

RESEARCH ARTICLE

WILEY

The impact of heterogeneity and geometry on the proof complexity of random satisfiability

Thomas Bläsius¹ | Tobias Friedrich² | Andreas Göbel² | Jordi Levy³ |
Ralf Rothenberger²

¹Karlsruhe Institute of Technology, Karlsruhe, Germany

²Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

³IIA, CSIC, Campus UAB, Bellaterra, Spain

Correspondence

Andreas Göbel, Hasso Plattner Institute, University of Potsdam, Potsdam, Germany.
Email: andreas.gobel@hpi.de

Funding information

German Research Foundation (DFG),
Grant/Award Number: 416061626

Abstract

Satisfiability is considered the canonical NP-complete problem and is used as a starting point for hardness reductions in theory, while in practice heuristic SAT solving algorithms can solve large-scale industrial SAT instances very efficiently. This disparity between theory and practice is believed to be a result of inherent properties of industrial SAT instances that make them tractable. Two characteristic properties seem to be prevalent in the majority of real-world SAT instances, heterogeneous degree distribution and locality. To understand the impact of these two properties on SAT, we study the proof complexity of random k -SAT models that allow to control heterogeneity and locality. Our findings show that heterogeneity alone does not make SAT easy as heterogeneous random k -SAT instances have superpolynomial resolution size. This implies intractability of these instances for modern SAT-solvers. In contrast, modeling locality with underlying geometry leads to small unsatisfiable subformulas, which can be found within polynomial time.

KEYWORDS

geometry/locality, heterogeneity, proof complexity, random satisfiability, voronoi diagrams

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Random Structures and Algorithms* published by Wiley Periodicals LLC.

1 | INTRODUCTION

Propositional satisfiability (SAT) is arguably among the most-studied problems for both theoretical and practical research. Nonetheless, the gap between theory and practice is huge. In theory, SAT is the prototypical hard problem and hardness of other problems is shown via reductions from SAT. Achieving even a running time of $O(2^{cn})$ for any $c < 1$ and n variables would be a major breakthrough and a somewhat surprising one at that. On the contrary, reductions to SAT are used to solve various problems appearing in practice, as state-of-the-art SAT solvers can easily handle industrial instances with millions of variables.

This theory–practice gap does not come from the lack of a sufficiently precise theoretical analysis of modern SAT solvers. They are actually provably slow on most instances, that is, drawing an instance uniformly at random yields a hard instance with probability tending to 1 for $n \rightarrow \infty$, if the clause-variable ratio is not too low or way too high [9,21]. Instead, the discrepancy comes from the fact that industrial instances have properties that make them easier than worst-case instances. In 2014, Vardi [59] wrote that “we have no understanding of why the specific sets of heuristics employed by modern SAT solvers are so effective in practice” and that we need this understanding to successfully advance SAT solving further.

In recent years, scientists have been studying properties of industrial SAT instances to gain this understanding. By modeling SAT instances as graphs, for example, with edges indicating inclusion of variables in clauses, one can benefit from the extensive research conducted in the field of network science. Two properties commonly observed in real-world networks are heterogeneity and locality. *Heterogeneity* refers to the degree distribution, meaning that vertices have strongly varying degrees. In fact, one usually observes a heavy-tailed distribution with many vertices of low degree and few vertices of high degree. A common assumption is a power-law distribution [60], where the number of vertices of degree k is roughly proportional to $k^{-\beta}$. The constant β is called the *power-law exponent*. *Locality* refers to the fact that edges tend to connect vertices that are close in the sense that they remain well connected even when ignoring their direct connection. This can also be seen as having strong community structures, with high connectivity within communities and loose ties between communities.

With respect to these two properties, industrial SAT instances are similar to real-world networks. In many cases, the variable frequencies are heterogeneous [1] and there is a high level of locality [2]. The latter is often measured in terms of modularity. Inspired by network science, researchers have studied models that resemble industrial instances with respect to these properties. Particularly, Ansótegui et al. [4] introduced a power-law SAT model for heterogeneous instances, which has been theoretically studied in terms of satisfiability thresholds [33–35]. A different model with heterogeneous degree distributions has been studied by Cooper et al. [22], Ansótegui et al. [3], and Omelchenko and Bulatov [52]. Moreover, Giráldez-Cru and Levy [40] introduced a model in which variable weights lead to heterogeneity while an underlying geometry facilitates locality. Comparing this to network models, the former model [4] is the SAT-variant of Chung-Lu graphs [19,20]. The latter [40] is based on the popularity-similarity model [53], which is closely related to hyperbolic random graphs [44] and geometric inhomogeneous random graphs [18].

Besides serving as somewhat realistic benchmarks for SAT competitions [39], these SAT models can be used to study solver behavior depending on heterogeneity and locality. One can experimentally observe that a high level of heterogeneity improves the performance of SAT solvers that also perform well on industrial instances [4,13]. Moreover, locality seems very beneficial as solvers appear to implicitly use the locality of a given instance [40]. This coincides with the findings of experiments on actual industrial instances that show that the locality (measured using modularity) of an instance is a good predictor for solver performance [51,62,63].

Up to date, there are no theoretical results supporting these experimental observations. On the contrary, it has been shown that instances generated by the community attachment model [38], which enforces a community structure, are hard for modern SAT solvers [48]. With this article, we provide a theoretical foundation that matches the observations in practice by studying the proof complexity of k -SAT instances (for constant k) drawn from the power-law SAT model, and from a very general model with underlying geometry. The former was introduced by Ansótegui et al. [4], the latter is a generalization of the geometric model by Giráldez-Cru and Levy [40] in the same way as geometric inhomogeneous random graphs [18] are a generalization of hyperbolic random graphs [44]. Our findings are that heterogeneous instances are hard asymptotically almost surely¹ in that modern SAT solvers require superpolynomial or even exponential running time to refute unsatisfiable instances. On the contrary, instances with a high level of locality facilitated by an underlying geometry are a. a. s. easy to solve. Our results focus on unsatisfiable instances, that is, on the case where a solver has to prove that no satisfying assignment exists. This is typically much harder than finding a satisfying assignment, making the unsatisfiable regime arguably more relevant. Besides these results on SAT, we provide insights on the complexity of weighted higher-order Voronoi diagrams in higher dimensions, which is of independent interest.

The power-law and geometric models both mimic specific properties observed in industrial instances while trying to make as little additional assumptions as possible. Though this makes the resulting instances arguably more realistic than, for example, instances drawn uniformly at random, we want to stress that even the geometric model is far from a perfect representation of industrial instances. Thus, our results do not claim to completely explain the efficiency of modern SAT solvers on industrial instances. However, to the best of our knowledge, we provide the first theoretical result that links a high level of locality to provably more tractable instances, which we believe to be a first step towards closing the theory–practice gap.

Outline

We state and discuss our main results and technical contributions in Section 2. Formal definitions are in Section 3. A short outline of our core arguments is in Section 4, followed by the formal proofs: lower bounds for the power-law model in Section 5, upper bounds on the complexity of Voronoi diagrams in Section 6, and upper bounds for the geometric SAT model in Section 7. To not distract from the core arguments, results we use that were either known before or are straight-forward to prove are outsourced to Appendix A.

2 | RESULTS, TECHNICAL CONTRIBUTION, DISCUSSION

In this section, we state our results and discuss the contribution, also in context to previous results. To make the results understandable, we briefly discuss, for example, the probability distributions over SAT formulas we study. These are short and not meant to be formal definitions. For complete definitions, see Section 3.

¹Asymptotically almost surely (a. a. s.) refers to a probability that tends to 1 for $n \rightarrow \infty$. With high probability (w. h. p.) refers to the stronger requirement that the probability is in $1 - O(1/n)$. Additionally, we say that an event holds with overwhelming probability, if for every $c > 0$ it holds with probability at least $1 - O(n^{-c})$.

2.1 | Power-law SAT

The power-law SAT model has four parameters: the number of variables n , the number of clauses m , the number k of variables appearing in each clause, and a power-law exponent β . To draw a formula, power-law weights with exponent β are assigned to the variables and then each clause is generated independently by drawing k variables without repetition using probabilities proportional to the weights. Each literal is negated with probability $1/2$.

To discuss our first main contribution, let Φ be a formula drawn from the power-law model with density at or above the satisfiability threshold, that is, Φ is unsatisfiable at least with constant probability. We show that, although it is likely that Φ is unsatisfiable, it is highly unlikely that modern SAT solvers can figure that out in polynomial time. We prove this using resolution proof complexity.

Resolution is a refutation technique for propositional and first-order logic introduced by [24]. If an application of resolution steps leads to a contradiction, the formula is unsatisfiable. The sequence of resolved clauses then serves as a proof for unsatisfiability, also called a refutation of the formula. The resolution proof system exhibits a strong connection to modern Davis–Putnam–Logemann–Loveland (DPLL) and conflict-driven clause learning (CDCL) SAT solvers: DPLL is polynomially equivalent to tree-like resolution [58] and CDCL with unlimited restarts is polynomially equivalent to resolution [7,54]. Thus, the minimum number of steps necessary to derive a contradiction also yields a lower bound on the running time of solvers simulating the same process. This number of steps is also called the *resolution size* of a formula, that is, the minimum number of resolution steps necessary to arrive at a contradiction. Equivalently, the width of a resolution proof is the size of the largest clause appearing in the proof and the *resolution width* of a formula is the smallest width of any proof refuting that formula. Interestingly, a lower bound w on the resolution width of a formula also implies a lower bound on its resolution size [9]: every resolution proof of a formula in k -CNF has size $\exp(\Omega((w-k)^2/n))$ and every tree-like resolution proof has size 2^{w-k} .

We will show a lower bound for the resolution width of unsatisfiable formulas drawn from the power-law model. Our results translate to lower bounds on the resolution size and thus to matching lower bounds on the running time of CDCL solvers. For DPLL solvers, which use tree-like resolution, the bounds are even stronger. We only consider the resolution width of unsatisfiable instances. Thus, the probability bound we get is actually a conditional probability conditioned on instances being unsatisfiable. Note that our bound does not only hold above the satisfiability threshold, where a random formula Φ is a. a. s. unsatisfiable, but also *at* the threshold, where it is unsatisfiable with constant probability.

Theorem 5.8. *Let Φ be an unsatisfiable random power-law k -SAT formula with n variables, $m \in \Omega(n)$ clauses, $k \geq 3$, and power-law exponent $\beta > \frac{2k-1}{k-1}$. Let $\Delta = m/n$ be large enough so that Φ is unsatisfiable at least with constant probability. Let $\epsilon, \epsilon_1, \dots, \epsilon_3$ be constants with $\epsilon > 0$, $\epsilon_1 = \frac{k-\epsilon}{2} - 1 > 0$, $\epsilon_2 = (k-\epsilon) \cdot \frac{\beta-2}{\beta-1} - 1 > 0$, and $0 < \epsilon_3 < (\frac{k}{2} - 1) \cdot \frac{\beta-2}{\beta-1} - 1$. For the resolution width w of Φ , it holds a. a. s. that:*

- (i) *If $\beta \in (\frac{2k-1}{k-1}, 3)$ and $\Delta \in o(n^{\epsilon_2})$, then $w \in \Omega(n^{\epsilon_2/\epsilon_1} \Delta^{-1/\epsilon_1})$.*
- (ii) *If $\beta = 3$ and $\Delta \in o(n^{\epsilon_1/\log^{1+\epsilon_1} n})$, then $w \in \Omega(n \cdot \Delta^{-1/\epsilon_1} \log^{1+\epsilon_1} n)$.*
- (iii) *If $\beta > 3$ and $\Delta \in o(n^{\epsilon_1})$, then $w \in \Omega(n \cdot \Delta^{-1/\epsilon_1})$.*
- (iv) *If $\beta > \frac{2k-2}{k-2}$ and $\Delta \in o(n^{\epsilon_3/\log^{\epsilon_3} n})$, then $w \in \Omega(n \cdot \Delta^{-1/\epsilon_3})$.*

The above lower bounds allow the density Δ to be super-constant (even polynomial), which is asymptotically above the satisfiability threshold. For the sake of simplicity, assume Δ to be constant in the following. Starting at the bottom (iii, iv), we get a linear bound for w if β is sufficiently large,

that is, greater than 3 or $(2k-2)/(k-2)$. For $\beta = 3$ (ii), the bound is still almost linear. Note that these results in particular imply exponential lower bounds on the resolution size and thus on the running time of CDCL and DPLL. For smaller β (i), we get a polynomial bound for the width with exponent $\varepsilon_2/\varepsilon_1$; see Figure 1 for a plot with ε close to 0.

Interestingly enough, our bounds only hold for power law exponents $\beta > \frac{2k-1}{k-1}$. This is complemented by a previous result [34], which shows that the satisfiability threshold of power-law random k -SAT is at density $\Delta = \Theta(1)$ for power law exponents $\beta > \frac{2k-1}{k-1}$ and that asymptotically almost surely instances with constant constraint densities are trivially unsatisfiable for power law exponents $\beta < \frac{2k-1}{k-1}$. Thus, the resolution width is constant in the latter case.

Part iv of Theorem 5.8 is derived via lower bounds on the bipartite expansion of the clause-variable incidence graph of these instances. These results can be of independent interest for hypergraphs with edge size k and for random $(0, 1)$ -matrices. Additionally, these expansion properties yield lower bounds for the clause space complexity, which in turn gives lower bounds on the tree-like resolution size of such formulas (Section 5.2). More precisely, this results in an exponential lower bound on the tree-like resolution size for $\beta > \frac{2k-3}{k-2}$. This is an improvement of the bound obtained via resolution width.

It is interesting to note that this result on the non-geometric model supports the claim that locality is a crucial factor for easy SAT instances. The lower bounds for the power-law model are solely based on the fact that every set of clauses covers a comparatively large set of variables. In other words, we only use that there are no clusters of clauses with similar variables, that is, we explicitly use the lack of locality.

2.2 | Geometric SAT

The geometric model has the following parameters: n , m , and k have the same meaning as for the power-law model. Moreover, w is a weight function assigning each variable v a weight w_v and T is the so-called temperature that controls the strength of locality by varying the impact of the geometry. As underlying geometric space, we use the d -dimensional torus $\mathbb{T}^d = \mathbb{R}^d/\mathbb{Z}^d$ (see Section 3) equipped with a \mathfrak{p} -norm with $\mathfrak{p} \in \mathbb{N}^+ \cup \infty$. To draw a formula, the variables and clauses are assigned random positions in \mathbb{T}^d . Then, for each clause, k variables are drawn without repetition with probabilities depending on the variable weight and on the geometric distance between clause and variable. In the extreme case of $T = 0$, each clause deterministically includes the k closest variables (where closeness is a combination of geometric distance and weight), while increasing the temperature T increases the probability for the inclusion of more distant variables. For $T \rightarrow \infty$, the model converges to uniform random SAT. Note that the weights are a parameter of the model and not drawn randomly. We have the following theorem, where W denotes the sum of all variable weights. The condition on the weights is in particular satisfied by power-law distributed weights.




FIGURE 1 Exponent of the bound (i) in Theorem 5.8. Dashed vertical lines show where the bound (iv) takes over.

Theorem 7.12. *Let Φ be a formula with n variables and $m \in \Theta(n)$ clauses drawn from the weighted geometric model with ground space \mathbb{T}^d equipped with a p -norm, temperature $T < 1$, $W \in O(n)$, and $w_v \in O(n^{1-\epsilon})$ for every $v \in V$ and any constant $\epsilon > 0$. Then, Φ contains a. a. s. an unsatisfiable subformula of constant size, which can be found in $O(n \log n)$ time.*

To briefly explain how we prove this, consider a simplified version where variables and clauses are points in the Euclidean plane and each clause contains the k variables geometrically closest to it (temperature $T = 0$). Now consider the equivalence relation obtained by defining two points of the plane equivalent if and only if they have the same set of k closest variables. The equivalence classes of this relation are the regions of the order- k Voronoi diagram of the variable positions. With this connection, we can use upper bounds on the complexity of order- k Voronoi diagrams [46] to prove the existence of small and easy to find unsatisfiable subformulas. We note that this result is of asymptotic nature. In particular for small densities, the number of variables n has to be very large before the instances actually get as easy as stated in Theorem 7.12. Nevertheless, this result strongly suggests that an underlying geometry makes SAT instances more tractable.

