda(\alpha)$, and all open sets $S \subseteq \mathcal{R}^{(\ell)}$,*

$$\mathbb{E} \max_{\mathbf{x} \in U_n^{(\ell)}(S)} \frac{1}{\ell n} \mathcal{H}(\mathbf{x}) = \mathbb{E} [f] + \mathbb{E} \max_{f \sim \Lambda} \frac{1}{\sqrt{\alpha}} \frac{1}{\ell n} \mathcal{G}(\mathbf{x}) + O(\alpha^{-2/3}) + o_n(1). \quad (76)$$

where the second-to-last term (which may depend on n) satisfies $\left| O(\alpha^{-2/3}) \right| \leq C \cdot \alpha^{-2/3}$ for all $\alpha \geq C$ for a constant C depending on ℓ .

Furthermore, by the concentration arguments in [Section 2.5](#), we may drop the expectations over \mathcal{G}, \mathcal{H} with the qualitative change that the last term is random and is $o_n(1)$ with high probability.

Proof. Using [Theorem 3.5](#), we have

$$\mathbb{E} \phi_{\mathcal{H}, S}(\beta) = \mathbb{E} [f] + \frac{1}{\sqrt{\alpha}} \mathbb{E} \phi_{\mathcal{G}, S}(\beta) + O\left(\frac{\beta^2}{\alpha^2}\right) + o_n(1). \quad (77)$$

Recall that there are at most $2^{\ell n}$ choices of \mathbf{x} . Applying [Fact 2.22](#) for \mathcal{G} and for \mathcal{H} , we get

$$\left| \phi_{\mathcal{G}, S}(\beta) - \max_{\mathbf{x} \in U_n(S)} \frac{1}{\ell n} \mathcal{G}(\mathbf{x}) \right| \leq \frac{\log |U_n(S)|}{\beta \ell n} = O\left(\frac{1}{\beta}\right), \quad (78)$$

$$\left| \phi_{\mathcal{H}, S}(\beta) - \max_{\mathbf{x} \in U_n(S)} \frac{1}{\ell n} \mathcal{H}(\mathbf{x}) \right| = O\left(\frac{1}{\beta}\right). \quad (79)$$

By Jensen's inequality, the same holds when using expectations. Plugging in these bounds and choosing $\beta = \alpha^{2/3}$ proves the claim. \square

In the special case $\ell = \ell' = 1$, $(A, b) = (\mathbb{I}, (1))$, and $S = \mathcal{R}^{(\ell)} = [-1, 1]$, we conclude [Corollary 1.2](#).

4.1 Numerical calculations

We compute the value of the Parisi formula for the spin glasses associated with popular CSPs in [Table 1](#). This code can be run for any choice of spin glass and is available online⁶. Our code uses the zero-temperature representation of the Parisi functional from [\[AC17\]](#), which we restate below:

Definition 4.2 (Parisi functional at zero temperature [\[Par80, AC17\]](#)). *Given a function ζ in*

$$\mathcal{U} = \left\{ \zeta : [0, 1) \rightarrow \mathbb{R}_{\geq 0} : \zeta \text{ is right-continuous, non-decreasing, } \int_0^1 \zeta(t) dt < \infty \right\}, \quad (80)$$

the Parisi functional $\mathcal{P}_\infty^\xi(\zeta)$ is

$$\mathcal{P}_\infty^\xi(\zeta) = \Phi_\zeta(0, 0) - \frac{1}{2} \int_0^1 s \xi''(s) \zeta(s) ds, \quad (81)$$

where the function $\Phi_\zeta(x; t) : \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}$ is the solution of the Hamilton-Jacobi-Bellman equation

$$\partial_t \Phi_\zeta(x; t) + \frac{\xi''(t)}{2} (\partial_{xx} \Phi_\zeta(x; t) + \zeta(t) (\partial_x \Phi_\zeta(x; t))^2) = 0, \quad (82)$$

with initial condition $\Phi_\zeta(x; 1) = |x|$.

As stated in the introduction, the optimum value of a spin glass is $\inf_{\zeta \in \mathcal{U}} \mathcal{P}_\infty^\xi(\zeta)$. The value ALG achievable by efficient algorithms is given by $\inf_{\zeta \in \mathcal{L}} \mathcal{P}_\infty^\xi(\zeta)$ for an extended class of functions $\mathcal{L} \supseteq \mathcal{U}$ defined below. See [\[HS21, Section 6.1\]](#) and references therein for well-posedness considerations.

Definition 4.3 (ALG). *Let $ALG = \inf_{\zeta \in \mathcal{L}} \mathcal{P}_\infty^\xi(\zeta)$, where the class of functions \mathcal{L} is defined by*

$$\mathcal{L} = \left\{ \begin{array}{l} \zeta : [0, 1) \rightarrow \mathbb{R}_{\geq 0} : \zeta \text{ is right-continuous,} \\ \|\xi'' \cdot \zeta\|_{TV[0,1]} < \infty \text{ for all } t \in [0, 1), \int_0^1 \xi''(t) \zeta(t) dt < \infty \end{array} \right\} \quad (83)$$

where the total variation $\|f\|_{TV(J)}$ on an interval J is

$$\|f\|_{TV(J)} = \sup_{n \in \mathbb{N}} \sup_{t_0 < \dots < t_n, t_i \in J} \sum_{i=1}^n |f(t_i) - f(t_{i-1})|. \quad (84)$$

Theorem 4.4 ([\[CGPR19, Lemma 5.4\]](#)). *For even mixture polynomials ξ such that $c_2 = 0$, $ALG < GSED(SG_\xi)$ holds with strict inequality.*

5 Overlap gaps in a random Max-CSP

[Theorem 3.5](#) implies that many quantities are equivalent for $\text{CSP}_\Lambda(\alpha)$ at large α and its associated spin glass model. For example:

1. Free energy of a single instance (proving [Theorem 1.1](#)).
($\ell = 1, \ell' = 1, A = [[1]], b = [1], S = \mathcal{R}$)

⁶<https://github.com/marwahaha/csp-parisi/>

k	Max 1-in- k SAT	Max k NAESAT	Max k SAT	Max k XOR
2	$\frac{1}{2} + \frac{0.54}{\sqrt{\alpha}}$	$\frac{1}{2} + \frac{0.54}{\sqrt{\alpha}}$	$\frac{3}{4} + \frac{0.40}{\sqrt{\alpha}}$	$\frac{1}{2} + \frac{0.54}{\sqrt{\alpha}}$
3	$\frac{3}{8} + \frac{0.54}{\sqrt{\alpha}}$	$\frac{3}{4} + \frac{0.47}{\sqrt{\alpha}}$	$\frac{7}{8} + \frac{0.33}{\sqrt{\alpha}}$	$\frac{1}{2} + \frac{0.58}{\sqrt{\alpha}}$
4	$\frac{1}{4} + \frac{0.48}{\sqrt{\alpha}}$	$\frac{7}{8} + \frac{0.37}{\sqrt{\alpha}}$	$\frac{15}{16} + \frac{0.26}{\sqrt{\alpha}}$	$\frac{1}{2} + \frac{0.58}{\sqrt{\alpha}}$
5	$\frac{5}{32} + \frac{0.41}{\sqrt{\alpha}}$	$\frac{15}{16} + \frac{0.28}{\sqrt{\alpha}}$	$\frac{31}{32} + \frac{0.20}{\sqrt{\alpha}}$	$\frac{1}{2} + \frac{0.59}{\sqrt{\alpha}}$

Table 1: Optimal value of a random k -CSP with n variables and an clauses, as $n \rightarrow \infty$. The calculation uses the zero-temperature Parisi functional in Definition 4.2. All values are rounded to two decimal places. The values for k XOR match those in [MH22]. We expect the values to be accurate to two significant figures, based on consistency with independently calculated values for 2XOR and 3XOR [AMS21b].