To extend the above argument to the general statement in Theorem 7.12, we extend the complexity bounds for order- k Voronoi diagrams in various ways; see next section for more details. Moreover, for non-zero temperatures, clauses no longer include exactly the k closest variables but can, in principle, consist of any set of k variables. However, we can show that, with high probability, a linear fraction of clauses behaves as in the $T = 0$ case. We note that analyses of similar structures, such as hyperbolic random graphs, are often restricted to the simpler but less realistic $T = 0$ case, for example, [11,12,14,50]. We believe that our analysis provides insights on the non-zero temperature case that can be helpful for such related questions.

We note that our results seem to contradict the results of Mull et al. [48], stating that (i) a strong community structure is not sufficient to have tractable SAT instances and that (ii) the community attachment model [38], which enforces a community structure, generates hard instances. However, at a closer look, this is not a contradiction at all. Though measuring the community structure, for example, via modularity, is a good indicator for locality, the concept of locality goes deeper. If the instance can be partitioned such that there are strong ties within each partition and loose ties between partitions, then the instance has a strong community structure. However, to have a high level of locality, this concept has to hierarchically repeat on different levels of magnitude, that is, there needs to be community structure within each partition and between the partitions. To state this slightly differently, consider locality based on a notion of similarity between objects (here: variables or clauses). In this article, we use distances between random points in a geometric space as a measure for similarity, which gives us a continuous range of more or less similar objects. In contrast to that, in the above mentioned papers focusing on a flat community structure [38,48], similarity is a binary equivalence relation: two objects are either similar or they are not.

2.3 | Voronoi diagrams

Consider a finite set of points, called sites, in a geometric space. The most commonly studied type of Voronoi diagram assumes the 2-dimensional Euclidean plane as ground space and has one Voronoi region for each site, containing all points closer to this site than to any other site. We deviate from this default setting in four ways: (i) We allow an arbitrary constant dimension d , where the ground space is the torus or a hypercube in \mathbb{R}^d . (ii) We consider the order- k Voronoi diagram, which has for every subset A of sites with $|A| = k$ a (possibly empty) Voronoi region containing all points for which




FIGURE 2 (A) Weighted Voronoi diagram (order-1) of the colored sites. Continuing the construction with $n/2$ high-weight sites on the left and $n/2$ low-weight sites towards the right yields $\Theta(n^2)$ vertices (small black dots). Note that each vertex lies on the boundary of three regions and has thus equal weighted distance to its three closest sites. (B) The order-3 Voronoi diagram for the same sites (excluding one). The colored boxes indicate the three closest sites. The order-1 diagram is shown in the background. Each order-1 vertex lies in the interior of an order-3 region as it has equal weighted distance to its three closest sites. As at most two order-1 vertices share an order-3 region, we get $\Omega(n^2)$ order-3 regions. Theorem 6.2 generalizes this observation.

A are the k nearest sites. The number of non-empty order- k Voronoi regions is called the *complexity* of the diagram. (iii) The sites have multiplicative weights that scale the influence of the different sites. Without loss of generality, we assume the weights to be scaled such that the minimum is 1. (iv) We allow the \mathfrak{p} -norm for arbitrary $\mathfrak{p} \in \mathbb{N}^+ \cup \infty$.

Theorem 6.9. *Let S be a set of n sites with minimum weight 1, total weight W , and random positions on the d -dimensional torus equipped with a \mathfrak{p} -norm, for constant d . For every fixed k , the expected number of regions of the weighted order- k Voronoi diagram of S is in $O(W)$. The same holds for random sites in a hypercube.*

To set this result into context, we briefly discuss previous work on the complexity of Voronoi diagrams in different settings. See the book by Aurenhammer et al. [6] for a general overview on Voronoi diagrams. To this end, we use the following theorem that relates the complexity in terms of Voronoi regions (which is what we are concerned with in this article) with the complexity in terms of vertices.²

Theorem 6.2. *Let S be a set of n weighted sites in general position in \mathbb{R}^d equipped with a \mathfrak{p} -norm. If the order- k Voronoi diagram has ℓ vertices, then the order- $(k + d)$ Voronoi diagram has $\Omega(\ell)$ non-empty regions.*

We note that, using insights from previous work, this theorem is not hard to prove. One basically has to generalize the result by Lê [45] bounding the number of d -spheres going through $d + 1$ points in d -dimensional space to weighted sites, and then observe how the Voronoi diagram changes in the construction by Lee [46] for $d = 2$, when going from order- k to order- $(k + 1)$. However, we are not aware of previous work stating this connection between vertices and non-empty regions in higher orders explicitly.

The four above-mentioned generalizations of the basic Voronoi diagram (higher dimension, higher order, multiplicative weights, and different \mathfrak{p} -norms) have all been considered before. However, to the best of our knowledge, not all of them together.

²Although the Voronoi regions are not necessarily polytopes in the weighted setting, we adopt the notion for polytopes and call the corners of Voronoi regions *vertices*. That is, vertices are the 0-dimensional elements (a.k.a. points) of the boundary, where higher-dimensional elements (a.k.a. edges, faces, etc.) intersect. They are represented as small black dots in Figure 2.

Higher-order Voronoi diagrams have been introduced by Shamos and Hoey [57]. Lee [46] showed that the order- k Voronoi diagram in the plane (unweighted with Euclidean metric) has complexity $O(k(n-k))$ (in terms of number of regions), which is linear for constant k . For the 1- and ∞ -norm, Liu et al. [47] improved this bound to $O(\min\{k(n-k), (n-k)^2\})$. Closely related to the 1-norm, Gemsa et al. [37] showed similar complexity bounds for higher-order Voronoi diagrams on transportation networks of axis-parallel line segments. Bohler et al. [15] show an upper bound of $2k(n-k)$ for the much more general setting of abstract Voronoi diagrams. There, the metric is replaced by curves separating pairs of sites such that certain natural (but rather technical) conditions are satisfied. One obtains normal Voronoi diagrams when using perpendicular bisectors for these curves. This in particular shows that the $2k(n-k)$ bound on the number of regions in the order- k Voronoi diagram holds for arbitrary p -norms in 2-dimensional space and for the hyperbolic plane. As the hyperbolic plane is closely related to 1-dimensional space with sites having multiplicative power-law weights [18], we suspect that the bound by Bohler et al. [15] also covers this case.

In general one can say that higher-order Voronoi diagrams of unweighted sites in 2-dimensional space are well-behaved in that they have linear complexity. This still holds true for arbitrary p -norms. However, this picture changes for weighted sites or higher dimensions.

Voronoi diagrams with multiplicative weights were first considered by Boots [17]³ due to applications in economics. Beyond that, multiplicatively weighted Voronoi diagrams have applications in sensor networks [23], logistics [36] and the growth of crystals [25]. However, even in the most basic setting of 2-dimensional Euclidean space and order 1, weighted Voronoi diagrams can have quadratic complexity [5] (in terms of number of vertices). This comes from the fact that Voronoi cells are not necessarily connected; see Figure 2A for the construction of Aurenhammer and Edelsbrunner [5] that proves the lower bound. With Theorem 6.2, and as illustrated in Figure 2, this implies that even the order-3 Voronoi diagram of weighted sites in 2-dimensional Euclidean space has a quadratic number of non-empty regions. As a special case, Theorem 6.9 shows that this complexity is only linear in the total weight for sites positioned randomly in the unit square. Moreover, this also implies that the number of vertices of the corresponding order-1 Voronoi diagram is linear. This nicely complements the result by Har-Peled and Raichel [42], who show that the expected complexity of order-1 Voronoi diagrams of sites in 2-dimensional Euclidean space with random weights is $O(n \text{ polylog } n)$. Only recently, Fan and Raichel [32] showed that sites with weights chosen randomly form a constant-sized set of possible weights yield Voronoi diagrams with linear complexity. Moreover, more closely related, they show that the Voronoi diagram of sites with arbitrary weights and with random positions chosen in the unit square has linear complexity in expectation. We are not aware of any results concerning the complexity of Voronoi diagrams when combining multiplicative weights with higher dimension, higher order or other norms.

For higher dimensions, even normal (first order, unweighted) Voronoi diagrams in 3-dimensional Euclidean space can have $\Theta(n^2)$ [43,56] vertices. Theorem 6.2 thus implies that the order-4 Voronoi diagram has a quadratic number of non-empty regions. Moreover, the complexity of higher-order Voronoi diagrams in higher dimensions has been considered before by Mulmuley [49], who obtains polynomial bounds with the degree of the polynomial depending on the dimension. Our Theorem 6.9 in particular shows that this complexity is much lower, namely linear, for the hypercube with randomly positioned sites. Moreover, via Theorem 6.2 this gives a linear bound on number of vertices in the normal order-1 Voronoi diagram in higher dimensions. We note that this special case of our result coincides with a previous result by Bienkowski et al. [10]. Similarly, Dwyer [28] showed that sites drawn uniformly from a higher dimensional unit sphere (instead of a hypercube) yield Voronoi

³In this article, Voronoi regions are called *Thiessen polygons*.

diagrams of linear complexity in expectation. Moreover, due to Golin and Na [41] and Driemel et al. [26], the same is true for random sites on 3-dimensional polytopes and random sites on polyhedral terrains, respectively. Thus, though higher dimensional Voronoi diagrams can be rather complex in the worst case, these results indicate that one can expect most instances to be rather well behaved. An alternative explanation of why the complexity of practical instances is lower than the worst-case indicates is given by Erickson [29,30], who studies the complexity of 3-dimensional Voronoi diagrams depending on the so-called spread of the sites.

The above results for higher dimensional Voronoi diagrams consider the Euclidean norm. For general p -norms, L  [45] showed that the complexity of the Voronoi diagram is bounded by $O(n^c)$, where c is a constant independent of p but dependent on the dimension d . With the same argument as above, Theorem 6.9 together with Theorem 6.2 implies a linear bound for this complexity that holds in expectation. Moreover, Boissonnat et al. [16] show more precise bounds of $\Theta(n^{\lceil d/2 \rceil})$ and $\Theta(n^2)$ for the ∞ - and the 1-norm, respectively. Again, our result implies linear bounds for random sites in this setting.

3 | FORMAL DEFINITIONS

Here, we provide formal definitions for all concepts we use throughout the article, including the power-law and geometric random SAT models, resolution, and Voronoi diagrams.

k -SAT

We let x_1, x_2, \dots, x_n denote Boolean variables that can be either true or false. A clause is a disjunction of literals $\ell_1 \vee \dots \vee \ell_k$, where each literal assumes a (possibly negated) variable. For a literal ℓ_i let $|\ell_i|$ denote the variable of the literal. A formula Φ in conjunctive normal form (CNF) is a conjunction of clauses $c_1 \wedge \dots \wedge c_m$ and a formula in k -CNF is a conjunction of clauses, where each clause contains exactly three distinct literals. We conveniently interpret a Boolean formula in CNF as a set of clauses and a clause c both as a Boolean formula and as a set of literals. We say that Φ is satisfiable if there exists an assignment of variables x_1, \dots, x_n such that the formula evaluates to true.

Power-law random k -SAT

The power-law model can be defined via the more general *non-uniform model*. To draw a k -SAT formula from the non-uniform model, let n and m be the number of variables and clauses, respectively, and let w_1, \dots, w_n be variable weights. We sample m clauses independently at random. Each clause is sampled by drawing k variables without repetition with probabilities proportional to their weights. Then each of the k variables is negated independently at random with probability $1/2$.

The *power-law model* for a power-law exponent $\beta > 2$ is an instantiation of the non-uniform model with discrete power-law weights

$$w_i = i^{-\frac{1}{\beta-1}}.$$

Resolution

The resolution proof system uses two rules, the resolution rule and the weakening rule. Given two clauses $a \vee x$ and $b \vee \bar{x}$, where a and b are clauses and x is a Boolean variable, the resolution rule states

$$\frac{a \vee x \quad b \vee \bar{x}}{a \vee b},$$

that is, the clause $a \vee b$ is a logical consequence of the two given clauses. The weakening rule states that for any two clauses a and b it holds that

$$\frac{a}{a \vee b},$$

that is, if a holds, then $a \vee b$ holds as well. For a formula $\Phi = \{c_1, c_2, \dots, c_m\}$ in CNF a resolution *derivation* of a clause c from Φ is a sequence of clauses (d_1, d_2, \dots, c) such that each clause d_i is either one of the initial clauses c_1, \dots, c_m or derived from previous clauses with either the resolution rule or the weakening rule. A resolution *refutation* is a resolution derivation of the empty clause. The *size* of a derivation is the number of clauses it contains. The size of a formula in CNF is the size of a smallest refutation for it. The *width* of a derivation is the size of the largest clause in it. The width of a formula in CNF is the smallest width of any refutation for it.

Graph representation and expansion

Let Φ be a SAT-formula with variable set V and clause set C . The *clause-variable incidence graph* $G(\Phi)$ of Φ has vertex set $C \cup V$, with an edge between a clause and a variable if and only if the clause contains the variable. Clearly, $G(\Phi)$ is bipartite. It is an (r, c) -*bipartite expander* if for all $C' \subset C$ with $|C'| \leq r$ it holds that $|N(C')| \geq (1 + c) \cdot |C'|$, where $N(C')$ is the neighborhood of C' .

Geometric ground space

We regularly deal with points with random positions in some geometric space. With *random point*, we refer to the uniform distribution in the sense that the probability for a point to lie in a region A is proportional to its volume $\text{vol}(A)$. For this to work, the volume of the ground space has to be bounded. Canonical options are, for example, a unit-hypercube or a unit-ball. These, however, lead to the necessity of special treatment for points close to the boundary, which makes the analysis more tedious without giving additional insights. To circumvent this, we use a torus as ground space, which is completely symmetric.

The d -dimensional torus \mathbb{T}^d is defined as the d -dimensional hypercube $[0, 1]^d$ in which opposite borders are identified, that is, a coordinate of 0 is identical to a coordinate of 1.⁴ It is equipped with the \mathfrak{p} -norm as metric, for arbitrary but fixed $\mathfrak{p} \in \mathbb{N}^+ \cup \{\infty\}$. To define it formally for the torus, let $\mathbf{p} = (p_1, \dots, p_d)$ and $\mathbf{q} = (q_1, \dots, q_d)$ be two points in \mathbb{T}^d . The circular difference between the i th coordinates is $|p_i - q_i|_o = \min\{|p_i - q_i|, 1 - |p_i - q_i|\}$. With this, the *distance* between \mathbf{p} and \mathbf{q} is

$$\|\mathbf{p} - \mathbf{q}\| = \begin{cases} \sqrt[\mathfrak{p}]{\sum_{i \in [d]} |p_i - q_i|_o^{\mathfrak{p}}} & \text{for } \mathfrak{p} \neq \infty, \\ \max_{i \in [d]} \{|p_i - q_i|_o\} & \text{for } \mathfrak{p} = \infty. \end{cases}$$

⁴For convenience reasons, we sometimes work with $[-0.5, 0.5]^d$ instead of $[0, 1]^d$.

Random points

We obtain the uniform distribution for a point $\mathbf{p} = (p_1, \dots, p_d)$ by drawing each coordinate p_i uniformly at random from $[0, 1]$. For two random points \mathbf{p} and \mathbf{q} , their distance $\|\mathbf{p} - \mathbf{q}\|$ is a random variable. Let $F_{\text{dist}}(x)$ be its *cumulative distribution function (CDF)*, that is, $F_{\text{dist}}(x) = \Pr[\|\mathbf{p} - \mathbf{q}\| \leq x]$. To determine $F_{\text{dist}}(x)$, fix the position of \mathbf{p} . Then, for $x \leq 0.5$, the set of points of distance at most x to \mathbf{p} is simply the ball $B_{\mathbf{p}}(x)$ of radius x around \mathbf{p} , yielding

$$\begin{aligned} F_{\text{dist}}(x) &= \text{vol}(B_{\mathbf{p}}(x)) \\ &= \Pi_{d,\mathbf{p}} \cdot x^d \quad \text{for } 0 \leq x \leq 0.5, \\ \text{with } \Pi_{d,\mathbf{p}} &= \frac{(2\Gamma(1/\mathbf{p} + 1))^d}{\Gamma(d/\mathbf{p} + 1)}, \end{aligned} \quad (1)$$

where Γ is the gamma function. Note that $\Pi_{d,\mathbf{p}}$ only depends on d and \mathbf{p} but is constant in x . Moreover $\Pi_{2,2} = \pi$ (thus the name Π), and $\Pi_{d,\infty} = \lim_{\mathbf{p} \rightarrow \infty} \Pi_{d,\mathbf{p}} = 2^d$. For distances $x > 0.5$, the formula for $F_{\text{dist}}(x)$ is more complicated (we basically have to subtract the parts reaching out of the hypercube). However, for our purposes, it suffices to know $F_{\text{dist}}(x)$ for $x \leq 0.5$ and use the obvious bound $F_{\text{dist}}(x) \leq 1$ for $x > 0.5$.

Weighted points and distances

We regularly consider a fixed set of n points equipped with weights, which we call *sites*. For a site s_i with weight w_i , the *weighted distance* of a point \mathbf{p} to s_i is $\|s_i - \mathbf{p}\|/w_i^{1/d}$. For a fixed value x , the set of points with weighted distance at most x are the points with $\|s_i - \mathbf{p}\| \leq xw_i^{1/d}$. Note that the volume of this set is proportional to w_i . Intuitively, the region of influence of a site is thus proportional to its weight. To simplify notation in some places, we define *normalized weights* $\omega_i = w_i^{1/d}$.⁵

Geometric random k-SAT

In the *geometric model*, we sample positions for the variables and clauses uniformly at random in the d -dimensional torus \mathbb{T}^d . For $v \in V$ and $c \in C$, we use \mathbf{v} and \mathbf{c} to denote their positions, respectively. Let w_1, \dots, w_n be *variable weights* that are normalized such that the smallest weight is 1. Moreover, let $W = \sum_{v=1}^n w_v$. For a clause c and a variable v , define the *connection weight*

$$X(c, v) = \left(\frac{w_v}{\|\mathbf{c} - \mathbf{v}\|^d} \right)^{1/T}.$$

This is the reciprocal of the weighted distance between \mathbf{v} and \mathbf{c} raised to the power d/T . The k variables for the clause c are drawn without repetition with probabilities proportional to $X(c, v)$. Among all possible combinations, we choose which of the k variables to negate uniformly at random, without repetition if possible, that is, we only get the same clause twice if we have more than 2^k clauses with the same variable set. For $T \rightarrow 0$ the model converges to the threshold case where c contains the k variables with smallest weighted distance.

⁵We note, in the context of weighted Voronoi diagrams, it is common to only use the normalized weights (just calling them “weights”). In the context of random networks, however, the non-normalized weights are more common. As both notions have their advantages in different situations, we use both.

The connection weight $X(c, v)$ is a random variable. We denote the CDF of $X(c, v)$ with $F_X(x)$. With the CDF for the distance between two random points in Equation (1), we obtain the following; see Lemma A.2 for a proof:

$$F_X(x) = 1 - \Pi_{d, \mathbf{p}} w_v x^{-T} \quad \text{for } x \geq (2^d w_v)^{1/T}. \quad (2)$$

Voronoi diagrams

Let $S = \{s_1, \dots, s_n\}$ be a set of sites with weights w_1, \dots, w_n . A point \mathbf{p} belongs to the (*open*) *Voronoi region* of a site s_i if its weighted distance to s_i is smaller than its weighted distance to any other site. The collection of all Voronoi regions is the *Voronoi diagram* of S . *Order- k Voronoi regions* are defined analogously for subsets $A \subseteq S$ with $|A| = k$, that is, the region of A contains a point \mathbf{p} if and only if the weighted distances of \mathbf{p} to all sites in A is smaller than the weighted distance to any site not in A . More formally, \mathbf{p} belongs to the order- k Voronoi region of A if there exists a radius r such that $\|s_i - \mathbf{p}\| \leq \omega_i r$ for $s_i \in A$ and $\|s_i - \mathbf{p}\| > \omega_i r$ for $s_i \notin A$. Note that the order- k Voronoi region of A is potentially empty. The *order- k Voronoi diagram* is the collection of all non-empty order- k Voronoi regions. Its *complexity* is the number of such non-empty regions.

4 | CORE ARGUMENTS

Before delving into the technical details of our proofs in the subsequent sections, we briefly discuss the core arguments.

4.1 | Power-law SAT

We use a framework that Ben-Sasson and Wigderson [9] introduced for the uniform SAT model. We prove lower bounds for the resolution width, which imply lower bounds for the resolution size and the tree-like resolution size, which then imply lower bounds for the running times of CDCL and DPLL solvers, respectively.

To bound the resolution width, we essentially have to show that different clauses do not overlap too heavily. Specifically, a formula has resolution width $\Omega(w)$ if (1) every set S of at most w clauses contains at least $|S|$ different variables and (2) every set S of $\frac{1}{3}w \leq |S| \leq \frac{2}{3}w$ clauses contains at least a constant fraction of unique variables.

We achieve the bounds in Theorem 5.8 (i–iii) by showing the above two properties directly. For the bound in Theorem 5.8 (iv), we first observe that both properties are fulfilled if the clause-variable incidence graph of a k -CNF formula Φ has high enough bipartite expansion. Recall the definition of bipartite expansion from Section 3 and note how the requirement that the neighborhood of clause vertices is large resembles the requirement that clauses do not overlap too heavily. We show that $G(\Phi)$ is a bipartite expander asymptotically almost surely if Φ is drawn from the power-law model, which yields the lower bound of Theorem 5.8 (iv).

Compared to the uniform case, the weights make the properties required for the lower bounds less likely. Variables with high weight appear in many clauses, making the clauses less diverse. Thus, it is less likely that every clause set covers a large variety of variables.

4.2 | Geometric SAT

To explain the core idea of our proof, consider the following simplified geometric model. Map n variables and m clauses to distinct points in the 2-dimensional Euclidean plane (randomly or deterministically). Build the SAT instance by including in each clause c the k variables with the smallest geometric distance to c . Now consider the order- k Voronoi diagram defined by the positions of the n variables. As a clause c contains the k closest variables, the k variables contained in c are exactly the k variables defining the Voronoi region of c 's position. Independent of the positions of the n variables, there are only at most $2k(n - k)$ regions in the order- k Voronoi diagram [15]. Thus, if we have at least $2^k 2k(n - k)$ clauses, then, by the pigeonhole principle, at least one Voronoi region contains 2^k clauses. As k is considered to be a constant, this number of clauses is linear in n , that is, we still have constant density. Moreover, as repeating the same clause (with the same variable negations) is avoided whenever possible, there is a set of k variables that has a clause for every combination of literals. Thus, we have an unsatisfiable subformula of constant size 2^k , which implies low proof complexity.

This result can be varied and strengthened in multiple ways, for example, by allowing weighted variables, a higher dimensional ground space, or by softening the requirement that every clause contains the k closest variables (model with higher temperature). In the following, we briefly discuss how these generalizations can be achieved.

Abstract geometric spaces

The result by Bohler et al. [15] on the complexity of order- k Voronoi diagrams is very general in the sense that it holds for abstract Voronoi diagrams. Roughly speaking, abstract Voronoi diagrams are based on separating curves between pairs of points that take the role of perpendicular bisectors. In this way, one can abstract from the specific geometric ground space. Whether a point p is closer to site s_1 or to site s_2 is no longer determined by comparing distances $\|s_1 - p\|$ and $\|s_2 - p\|$ but by the curve separating s_1 from s_2 . For this to work, the separating curves have to satisfy a handful of basic axioms. These are for example satisfied by perpendicular bisectors in the Euclidean or the hyperbolic plane. Thus, the above argumentation for low proof complexity directly carries over to the hyperbolic plane, or more generally, to any abstract geometric space satisfying the axioms.

Lower density via random clause positions

Assume the variable positions are fixed. Now choose random positions for the clauses and observe in which regions of the order- k Voronoi diagram they end up. We want to know whether there is a region that contains at least 2^k clauses. This comes down to a balls into bins experiment. Each Voronoi region is a bin and each clause is a ball. Thus, there are $O(n)$ bins and m balls. Moreover, we are interested in the maximum load, that is, the maximum number of balls that land in a single bin. Due to a result by Raab and Steger [55], the maximum load is a. a. s. in $\Omega(\frac{\log n}{\log \log n})$ if we throw $\Omega(\frac{n}{\text{polylog } n})$ balls. Thus, even for a slightly sublinear number of balls, the maximum load is superconstant. We note that this result holds for uniform bins. In our case, we have non-uniform bins, as the probability for a clause to end up in a particular Voronoi region is proportional to the area of the region. However, it is not hard to see that the result by Raab and Steger [55] remains true for non-uniform bins; see Section A.5. Thus, even if the number of clauses m is slightly sublinear in the number of variables n , we get a small unsatisfiable subformula asymptotically almost surely if the Voronoi diagram has low complexity.

Positive or negative literals with repetition

Above we assumed that we get the exact same clause with coinciding negations twice only if we already have more than 2^k clauses with the same set of k variables. Although this is arguably a reasonable assumption for the model, we can make a similar argument without it. Assume instead that for each variable, we choose the positive and negative literal uniformly at random, independently of all other choices. Moreover, assume for an increasing function f , that there are $f(n)$ clauses that have the same set of k variables. With the above balls into bins argument, we, for example, have $f(n) \in \Omega(\frac{\log n}{\log \log n})$. Then the probability that there is a combination of positive and negative literals that we did not see at least once is at most $2^k(1 - 2^{-k})^{f(n)}$. This probability goes to 0 for $n \rightarrow \infty$, that is, a. a. s., there is an unsatisfiable subformula of constant size 2^k .

Higher dimension and weighted variables

At the core of our argument lies the fact that order- k Voronoi diagrams have linear complexity in the plane. As already mentioned in Section 2.3, this is no longer true for order- k Voronoi diagrams in higher dimensions or if the variables have multiplicative weights. A formal argument for why this property breaks is in Section 6.1. However, for sites distributed uniformly at random, we show in Section 6.2 that the complexity can be expected to be linear in the total weight, even in the more general setting. Thus, using that the variables have random positions (a requirement we did not need before), we can apply the above argument to obtain low proof complexity.

Non-zero temperature

Non-zero temperatures make it so that clauses do not necessarily contain the k closest variables. Instead, variables are included with probabilities depending on the distance. Thus, we cannot simply look at the order- k Voronoi diagram to determine which variables are contained in a given clause. We resolve this issue in Section 7. For this, we call a clause *nice*, if it behaves as it would in the $T = 0$ case, that is, if it includes the k closest variables. In Section 7.1, we show that, in expectation, a constant fraction of clauses is actually nice. Moreover, in Section 7.2, we show that the number of nice clauses is concentrated around its expectation. With this, we can apply the same arguments as before to only the nice clauses, of which we have linearly many, to obtain a low proof complexity.

4.3 | Voronoi diagrams

The worst-case lower bounds for the complexity of order- k Voronoi diagrams follow from existing lower bounds on the number of vertices together with Theorem 6.2, which connects the complexity in terms of regions with the complexity in terms of vertices. This connection is obtained by observing how the order- k Voronoi diagram changes when increasing k .

For the average-case linear upper bound on the number of regions, the argument works roughly as follows, assuming the unweighted case for the sake of simplicity. For each size- k subset A of the sites, we devise an upper bound on the probability that A has a non-empty order- k Voronoi region. This region is non-empty if and only if there are points that have A as the k closest sites, that is, if there is a ball that contains the sites of A and no other sites. With this observation, we can use a win-win-style argument. Either the radius of this ball is small, which makes it unlikely that all sites of A lie in the ball, or the ball is large, which makes it unlikely that it contains no other sites.

5 | RESOLUTION SIZE OF POWER-LAW RANDOM K-SAT

5.1 | The direct approach

As stated in Section 4.1, a formula has resolution width $\Omega(w)$ if (1) every set S of at most w clauses contains at least $|S|$ different variables and (2) every set S of $\frac{1}{3}w \leq |S| \leq \frac{2}{3}w$ clauses contains at least a constant fraction of unique variables. In this section, we are going to show that both conditions are satisfied for power-law exponents $\beta > \frac{2k-1}{k-1}$ and clause-variable ratios $\Delta \in \Omega(1)$. The first condition can also be interpreted in terms of bipartite expansion. It states that the clause-variable incidence graph $G(\Phi)$ is a $(w, 0)$ -bipartite expander. The following lemma states bounds on w for which $G(\Phi)$ is a $(w, 0)$ -bipartite expander asymptotically almost surely. These bounds depend on the power-law exponent β as well as on the clause-variable ratio Δ . Note that our choices of k and β in the lemma ensure $\varepsilon_1, \varepsilon_2 > 0$.

Lemma 5.1. *Let Φ be a random power-law k -SAT formula with n variables, $\Delta \cdot n = m \in \Omega(n)$ clauses, $k \geq 3$, and power-law exponent $\beta > \frac{2k-1}{k-1}$. Let $\varepsilon_1 = k \cdot \frac{\beta-2}{\beta-1} - 1 > 0$ and $\varepsilon_2 = (k-2) \frac{\beta-2}{\beta-1} > 0$. Then $G(\Phi)$ has $(w, 0)$ -bipartite expansion a. a. s. if*

- (i) $\beta \in \left(\frac{2k-1}{k-1}, 3\right)$, $\Delta \in o(n^{\varepsilon_1}/\log^{\varepsilon_2}(n))$, and $w \in O(n^{\varepsilon_1/\varepsilon_2} \cdot \Delta^{-1/\varepsilon_2})$.
- (ii) $\beta = 3$, $\Delta \in o(n^{(k-2)/2}/\log^{1+(k-2)/2}(n))$, and $w \in O(n \cdot (\Delta \cdot \ln n)^{-2/(k-2)})$.
- (iii) $\beta > 3$, $\Delta \in o(n^{\varepsilon_2}/\log^{\varepsilon_2} n)$, and $w \in O(n \cdot \Delta^{-1/\varepsilon_2})$.

Proof. We are interested in showing $|N(C')| \geq |C'|$ for all $C' \subseteq C$ with $|C'| \leq w$. We consider a smallest C' such that $|N(C')| \leq |C'| - 1$ and denote it by \hat{C} . Let \mathcal{E}_i be the event that $|\hat{C}| = i$. Thus, \mathcal{E}_i implies that for all $C' \subseteq C$ with $|C'| < i$ it holds that $|N(C')| \geq |C'|$. This implies that every variable in $N(\hat{C})$ has to appear at least twice. Otherwise, one could delete a clause with a unique variable from \hat{C} to get a set \hat{C}' with $|\hat{C}'| = i - 1$ and $|N(\hat{C}')| \leq i - 2$. This would violate the minimality of \hat{C} . Also, \hat{C} must contain exactly $i - 1$ different variables. Otherwise, we could remove any clause from \hat{C} and violate minimality.