2. Free energy of a single instance, restricting the Hamming weight to $W \subseteq [-1, 1]$ (related to [JS20]).
($\ell = 1, \ell' = 1, A = [[1]], b = [1], S = \{r \in \mathcal{R} : r_{\{1\}} \in W\}$)
3. Existence (or non-existence) of an overlap gap property in the overlap range (s, t) .
($\ell = 2, \ell' = 1, A = [[1], [1]], b = [1], S = \{r \in \mathcal{R} : s < r_{\{1,2\}} < t\}$)
4. Existence (or non-existence) of an overlap gap property in the overlap range (s, t) for η -coupled Hamiltonians [CGPR19].
($\ell = 2, \ell' = 3, A = [[1, 1, 0], [1, 0, 1]], b = [\eta, 1 - \eta, 1 - \eta], S = \{r \in \mathcal{R} : s < r_{\{1,2\}} < t\}$)
5. Existence (or non-existence) of the branching overlap gap property in [HS21].
(Choice of parameters in Section 5.1)

We use Theorem 3.5 to show how $\text{CSP}_\Lambda(\alpha)$ inherits the overlap gap property from a spin glass. This proof is a generalization of [CGPR19, Proof of Theorem 5] and [CLSS22, Lemma 8.14] to hold for arbitrary (A, b) -coupled instances. We use the same argument to transfer the branching OGP of [HS21], a hierarchical style of OGP described by a rooted tree.

Definition 5.1 (OGP [Gam21]). *A Hamiltonian $H : \{\pm 1\}^n \rightarrow \mathbb{R}$ exhibits the overlap gap property (OGP) at value v if there are $-1 \leq s < t \leq 1$ such that for all σ_1, σ_2 with $H(\sigma_1) \geq v, H(\sigma_2) \geq v$, we have*

$$R(\sigma_1, \sigma_2) \notin (s, t). \quad (85)$$

Definition 5.2 (Average-OGP, nonstandard definition). *A Hamiltonian H exhibits the average-OGP at value v if the same holds whenever $\frac{1}{2}(H(\sigma_1) + H(\sigma_2)) \geq v$.*

Remark 5.3. *The average-OGP implies the OGP. On the other hand, the OGP implies the average-OGP with a weakened (larger) value. The interpolation in this work transfers the average-OGP.*

Proposition 5.4. *If SG_ξ exhibits the average-OGP at value v with high probability, then for all $\varepsilon > 0$, for all sufficiently large α , $\text{CSP}_\Lambda(\alpha)$ exhibits the average-OGP at value $\mathbb{E}_{f \sim \Lambda}[f] + \frac{v+\varepsilon}{\sqrt{\alpha}}$ with high probability, when ξ is related to Λ as in Equation (1).*

Proof. Let $\ell = 2$, and let $S = (a, b)$ be the overlap gap for the spin glass model SG_ξ . Using

Corollary 4.1, for sufficiently large α , we have w.h.p

$$\max_{\mathbf{x} \in U_n^{(\ell)}(S)} \frac{1}{\ell n} \mathcal{H}(\mathbf{x}) \leq \mathbb{E}_{f \sim \Lambda} [f] + \frac{\max_{\mathbf{x} \in U_n^{(\ell)}(S)} \frac{1}{\ell n} \mathcal{G}(\mathbf{x}) + \varepsilon}{\sqrt{\alpha}} + o_n(1). \quad (86)$$

The overlap gap property implies that w.h.p

$$\max_{\mathbf{x} \in U_n^{(\ell)}(S)} \frac{1}{\ell n} \mathcal{G}(\mathbf{x}) \leq v. \quad (87)$$

Hence with high probability we also have

$$\max_{\mathbf{x} \in U_n^{(\ell)}(S)} \frac{1}{\ell n} \mathcal{H}(\mathbf{x}) \leq \mathbb{E}_{f \sim \Lambda} [f] + \frac{v + 2\varepsilon}{\sqrt{\alpha}}. \quad (88)$$

□

Analogously, we can transfer a generic version of the OGP on (A, b) -coupled models.

Definition 5.5 (Generic OGP). *For a relatively open subset $S \subseteq \mathcal{R}^{(\ell)}$, random Hamiltonians $H_1, \dots, H_\ell : \{\pm 1\}^n \rightarrow \mathbb{R}$ exhibit an S -OGP at value v if*

$$\max_{(\sigma_1, \dots, \sigma_\ell) \in U_n^{(\ell)}(S)} \frac{1}{\ell n} (H_1(\sigma_1) + \dots + H_\ell(\sigma_\ell)) \leq v. \quad (89)$$

The size of the OGP is ℓ .

Proposition 5.6. *If an (A, b) -coupled model of \mathbf{SG}_ξ exhibits an S -OGP at value v with high probability, then for all $\varepsilon > 0$, for all sufficiently large α , the (A, b) -coupled model of $\mathbf{CSP}_\Lambda(\alpha)$ exhibits an S -OGP at value $\mathbb{E}_{f \sim \Lambda} [f] + \frac{v + \varepsilon}{\sqrt{\alpha}}$ with high probability, when ξ is related to Λ as in Equation (1).*

5.1 The branching OGP

The branching OGP gives bounds *tight* algorithmic bounds for a class of spin glass models, matching the performance of certain approximate message passing algorithms [HS21]. We present the somewhat involved definition and show that it is captured by our framework.

Definition 5.7 (Tree-coupled ensemble). *A tree-coupled ensemble of Hamiltonians is defined by:*

1. *A rooted tree of height D defined by the vector $\vec{k} \in \mathbb{Z}_+^D$, so that every node at depth d has k_{d+1} children. We describe the nodes of the tree as $\mathbb{T}(\vec{k})$ and the leaves of the tree as $\mathbb{L}(\vec{k})$.*
2. *A coupling vector $\vec{p} \in \mathbb{Z}_+^{D+1}$, so that*

$$0 = p_0 \leq p_1 \leq \dots \leq p_D = 1. \quad (90)$$

Given these parameters, we generate a family of Hamiltonians $(H^{(u)})_{u \in \mathbb{L}(\vec{k})}$ as follows. Generate independent instances of the Hamiltonian $\tilde{H}^{(v)}$ for every non-root $v \in \mathbb{T}(\vec{k})$, and scale⁷ each by a factor of $p_d - p_{d-1}$ where

⁷For the spin glass, scale the variance of each Gaussian by $p_d - p_{d-1}$; for the CSP, scale the number of clauses by $p_d - p_{d-1}$.

d is the depth of v in the tree. Use these to construct the Hamiltonians $H^{(u)}$ for each $u \in \mathbb{L}(\vec{k})$ defined by

$$H^{(u)} = \sum_{d=0}^{D-1} \tilde{H}^{(a(u,d))}, \quad (91)$$

where $a(u,d)$ is the d th ancestor of u . The size of the ensemble is $|\mathbb{L}(\vec{k})| = \prod_{i=1}^D \vec{k}_i$.

Definition 5.8 (Branching OGP). A tree-coupled ensemble of Hamiltonians exhibits a branching OGP with gap η at value v if:

- there is $\vec{q} \in \mathbb{Z}_+^{D+1}$ with $0 \leq q_0 < q_1 < \dots < q_D = 1$,
- define $Q \in \mathbb{R}^{\mathbb{L}(\vec{k}) \times \mathbb{L}(\vec{k})}$ by $Q_{u,v} = q_{\text{lca}(u,v)}$, where $\text{lca}(u,v)$ is the depth of the least common ancestor of u and v ,
- define

$$\mathcal{Q}(\eta) = \left\{ \vec{\sigma} \in (\{\pm 1\}^n)^{\mathbb{L}(\vec{k})} : \forall u, v \in \mathbb{L}(\vec{k}). |R(\sigma^{(u)}, \sigma^{(v)}) - Q_{u,v}| \leq \eta \right\}$$

- it holds that for all $(\sigma^{(u)})_{u \in \mathbb{L}(\vec{k})} \in \mathcal{Q}(\eta)$, the average value over $u \in \mathbb{L}(\vec{k})$ of $H^{(u)}(\sigma^{(u)})$ is at most v .

The size of the OGP is $|\mathbb{L}(\vec{k})| = \prod_{i=1}^D \vec{k}_i$.

Remark 5.9. Our definition is slightly simplified from [HS21]. We always fix (using their notation) $\mathbf{m} = 0$, because as we show below in [Appendix C](#), it suffices to obstruct algorithms which have mean zero.

A tree-coupled ensemble is an instance of an (A, b) -coupled model, [Definition 3.1](#), using the following choice of parameters:

- $\ell = |\mathbb{L}(\vec{k})|$.
- $\ell' = |\mathbb{T}(\vec{k})|$.
- $A_{u,v} = 1$ if $u = v$, or $v \in \mathbb{T}(\vec{k})$ is a non-root ancestor of $u \in \mathbb{L}(\vec{k})$; otherwise $A_{u,v} = 0$.
- $b_v = p_d - p_{d-1}$ for every $v \in \mathbb{T}(\vec{k})$ at the d th level of the tree.