Now we bound

$$\sum_{i=1}^w \Pr(\mathcal{E}_i) \leq \sum_{i=1}^w \binom{m}{i} P_i,$$

where P_i is the probability to draw i clauses which contain at most $i - 1$ different variables and all of them at least twice. We can now imagine the $k \cdot i$ variables of the i clauses to be drawn independently with replacement. This would only increase the probability that the i clauses contain at most $i - 1$ different variables and all of them at least twice. Thus, the probability we consider is an upper bound. Now, we consider the $i - 1$ different variables drawn. Then, we choose the $i - 1$ pairs of positions where each variable appears for the first and second time. As a rough upper bound we have at most

$$\binom{\binom{k \cdot i}{2}}{i-1} \leq \left(\frac{(k \cdot i)^2 \cdot e}{2 \cdot (i-1)} \right)^{i-1}$$

many possibilities by simply choosing $i - 1$ from all $\binom{k \cdot i}{2}$ possible pairs. Now we bound the probability that at these pairs of positions the same variables do appear. This is at most $\sum_{j=1}^n p_j^2$ per pair of positions, where $p_j = \frac{j^{-1/(\beta-1)}}{\sum_{\ell=1}^n \ell^{-1/(\beta-1)}}$ is the probability that the j th variable

is drawn. At the remaining $k \cdot i - 2 \cdot (i - 1)$ positions we can only choose from at most those $i - 1$ variables. Thus, the probabilities at all other positions are the sum of the $i - 1$ variable probabilities, which is at most the sum of the $i - 1$ highest variable probabilities. Let $F(i)$ be the sum of the i highest variable probabilities. Then it holds that

$$\begin{aligned} P_i &\leq \left(\frac{(k \cdot i)^2 \cdot e}{2 \cdot (i - 1)} \right)^{i-1} \cdot \left(\sum_{j=1}^n p_j^2 \right)^{i-1} \cdot F(i-1)^{k \cdot i - 2 \cdot (i-1)} \\ &\leq \kappa^{i-1} \cdot \left(\frac{i^2}{i-1} \right)^{i-1} \cdot \left(\sum_{j=1}^n p_j^2 \right)^{i-1} \cdot \left(\frac{i-1}{n} \right)^{(k \cdot i - 2 \cdot (i-1)) \frac{\beta-2}{\beta-1}} \end{aligned}$$

for a constant $\kappa = \kappa(k, \beta) > 0$ that might depend on other parameters, which are fixed to constants as well. We will use κ to collect all constant factors. According to Lemma A.1

$$\sum_{j=1}^n p_j^2 = \begin{cases} \Theta \left(n^{-2 \frac{\beta-2}{\beta-1}} \right), & \beta < 3; \\ \Theta (\ln n / n), & \beta = 3; \\ \Theta (n^{-1}), & \beta > 3. \end{cases}$$

Thus, our result depends on the power law exponent β . For $\beta < 3$ we get

$$\begin{aligned} P_i &\leq \kappa^i \cdot \left(\frac{i^2}{i-1} \right)^{i-1} \cdot n^{-(i-1) \cdot 2 \frac{\beta-2}{\beta-1}} \cdot \left(\frac{i-1}{n} \right)^{(k \cdot i - 2 \cdot (i-1)) \frac{\beta-2}{\beta-1}} \\ &\leq \kappa^i \cdot n^{-k \cdot i \cdot \frac{\beta-2}{\beta-1}} \cdot i^{i-1} \cdot (i-1)^{(k \cdot i - 2 \cdot (i-1)) \frac{\beta-2}{\beta-1}} \\ &\leq \kappa^i \cdot n^{-k \cdot i \cdot \frac{\beta-2}{\beta-1}} \cdot i^{(k \cdot i - 2 \cdot (i-1)) \frac{\beta-2}{\beta-1} + (i-1)} \\ &= \kappa^i \cdot n^{-k \cdot i \cdot \frac{\beta-2}{\beta-1}} \cdot i^{(k \cdot i - 2 \cdot i) \frac{\beta-2}{\beta-1} + i + 2 \frac{\beta-2}{\beta-1} - 1} \\ &\leq \kappa^i \cdot n^{-k \cdot i \cdot \frac{\beta-2}{\beta-1}} \cdot i^{(k \cdot i - 2 \cdot i) \frac{\beta-2}{\beta-1} + i} = \kappa^i \cdot n^{-i \cdot (\varepsilon_1 + 1)} \cdot i^{i \cdot (\varepsilon_2 + 1)}, \end{aligned}$$

where we used $\left(\frac{i^2}{i-1} \right)^{i-1} \leq e \cdot i^{i-1}$ in the second line and upper-bounded $i - 1$ in the base by i , which we can do since $(k \cdot i - 2 \cdot (i - 1)) \frac{\beta-2}{\beta-1} > 0$ due to $k \geq 3$ and $\beta > 2$. In the third line, we used $2 \frac{\beta-2}{\beta-1} - 1 < 0$, which holds since $\beta < 3$.

We can now see that

$$\begin{aligned} \sum_{i=1}^w \binom{m}{i} P_i &\leq \sum_{i=1}^w \binom{m}{i} \cdot \kappa^i \cdot n^{-i \cdot (\varepsilon_1 + 1)} \cdot i^{i \cdot (\varepsilon_2 + 1)} \\ &\leq \sum_{i=1}^w \kappa^i \cdot \Delta^i \cdot n^{-i \cdot \varepsilon_1} \cdot i^{i \cdot \varepsilon_2}, \end{aligned} \quad (3)$$

which holds since we assume $m = \Delta \cdot n$ and $\binom{m}{i} \leq \left(\frac{e \cdot m}{i} \right)^i$. In order to have a sum which is $o(1)$ we want to ensure that

$$\kappa \cdot \Delta \cdot n^{-\varepsilon_1} \cdot i^{\varepsilon_2},$$

is at most a constant smaller than 1. It is easy to check that this holds for

$$i \in O\left(n^{\varepsilon_1/\varepsilon_2} \cdot \Delta^{-1/\varepsilon_2}\right).$$

Thus, we can set w to this value. If we split the sum in Equation (3) at $i_0 = \lfloor \varepsilon_1 \log n \rfloor$, the part with $i \leq i_0$ is upper-bounded by $O\left(\Delta \cdot n^{-\varepsilon_1} \cdot i_0^{\varepsilon_2}\right) \in O\left(\Delta (\log^{\varepsilon_2}(n)/n^{\varepsilon_1})\right)$ via a geometric series. The part with $i > i_0$ is upper-bounded by the first term. If we chose $w \in \Theta(n^{\varepsilon_1/\varepsilon_2} \cdot \Delta^{-1/\varepsilon_2})$ so that $\Delta \cdot n^{-\varepsilon_1} \cdot i^{\varepsilon_2} \leq c$ for a constant $c \in (0, 1)$, the second term yields at most $c^{\Theta(\log n)} = o(1)$. Thus, we get $\left(\Theta\left(n^{\varepsilon_1/\varepsilon_2} \cdot \Delta^{-1/\varepsilon_2}\right), 0\right)$ -expansion with probability at least $1 - \Theta\left(\Delta (\log^{\varepsilon_2}(n)/n^{\varepsilon_1})\right)$ or a. a. s. if $\Delta \in o(n^{\varepsilon_1}/(\log(n))^{\varepsilon_2})$.

For $\beta > 3$ we get

$$\begin{aligned} P_i &\leq \kappa^i \cdot \left(\frac{i^2}{i-1}\right)^{i-1} \cdot n^{-(i-1)} \cdot \left(\frac{i-1}{n}\right)^{(k \cdot i - 2 \cdot (i-1)) \frac{\beta-2}{\beta-1}} \\ &\leq \kappa^i \cdot \left(\frac{i}{n}\right)^{i-1+(k \cdot i - 2 \cdot (i-1)) \frac{\beta-2}{\beta-1}}. \end{aligned}$$

Thus,

$$\begin{aligned} \sum_{i=1}^r \binom{m}{i} P_i &\leq \sum_{i=1}^w \left(\frac{n}{i}\right)^i \cdot \kappa^i \cdot \Delta^i \cdot \left(\frac{i}{n}\right)^{i-1+(k \cdot i - 2 \cdot (i-1)) \frac{\beta-2}{\beta-1}} \\ &= \sum_{i=1}^w \kappa^i \cdot \Delta^i \cdot \left(\frac{i}{n}\right)^{(k \cdot i - 2 \cdot i) \frac{\beta-2}{\beta-1} + 2 \frac{\beta-2}{\beta-1} - 1} \\ &\leq \sum_{i=1}^w \kappa^i \cdot \Delta^i \cdot \left(\frac{i}{n}\right)^{i \cdot (k-2) \frac{\beta-2}{\beta-1}} \\ &\leq \sum_{i=1}^w \kappa^i \cdot \Delta^i \cdot \left(\frac{i}{n}\right)^{i \cdot \varepsilon_2}, \end{aligned}$$

which holds since $\frac{i}{n} \leq 1$ and $2 \cdot \frac{\beta-2}{\beta-1} - 1 \geq 0$ for $\beta \geq 3$. It is now easy to show that $\kappa \cdot \Delta \cdot (i/n)^{\varepsilon_2}$ is at most a small constant for $w \in \Theta(n \cdot \Delta^{-1/\varepsilon_2})$ sufficiently small. By splitting the sum as before, we can show $((n \cdot \Delta^{-1/\varepsilon_2}), 0)$ -expansion with probability at least $1 - \Theta(\Delta \cdot \log^{\varepsilon_2} n / n^{\varepsilon_2})$ or a. a. s. for $\Delta \in o(n^{\varepsilon_2} / \log^{\varepsilon_2} n)$.

For $\beta = 3$ we get the same result as for $\beta > 3$, except for an additional factor of $(\ln n)^{i-1}$. Thus,

$$\begin{aligned} \sum_{i=1}^w \binom{m}{i} P_i &\leq \sum_{i=1}^w \kappa^i \cdot \Delta^i \cdot \left(\frac{i}{n}\right)^{i \cdot (k-2) \frac{\beta-2}{\beta-1}} \ln^{i-1} n \\ &\leq \sum_{i=1}^w \kappa^i \cdot \Delta^i \cdot \left(i \cdot \frac{\ln^{2/(k-2)} n}{n}\right)^{i \cdot \frac{k-2}{2}}. \end{aligned}$$

By assuming

$$w \in \Theta\left(n \cdot (\Delta \cdot \log n)^{-2/(k-2)}\right)$$

small enough, we can ensure that this sum is at most $O(\Delta \cdot \log n \cdot (\log(n)/n)^{(k-2)/2})$ by splitting the expression at $\lfloor i_0 = \ln n \rfloor$ again. Hence, we get $(\Theta(n \cdot (\Delta \cdot \log n)^{-2/(k-2)}), 0)$ -expansion with probability at least $1 - O(\Delta \cdot \ln n \cdot (\log(n)/n)^{(k-2)/2})$ or a. a. s. for $\Delta \in o(n^{(k-2)/2} / \log^{(k-2)/2+1}(n))$. ■

Now, we want to show the second requirement of Theorem 5.3, that every set S of $\frac{1}{3}w \leq |S| \leq \frac{2}{3}w$ clauses contains at least a constant fraction of unique variables. Again, our choices of k and β in the lemma ensure that we can always choose an $\varepsilon > 0$ with $\varepsilon_1, \varepsilon_2 > 0$.

Lemma 5.2. *Let Φ be a random power-law k -SAT formula with n variables, $\Delta \cdot n = m \in \Omega(n)$ clauses, $k \geq 3$, and power-law exponent $\beta > \frac{2k-1}{k-1}$. Let $\varepsilon, \varepsilon_1, \varepsilon_2$ be constant such that $\varepsilon > 0$, $\varepsilon_1 = \frac{k-\varepsilon}{2} - 1 > 0$, and $\varepsilon_2 = (k-\varepsilon) \cdot \frac{\beta-2}{\beta-1} - 1 > 0$. There is a W such that for all $w \in \omega(1)$ with $w \leq W$ a. a. s. all sets C' of clauses from Φ with $\frac{1}{3}w \leq |C'| \leq \frac{2}{3}w$ contain at least $\varepsilon \cdot |C'|$ unique variables. It holds that:*

- (i) If $\beta \in \left(\frac{2k-1}{k-1}, 3\right)$ and $\Delta \in o(n^{\varepsilon_2})$, then $W \in \Theta(n^{\varepsilon_2/\varepsilon_1} \cdot \Delta^{-1/\varepsilon_1})$.
- (ii) If $\beta = 3$ and $\Delta \in o(n^{\varepsilon_1/\ln^{\varepsilon_1+1}n})$, then $W \in \Theta\left(n \cdot \Delta^{-1/\varepsilon_1} / \ln^{1+\frac{1}{\varepsilon_1}}n\right)$.
- (iii) If $\beta > 3$ and $\Delta \in o(n^{\varepsilon_1})$, then $W \in \Theta(n \cdot \Delta^{-1/\varepsilon_1})$.

Proof. Let $\varepsilon \in (0, \min\{k - \frac{\beta-1}{\beta-2}, k-2\})$ be a constant. The upper bounds on ε ensure $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$. We want to bound the probability that there is a set of clauses C' with $\frac{1}{3}w \leq |C'| \leq \frac{2}{3}w$ and at most $\varepsilon \cdot |C'|$ many unique variables. Let P_i be the probability that there is a set C' of size i with that property. We assume the $k \cdot i$ Boolean variables to be drawn independently at random, that is, we allow duplicate variables inside clauses. This only decreases the probability of having unique variables. Additionally, we split the probability into parts depending on the number j of different variables that appear in C' in addition to the $\varepsilon \cdot i$ unique ones. It holds that

$$\begin{aligned}
 P_i &\leq \underbrace{\binom{m}{i}}_{\text{choices of clauses}} \cdot \sum_{j=1}^{\frac{k-\varepsilon}{2} \cdot i} \underbrace{\binom{k \cdot i}{\varepsilon \cdot i}}_{\text{possible positions for the } \varepsilon \cdot i \text{ unique variables}} \underbrace{\binom{\binom{(k-\varepsilon) \cdot i}{2}}{j}}_{\text{possible positions for the first two appearances of the } j \text{ other variables}} \cdot \underbrace{1^{\varepsilon \cdot i}}_{\text{probability to draw a new variable}} \underbrace{\left(\sum_{x=1}^n p_x^2\right)^j}_{\text{probability that variables are same at positions for first two appearances}} \cdot \underbrace{F(j)^{k-i-\varepsilon \cdot i-2j}}_{\text{upper bound on probability to draw } j \text{ chosen variables again}} \\
 &\leq \kappa^i \cdot \Delta^i \left(\frac{n}{i}\right)^i \cdot \sum_{j=1}^{\frac{k-\varepsilon}{2} \cdot i} \left(\frac{k \cdot i}{\varepsilon \cdot i}\right)^{\varepsilon \cdot i} \left(\frac{(k-\varepsilon)^2 \cdot i^2}{j}\right)^j \cdot \left(\sum_{x=1}^n p_x^2\right)^j \cdot F(j)^{k-i-\varepsilon \cdot i-2j} \\
 &\leq \kappa^i \cdot \Delta^i \left(\frac{n}{i}\right)^i \cdot \sum_{j=1}^{\frac{k-\varepsilon}{2} \cdot i} \left(\frac{i^2}{j}\right)^j \cdot \left(\sum_{x=1}^n p_x^2\right)^j \cdot \left(\frac{j}{n}\right)^{(k-i-\varepsilon \cdot i-2j) \frac{\beta-2}{\beta-1}},
 \end{aligned}$$

where $\kappa = \kappa(k, \varepsilon, \beta) > 0$ is a constant that might depend on other parameters, which are fixed to constants. Note that we estimated the probability to draw a new (unique) variable with 1. Thus, this also accounts for the probability to draw a variable that is not actually new. Especially, it accounts for the probability to draw one of the j non-unique variables. This means, the expression we have is an upper bound for the probability to draw *at most*

$\varepsilon \cdot i$ unique variables. As in the proof of Lemma 5.1 we have to distinguish three cases depending on the power law exponent β . Using Lemma A.1 we see that for $\beta < 3$

$$\begin{aligned}
 P_i &\leq \kappa^i \cdot \Delta^i \left(\frac{n}{i}\right)^i \cdot \sum_{j=1}^{\frac{k-\varepsilon}{2} \cdot i} \left(\frac{i^2}{j}\right)^j \cdot \left(\sum_{x=1}^n p_x^2\right)^j \cdot \left(\frac{j}{n}\right)^{(k-i-\varepsilon-i-2j)\frac{\beta-2}{\beta-1}} \\
 &\leq \kappa^i \cdot \Delta^i \left(\frac{n}{i}\right)^i \cdot \sum_{j=1}^{\frac{k-\varepsilon}{2} \cdot i} \left(\frac{i^2}{j}\right)^j \cdot n^{-2j\frac{\beta-2}{\beta-1}} \cdot \left(\frac{j}{n}\right)^{(k-i-\varepsilon-i-2j)\frac{\beta-2}{\beta-1}} \\
 &= \kappa^i \cdot \Delta^i \cdot n^{i(1-(k-\varepsilon)\frac{\beta-2}{\beta-1})} \cdot i^{-i} \cdot \sum_{j=1}^{\frac{k-\varepsilon}{2} \cdot i} \left(\frac{i^2}{j}\right)^j \cdot j^{(k-i-\varepsilon-i-2j)\frac{\beta-2}{\beta-1}}. \quad (4)
 \end{aligned}$$

Now, it remains to bound the inner sum. In order to do so, we will split it at $j_0 = \frac{3-\beta}{4}(k-\varepsilon) \cdot i$. It is easy to see that $0 < \frac{3-\beta}{4} < \frac{1}{4}$ for $2 < \beta < 3$, thus this choice of j is valid. For the first part of the sum it holds that

$$\begin{aligned}
 \sum_{j=1}^{\frac{3-\beta}{4}(k-\varepsilon) \cdot i} \left(\frac{i^2}{j}\right)^j \cdot j^{(k-i-\varepsilon-i-2j)\frac{\beta-2}{\beta-1}} &\leq \kappa^i \sum_{j=1}^{\frac{3-\beta}{4}(k-\varepsilon) \cdot i} i^{2j} \cdot j^{-j} \cdot i^{((k-\varepsilon)-i-2j)\frac{\beta-2}{\beta-1}} \\
 &\leq \kappa^i \cdot i^{(k-\varepsilon) \cdot i \cdot \frac{\beta-2}{\beta-1}} \sum_{j=1}^{\frac{3-\beta}{4}(k-\varepsilon) \cdot i} i^{\frac{2j}{\beta-1}} \\
 &\leq \kappa^i \cdot i^{(k-\varepsilon) \cdot i \cdot \frac{\beta-2}{\beta-1}} \cdot i^{2 \cdot \frac{3-\beta}{4} \cdot \frac{k-\varepsilon}{\beta-1} \cdot i} \\
 &= \kappa^i \cdot i^{\frac{k-\varepsilon}{2} \cdot i},
 \end{aligned}$$

where we used $j \leq \frac{3-\beta}{4}(k-\varepsilon) \cdot i$ and $((k-\varepsilon) \cdot i - 2j) \geq 0$ in the first line. The derived sum in the second line is a geometric series with base $i^{\frac{2}{\beta-1}} \geq 1$. This series is dominated by the term with $j = \frac{3-\beta}{4}(k-\varepsilon) \cdot i$. Additional factors of at most c^i for positive constants c are hidden in κ^i . For the second part of the sum it holds that

$$\begin{aligned}
 \sum_{j=\frac{3-\beta}{4}(k-\varepsilon) \cdot i}^{\frac{k-\varepsilon}{2} \cdot i} \left(\frac{i^2}{j}\right)^j \cdot j^{(k-i-\varepsilon-i-2j)\frac{\beta-2}{\beta-1}} &\leq \kappa^i \sum_{j=\frac{3-\beta}{4}(k-\varepsilon) \cdot i}^{\frac{k-\varepsilon}{2} \cdot i} i^{2j} \cdot j^{-j} \cdot i^{((k-\varepsilon)-i-2j)\frac{\beta-2}{\beta-1}} \\
 &\leq \kappa^i \cdot i^{(k-\varepsilon) \cdot i \cdot \frac{\beta-2}{\beta-1}} \sum_{j=\frac{3-\beta}{4}(k-\varepsilon) \cdot i}^{\frac{k-\varepsilon}{2} \cdot i} i^{j-2j\frac{\beta-2}{\beta-1}} \\
 &\leq \kappa^i \cdot i^{(k-\varepsilon) \cdot i \cdot \frac{\beta-2}{\beta-1}} \cdot i^{\frac{3-\beta}{4} \cdot \frac{k-\varepsilon}{\beta-1} \cdot i} \\
 &= \kappa^i \cdot i^{\frac{k-\varepsilon}{2} \cdot i},
 \end{aligned}$$

where we used $j \geq \frac{3-\beta}{4}(k-\varepsilon) \cdot i$ in the second and a geometric series in the third line. The base of the series is $i^{\frac{1}{\beta-1}} \geq 1$. Thus, the last term with $j = \frac{k-\varepsilon}{2} \cdot i$ dominates and we get the shown estimate with factors c^i for positive constants c hidden in κ^i again.

Thus,

$$\sum_{j=1}^{\frac{k-\varepsilon}{2}i} \left(\frac{i^2}{j}\right)^j \cdot j^{(k-i-\varepsilon-i-2j)\frac{\beta-2}{\beta-1}} \leq \kappa^i \cdot i^{\frac{k-\varepsilon}{2}i},$$

and plugging this into Equation (4) yields

$$P_i \leq \kappa^i \cdot \Delta^i \cdot n^{i(1-(k-\varepsilon)\frac{\beta-2}{\beta-1})} \cdot i^{i(\frac{k-\varepsilon}{2}-1)} = \kappa^i \cdot \Delta^i \cdot n^{-\varepsilon_2 \cdot i} \cdot i^{\varepsilon_1 \cdot i}.$$

Since, we want to sum overall P_i with $\frac{1}{3}w \leq i \leq \frac{2}{3}w$ for some w , it holds that

$$\begin{aligned} \sum_{i=\frac{1}{3}w}^{\frac{2}{3}w} P_i &\leq \sum_{i=\frac{1}{3}w}^{\frac{2}{3}w} \kappa^i \cdot \Delta^i \cdot n^{-\varepsilon_2 \cdot i} \cdot i^{\varepsilon_1 \cdot i} \\ &\leq \sum_{i=\frac{1}{3}w}^{\frac{2}{3}w} (\kappa \cdot \Delta \cdot n^{-\varepsilon_2} \cdot w^{\varepsilon_1})^i. \end{aligned}$$

This sums up to $o(1)$ as soon as $\kappa \cdot \Delta \cdot n^{-\varepsilon_2} \cdot w^{\varepsilon_1}$ is a suitably small constant and w is super-constant. In our case, we see that this holds for some

$$w \in O(n^{\varepsilon_2/\varepsilon_1} \Delta^{-1/\varepsilon_1}).$$

For $\beta = 3$ we get

$$\begin{aligned} P_i &\leq \kappa^i \cdot \Delta^i \left(\frac{n}{i}\right)^i \cdot \sum_{j=1}^{\frac{k-\varepsilon}{2}i} \left(\frac{i^2}{j}\right)^j \cdot \left(\sum_{x=1}^n p_x^2\right)^j \cdot \left(\frac{j}{n}\right)^{((k-\varepsilon)-i-2j)\frac{\beta-2}{\beta-1}} \\ &\leq \kappa^i \cdot \Delta^i \left(\frac{n}{i}\right)^i \cdot \sum_{j=1}^{\frac{k-\varepsilon}{2}i} \left(\frac{i^2}{j}\right)^j \cdot \left(\frac{\ln n}{n}\right)^j \cdot \left(\frac{j}{n}\right)^{((k-\varepsilon)-i-2j)\frac{1}{2}} \\ &= \kappa^i \cdot \Delta^i \cdot \left(\frac{n}{i}\right)^i \cdot n^{-\frac{k-\varepsilon}{2}i} \sum_{j=1}^{\frac{k-\varepsilon}{2}i} i^{2j} \cdot j^{\frac{k-\varepsilon}{2}i-2j} \cdot (\ln n)^j. \end{aligned} \quad (5)$$

We want to show that this inner sum is at most $\kappa^i \cdot (i \cdot \ln n)^{\frac{k-\varepsilon}{2}i}$. As before, we can split the sum. This time we split it at $j_0 = \frac{k-\varepsilon}{4}i$. For the first part we get

$$\begin{aligned} \sum_{j=1}^{\frac{k-\varepsilon}{4}i} i^{2j} \cdot j^{\frac{k-\varepsilon}{2}i-2j} \cdot (\ln n)^j &\leq \kappa^i \cdot \sum_{j=1}^{\frac{k-\varepsilon}{4}i} i^{2j} \cdot i^{\frac{k-\varepsilon}{2}i-2j} \cdot (\ln n)^j \\ &\leq \kappa^i \cdot i^{\frac{k-\varepsilon}{2}i} \cdot \sum_{j=1}^{\frac{k-\varepsilon}{4}i} (\ln n)^j \\ &\leq \kappa^i \cdot i^{\frac{k-\varepsilon}{2}i} \cdot (\ln n)^{\frac{k-\varepsilon}{4}i}, \end{aligned}$$

where we used that $\frac{k-\epsilon}{2}i - 2j \geq 0$ in the first line. The second line contains a geometric series with base $\ln n \geq 1$ again that we estimated by its dominating term $(\ln n)^{\frac{k-\epsilon}{4}i}$. The second part of the sum yields

$$\begin{aligned} \sum_{j=\frac{k-\epsilon}{4}i}^{\frac{k-\epsilon}{2}i} i^{2j} \cdot j^{\frac{k-\epsilon}{2}i-2j} \cdot (\ln n)^j &\leq \kappa^i \cdot \sum_{j=\frac{k-\epsilon}{4}i}^{\frac{k-\epsilon}{2}i} i^{2j} \cdot i^{\frac{k-\epsilon}{2}i-2j} \cdot (\ln n)^j \\ &\leq \kappa^i \cdot i^{\frac{k-\epsilon}{2}i} \sum_{j=\frac{k-\epsilon}{4}i}^{\frac{k-\epsilon}{2}i} (\ln n)^j \leq \kappa^i \cdot i^{\frac{k-\epsilon}{2}i} (\ln n)^{\frac{k-\epsilon}{2}i}, \end{aligned}$$

since $j \in \Theta(i)$. Plugging this into Equation (5) gives us

$$P_i \leq \kappa^i \cdot \Delta^i \cdot \left(\frac{n}{i}\right)^{i\left(1-\frac{k-\epsilon}{2}\right)} \cdot (\ln n)^{\frac{k-\epsilon}{2}i} = \kappa^i \cdot \Delta^i \cdot \left(\frac{n}{i}\right)^{-\epsilon_1 i} \cdot (\ln n)^{(\epsilon_1+1)i}.$$

As before, we can see that this is at most κ^i for some constant $\kappa \in (0, 1)$ if

$$w \in O\left(n/\ln^{1+\frac{1}{\epsilon_1}}n \cdot \Delta^{-1/\epsilon_1}\right)$$

is small enough.

For $\beta > 3$ we get

$$\begin{aligned} P_i &\leq \kappa^i \cdot \Delta^i \left(\frac{n}{i}\right)^i \cdot \sum_{j=1}^{\frac{k-\epsilon}{2}i} \left(\frac{i^2}{j}\right)^j \cdot \left(\sum_{x=1}^n p_x^2\right)^j \cdot \left(\frac{j}{n}\right)^{((k-\epsilon)\cdot i-2j)\frac{\beta-2}{\beta-1}} \\ &\leq \kappa^i \cdot \Delta^i \left(\frac{n}{i}\right)^i \cdot \sum_{j=1}^{\frac{k-\epsilon}{2}i} \left(\frac{i^2}{j}\right)^j \cdot n^{-j} \cdot \left(\frac{j}{n}\right)^{((k-\epsilon)\cdot i-2j)\frac{\beta-2}{\beta-1}} \\ &= \kappa^i \cdot \Delta^i \cdot \left(\frac{n}{i}\right)^i \cdot n^{-\frac{(k-\epsilon)\cdot i}{\beta-1}} \sum_{j=1}^{\frac{k-\epsilon}{2}i} i^{2j} \cdot n^{j\left(2\frac{\beta-2}{\beta-1}-1\right)} \cdot j^{((k-\epsilon)\cdot i-2j)\frac{\beta-2}{\beta-1}-j}. \end{aligned} \quad (6)$$

This time we are going to show that the inner sum is bounded by $i^{(k-\epsilon)\frac{\beta-2}{\beta-1}i} \cdot \left(\frac{n}{i}\right)^{\frac{k-\epsilon}{2}i\left(2\frac{\beta-2}{\beta-1}-1\right)}$.

Again, we split the sum. This time at

$$j_0 = \frac{(k-\epsilon)\frac{\beta-2}{\beta-1}i}{1+2\frac{\beta-2}{\beta-1}}i.$$

Our choice ensures $((k-\epsilon)\cdot i-2j)\frac{\beta-2}{\beta-1}-j \geq 0$ for $j \leq j_0$. Thus, in the first part of the sum all exponents are positive. It now holds that

$$j^{((k-\epsilon)\cdot i-2j)\frac{\beta-2}{\beta-1}-j} \leq j_0^{((k-\epsilon)\cdot i-2j)\frac{\beta-2}{\beta-1}-j} \leq \kappa^i \cdot i^{((k-\epsilon)\cdot i-2j)\frac{\beta-2}{\beta-1}-j},$$

for some constant κ that we can incorporate in the κ we already have. In the second part of the sum the exponent $((k - \epsilon) \cdot i - 2j) \frac{\beta-2}{\beta-1} - j$ is negative. However, we know that the base is $j \geq j_0 = \frac{(k-\epsilon) \frac{\beta-2}{\beta-1}}{1+2 \frac{\beta-2}{\beta-1}} i$. Thus,

$$j^{((k-\epsilon) \cdot i - 2j) \frac{\beta-2}{\beta-1} - j} \leq j_0^{((k-\epsilon) \cdot i - 2j) \frac{\beta-2}{\beta-1} - j} \leq \kappa^i \cdot i^{((k-\epsilon) \cdot i - 2j) \frac{\beta-2}{\beta-1} - j},$$

as well. This yields

$$\begin{aligned} & \sum_{j=1}^{\frac{k-\epsilon}{2} \cdot i} i^{2j} \cdot n^{j \left(2 \frac{\beta-2}{\beta-1} - 1 \right)} \cdot j^{((k-\epsilon) \cdot i - 2j) \frac{\beta-2}{\beta-1} - j} \\ & \leq \kappa^i \cdot \sum_{j=1}^{\frac{k-\epsilon}{2} \cdot i} i^{2j} \cdot n^{j \left(2 \frac{\beta-2}{\beta-1} - 1 \right)} \cdot i^{((k-\epsilon) \cdot i - 2j) \frac{\beta-2}{\beta-1} - j} \\ & = \kappa^i \cdot i^{(k-\epsilon) \frac{\beta-2}{\beta-1} i} \cdot \sum_{j=1}^{\frac{k-\epsilon}{2} \cdot i} \left(\frac{n}{i} \right)^{j \left(2 \frac{\beta-2}{\beta-1} - 1 \right)} \\ & \leq \kappa^i \cdot i^{(k-\epsilon) \frac{\beta-2}{\beta-1} i} \cdot \left(\frac{n}{i} \right)^{\frac{k-\epsilon}{2} i \left(2 \frac{\beta-2}{\beta-1} - 1 \right)}, \end{aligned}$$

where the last line holds, since $2 \frac{\beta-2}{\beta-1} - 1 > 0$, which implies that we have a geometric series with base at least one again, that we estimate by its dominating term, that is, the term with $j = \frac{k-\epsilon}{2} \cdot i$. If we plug our estimate into Equation (6) this gives us

$$P_i \leq \kappa^i \cdot \Delta^i \left(\frac{n}{i} \right)^{\left(1 - \frac{k-\epsilon}{2} \right) i} = \kappa^i \cdot \Delta^i \left(\frac{n}{i} \right)^{-\epsilon_1 \cdot i}.$$

We can now find a $w \in \Theta(n \cdot \Delta^{-1/\epsilon_1})$ small enough such that the property holds as desired.

In all three cases, we can choose w in such a way that the probability for the property not to hold is at most $\kappa^{\frac{w}{3}}$ for some constant $\kappa \in (0, 1)$. This means, the property holds a. a. s. for $w \in \omega(1)$. ■

The two properties, we showed in Lemmas 5.1 and 5.2 can be used to derive lower bounds on resolution width via the following theorem by Ben-Sasson and Wigderson [9].

Theorem 5.3 ([9]). *Let Φ be an unsatisfiable k -CNF formula with $k \geq 3$. If there is a $w \in \mathbb{N}$ such that*

- (i) *for all sets of clauses C' with $|C'| \leq w$ it holds that C' contains at least $|C'|$ different Boolean variables and;*
- (ii) *for all sets of clauses C' with $\frac{1}{3}w \leq |C'| \leq \frac{2}{3}w$ it holds that C' contains at least $\epsilon \cdot |C'|$ unique variables for some constant $\epsilon > 0$;*

then the resolution width of Φ is $\Omega(w)$.