The presence of the branching OGP is determined by the set

$$S = \{r \in \mathcal{R} : \forall u, v \in \mathbb{L}(\vec{k}). |r_{\{u,v\}} - Q_{u,v}| < \eta\}. \quad (92)$$

Therefore, by [Proposition 5.6](#), the branching OGP transfers from the spin glass to the CSP.

6 Implications for algorithmic hardness

We prove that the OGP has consequences for the performance of many Max-CSP algorithms, specifically those which are *overlap-concentrated*. This section follows [HS21, Section 3] which proves a similar result for spin glasses, with the changes that (1) an instance of (the Poisson version of) the CSP model consists of i.i.d Poisson random variables whereas a spin glass consists of i.i.d Gaussian random variables; (2) we debias the algorithms first to simplify the proof. We do not enforce any

time or space constraints on algorithms and the proofs apply equally well to arbitrary functions satisfying the assumptions.

Definition 6.1 (*t*-correlated instances). For $t \in [0, 1]$, a *t*-correlated pair of instances \mathcal{I}_t^1 and \mathcal{I}_t^2 of the model $\text{CSP}_\Lambda(\alpha)$ is created by:

1. Generate $\text{Pois}(\alpha n)$ random clauses of the CSP. Let \mathcal{I}_t^1 and \mathcal{I}_t^2 both be initially equal to this instance.
2. Independently generate two additional sets of $\text{Pois}((1-t)\alpha n)$ clauses. Add the first set of clauses to \mathcal{I}_t^1 and the second set to \mathcal{I}_t^2 .

Equivalently, we let \mathcal{I}_t^1 be an instance with $\text{Pois}(\alpha n)$ clauses, then independently with probability $1-t$ for each possible clause and sign pattern, we re-sample the multiplicity $\text{Pois}\left(\frac{\alpha}{2^k n^{k-1}}\right)$ of the clause and sign pattern to generate \mathcal{I}_t^2 .

Equivalently, sample $\mathcal{I}_t^1, \mathcal{I}_t^2$ from the (A, b) -coupled model with $A = [[1, 1, 0], [1, 0, 1]]$, $b = [t, 1-t, 1-t]$.

Definition 6.2 (Overlap-concentrated algorithm, [HS21, Definition 2.1]). A deterministic algorithm \mathcal{A} is overlap-concentrated if for every $t \in [0, 1]$ and constant $\delta > 0$, the event

$$\left| R(\mathcal{A}(H_t^1), \mathcal{A}(H_t^2)) - \mathbb{E} R(\mathcal{A}(H_t^1), \mathcal{A}(H_t^2)) \right| \leq \delta \quad (93)$$

occurs w.h.p, where H_t^1, H_t^2 are *t*-correlated instances.

Our results also apply to randomized algorithms which are overlap-concentrated w.h.p over the random seed. For the remainder of the section, we assume the algorithm is deterministic.

By a debiasing procedure, we may assume without loss of generality that an overlap-concentrated algorithm \mathcal{A} satisfies $\mathbb{E}[\mathcal{A}(H)] = 0$. See [Appendix C](#) for details.

Fix an algorithm \mathcal{A} . For $t \in [0, 1]$, let H_t^1 and H_t^2 be a *t*-correlated pair of Hamiltonians for $\text{CSP}_\Lambda(\alpha)$. We define the correlation function $\chi : [0, 1] \rightarrow \mathbb{R}$ for \mathcal{A} by

$$\chi(t) = \mathbb{E} \left[R(\mathcal{A}(H_t^1), \mathcal{A}(H_t^2)) \right]. \quad (94)$$

The correlation function χ satisfies the following properties.

Proposition 6.3. For any algorithm \mathcal{A} run on an instance \mathcal{I} of $\text{CSP}_\Lambda(\alpha)$ such that $\mathbb{E}[\mathcal{A}(\mathcal{I})] = 0$, $\chi : [0, 1] \rightarrow [0, 1]$ satisfies

1. χ is continuous,
2. $\chi(t)$ is strictly increasing,
3. $\chi(0) = 0, \chi(1) = 1$
4. for all $t \in [0, 1]$, $\chi(t) \leq t$.

Proof sketch. As we show in [Appendix B](#), the correlation function $\chi(p)$ is equal to the noise stability of \mathcal{A} viewed as a function of independent Poisson random variables, from which the listed properties easily follow. See the appendix for the full proof. \square

Definition 6.4 (Correlation function, [HS21, Proposition 3.1]). A function $\chi : [0, 1] \rightarrow [0, 1]$ that satisfies the properties of [Proposition 6.3](#) is called a correlation function.

Definition 6.5 (η -forbid an algorithm). An OGP with gap (s, t) η -forbids an algorithm \mathcal{A} if $\mathbb{E} R(\mathcal{A}(H_r^1), \mathcal{A}(H_r^2))$ is in the η -interior of the interval (s, t) for H_r^1 and H_r^2 a pair of r -correlated instances for some $r \in [0, 1]$. In other words, the open ball with radius η around the point is contained in (s, t) .

More generally, an S -OGP η -forbids an algorithm \mathcal{A} if $\mathbb{E} Q(\mathcal{A}(H_1), \dots, \mathcal{A}(H_\ell))$ is in the η -interior of S .

Definition 6.6 (η -forbid a correlation function). An S -OGP η -forbids a correlation function χ if $S \subseteq [-1, 1]^{\binom{\ell}{2}}$ only depends on pairwise overlaps, and $(\chi(\rho_{i,j}))_{i,j}$ is in the η -interior of S where $\rho_{i,j} = \mathbb{E} H_i(\sigma) H_j(\sigma)$ is the correlation of Hamiltonians H_i and H_j (for any fixed arbitrary σ).

Observe that if an OGP forbids the correlation function of an algorithm, then it also forbids the algorithm.

Theorem 6.7 (Existence of even spin glass OGP, [HS21, Proposition 3.2]). For all even functions ξ and $v > \text{ALG}$, there is $\eta > 0$ and an integer ℓ such that for all correlation functions χ there is a coupled model for SG_ξ on at most ℓ Hamiltonians and a set S η -forbidding χ such that with high probability the coupled model exhibits an S -OGP at value v . Specifically, the coupled model is a tree-coupled ensemble and the OGP is a branching OGP.

We now show that the presence of a branching OGP on a random Max-CSP upper-bounds the performance of any overlap-concentrated algorithm.

Theorem 6.8 (Formal version of [Theorem 1.4](#)). Suppose that $\text{CSP}_\Lambda(\alpha)$ exhibits a size- ℓ branching OGP at value v for a constant ℓ . Suppose that an overlap-concentrated algorithm \mathcal{A} is η -forbidden by the OGP for a constant $\eta > 0$. Then w.h.p the output of \mathcal{A} has value at most v .

Proof. Let p be the probability that \mathcal{A} outputs a solution of value at least v when run on a single instance of $\text{CSP}_\Lambda(\alpha)$, so that we want to prove $p \leq o_n(1)$. Fix the correlation structure of the Hamiltonians for which the branching OGP holds, and let $H^{(u)}$ be the correlated Hamiltonians. Consider running algorithm \mathcal{A} on each of the instances $H^{(u)}$. Let \mathcal{E} be the event that the average value of the output solutions is at least v . On the one hand, by [HS21, Proof of Proposition 3.6(a)],

$$\Pr[\mathcal{E}] \geq p^\ell. \quad (95)$$

On the other hand, \mathcal{E} is unlikely, because the expected output of \mathcal{A} lies in the forbidden region of the OGP, and if the output of \mathcal{A} lies in the forbidden region, then by definition \mathcal{E} does not hold. Specifically,

$$\Pr[\mathcal{E}] \leq \Pr[\mathcal{A}(H^{(u)}) \text{ outside of forbidden region}]. \quad (96)$$

The right-hand side can be union-bounded across the $\binom{\ell}{2}$ pairwise overlaps describing the forbidden region. In order to lie outside the region, the deviation from the expectation must be at least η in one of the coordinates. By overlap concentration, the deviation probability is at most $\binom{\ell}{2} \cdot o_n(1) = o_n(1)$. Combining this with [Equation \(95\)](#) and taking the ℓ th root yields the claim. \square

Remark 6.9. Note that [HS21] obstruct a broader class of algorithms that can output solutions in the convex hull of the solution space; since we restrict to algorithms that output valid values, we do not use the technical considerations of [HS21, Section 3.3].