Lemmas 5.1 and 5.2 together with Theorem 5.3 imply Corollary 5.4. However, Theorem 5.3 only works for unsatisfiable instances. Since the two lemmas do not condition on instances being unsatisfiable, we also need to make sure that the probability for having unsatisfiable instances is large

enough. In particular, we have to guarantee that this probability is larger than the error probabilities of Lemmas 5.1 and 5.2. If the probability of generating unsatisfiable instances is asymptotically larger than those error probabilities, the conditional probability of our width lower bounds to hold conditioned on instances being unsatisfiable will be approaching one. Since the error probabilities of the two lemmas are $o(1)$, we want the clause-variable ratio Δ to be high enough for instances to be unsatisfiable with at least constant probability. The resulting corollary is stated below. It only holds for unsatisfiable instances as well, that is, the probability bound on resolution width is actually a conditional probability conditioned on instances being unsatisfiable.

Corollary 5.4. *Let Φ be an unsatisfiable random power-law k -SAT formula with n variables, $m \in \Omega(n)$ clauses, $k \geq 3$, and power-law exponent $\beta > \frac{2k-1}{k-1}$ constant. Let $\Delta = m/n$ be large enough so that Φ is unsatisfiable at least with constant probability. Let $\varepsilon, \varepsilon_1, \varepsilon_2$ be constants with $\varepsilon > 0$, $\varepsilon_1 = \frac{k-\varepsilon}{2} - 1 > 0$, and $\varepsilon_2 = (k - \varepsilon) \cdot \frac{\beta-2}{\beta-1} - 1 > 0$. For the resolution width w of Φ , it holds a. a. s. that:*

- (i) *If $\beta \in \left(\frac{2k-1}{k-1}, 3\right)$ and $\Delta \in o(n^{\varepsilon_2})$, then $w \in \Omega(n^{\varepsilon_2/\varepsilon_1} \cdot \Delta^{-1/\varepsilon_1})$.*
- (ii) *If $\beta = 3$ and $\Delta \in o(n^{\varepsilon_1/\log^{\varepsilon_1+1} n})$, then $w \in \Omega\left(n \cdot \Delta^{-1/\varepsilon_1} / \log^{1+\frac{1}{\varepsilon_1}} n\right)$.*
- (iii) *If $\beta > 3$ and $\Delta \in o(n^{\varepsilon_1+1})$, then $w \in \Omega(n \cdot \Delta^{-1/\varepsilon_1})$.*

Proof. If both Lemmas 5.1 and 5.2 hold, we can use Theorem 5.3 to get the desired bound on resolution width. As stated before, Theorem 5.3 only holds for unsatisfiable instances. Thus, if a random formula Φ is unsatisfiable at least with constant probability, it holds that the conditional probability for the bounds stated in the corollary to hold is at least

$$\frac{\Pr(\Phi \text{ unsat}) - o(1)}{\Pr(\Phi \text{ unsat})} = 1 - o(1),$$

conditioned on Φ being unsatisfiable, where the $o(1)$ term is the error probability from Lemmas 5.1 and 5.2. We are going to show that the values of w from Lemma 5.2 are smaller than those from Lemma 5.1. The expansion bound from Lemma 5.1 also holds for those smaller values of w due to the definition of bipartite expansion. Thus, the bound from Lemma 5.2 gives us the maximum w we can achieve.

First, consider the case $\beta \in \left(\frac{2k-1}{k-1}, 3\right)$. Let $\varepsilon_3 = k \frac{\beta-2}{\beta-1} - 1$ and $\varepsilon_4 = (k-2) \frac{\beta-2}{\beta-1}$. We want to show that

$$n^{\varepsilon_2/\varepsilon_1} \cdot \Delta^{-1/\varepsilon_1} \leq n^{\varepsilon_3/\varepsilon_4} \cdot \Delta^{-1/\varepsilon_4}. \quad (7)$$

Both bounds only hold for

$$\Delta \in o(n^{\varepsilon_2}) \subseteq o(n^{\varepsilon_3}/\log^{\varepsilon_4}(n)),$$

since $\varepsilon_2 = (k - \varepsilon) \frac{\beta-2}{\beta-1} - 1 < k \frac{\beta-2}{\beta-1} - 1 = \varepsilon_3$. It holds that

$$\varepsilon_2/\varepsilon_1 = \frac{(k - \varepsilon) \cdot \frac{\beta-2}{\beta-1} - 1}{\frac{k-\varepsilon}{2} - 1} < \frac{k \cdot \frac{\beta-2}{\beta-1} - 1}{\frac{k}{2} - 1} < \frac{k \frac{\beta-2}{\beta-1} - 1}{(k-2) \cdot \frac{\beta-2}{\beta-1}} = \varepsilon_3/\varepsilon_4.$$

We can now distinguish four cases. First, assume $\Delta \geq 1$. If $\varepsilon_1 \leq \varepsilon_4$, then $\Delta^{-1/\varepsilon_1} \leq \Delta^{-1/\varepsilon_4}$, which implies Inequality (7). If $\varepsilon_1 > \varepsilon_4$, we need to ensure

$$\Delta \leq n^{\left(\frac{\varepsilon_3}{\varepsilon_4} - \frac{\varepsilon_2}{\varepsilon_1}\right) / \left(\frac{1}{\varepsilon_4} - \frac{1}{\varepsilon_1}\right)}.$$

This is already the case, since we assume $\Delta \in o(n^{\varepsilon_2})$ and $\varepsilon_2 \leq \left(\frac{\varepsilon_3}{\varepsilon_4} - \frac{\varepsilon_2}{\varepsilon_1}\right) / \left(\frac{1}{\varepsilon_4} - \frac{1}{\varepsilon_1}\right)$ due to $\varepsilon_1 > \varepsilon_4$ and $\varepsilon_3 \geq \varepsilon_2$. Thus, Inequality (7) holds.

Now assume $\Delta < 1$. If $\varepsilon_1 \leq \varepsilon_4$, we need to ensure that

$$\Delta \geq n^{\left(\frac{\varepsilon_2}{\varepsilon_1} - \frac{\varepsilon_3}{\varepsilon_4}\right) / \left(\frac{1}{\varepsilon_1} - \frac{1}{\varepsilon_4}\right)}.$$

This already holds, since we assume $\Delta \in \Omega(1)$ and $\left(\frac{\varepsilon_2}{\varepsilon_1} - \frac{\varepsilon_3}{\varepsilon_4}\right) / \left(\frac{1}{\varepsilon_1} - \frac{1}{\varepsilon_4}\right) \leq 0$ due to $\varepsilon_1 \leq \varepsilon_4$ and $\varepsilon_2/\varepsilon_1 \leq \varepsilon_3/\varepsilon_4$. Thus, Inequality (7) holds. If $\varepsilon_4 \leq \varepsilon_1$, then $\Delta^{-1/\varepsilon_1} \leq \Delta^{-1/\varepsilon_4}$ and Inequality (7) holds as well.

Now consider $\beta = 3$. We need to show that

$$n/(\ln n)^{\frac{\varepsilon_1+1}{\varepsilon_1}} \Delta^{-1/\varepsilon_1} = n/(\ln n)^{\frac{k-\varepsilon}{k-\varepsilon-2}} \Delta^{-2/(k-\varepsilon-2)} \in O(n \cdot (\Delta \cdot \ln n)^{-2/(k-2)}).$$

Again, the left-hand side is from Lemma 5.2 and the right-hand side is from Lemma 5.1. This holds, due to our assumption $\Delta \in \Omega(1)$ and since $\varepsilon_1 = \frac{k-\varepsilon}{2} - 1 > 0$ implies $0 < \varepsilon < k - 2$ and thus $\frac{k-\varepsilon}{k-\varepsilon-2} > \frac{2}{k-\varepsilon-2} > \frac{2}{k-2}$. Additionally, the bound only holds up to $\Delta \in o\left(n^{(k-\varepsilon-2)/2} / \ln^{(k-\varepsilon)/2}(n)\right) \subseteq o\left(n^{(k-2)/2} / \log^{(k-2)/2+1}(n)\right)$.

For $\beta > 3$ we have to show

$$n \cdot \Delta^{-1/\varepsilon_1} \in O(n \cdot \Delta^{-1/\varepsilon_4}),$$

as well as $\Delta \in o(n^{\varepsilon_1}) \subseteq o((n/\log n)^{\varepsilon_4})$. This holds since $\varepsilon_1 = \frac{k-\varepsilon}{2} - 1 \leq (k-2)\frac{\beta-2}{\beta-1} = \varepsilon_4$ due to $\beta > 3$. This shows that in all three cases the bounds from Lemma 5.2 are smaller, thus giving us the lower bounds on resolution width as stated in the corollary. ■

This is nearly the statement of Theorem 5.8. However, via bipartite expansion we can already show linear resolution width at constant clause-variable ratios for $\beta > \frac{2k-2}{k-2}$ instead of $\beta > 3$. This gives a better bound for $k \geq 5$. The bounds on bipartite expansion and the resulting bounds on resolution width will be derived in the next section.

5.2 | A lower bound on bipartite expansion

In this section, we show an improved bound on the bipartite expansion. We will use it to obtain a linear lower bound on resolution width for $\beta > \frac{2k-2}{k-2}$, which is potentially smaller than three, and therefore improves the previous bound. Recall that linear resolution width implies exponential resolution size, and thus also exponential tree-like resolution size. Moreover, our bound on the bipartite expansion can also be used to bound the so-called resolution clause space, which additionally yields an exponential lower bound on tree-like resolution size for $\beta > \frac{2k-3}{k-2}$ as we will see at the end of this section. The following lemma shows the bipartite expansion property.

Lemma 5.5. *Let Φ be a random power-law k -SAT formula with n variables, m clauses, $k \geq 3$, power-law exponent $\beta > \frac{2k-3}{k-2}$, and let $\varepsilon \in (0, (k-1) \cdot \frac{\beta-2}{\beta-1} - 1)$ constant. If $\Delta = m/n \in o(n^\varepsilon / \log^\varepsilon n)$, then there exists an $r \in \Theta(n \cdot \Delta^{-1/\varepsilon})$ such that the clause-variable incidence graph $G(\Phi)$ is an (r, c) -bipartite expander a. a. s. for $c = (k-1) - (1+\varepsilon) \cdot \frac{\beta-1}{\beta-2}$.*

Proof. First, note that our choice of $\beta > \frac{2k-3}{k-2}$ guarantees that the interval $(0, (k-1) \cdot \frac{\beta-2}{\beta-1} - 1)$ from which we choose ε is not empty. This interval is chosen in such a way that $c > 0$ is guaranteed. As in the proof of [8, Lemma 5.1], we define a bad event \mathcal{E} , that $G(\Phi)$ is not an (r, c) -bipartite expander. If \mathcal{E} happens, then there is a set $C' \subseteq C$ with $1 \leq |C'| \leq r$ such that $|N(C')| < (1+c) \cdot |C'|$. Given a set $C' \subseteq C = [m]$ of clause indices with $|C'| = i$ we want to bound the probability P_i that the $k \cdot i$ indices of variables appearing in those clauses contain at most $(1+c) \cdot i$ different variables. Since clauses contain variables without repetition, it holds that P_i is dominated by the probability to draw at most $(1+c) \cdot i$ different variables when drawing $k \cdot i$ Boolean variables independently at random. Now imagine sampling these $k \cdot i$ variables in some arbitrary, but fixed order. It holds that the probability to draw a new variable is at most 1, while the probability to draw an old variable is at most the probability to draw one of the $(1+c) \cdot i$ variables of maximum probability. As before, the sum of these probabilities is denoted by $F((1+c) \cdot i)$. This gives us

$$P_i \leq \binom{m}{i} \cdot \binom{k \cdot i}{(1+c) \cdot i} \cdot 1^{(1+c) \cdot i} \cdot F((1+c) \cdot i)^{k \cdot i - (1+c) \cdot i}.$$

Note that this expression also captures the case that we draw fewer than $(1+c) \cdot i$ different variables, since the probability to draw a new variable is bounded by one and thus also captures the probability that this new variable is in fact an old one. In the case of a power-law distribution, we have

$$F((1+c) \cdot i) \sim \left(\frac{(1+c) \cdot i}{n} \right)^{\frac{\beta-2}{\beta-1}},$$

due to Lemma A.1 and thus

$$\begin{aligned} P_i &\leq \binom{m}{i} \cdot \binom{k \cdot i}{(1+c) \cdot i} \cdot \left(\frac{(1+c) \cdot i}{n} \right)^{(k-(1+c)) \cdot \frac{\beta-2}{\beta-1} \cdot i} \\ &\leq \left(\frac{e \cdot m}{i} \right)^i \cdot \left(\frac{e \cdot k}{1+c} \right)^{(1+c) \cdot i} \cdot \left(\frac{(1+c) \cdot i}{n} \right)^{(k-(1+c)) \cdot \frac{\beta-2}{\beta-1} \cdot i} \\ &= \kappa(c, \beta, k)^i \cdot \Delta^i \left(\frac{i}{n} \right)^{i((k-(1+c)) \cdot \frac{\beta-2}{\beta-1} - 1)} \\ &= \kappa(c, \beta, k)^i \cdot \Delta^i \left(\frac{i}{n} \right)^{i \cdot \varepsilon}, \end{aligned}$$

for some constant $\kappa(c, \beta, k) > 0$, $m = \Delta \cdot n$, and $c = (k-1) - (1+\varepsilon) \cdot \frac{\beta-1}{\beta-2}$.

Summing overall $i \geq 1$ now yields

$$\Pr[\mathcal{E}] \leq \sum_{i=1}^r \kappa(c, \beta, k)^i \cdot \Delta^i \cdot \left(\frac{i}{n} \right)^{i \cdot \varepsilon}.$$

We split this sum into two parts, the first part from $i = 1$ to $\lfloor \varepsilon \cdot \log n \rfloor$ and the second part from $\lfloor \varepsilon \cdot \log n \rfloor$ to r . For the first part we get

$$\begin{aligned} \sum_{i=1}^{\lfloor \varepsilon \cdot \log n \rfloor} \kappa(c, \beta, k)^i \cdot \Delta^i \cdot \left(\frac{i}{n}\right)^{i \cdot \varepsilon} &\leq \sum_{i=1}^{\lfloor \varepsilon \cdot \log n \rfloor} \kappa(c, \beta, k)^i \cdot \Delta^i \cdot \left(\frac{\varepsilon \cdot \log n}{n}\right)^{i \cdot \varepsilon} \\ &\leq 2 \cdot \kappa(c, \beta, k) \cdot \Delta \cdot \left(\frac{\varepsilon \cdot \log n}{n}\right)^{\varepsilon} \\ &\in O\left(\Delta \left(\frac{\log n}{n}\right)^{\varepsilon}\right), \end{aligned}$$

which holds, since $\sum_{i=1}^m \alpha^i \leq 2 \cdot \alpha$ for all $m \geq 1$ and $\alpha < \frac{1}{2}$. This holds for big enough values of n and for $\Delta \in o(n^{\varepsilon} / \log^{\varepsilon} n)$. For the second part we get

$$\sum_{i=\lfloor \varepsilon \cdot \log n \rfloor}^r \kappa(c, \beta, k)^i \cdot \Delta^i \cdot \left(\frac{i}{n}\right)^{i \cdot \varepsilon} \leq \sum_{i=\lfloor \varepsilon \cdot \log n \rfloor}^r 2^{-i} \in O\left(\left(\frac{1}{n}\right)^{\varepsilon}\right),$$

which holds if we choose

$$r \in O(n \cdot \Delta^{-1/\varepsilon})$$

small enough so that $\Delta \cdot \left(\frac{r}{n}\right)^{\varepsilon} < \frac{1}{2 \cdot \kappa(c, \beta, k)}$. ■

This notion of bipartite expansion is connected to the resolution width of a formula. The following corollary, implicitly stated by Ben-Sasson and Wigderson [9], formalizes this connection.