Combining Theorem 6.7, Proposition 5.6, Theorem 6.8, we conclude the following:

Corollary 6.10. Let \mathcal{I} be an instance of $\text{CSP}_\Lambda(\alpha)$ such that every predicate f in the support of Λ is even. For all $\varepsilon > 0$ there is $\alpha_0 > 0$ such that for all $\alpha \geq \alpha_0$ and all overlap-concentrated algorithms, w.h.p it holds that

$$\frac{1}{n}H^\alpha(\mathcal{A}(\mathcal{I})) \leq \mathbb{E}_{f \sim \Lambda} [f] + \frac{ALG + \varepsilon}{\sqrt{\alpha}}. \quad (97)$$

where ALG is defined in Definition 4.3 for ξ related to Λ as in Equation (1).

Using Theorem 4.4, which states that ALG is strictly less than the ground state energy for even spin glasses with no degree-2 part,

Corollary 6.11 (Formal version of Corollary 1.5). Let \mathcal{I} be an instance of $\text{CSP}_\Lambda(\alpha)$ such that for every predicate f in the support of Λ , the only nonzero Fourier coefficients of f have even degree $j \geq 4$. There are constants $\varepsilon > 0, \alpha_0 > 0$ such that for all $\alpha \geq \alpha_0$ and all overlap-concentrated algorithms \mathcal{A} , w.h.p it holds that

$$\frac{1}{n}H^\alpha(\mathcal{A}(\mathcal{I})) \leq v_{\mathcal{I}} - \frac{\varepsilon}{\sqrt{\alpha}}. \quad (98)$$

In particular, for $\alpha = \alpha_0$, the value is bounded away from $v_{\mathcal{I}}$ by a constant factor.

Remark 6.12. As in Remark 2.31, assuming:

1. the OGP holds with probability at least $1 - \exp(-C_1 \cdot n)$, which is shown by [HS21] in their proof of Theorem 6.7
2. \mathcal{A} is overlap-concentrated with probability at least $1 - \exp(-C_2 \cdot n)$ for $C_2 = C_2(\delta)$

then the conclusions of Theorem 6.8, Corollary 6.10, Corollary 6.11 hold with probability at least $1 - \exp(-C \cdot n)$, where C depends on C_1, C_2 .

7 Discussion

We establish a formal average-case link between spin glasses and Max-CSPs of certain minimum clause densities. This link shows an equivalence of optimal value, overlap gaps, and hardness for a large class of algorithms. Curiously, every Max-CSP with the same noise stability polynomial is linked to the same spin glass.

As part of this work, we extend the list of Max-CSPs known to have an OGP. It is an open question to completely classify which spin glasses (and thus which Max-CSPs) have an OGP. In the spherical spin glass setting [Sub18, Proposition 1], the weight of the quadratic terms exactly determine the presence of an OGP; the same may be true on the hypercube. There is a technical hurdle to proving the existence of OGPs on spin glasses when the mixture polynomial is not even. For example, the associated spin glass to Max-3XOR is not proven to have an OGP, although it is expected to have one [AMS21b].

It is also possible that the onset of the branching OGP of [HS21] marks the hardness threshold for all efficient algorithms for spin glasses and Max-CSPs, and not just overlap-concentrated algorithms.

Furthermore, it is possible that a single algorithm (namely, suitably-applied message-passing) is the optimal algorithm. This is remarkably similar to the situation for worst-case analysis vis-à-vis the *Unique Games Conjecture (UGC)* [Kho02]. The UGC implies that the standard SDP is the optimal approximation algorithm for any CSP [Rag08]. It is interesting to investigate whether other algorithms, such as Sum-of-Squares algorithms, can be designed in a way that has equivalent performance to approximate message-passing, although there are existing certification lower bounds for a variety of models [BGL17, GJJ⁺20]. Our curiosity is heightened by the observation that the Parisi formula, and algorithms for spin glasses (and hence average-case CSPs), only use the “degree-2” part of the overlap distribution on $\mathcal{R}^{(\ell)}$, analogous to how the basic SDP is also a “degree-2” algorithm.

How general is overlap concentration? [HS21] show that Langevin dynamics and certain families of approximate message-passing algorithms are overlap-concentrated on spin glasses, and therefore obstructed in the presence of a branching OGP to a constant given by the extended Parisi formula. Algorithms representable as low-degree polynomials are known to be *stable* on spin glasses, and obstructed for some spin glasses [GJW20] and CSPs [BH22]. However, it remains a technical challenge to show that low-degree polynomials are overlap-concentrated:

Conjecture 7.1 (Low-degree polynomials are overlap-concentrated on CSPs and spin glasses). *Any low-degree polynomial algorithm that solves a typical instance of $\text{CSP}_\Lambda(\alpha)$ or SG_ξ as defined in [GJW20, Definition 2.3] is overlap-concentrated.*

There may be other forms of equivalence between spin glasses and Max-CSPs. We conjecture that the two models have the same *distribution* of overlap vectors:

Conjecture 7.2 (Equivalence of distribution of overlap vector, informal). *Take any Max-CSP and consider the associated spin glass SG_ξ , where ξ is defined as in Equation (1). For any coupled ensemble of the models, the distribution of $Q(\sigma_1, \dots, \sigma_\ell)$ for random near-optimal solutions converges in some sense as $n \rightarrow \infty$. Furthermore, for $\alpha \rightarrow \infty$ the distribution for the CSP converges to the distribution for the spin glass.*

It is also likely that algorithms beyond the QAOA have identical average-case performance on a random instance of Max-CSP and that of the corresponding spin glass. In fact, we suspect that every Max-CSP has an optimal algorithm related to message-passing that obeys this equivalence.

Spin glasses may also be related to more classes of CSPs, such as those with non-Boolean inputs and those without random literal signs. Also of interest is the problem of refuting the CSP when it is unsatisfiable, which typically requires α superconstant. It is not known how to use statistical physics methods to study refutation [AOW15, HMX21]. Another open problem (related to studying CSPs at small constraint density α) is to determine the optimal average-case value to higher-order terms in α .

Acknowledgements

We thank Antares Chen for collaborating during the early stages of this project.

We also thank Peter J. Love for helpful feedback on a previous version of this manuscript.

KM thanks Ryan Robinett for a tip on improving the numerical integration scheme. JSS did some of this work as a visiting student at Bocconi University. CJ thanks Tim Stumpf for help with the probabilistic arguments.

We thank the anonymous reviewers for many suggestions to improve the text, for pointing out an error in the proof of [Theorem 6.8](#), and for the reference [\[BP22\]](#).

References

- [AC17] 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. [i](#), [7](#), [20](#)
- [ACO08] 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. [4](#)
- [ACZ20] Antonio Auffinger, Wei-Kuo Chen, and Qiang Zeng. The SK model is infinite step replica symmetry breaking at zero temperature. *Communications on Pure and Applied Mathematics*, 73(5), 2020. [3](#), [5](#)
- [AM02] 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. [5](#)
- [AM20] Ahmed El Alaoui and Andrea Montanari. Algorithmic Thresholds in Mean Field Spin Glasses. *arXiv preprint arXiv:2009.11481*, 2020. [5](#)
- [AMS21a] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Local algorithms for Maximum Cut and Minimum Bisection on locally treelike regular graphs of large degree. *arXiv preprint arXiv:2111.06813*, 2021. [4](#), [5](#)
- [AMS21b] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Optimization of mean-field spin glasses. *The Annals of Probability*, 49(6), November 2021. [1](#), [4](#), [5](#), [21](#), [26](#)
- [AOW15] Sarah R. Allen, Ryan O’Donnell, and David Witmer. How to Refute a Random CSP. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 689–708, 2015. [27](#)
- [ART06] 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. [5](#)
- [AWZ20] Gérard Ben Arous, Alexander S Wein, and Ilias Zadik. Free Energy Wells and Overlap Gap Property in Sparse PCA. In *Conference on Learning Theory*, pages 479–482. PMLR, 2020. [5](#)

- [BFM⁺22] Joao Basso, Edward Farhi, Kunal Marwaha, Benjamin Villalonga, and Leo Zhou. The Quantum Approximate Optimization Algorithm at High Depth for MaxCut on Large-Girth Regular Graphs and the Sherrington-Kirkpatrick Model. In *17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022)*, pages 7:1–7:21, 2022. [5](#)
- [BGL17] Vijay Bhattiprolu, Venkatesan Guruswami, and Euiwoong Lee. Sum-of-Squares Certificates for Maxima of Random Tensors on the Sphere. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017)*, pages 31:1–31:20, 2017. [27](#)
- [BGMZ22] Joao Basso, David Gamarnik, Song Mei, and Leo Zhou. Performance and limitations of the QAOA at constant levels on large sparse hypergraphs and spin glass models. *FOCS 2022; arXiv:2204.10306*, 2022. [5](#)
- [BH22] 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. [3](#), [5](#), [27](#)
- [BM21] Sami Boulebnane and Ashley Montanaro. Predicting parameters for the Quantum Approximate Optimization Algorithm for MAX-CUT from the infinite-size limit. *arXiv preprint arXiv:2110.10685*, 2021. [5](#)
- [BMZ05] Alfredo Braunstein, Marc Mézard, and Riccardo Zecchina. Survey propagation: An algorithm for satisfiability. *Random Structures & Algorithms*, 27(2):201–226, 2005. [3](#)
- [BP22] Jean Barbier and Dmitry Panchenko. Strong replica symmetry in high-dimensional optimal Bayesian inference. *Communications in Mathematical Physics*, pages 1–41, 2022. [10](#), [28](#)
- [CGPR19] Wei-Kuo Chen, David Gamarnik, Dmitry Panchenko, and Mustazee Rahman. Suboptimality of local algorithms for a class of max-cut problems. *Annals of Probability*, 47(3):1587–1618, 2019. [3](#), [5](#), [12](#), [20](#), [21](#)
- [CLSS22] Chi-Ning Chou, Peter J. Love, Juspreet Singh Sandhu, and Jonathan Shi. Limitations of Local Quantum Algorithms on Random MAX-k-XOR and Beyond. In *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, pages 41:1–41:20, 2022. [i](#), [3](#), [5](#), [12](#), [21](#), [33](#)
- [CR02] Andrea Crisanti and Tommaso Rizzo. Analysis of the ∞ -replica symmetry breaking solution of the Sherrington-Kirkpatrick model. *Physical Review E*, 65(4), Apr 2002. [4](#)
- [CvD21] Jahan Claes and Wim van Dam. Instance Independence of Single Layer Quantum Approximate Optimization Algorithm on Mixed-Spin Models at Infinite Size. *Quantum*, 5:542, Sep 2021. [5](#)
- [DMS17] Amir Dembo, Andrea Montanari, and Subhabrata Sen. Extremal cuts of sparse random graphs. *The Annals of Probability*, 45(2):1190–1217, 2017. [1](#), [2](#), [5](#)

- [DSS16a] Jian Ding, Allan Sly, and Nike Sun. Maximum independent sets on random regular graphs. *Acta Mathematica*, 217(2):263–340, 2016. [1](#), [5](#)
- [DSS16b] Jian Ding, Allan Sly, and Nike Sun. Satisfiability Threshold for Random Regular NAE-SAT. *Communications in Mathematical Physics*, 341(2):435–489, Jan 2016. [1](#), [5](#)
- [DSS22] Jian Ding, Allan Sly, and Nike Sun. Proof of the satisfiability conjecture for large k . *Annals of Mathematics*, 196(1):1–388, 2022. [1](#), [5](#)
- [FA86] Yaotian Fu and Philip W. Anderson. Application of statistical mechanics to NP-complete problems in combinatorial optimisation. *Journal of Physics A: Mathematical and General*, 19(9):1605–1620, Jun 1986. [5](#)
- [FGG14] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014. [5](#)
- [FGGZ22] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Leo Zhou. The Quantum Approximate Optimization Algorithm and the Sherrington-Kirkpatrick Model at Infinite Size. *Quantum*, 6:759, July 2022. [5](#)
- [Gam21] David Gamarnik. The overlap gap property: A topological barrier to optimizing over random structures. *Proceedings of the National Academy of Sciences*, 118(41):e2108492118, 2021. [1](#), [3](#), [5](#), [21](#)
- [GJ21] 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. [1](#), [5](#)
- [GJJ⁺20] Mrinalkanti Ghosh, Fernando Granha Jeronimo, Chris Jones, Aaron Potechin, and Goutham Rajendran. Sum-of-Squares Lower Bounds for Sherrington-Kirkpatrick via Planted Affine Planes. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 954–965. IEEE, 2020. [27](#)
- [GJW20] 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. [1](#), [5](#), [27](#)
- [GL18] David Gamarnik and Quan Li. Finding a large submatrix of a Gaussian random matrix. *The Annals of Statistics*, 46(6A):2511–2561, 2018. [5](#)
- [GS14] David Gamarnik and Madhu Sudan. Limits of local algorithms over sparse random graphs. In *Proceedings of the 5th conference on Innovations in Theoretical Computer Science*, pages 369–376, 2014. [3](#), [5](#)
- [GT04] Francesco Guerra and Fabio Lucio Toninelli. The high temperature region of the Viana–Bray diluted spin glass model. *Journal of Statistical Physics*, 115(1):531–555, 2004. [5](#)
- [Gue03] Francesco Guerra. Broken replica symmetry bounds in the mean field spin glass model. *Communications in Mathematical Physics*, 233(1):1–12, 2003. [4](#), [7](#)

- [HMX21] Jun-Ting Hsieh, Sidhanth Mohanty, and Jeff Xu. Certifying solution geometry in random CSPs: counts, clusters and balance. *arXiv preprint arXiv:2106.12710*, 2021. [27](#)
- [HS21] Brice Huang and Mark Sellke. Tight Lipschitz hardness for optimizing mean field spin glasses. *arXiv preprint arXiv:2110.07847*, 2021. [i](#), [1](#), [3](#), [4](#), [5](#), [20](#), [21](#), [22](#), [23](#), [24](#), [25](#), [26](#), [27](#), [36](#)
- [Hua22] Brice Huang. *Computational Hardness in Random Optimization Problems from the Overlap Gap Property*. Master’s thesis, Massachusetts Institute of Technology, 2022. [5](#)
- [JS20] Aukosh Jagannath and Subhabrata Sen. On the unbalanced cut problem and the generalized Sherrington–Kirkpatrick model. *Annales de l’Institut Henri Poincaré D*, 8(1):35–88, 2020. [12](#), [21](#)
- [Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the thirty-fourth annual ACM Symposium on Theory of Computing*, pages 767–775, 2002. [27](#)
- [KLLM21] Peter Keevash, Noam Lifshitz, Eoin Long, and Dor Minzer. Global hypercontractivity and its applications. *arXiv preprint arXiv:2103.04604*, 2021. [36](#)
- [MH22] Kunal Marwaha and Stuart Hadfield. Bounds on approximating Max k XOR with quantum and classical local algorithms. *Quantum*, 6:757, 2022. [5](#), [21](#)
- [Mon21] Andrea Montanari. Optimization of the Sherrington-Kirkpatrick Hamiltonian. *SIAM Journal on Computing*, pages FOCS19–1–FOCS19–38, 2021. [4](#)
- [MPV87] Marc Mézard, Giorgio Parisi, and Miguel Angel Virasoro. *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, volume 9. World Scientific Publishing Company, 1987. [5](#)
- [O’D14] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014. [10](#), [11](#), [36](#)
- [Pan05] Dmitry Panchenko. Free energy in the generalized Sherrington–Kirkpatrick mean field model. *Reviews in Mathematical Physics*, 17(07):793–857, 2005. [12](#)
- [Pan13] Dmitry Panchenko. *The Sherrington-Kirkpatrick model*. Springer Science & Business Media, 2013. [4](#), [7](#), [9](#), [11](#), [17](#)
- [Pan14] Dmitry Panchenko. The Parisi formula for mixed p -spin models. *The Annals of Probability*, 42(3):946–958, 2014. [7](#)
- [Pan18] Dmitry Panchenko. On the K-sat model with large number of clauses. *Random Structures & Algorithms*, 52(3):536–542, 2018. [2](#), [5](#), [12](#)
- [Par80] Giorgio Parisi. A sequence of approximated solutions to the SK model for spin glasses. *Journal of Physics A: Mathematical and General*, 13(4):L115, 1980. [i](#), [4](#), [7](#), [20](#)
- [Pas13] Vangelis Th Paschos. *Paradigms of Combinatorial Optimization: Problems and New Approaches*, volume 2. John Wiley & Sons, 2013. [1](#)
- [PT04] Dmitry Panchenko and Michel Talagrand. Bounds for diluted mean-fields spin glass models. *Probability Theory and Related Fields*, 130(3):319–336, 2004. [5](#)