Corollary 5.6 ([9]). *Let $k \geq 3$ integer and constant, let $\varepsilon > 0$ constant, and let Φ be an unsatisfiable Boolean formula in k -CNF. If there is a constant $\varepsilon > 0$ such that $G(\Phi)$ is a $\left(r, \frac{k+\varepsilon}{2} - 1\right)$ -bipartite expander, then Φ has resolution width at least $\Omega(r)$.*

Proof. Due to the definition of bipartite expansion, $\frac{k+\varepsilon}{2} > 1$ ensures the first condition of Theorem 5.3. We will show that the second condition is fulfilled as well. Let $G(\Phi) = (C, V, E)$ and let $C' \subseteq C$ with $\frac{1}{3}r \leq |C'| \leq \frac{2}{3}r$. Let $\delta C'$ denote the set of unique variables from C' , that is, $\delta C' = \{v \in N(C') \mid |N(v) \cap C'| = 1\}$. As Ben-Sasson and Wigderson state in [9, Proof of Theorem 6.5] it holds that:

$$|N(C')| - |\delta C'| \leq (k \cdot |C'| - |\delta C'|)/2,$$

which implies

$$|\delta C'| \geq 2|N(C')| - k \cdot |C'| \geq \varepsilon \cdot |C'|,$$

due to the $\left(r, \frac{k+\varepsilon}{2} - 1\right)$ -bipartite expansion. These two properties imply a resolution width of $\Omega(r)$. ■

This result on the bipartite expansion of power-law random k -SAT allows us to derive the following corollary on resolution width. Again, we require the clause-variable ratio Δ to be high enough for instances to be unsatisfiable with at least constant probability.

Corollary 5.7. *Let Φ be an unsatisfiable random power-law k -SAT formula with n variables, $m \in \Omega(n)$ clauses, $k \geq 3$, and power-law exponent $\beta > \frac{2k-2}{k-2}$. Let $\Delta = m/n$ be large enough so that Φ is unsatisfiable at least with constant probability. For $0 < \varepsilon < \frac{k}{2} \cdot \frac{\beta-2}{\beta-1} - 1$ constant and $\Delta \in o(n^\varepsilon / \log^\varepsilon n)$ it holds a. a. s. that Φ has resolution width $w \in \Omega(n \cdot \Delta^{-1/\varepsilon})$.*

Proof. Due to Lemma 5.5 $G(\Phi)$ is a $(\Omega(n \cdot \Delta^{-1/\varepsilon}), c)$ -bipartite expander for $c = (k-1) - (1+\varepsilon)\frac{\beta-1}{\beta-2}$. With $\beta > \frac{2k-2}{k-2}$, it holds that we can choose an $\varepsilon > 0$ so that $c > \frac{k}{2} - 1$. This means, the requirement of Corollary 5.6 is fulfilled and implies the statement. ■

Together with Corollary 5.4 the former corollary implies Theorem 5.8.

Theorem 5.8. *Let Φ be an unsatisfiable random power-law k -SAT formula with n variables, $m \in \Omega(n)$ clauses, $k \geq 3$, and power-law exponent $\beta > \frac{2k-1}{k-1}$. Let $\Delta = m/n$ be large enough so that Φ is unsatisfiable at least with constant probability. Let $\varepsilon, \varepsilon_1, \dots, \varepsilon_3$ be constants with $\varepsilon > 0$, $\varepsilon_1 = \frac{k-\varepsilon}{2} - 1 > 0$, $\varepsilon_2 = (k-\varepsilon) \cdot \frac{\beta-2}{\beta-1} - 1 > 0$, and $0 < \varepsilon_3 < (\frac{k}{2} - 1) \cdot \frac{\beta-2}{\beta-1} - 1$. For the resolution width w of Φ , it holds a. a. s. that:*

- (i) *If $\beta \in (\frac{2k-1}{k-1}, 3)$ and $\Delta \in o(n^{\varepsilon_2/\varepsilon_1})$, then $w \in \Omega(n^{\varepsilon_2/\varepsilon_1} \Delta^{-1/\varepsilon_1})$.*
- (ii) *If $\beta = 3$ and $\Delta \in o(n^{\varepsilon_1/\log^{1+\varepsilon_1} n})$, then $w \in \Omega(n \cdot \Delta^{-1/\varepsilon_1} / \log^{1+1/\varepsilon_1} n)$.*
- (iii) *If $\beta > 3$ and $\Delta \in o(n^{\varepsilon_1})$, then $w \in \Omega(n \cdot \Delta^{-1/\varepsilon_1})$.*
- (iv) *If $\beta > \frac{2k-2}{k-2}$ and $\Delta \in o(n^{\varepsilon_3/\log^{\varepsilon_3} n})$, then $w \in \Omega(n \cdot \Delta^{-1/\varepsilon_3})$.*

Additionally, Ben-Sasson and Galesi [8] state a theorem that directly connects bipartite expansion and tree-like resolution size. An application of this theorem yields a slightly better bound on tree-like resolution size than the ones derived from resolution width.

Theorem 5.9 ([8]). *Let Φ be an unsatisfiable CNF and let $G(\Phi) = (U \cup V, E)$ be the clause-variable incidence graph of Φ . If $G(\Phi)$ is a (r, c) -bipartite expander then Φ has resolution clause space of at least $\frac{c \cdot r}{2+c}$ and tree-like resolution size of at least $\exp\left(\Omega\left(\frac{c \cdot r}{2+c}\right)\right)$.*

Proof. ([8, Theorems 4.2 and 3.3]). State together that any bipartite graph that is an (r, c) -bipartite expander has a resolution clause space of at least $\frac{c \cdot r}{2+c}$. Thus, with [31, Theorem 1.6], it holds that the resolution size for formulas whose clause-variable incidence graph is an (r, c) -bipartite expander, is at least $\exp\left(\frac{c \cdot r}{2+c}\right)$. ■

This leads to the following corollary, which already asserts exponential tree-like resolution size for constant clause-variable ratios at $\beta > \frac{2k-3}{k-2}$.

Corollary 5.10. *Let Φ be an unsatisfiable random power-law k -SAT formula with n variables, $m = \Omega(n)$ clauses, $k \geq 3$, and power-law exponent $\beta > \frac{2k-3}{k-2}$. Let $\Delta = m/n$ be large enough so that Φ is unsatisfiable at least with constant probability. For $0 < \varepsilon < (k-1) \cdot \frac{\beta-2}{\beta-1} - 1$ constant and $\Delta \in o((n/\log n)^\varepsilon)$, it holds that Φ has tree-like resolution size $\exp(\Omega(n \cdot \Delta^{-1/\varepsilon}))$.*

Proof. Using Lemma 5.5 we see that for $\beta > \frac{2k-3}{k-2}$ the clause-variable incidence graph a. a. s. is a $(\Theta(n \cdot \Delta^{-1/\varepsilon}), c)$ -bipartite expander for some constant $c > 0$. Thus, Theorem 5.9 implies the statement. ■

6 | THE COMPLEXITY OF VORONOI DIAGRAMS

We first show quadratic lower bounds on the complexity (number of non-empty regions) of order- k Voronoi diagrams that already hold in rather basic settings. Afterwards, we consider random point sets and prove a linear upper bound.

6.1 | Worst-case lower bounds

In this section, we show worst-case lower bounds on the number of non-empty regions of higher-order Voronoi diagrams. As already mentioned in Section 2.3, our lower bounds are based on previously known lower bounds on the number of vertices of Voronoi diagrams, in conjunction with a new theorem connecting the number of vertices with the number of regions in higher orders. This theorem relies on the fact that there are not too many different points with equal distance to a set of $d + 1$ sites in d -dimensional space. For the unweighted case and for $\mathfrak{p} \neq \infty$, the result in the next lemma was shown by Lê [45]. We extend it to weighted sites and $\mathfrak{p} = \infty$, following along the lines of Lê's proof [45] (at least for $\mathfrak{p} \neq \infty$): (i) Observe that the points with equal distance to the $d + 1$ sites is the set of solutions to a system of polynomial equations. (ii) Show that the so-called *additive complexity* of these polynomial equations is bounded by a constant only depending on d . (iii) Apply [45, Proposition 3], giving an upper bound on the number of solutions to a system of equations that only depends on d and on the additive complexities of the equations.

Lemma 6.1. *Let A be a set of $d + 1$ weighted sites in general position⁶ in \mathbb{R}^d equipped with a \mathfrak{p} -norm. Then, the number of points with equal weighted distance to all sites in A only depends on d .*

Proof. Assume $\mathfrak{p} \neq \infty$, and let s_0, \dots, s_d be $d + 1$ sites with normalized weights $\omega_0, \dots, \omega_d$. Recall that the weighted distance between s_i and a point \mathbf{p} is $\|s_i - \mathbf{p}\|/\omega_i$. Thus, \mathbf{p} has the same distance to all $d + 1$ sites if, for all $i \in [d]$, it satisfies

$$\frac{\|s_0 - \mathbf{p}\|}{\omega_0} - \frac{\|s_i - \mathbf{p}\|}{\omega_i} = 0. \quad (8)$$

We note that this polynomial has the same form in the unweighted case [45, eq. 10], except we have the additional factors $1/\omega_0$ and $1/\omega_i$.

Concerning (ii), it thus suffices to note that these additional factors do not significantly increase the so-called additive complexity. We do not fully define the additive complexity here, but rather cite the properties crucial for this proof. The additive complexity $L_+(P)$ of a polynomial P is defined to be 0 if P is a monomial. Moreover, by [45, Lemma 4], it holds that

$$\begin{aligned} L_+(P_1 + \dots + P_n) &\leq n - 1 + L_+(P_1) + \dots + L_+(P_n), \\ L_+(P^m) &\leq L_+(P), \text{ for any } m \in \mathbb{N}, \text{ and} \\ L_+(PQ) &\leq L_+(P) + L_+(Q), \end{aligned}$$

⁶For a formal definition what general position means in this context, see [45]. As usual, the configurations excluded by the assumption of general position have measure 0.

where all P_i , P , and Q are polynomials. With this, it is easy to see that the additive complexity of the polynomial in Equation (8) is bounded by a constant only depending on d . In fact, the last bound, $L_+(PQ) \leq L_+(P) + L_+(Q)$ in conjunction with the property that constants are monomials with additive complexity 0, makes it so that the additional constant factors ω_0 and ω_i do not increase the additive complexity at all. Thus, the additive complexity is bounded by $4d - 1$ [45, Lemma 5].

Finally, applying [45, Proposition 3] directly yields the claim, which concludes the proof for $\mathfrak{p} \neq \infty$.

For $\mathfrak{p} = \infty$, we cannot use the same argument, as Equation (8) is not polynomial: $\|s_i - \mathfrak{p}\|$ involves the maximum over all coordinates. However, for each s_i , there are only d possibilities to which coordinate the maximum is evaluated, leading to d^{d+1} combinations. For each of these combinations, we consider its own system of equations. Denote the resulting set of systems of equations with \mathcal{E} . Clearly, every solution for the system of equations in (8) is a solution to at least one system in \mathcal{E} . Thus, the number of solutions to (8) is bounded by the total number of solutions to systems in \mathcal{E} . Clearly, with the same argument as above, the number of solutions to each system of equations in \mathcal{E} is bounded by a constant only depending on d . As \mathcal{E} contains only d^{d+1} systems, this bounds the number of solutions to (8) by a constant only depending on d . ■

With this, we can now prove the theorem establishing the connection between vertices and non-empty regions.

Theorem 6.2. *Let S be a set of n weighted sites in general position in \mathbb{R}^d equipped with a \mathfrak{p} -norm. If the order- k Voronoi diagram has ℓ vertices, then the order- $(k + d)$ Voronoi diagram has $\Omega(\ell)$ non-empty regions.*

Proof. We first show that a vertex of the order- k Voronoi diagram is an interior point of a non-empty region of the order- $(k + d)$ Voronoi diagram. Afterwards, we show that only a constant number of different vertices can end up in the same region.

Let $\mathfrak{p} \in \mathbb{R}^d$ be a vertex of the order- k Voronoi diagram. Then \mathfrak{p} has equal weighted distance to exactly $d + 1$ sites (the sites are in general position). Let $\{s_1, \dots, s_{d+1}\} = A \subseteq S$ be these sites and let P be the ε -environment of \mathfrak{p} , that is, a ball with sufficiently small radius ε centered at \mathfrak{p} . For a point $\mathfrak{p}' \in P$, sort all sites in S by weighted distance from \mathfrak{p}' . Then all sites in A appear consecutive in this order. Moreover, we obtain almost the same order of S for every $\mathfrak{p}' \in P$. The only difference is that the sites of A might be reordered. Also, as \mathfrak{p} is a vertex of the order- k Voronoi diagram, at least one site from A belongs to the k sites with smallest weighted distance to \mathfrak{p} . It follows that the first $k + d$ sites in this order completely include all sites from A . Thus, the $k + d$ closest sites are the same for all points in the ε -environment P around \mathfrak{p} ; let B be the set of these sites. It follows that B has non-empty Voronoi region in the order- $(k + d)$ Voronoi diagram as this region has \mathfrak{p} in its interior.

It remains to show that only a constant number of vertices of the order- k Voronoi diagram can be contained in the same region of the order- $(k + d)$ Voronoi diagram, that is, the order- $(k + d)$ region belonging to B includes only a constant number of order- k vertices. As stated above, every order- k vertex belongs to a subset $A \subseteq B$ with $|A| = d + 1$. There are only $\binom{|B|}{|A|} \leq \binom{k+d}{d+1}$ such subsets A , which is constant for constant k and d . Moreover, every fixed subset A of $d + 1$ sites is responsible for only a constant number of

vertices due to Lemma 6.1. Thus, only a constant number of order- k vertices end up in the same order- $k + d$ region, which concludes the proof. ■

Theorem 6.2 transfers some known lower bounds on the number of vertices of Voronoi diagrams to lower bounds on the number of non-empty regions of order- k Voronoi diagrams. In particular, we get the following corollaries.

Corollary 6.3. *In the worst case, the order-4 Voronoi diagram of n (unweighted) sites in 3-dimensional Euclidean space has $\Omega(n^2)$ non-empty regions.*

Proof. In the worst case, the ordinary (order-1, unweighted) Voronoi diagram of n sites in 3-dimensional Euclidean space has $\Omega(n^2)$ vertices [43,56]. Applying Theorem 6.2 yields the claim. ■

Corollary 6.4. *In the worst case, the order-3 Voronoi diagram of n weighted sites in 2-dimensional Euclidean space has $\Omega(n^2)$ non-empty regions.*

Proof. In the worst case, the order-1 Voronoi diagram of n weighted sites in 2-dimensional Euclidean space has $\Omega(n^2)$ vertices [5]; also see Figure 2. Applying Theorem 6.2 yields the claim. ■

6.2 | Upper bounds for sites with random positions

Let $S = \{s_1, \dots, s_n\} \subseteq \mathbb{T}^d$ be n randomly positioned sites with weights w_1, \dots, w_n . In the following, we bound the complexity of the weighted order- k Voronoi diagram in terms of non-empty regions. Recall from Section 3 that the torus \mathbb{T}^d is the hypercube $[0, 1]^d$ that wraps around in every dimension in the sense that opposite sides are identified. However, the following arguments do not require this property. Thus, the exact same results hold for Voronoi diagrams in hypercubes.

For the normalized weights $\omega_1, \dots, \omega_n$, recall from Section 3, that the point $p \in \mathbb{T}^d$ belongs to the Voronoi region corresponding to $A \subseteq S$ with $|A| = k$ if there exists a radius r such that $\|p - s_i\| \leq \omega_i r$ if $s_i \in A$ and $\|p - s_i\| > \omega_i r$ if $s_i \notin A$. Thus, A has non-empty order- k Voronoi region if and only if there exists such a point p . Our goal in the following is to bound the probability for its existence.

Our general approach to achieve such a bound is the following. The condition $\|p - s_i\| \leq \omega_i r$ for $s_i \in A$ basically tells us the sites in A are either close together or that r has to be large. In contrast to that, the condition $\|p - s_i\| > \omega_i r$ for $s_i \notin A$ tells us that many sites (namely all $n - k$ sites in $S \setminus A$) have to lie sufficiently far away from p , which is unlikely if r is large. How unlikely this is of course depends on r and thus on how close the sites in A lie together. Therefore, to follow this approach, we first condition on how close the sites in A lie together.