- [Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the fortieth annual ACM Symposium on Theory of Computing*, pages 245–254, 2008. [27](#)
- [RV17] Mustazee Rahman and Balint Virag. Local algorithms for independent sets are half-optimal. *The Annals of Probability*, 45(3):1543–1577, 2017. [5](#)
- [Sel21] Mark Sellke. Optimizing Mean Field Spin Glasses with External Field. *arXiv preprint arXiv:2105.03506*, 2021. [4](#), [5](#)
- [Sen18] Subhabrata Sen. Optimization on sparse random hypergraphs and spin glasses. *Random Structures & Algorithms*, 53(3):504–536, 2018. [1](#), [2](#), [5](#)
- [SK75] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical Review Letters*, 35(26):1792, 1975. [1](#), [4](#)
- [Sub18] Eliran Subag. Following the ground-states of full-RSB spherical spin glasses, 2018. [3](#), [26](#)
- [Tal06] Michel Talagrand. The parisi formula. *Annals of Mathematics*, pages 221–263, 2006. [i](#), [4](#), [7](#)
- [Zha21] Yu Zhao. *Generalizations and Applications of Hypercontractivity and Small-Set Expansion*. PhD thesis, Carnegie Mellon University, 2021. [36](#)

Appendix

A Concentration Proofs

In this section, we furnish the proofs of concentration inequalities that we use. We start with [Corollary 2.33](#). The exponent in this corollary can be improved to $\exp\left(-\frac{x^2}{2n}\right)$, which is the statement of the Borell-TIS inequality, but we utilize this argument since it analogously proves [Corollary 2.35](#).

Corollary 2.33. *Let $\{H(\sigma)\}_{\sigma \in \{\pm 1\}^n}$ be a Gaussian process such that $\mathbb{E} H(\sigma)^2 \leq s^2$ for all $\sigma \in \{\pm 1\}^n$. Then for all $x \geq 0$,*

$$\Pr\left[\left|\max_{\sigma \in \{\pm 1\}^n} H(\sigma) - \mathbb{E} \max_{\sigma \in \{\pm 1\}^n} H(\sigma)\right| \geq x\right] \leq 2 \exp\left(\frac{-x^2}{4s^2}\right). \quad (24)$$

Proof. Take $f(\sigma) = 1$ in [Lemma 2.32](#) to obtain the free energy $X = Z_H(\beta)$. By the lemma, we have

$$\Pr\left[\left|\frac{1}{\beta} Z_H(\beta) - \mathbb{E} \frac{1}{\beta} Z_H(\beta)\right| \geq x\right] \leq 2 \exp\left(\frac{-x^2}{4s^2}\right). \quad (99)$$

On the other hand, by [Fact 2.22](#), $\lim_{\beta \rightarrow \infty} \frac{1}{\beta} Z_H(\beta) \stackrel{\text{a.s.}}{=} \max_{\sigma \in \{\pm 1\}^n} H(\sigma)$, and since the right-hand side concentration is uniform in β , we may take the limit $\beta \rightarrow \infty$ on the left-hand side to conclude the claim. Formally, a.s. convergence implies $\lim_{\beta \rightarrow \infty} \mathbb{E} \frac{1}{\beta} Z_H(\beta) = \mathbb{E} \max_{\sigma \in \{\pm 1\}^n} H(\sigma)$. Slutsky's theorem then implies

$$\frac{1}{\beta} Z_H(\beta) - \mathbb{E} \frac{1}{\beta} Z_H(\beta) \stackrel{\text{a.s.}}{\rightarrow} \max_{\sigma \in \{\pm 1\}^n} H(\sigma) - \mathbb{E} \max_{\sigma \in \{\pm 1\}^n} H(\sigma). \quad (100)$$

This implies that the probability of lying in the test region $(-\infty, -x] \cup [x, \infty)$ is the same between the right side and the limit of the left-hand side. \square

We will now prove [Lemma 2.34](#). This proof applies to the Poisson model. The same proof works for the exact model, because McDiarmid's inequality is proven with a martingale argument, which still holds for non-independent random variables. We will use the following variant of McDiarmid's inequality, which applies to a Lipschitz function of i.i.d random variables which are highly biased.

Lemma A.1 ([\[CLSS22, Lemma 7.6 \(arXiv version\)\]](#)). *Suppose that X_1, \dots, X_N are sampled i.i.d from a distribution D over a finite set \mathcal{X} such that D assigns probability $1 - p$ to a particular outcome $x_0 \in \mathcal{X}$. Let $F : \mathcal{X}^N \rightarrow \mathbb{R}$ satisfy a bounded-differences inequality, so that*

$$\left|F(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_N) - F(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_N)\right| \leq c \quad (101)$$

for all $x_1, \dots, x_n, x'_i \in \mathcal{X}$. Then

$$\Pr\left[\left|F(X_1, \dots, X_N) - \mathbb{E} F(X_1, \dots, X_N)\right| \geq \varepsilon\right] \leq 2 \exp\left(\frac{-\varepsilon^2}{2Np(2-p)c^2 + 2c\varepsilon/3}\right). \quad (102)$$

We deduce the following corollary for Poisson random variables.

Corollary A.2. Suppose that X_1, \dots, X_N are sampled i.i.d from $\text{Pois}(p)$. Let $F(X_1, \dots, X_N)$ satisfy the bounded-differences inequality, so that

$$\left| F(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_N) - F(x_1, \dots, x_{i-1}, x_i + 1, x_{i+1}, \dots, x_N) \right| \leq c \quad (103)$$

for all $x_1, \dots, x_n \in \mathbb{N}$. Then

$$\Pr\left[\left| F(X_1, \dots, X_N) - \mathbb{E} F(X_1, \dots, X_N) \right| \geq \varepsilon \right] \leq 2 \exp\left(\frac{-\varepsilon^2}{2Np(2-p)c^2 + 2c\varepsilon/3} \right). \quad (104)$$

Proof. Fix $m \in \mathbb{N}$. We approximate $\text{Pois}(p) \approx \text{Bin}(m, p/m)$, apply the lemma to the corresponding (finitely many, 0/1-supported) Bernoulli random variables, then take the limit $m \rightarrow \infty$.

Formally, let $(X_{i,j})_{i \in [N], j \in [m]}$ be i.i.d $\text{Ber}(p/m)$, and let

$$F'(X_{i,j}) := F\left(\sum_{j=1}^m X_{1,j}, \dots, \sum_{j=1}^m X_{N,j} \right). \quad (105)$$

Applying the version of McDiarmid's inequality to F' , we conclude

$$\Pr\left[\left| F'(X_{i,j}) - \mathbb{E} F'(X_{i,j}) \right| \geq \varepsilon \right] \leq 2 \exp\left(\frac{-\varepsilon^2}{2Np(2-p)c^2 + 2c\varepsilon/3} \right). \quad (106)$$

Notice that the right-hand side is independent of m . Taking the limit $m \rightarrow \infty$ independent of all the other parameters, the left-hand side converges to $\Pr\left[\left| F(X_1, \dots, X_N) - \mathbb{E} F(X_1, \dots, X_N) \right| \geq \varepsilon \right]$ as desired. This is justified as follows. We have $\text{Bin}(m, p/m) \xrightarrow{d} \text{Pois}(p)$ as $m \rightarrow \infty$. Writing $F'(X_{i,j}) = F(X'_1, \dots, X'_N)$ where $X'_i \sim \text{Bin}(m, p/m)$, by the continuous mapping theorem (F is continuous since the domain is discrete) and Slutsky's theorem,

$$F(X'_1, \dots, X'_N) - \mathbb{E} F(X'_1, \dots, X'_N) \xrightarrow{d} F(X_1, \dots, X_N) - \mathbb{E} F(X_1, \dots, X_N).$$

Hence we may conclude that the probability of lying in the test region $(-\infty, -\varepsilon] \cup [\varepsilon, \infty)$ is the same between the right side and the limit of the left-hand side. \square

Lemma 2.34 (Concentration under CSP Gibbs distribution). Let $\beta > 0$ and $f : \{\pm 1\}^n \rightarrow \mathbb{R}_{\geq 0}$ be arbitrary. Let $\{H(\sigma)\}_{\sigma \in \{\pm 1\}^n}$ be a sample from $\text{CSP}_\Lambda(\alpha)$. Let $X = \log \sum_{\sigma \in \{\pm 1\}^n} \exp(\beta H(\sigma)) f(\sigma)$. Then for all $x \geq 0$,

$$\Pr\left[\left| X - \mathbb{E} X \right| \geq x \right] \leq 2 \exp\left(\frac{-x^2 \alpha}{4(n + \frac{x}{\beta}) \beta^2} \right). \quad (25)$$

Proof. We consider the n^k random variables X_e equal to the multiplicity of each edge $e \in [n]^k$ and the function $F = \frac{1}{\beta} X$. The X_e are i.i.d as $\text{Pois}(\alpha/n^{k-1})$. If a single edge is added, then since $H(\sigma)$ changes by at most

$$\frac{\max_{f \in \text{supp}(\Lambda), x \in \{\pm 1\}^k} |f(x)|}{\alpha} \leq \frac{1}{\alpha}, \quad (107)$$

we also have that F changes by only at most $\frac{1}{\alpha}$.

Therefore we may apply [Corollary A.2](#) with constant $c = \frac{1}{\alpha}$, $p = \frac{\alpha}{n^{k-1}}$, and $\varepsilon = \frac{x}{\beta}$. The denominator in the right-hand side exponent of [Equation \(104\)](#) is

$$\frac{-\varepsilon^2}{2Np(2-p)c^2 + 2c\varepsilon/3} \quad (108)$$

$$= \frac{-x^2}{\left(\frac{2n}{\alpha}\left(2 - \frac{\alpha}{n^{k-1}}\right) + \frac{2x}{3\alpha\beta}\right)\beta^2} \quad (109)$$

$$= \frac{-x^2\alpha}{\left(2n\left(2 - \frac{\alpha}{n^{k-1}}\right) + \frac{2x}{3\beta}\right)\beta^2} \quad (110)$$

$$\leq \frac{-x^2\alpha}{4\left(n + \frac{x}{\beta}\right)\beta^2}. \quad (111)$$

This completes the proof. \square

Lemma 3.6. Let $\phi_{\mathcal{H},S}^{(\text{Pois})}$ and $\phi_{\mathcal{H},S}^{(\text{exact})}$ denote the free energy density of a $\text{CSP}_\Lambda(\alpha)$ instance with $m \sim \text{Pois}(\alpha n)$ and $m = \alpha n$ clauses, respectively. Then

$$\left| \mathbb{E} \phi_{\mathcal{H},S}^{(\text{Pois})} - \mathbb{E} \phi_{\mathcal{H},S}^{(\text{exact})} \right| \leq \frac{1}{\sqrt{\alpha n}} = o_n(1). \quad (32)$$

Proof. Let m be the number of edges in the Poisson model and αn in the exact model. Couple the two models so that the first αn edges of the Poisson model equal the exact model. Comparing the Hamiltonians of the two instances, we find

$$\frac{1}{\alpha} \sum_{e \in E(\mathcal{I}^{(\text{exact})})} f_e(\sigma_e) - \frac{1}{\alpha} \cdot |m - \alpha n| \leq \frac{1}{\alpha} \sum_{e \in E(\mathcal{I}^{(\text{Pois})})} f_e(\sigma_e) \leq \frac{1}{\alpha} \sum_{e \in E(\mathcal{I}^{(\text{exact})})} f_e(\sigma_e) + \frac{1}{\alpha} \cdot |m - \alpha n|, \quad (112)$$

where we have used that $|f(x)| \leq 1$. This gives a corresponding bound on the difference in free energy density:

$$\left| \phi_{\mathcal{H},S}^{(\text{Pois})} - \phi_{\mathcal{H},S}^{(\text{exact})} \right| \leq \frac{|m - \alpha n|}{\alpha n}. \quad (113)$$

Taking the expectation,

$$\left| \mathbb{E} \phi_{\mathcal{H},S}^{(\text{Pois})} - \mathbb{E} \phi_{\mathcal{H},S}^{(\text{exact})} \right| \leq \mathbb{E} \left| \phi_{\mathcal{H},S}^{(\text{Pois})} - \phi_{\mathcal{H},S}^{(\text{exact})} \right| \quad (\text{Jensen's inequality}) \quad (114)$$

$$\leq \frac{\mathbb{E}_{m \sim \text{Pois}(\alpha n)} |m - \alpha n|}{\alpha n} \quad (\text{Equation (113)}) \quad (115)$$

$$\leq \frac{\left(\mathbb{E}_{m \sim \text{Pois}(\alpha n)} (m - \alpha n)^2\right)^{1/2}}{\alpha n} \quad (\text{Jensen's inequality}) \quad (116)$$

$$= \frac{\sqrt{\alpha n}}{\alpha n} = \frac{1}{\sqrt{\alpha n}}. \quad (117)$$

\square

B Correlation functions of sparse Hamiltonians

In this section, we prove [Proposition 6.3](#). The proof uses Fourier analysis of \mathcal{A} in the same way as [\[HS21, Proof of Proposition 3.1\]](#). We must use the more general Efron-Stein basis and a mild generalization of the noise operator defined in [Definition 2.27](#).

Definition B.1 (Efron-Stein decomposition). *For a function $f \in L^2(\Omega^m, \pi^{\otimes m})$, we define $f^{\subseteq S}(x) = \mathbb{E}_{x_{\bar{S}}} [f(x) \mid x_S]$, where the expectation means that the variables x_i for $i \in \bar{S}$ are independently resampled from π , and $f^{\subseteq S}(x) = \sum_{J \subseteq S} (-1)^{|S|-|J|} f^{\subseteq J}(x)$.*

A core property of this decomposition, as stated in [\[KLLM21, Section II.2.2\]](#), is that $f^{\subseteq S}$ depends only on x_i for $i \in S$, and furthermore $\langle f^{\subseteq S}, g \rangle = 0$ for any g that does not depend on all x_i for $i \in S$. Additionally, by [\[O'D14, Theorem 8.35 and Proposition 8.36\]](#),

$$f = \sum_{S \subseteq [m]} f^{\subseteq S} = \sum_{S \subseteq [m]} \left(\sum_{\alpha \in \mathbb{N}_{|\Omega|}^m, \text{supp}(\alpha) = S} \widehat{f}(\alpha) \phi_\alpha \right), \quad (118)$$

where $\mathbb{N}_{|\Omega|}^m = \{0, \dots, |\Omega| - 1\}^m$ and,

$$\phi_\alpha = \prod_{j=1}^m \phi_{\alpha_j}, \quad (119)$$

with $\phi_0 = \text{Id}$.

Using this decomposition, any $\mathcal{A} : \Omega^m \rightarrow [-1, 1]^n$ under the input measure $\pi^{\otimes m}$ can be written as

$$\mathcal{A} = (f_1, \dots, f_n), \quad (120)$$

where each f_i can be described as in [Definition B.1](#).

We now define the noise operator on a domain Ω^m with an arbitrary product measure $\pi^{\otimes m}$:

Definition B.2 (Noise operator, [\[KLLM21, Zha21\]](#)). *Given a probability space (Ω, π) and a noise parameter $\rho \in [0, 1]$, define the noise operator T_ρ , which acts on a function $f \in L_2(\Omega^m, \pi^{\otimes m})$ as*

$$T_\rho[f](\sigma_1, \dots, \sigma_m) = \mathbb{E}_{\tau \sim N_\rho(\sigma)} \left[f(\tau_1, \dots, \tau_m) \right], \quad (121)$$

where the distribution $N_\rho(\sigma)$ chooses each coordinate τ_i independently as

$$\begin{cases} \tau_i = \sigma_i & \text{with probability } \rho, \\ \tau_i \sim \pi & \text{with probability } 1 - \rho. \end{cases} \quad (122)$$

Proof of [Proposition 6.3](#). Define the Fourier weight on a level j as

$$W_j = \frac{1}{n} \sum_{i=1}^n \sum_{S \subseteq [m], |S|=j} \|\mathcal{A}_i^{\subseteq S}\|^2. \quad (123)$$

Then, when x, y are generated as in [Definition 6.1](#) with $t = p$,

$$\begin{aligned}\chi(p) &= \mathbb{E} [R(\mathcal{A}(x), \mathcal{A}(y))] = \frac{1}{n} \mathbb{E} [\langle \mathcal{A}(x), \mathcal{A}(y) \rangle] = \frac{1}{n} \mathbb{E} [\langle \mathcal{A}(x), T_p[\mathcal{A}](x) \rangle] \\ &= \frac{1}{n} \mathbb{E} [\langle T_{\sqrt{p}}[\mathcal{A}](x), T_{\sqrt{p}}[\mathcal{A}](x) \rangle] = \frac{1}{n} \sum_{i=1}^n \sum_{S \subseteq [m]} p^{|S|} \|f_i^{=S}\|^2 = \sum_{j \geq 0} p^{|S|} W_j.\end{aligned}\quad (124)$$

This follows for two reasons. First, the noise operator acts on a $p \cdot m$ -sized random subset of $[m]$, equivalent to the (average) number of shared clauses generated by [Definition 6.1](#). Second, the noise operator satisfies the semi-group property of $T_p = T_{\sqrt{p}} \circ T_{\sqrt{p}}$.

The desired properties for $\chi(p)$ now follow straightforwardly. $\chi(p)$ is continuous as it is a polynomial in p . It is strictly increasing as the coefficients are nonnegative and not all zero. $\chi(0) = 0$ and $\chi(1) = 1$ by computation. Finally we have

$$\chi(p) = \sum_{j \geq 1} p^{|S|} W_j \leq p \cdot \left(\sum_{j \geq 1} W_j \right) = p. \quad (125)$$

□

C Debiasing

Proposition C.1 (Debiasing spin glass algorithms). *Let \mathcal{A} be a deterministic algorithm for optimizing an instance H of SG_ξ such that*

1. \mathcal{A} is overlap-concentrated
2. $H(\mathcal{A}(H)) \geq v$ whp.

Then there is a deterministic algorithm \mathcal{A}' such that $\mathbb{E} \mathcal{A}'(H) = 0$ and

1. \mathcal{A}' is overlap-concentrated
2. $H(\mathcal{A}'(H)) \geq v - o(n)$ whp.

Proof. We use the first column of $J^{(2)}$ to symmetrize the algorithm. Since the first column is only a small fraction of the input, this will not significantly affect the output value. Our technique can be interpreted as a derandomization procedure for randomized spin glass algorithms that use a small number of random bits.

In the case of $\xi(s) = c_1^2 s$, the optimal algorithm is $\mathcal{A}(H) = \text{sgn}(J^{(1)})$ and we ignore this case.

Let $\tilde{C} \in \mathbb{R}^n$ be the first column of $J^{(2)}$ and let $C = \text{sgn}(\tilde{C}) \in \{\pm 1\}^n$ be the entrywise sign. Let

$$J'^{(p)} := J^{(p)} \circ C^{\otimes p} \quad \text{restricted to coordinates 2 through } n. \quad (126)$$

where \circ is the entrywise/Hadamard product. The algorithm \mathcal{A}' runs \mathcal{A} on $J'^{(p)}$ to produce a candidate solution $\sigma' \in \{\pm 1\}^{n-1}$ on coordinates 2 through n . Then it outputs $C \circ [1, \sigma']$.

To compute the expectation, observe that C and $J'^{(p)}$ are probabilistically independent. Therefore

$$\mathbb{E}[C \circ [1, \sigma']] = \mathbb{E}[C] \circ [1, \mathbb{E}[\sigma']] = 0 \circ [1, \mathbb{E}[\sigma']] = 0. \quad (127)$$

The value of the output is

$$H(C \circ [1, \sigma']) = \sum_{p \geq 1} \frac{c_p}{n^{(p-1)/2}} \langle J^{(p)}, C^{\otimes p} \circ [1, \sigma']^{\otimes p} \rangle \quad (128)$$

$$= \sum_{p \geq 1} \frac{c_p}{n^{(p-1)/2}} \left(\langle J^{(p)}, \sigma'^{\otimes p} \rangle + \sum_{\substack{(i_1, \dots, i_p) \in [n]^p: \\ \text{at least one equals 1}}} J^{(p)}[i_1, \dots, i_p] \sigma'_{i_1} \cdots \sigma'_{i_p} \right) \quad (129)$$

Since σ' is independent of $J^{(p)}[i_1, \dots, i_p]$ such that at least one index equals 1, the second term is a mean-zero Gaussian with variance $O(n^{p-1})$. Whp,

$$= \sum_{p \geq 1} \frac{c_p}{n^{(p-1)/2}} \left(\langle J^{(p)}, \sigma'^{\otimes p} \rangle \pm O(n^{(p-1)/2} \sqrt{\log n}) \right) \quad (130)$$

$$= \sum_{p \geq 1} \frac{c_p}{(n-1)^{(p-1)/2}} \langle J^{(p)}, \sigma'^{\otimes p} \rangle \pm o(n) \quad (131)$$

$$\geq v - o(n). \quad (132)$$

where the last inequality holds whp by assumption on \mathcal{A} .

It remains to show that \mathcal{A}' remains overlap-concentrated. Consider t -correlated Hamiltonians H_t^1, H_t^2 . Let $C^1, C^2 \in \{\pm 1\}^n$ be the signs of the respective first columns of $J^{(2)}$ and let $q = \mathbb{E}[C_1^1 C_1^2]$ (note that all of the coordinates have the same distribution). We have

$$\mathbb{E}[R(\mathcal{A}'(H_t^1), \mathcal{A}'(H_t^2))] = q \cdot \mathbb{E}[R(\sigma^1, \sigma^2)]. \quad (133)$$

We also have

$$R(\mathcal{A}'(H_t^1), \mathcal{A}'(H_t^2)) = \frac{1}{n} C_1^1 C_1^2 + \frac{1}{n} \sum_{i=2}^n \sigma_i^1 \sigma_i^2 C_i^1 C_i^2. \quad (134)$$

As before, σ' and C^1, C^2 are independent. With σ' fixed, the terms $\sigma_i^1 \sigma_i^2 C_i^1 C_i^2$ are independent with deviation from the mean bounded by 2. Their mean over C^1, C^2 is $q \cdot R(\sigma^1, \sigma^2)$. Therefore by Bernstein's inequality,

$$\Pr_{C^1, C^2} \left[\left| R(\mathcal{A}'(H_t^1), \mathcal{A}'(H_t^2)) - q \cdot R(\sigma^1, \sigma^2) \right| \geq \frac{L}{n} \right] \leq \exp\left(\frac{-L^2/2}{n + 2L/3} \right).$$

By overlap concentration,

$$\Pr \left[\left| q \cdot R(\sigma^1, \sigma^2) - q \cdot \mathbb{E}[R(\sigma^1, \sigma^2)] \right| \geq \delta \right] \leq o_n(1). \quad (135)$$

Take $\frac{L}{n} = \delta$ and apply the triangle inequality and a union bound to conclude the claim. \square

Proposition C.2 (Debiasing random CSP algorithms). *The same statement holds for the model $\text{CSP}_\Lambda(\alpha)$.*

Proof. In the CSP model, we may use the first $\frac{n}{\log n}$ constraints to symmetrize the algorithm. Since the entropy of a single constraint is approximately $k \log n$, this will provide at least n bits of entropy for symmetrization.⁸ Any $\omega(1)$ bits of entropy are enough to maintain overlap concentration whp.

Let \mathcal{I} be the instance of $\text{CSP}_\Lambda(\alpha)$. Let $\tilde{\mathcal{C}}$ be the first $\frac{n}{\log n}$ clauses of the CSP and let $C \in \{\pm 1\}^n$ extract the bit labels of the first variable in each clause, so that C equals n independent random bits. Let \mathcal{I}' be the CSP instance consisting of the remaining clauses of \mathcal{I} , then flipping the sign of the variable i by C_i in each constraint.

The algorithm \mathcal{A}' runs \mathcal{A} on \mathcal{I}' to produce a candidate solution $\sigma' \in \{\pm 1\}^n$. Then it outputs $C \circ \sigma'$.

The analysis proceeds as in the previous proof. Since C and \mathcal{I}' are independent,

$$\mathbb{E}[C \circ \sigma'] = \mathbb{E}[C] \circ \mathbb{E}[\sigma'] = 0 \circ \mathbb{E}[\sigma'] = 0. \quad (136)$$

The number of clauses used in $\tilde{\mathcal{C}}$ is $o(n)$, hence the Hamiltonian value is affected by at most $o(n)$.

Overlap concentration of \mathcal{A}' follows exactly as in the prior proof. \square

⁸This is called *randomness extraction* in computer science literature.