To formalize this, consider a size- k subset $A \subseteq S$ and assume without loss of generality that $A = \{s_1, \dots, s_k\}$. The site in A with the lowest weight, without loss of generality s_1 , will play a special role. We define the random variable R_A to be

$$R_A = \max_{i \in [k]} \frac{\|s_1 - s_i\|}{\omega_1 + \omega_i}. \quad (9)$$

The intuition behind the definition of R_A is the following. The weighted center between s_1 and s_i is the point p on the line between them such that $\|s_1 - p\| = \omega_1 r$ and $\|s_i - p\| = \omega_i r$ for a radius $r \in \mathbb{R}$. Then R_A is the maximum value for r over $i \in [k]$. In the unweighted setting, R_A is just half the maximum

distance between s_1 and any other site s_i . In a sense, R_A describes how close the sites in A lie together. Thus, it provides a lower bound on r .

Based on R_A , we slightly relax the condition on A having non-empty Voronoi region. We call A *relevant* if there exists a point $\mathbf{p} \in \mathbb{T}^d$ and a radius $r \geq R_A$ such that $\|s_1 - \mathbf{p}\| \leq \omega_1 r$ and $\|s_i - \mathbf{p}\| > \omega_i r$ for $i > k$. The following lemma states that being relevant is in fact a weaker condition than having non-empty order- k Voronoi region. Thus, bounding the probability that a set is relevant from above also bounds the probability for a non-empty Voronoi region from above.

Lemma 6.5. *A subset of k sites that has a non-empty order- k Voronoi region is relevant.*

Proof. Assume $A = \{s_1, \dots, s_k\}$ has a non-empty order- k Voronoi region. Then there exists a point \mathbf{p} and a radius r such that $\|s_i - \mathbf{p}\| \leq \omega_i r$ if and only if $i \leq k$. Thus, $\|s_1 - \mathbf{p}\| \leq \omega_1 r$ and $\|s_i - \mathbf{p}\| > \omega_i r$ for $i > k$ clearly holds, and it remains to show $r \geq R_A$. From $\|s_i - \mathbf{p}\| \leq \omega_i r$ for $i \in [k]$ it follows that $\|s_1 - \mathbf{p}\| + \|s_i - \mathbf{p}\| \leq \omega_1 r + \omega_i r$ holds for any $i \in [k]$. Thus, by rearranging and applying the triangle inequality, we obtain $r \geq (\|s_1 - \mathbf{p}\| + \|s_i - \mathbf{p}\|) / (\omega_1 + \omega_i) \geq \|s_1 - s_i\| / (\omega_1 + \omega_i)$. This immediately yields $r \geq R_A$. ■

Now we proceed to bound the probability that a set A is relevant. The following lemma bounds this probability conditioned on the random variable R_A . At its core, we have to bound the probability of the event $\|s_i - \mathbf{p}\| > \omega_i r$ for $s_i \notin A$. For a fixed point \mathbf{p} and a fixed radius r , this is rather easy. Thus, most of the proof is concerned with eliminating the existential quantifiers for \mathbf{p} and r .

Lemma 6.6. *For constants c_1 and c_2 depending only on d and \mathfrak{p} , it holds that*

$$\Pr[A \text{ is relevant} | R_A] \leq c_1 \min_{s_i \in A} \{w_i\} \exp\left(-c_2 R_A^d \sum_{s_i \notin A} w_i\right).$$

Proof. As before, we assume that $A = \{s_1, \dots, s_k\}$ and that s_1 has minimum weight among sites in A , that is, $\min_{s_i \in A} \{w_i\} = w_1$. By definition, A is relevant conditioned on R_A , if and only if there exists a radius $r \geq R_A$ and point $\mathbf{p} \in \mathbb{T}^d$ such that $\|s_1 - \mathbf{p}\| \leq \omega_1 r$ and $\|s_i - \mathbf{p}\| > \omega_i r$ for $i > k$, that is, formally we have

$$\exists r \geq R_A \exists \mathbf{p} \in \mathbb{T}^d \forall i > k : \|s_1 - \mathbf{p}\| \leq \omega_1 r \wedge \|s_i - \mathbf{p}\| > \omega_i r. \quad (10)$$

The core difficulties of bounding the probability for this event are the existential quantifiers that quantify over the continuous variables r and \mathbf{p} . In both cases, we resolve this by using an appropriate discretization, for which we then apply the union bound.

We get rid of the existential quantifier for r by dividing the interval $[R_A, \infty)$, which covers the domain of r , into pieces of length at most R_A . More formally, we split the event $\exists r \geq R_A$ with the desired property into the disjoint events $\exists r \in [jR_A, (j+1)R_A)$ for $j \in \mathbb{N}^+$. For a fixed j , $r \geq jR_A$ and $\|s_i - \mathbf{p}\| > \omega_i r$ implies $\|s_i - \mathbf{p}\| > \omega_i jR_A$. Moreover, $r \leq (j+1)R_A$ and $\|s_1 - \mathbf{p}\| \leq \omega_1 r$ implies $\|s_1 - \mathbf{p}\| \leq \omega_1 (j+1)R_A$. Note that this completely eliminates the variable r from the event, which lets us drop the existential quantifier for r . Thus, the event in Equation (10) implies

$$\exists j \in \mathbb{N}^+ \exists \mathbf{p} \in \mathbb{T}^d \forall i > k : \|s_1 - \mathbf{p}\| \leq \omega_1 (j+1)R_A \wedge \|s_i - \mathbf{p}\| > \omega_i jR_A. \quad (11)$$

Note that the new existential quantifier for j is not an issue: as j is discrete, we can simply use the union bound and sum over the probabilities we obtain for the different values of j . We will later see that this sum is dominated by the first term corresponding to $j = 1$.

To deal with the existential quantifier for \mathbf{p} , assume $j \in \mathbb{N}^+$ to be a fixed number. First note that $\|s_1 - \mathbf{p}\| \leq \omega_1(j+1)R_A$ implies that \mathbf{p} lies somewhat close to s_1 . We discretize the space around s_1 using a grid such that the point \mathbf{p} is guaranteed to lie inside a grid cell. By choosing the distance between neighboring grid vertices sufficiently small, we guarantee that \mathbf{p} lies close to a grid vertex. Then, instead of considering \mathbf{p} itself, we deal with its closest grid vertex. To define the grid formally, let $\omega_{\min} = \min_{i=k+1}^n \{\omega_i\}$ be the minimum weight of sites not in A and let $x = \omega_{\min} j R_A / \sqrt[d]{d}$ (x will be the width of our grid cells). To simplify notation, assume that s_1 is the origin. Otherwise, we can simply translate the grid defined in the following to be centered at s_1 to obtain the same result. Let $\Gamma = \{\ell x \mid \ell \in \mathbb{Z} \wedge |(\ell-1)x| \leq \omega_1(j+1)R_A\}$ be the set of all multiples of x that are not too large. We use the grid defined by the Cartesian product Γ^d . Then the following three properties of Γ^d are easy to verify.

- (i) A point \mathbf{p} with $\|s_1 - \mathbf{p}\| \leq \omega_1(j+1)R_A$ lies in a grid cell.
- (ii) The maximum distance between a point in a grid cell and its closest grid vertex is $\sqrt[d]{d}x/2 = \omega_{\min} j R_A / 2$.
- (iii) Γ^d has at most $c'_1(\omega_1 j R_A / x)^d = c_1(\omega_1 / \omega_{\min})^d \leq c_1 \omega_1^d$ vertices for constants c_1 and c'_1 only depending on d and \mathbf{p} .

Going back to the event in Equation (11), let \mathbf{p} be a point with $\|s_1 - \mathbf{p}\| \leq \omega_1(j+1)R_A$ and $\|s_i - \mathbf{p}\| > \omega_i j R_A$ (for all $i > k$). By the first inequality and Property (i), \mathbf{p} lies in a grid cell of Γ^d . Let $\mathbf{p}' \in \Gamma^d$ be the grid vertex with minimum distance to \mathbf{p} . Then, by Property (ii), $\|\mathbf{p} - \mathbf{p}'\| \leq \omega_{\min} j R_A / 2$. Thus, using the triangle inequality and $\|s_i - \mathbf{p}\| > \omega_i j R_A$, we obtain

$$\|s_i - \mathbf{p}'\| \geq \|s_i - \mathbf{p}\| - \|\mathbf{p} - \mathbf{p}'\| > \omega_i j R_A - \frac{\omega_{\min} j R_A}{2} \geq \frac{\omega_i j R_A}{2}.$$

It follows that the event in Equation (11) implies

$$\exists j \in \mathbb{N}^+ \exists \mathbf{p}' \in \Gamma^d \forall i > k : \|s_i - \mathbf{p}'\| > \frac{\omega_i j R_A}{2}.$$

For this event, we can now bound the probability. First note that $\|s_i - \mathbf{p}'\| > \omega_i j R_A / 2$ implies that the ball of radius $\omega_i j R_A / 2$ around \mathbf{p}' does not contain s_i . By Lemma A.3, the volume of this ball intersected with $[-0.5, 0.5]^d$ is $\min\{1, c_2(\omega_i j R_A)^d\}$ for a constant c_2 depending only on d and \mathbf{p} . As the s_i are chosen independently and using that $1 - x \leq \exp(-x)$ for $0 \leq x \leq 1$, we obtain

$$\begin{aligned} \Pr \left[\forall i > k : \|s_i - \mathbf{p}'\| > \frac{\omega_i j R_A}{2} \right] &= \prod_{i=k+1}^n \max \{0, 1 - c_2(\omega_i j R_A)^d\} \\ &\leq \prod_{i=k+1}^n \exp(-c_2(\omega_i j R_A)^d) \\ &= \exp \left(-c_2 j^d R_A^d \sum_{i=k+1}^n \omega_i^d \right). \end{aligned}$$

We resolve the two existential quantifiers for j and \mathbf{p}' using the union bound. Recall from Property (iii) that the grid Γ^d contains only $c_1 \omega_1^d$ vertices. Using that $\omega_i = w_i^{1/d}$, we obtain

$$\begin{aligned} \Pr[A \text{ is relevant} | R_A] &\leq \Pr \left[\exists j \in \mathbb{N}^+ \exists \mathbf{p}' \in \Gamma^d \forall i > k : \|s_i - \mathbf{p}'\| > \frac{\omega_j R_A}{2} \right] \\ &\leq \sum_{j=1}^{\infty} c_1 w_1 \exp \left(-c_2 j^d R_A^d \sum_{i=k+1}^n w_i \right). \end{aligned}$$

To conclude the proof, it remains to show that the sum over j is dominated by the first term corresponding to $j = 1$. For this, note that

$$\begin{aligned} \sum_{j=1}^{\infty} \exp(-x j^d) &= \exp(-x) \cdot \sum_{j=1}^{\infty} \frac{\exp(-x j^d)}{\exp(-x)} \\ &= \exp(-x) \cdot \sum_{j=1}^{\infty} \exp(-x(j^d - 1)) \\ &\leq \exp(-x) \cdot \sum_{j=1}^{\infty} (\exp(-x))^{j-1}. \end{aligned}$$

As x is positive in our case, the sum is bounded by a constant due to the convergence of the geometric series. This concludes the proof. ■

Now that we know the probability that $A \subseteq S$ is relevant conditioned on R_A , we want to understand how R_A is distributed. The following lemma gives an upper bound on its density function.

Lemma 6.7. *There exists a constant c depending only on k , d , and \mathbf{p} , such that the density function $f_{R_A}(x)$ of the random variable R_A satisfies*

$$f_{R_A}(x) \leq c x^{dk-d-1} \frac{1}{\min_{s_i \in A} \{w_i\}} \prod_{s_i \in A} w_i.$$

Proof. The density function $f_{R_A}(x)$ is the derivative of the distribution function $F_{R_A}(x) = \Pr[R_A \leq x]$. Thus, we have to upper bound the slope of $\Pr[R_A \leq x]$. As before, we assume that $A = \{s_1, \dots, s_k\}$ and that s_1 has minimum weight among sites in A , that is, $\min_{s_i \in A} \{w_i\} = w_1$. Recall the definition of R_A in Equation (9). It follows directly that $R_A \leq x$ if and only if $\|s_1 - s_i\|/(\omega_1 + \omega_i) \leq x$ for all $i \in [k]$. Note that this clearly holds for $i = 1$. For greater i , this is the case if and only if s_i lies in the ball $B_{s_1}((\omega_1 + \omega_i)x)$ of radius $(\omega_1 + \omega_i)x$ around s_1 . To simplify notation, we denote this ball with $B(x_i)$ in the following. Note that the volume $\text{vol}(B(x_i))$ is exactly the probability for s_i to lie sufficiently close to s_1 . As the positions of the different sites s_i are independent, we obtain

$$F_{R_A}(x) = \Pr[R_A \leq x] = \prod_{i=2}^k \text{vol}(B(x_i)).$$

To upper bound the derivative of this, we have to upper bound the growth of $\text{vol}(B(x_i))$ depending on x_i . For sufficiently small x_i , this volume is given by the volume of a ball

in \mathbb{R}^d . For larger x_i , due to the fact that our ground space⁷ is bounded, the growth of this volume slows down. Thus, to get an upper bound on the derivative, we can simply use the volume of a ball in \mathbb{R}^d . Thus, for appropriate constants c_1 and c_2 only depending on d and \mathfrak{p} , we obtain

$$\frac{d}{dx} \text{vol}(B(x_i)) \leq \frac{d}{dx} c_1 ((\omega_1 + \omega_i)x)^d \leq \frac{d}{dx} c_2 (\omega_i x)^d = \frac{d}{dx} c_2 w_i x^d.$$

With this, it follows that

$$f_{R_A}(x) = \frac{d}{dx} F_{R_A}(x) \leq \frac{d}{dx} \prod_{i=2}^k (c_2 w_i x^d),$$

which immediately yields the claimed bound. \blacksquare

By Lemma 6.6, we know the probability for a set A to be relevant conditioned on R_A and by Lemma 6.7 we know how R_A is distributed. Based on this, we can bound the unconditional probability that A is relevant.

Lemma 6.8. *Let $A \subseteq S$. For a constant c only depending on k , d , and \mathfrak{p} , the probability that A is relevant satisfies*

$$\Pr[A \text{ is relevant}] \leq c \frac{\prod_{s_i \in A} w_i}{\left(\sum_{s_i \notin A} w_i \right)^{k-1}}.$$

Proof. Let $A \subseteq S$ and let R_A be the random variable as defined before; see Equation (9). Note that $0 \leq R_A \leq \sqrt[k]{d}$. By the law of total probability, we have

$$\Pr[A \text{ is relevant}] = \int_0^{\sqrt[k]{d}} \Pr[A \text{ is relevant} | R_A = x] \cdot f_{R_A}(x) dx.$$

Using Lemmas 6.6 and 6.7, we obtain

$$\Pr[A \text{ is relevant} | R_A = x] \cdot f_{R_A}(x) \leq c_1 \prod_{s_i \in A} w_i \exp\left(-c_2 x^d \sum_{s_i \notin A} w_i\right) x^{dk-d-1},$$

for constants c_1 and c_2 only depending on k , d , and \mathfrak{p} . Ignoring the factors independent of x for now, this expression has the form

$$x^{\alpha d-1} \exp(-\beta x^d) \quad \text{with } \alpha = k-1 \text{ and } \beta = c_2 \sum_{s_i \notin A} w_i,$$

which lets us apply Lemma A.4 to bound the integral. We obtain

⁷Again, this is true for the torus as well as for the hypercube.

$$\begin{aligned}\Pr[A \text{ is relevant}] &= c_1 \prod_{s_i \in A} w_i \cdot \int_0^{\sqrt[d]{d}} x^{\alpha d - 1} \exp(-\beta x^d) dx \\ &\leq c_1 \prod_{s_i \in A} w_i \cdot \frac{\Gamma(\alpha)}{\beta^\alpha d}.\end{aligned}$$

As k is an integer, $\Gamma(\alpha) = \Gamma(k - 1) = (k - 2)!$, which is constant. Thus, substituting α and β by its corresponding values and aggregating all constant factors into c yields

$$\Pr[A \text{ is relevant}] \leq c \frac{\prod_{s_i \in A} w_i}{\left(\sum_{s_i \notin A} w_i\right)^{k-1}},$$

which is exactly the bound we wanted to prove. \blacksquare

Having bound the probability that a specific subset of sites $A \subseteq S$ of size k is relevant, we can now bound the expected total number of relevant subsets. By Lemma 6.5, this also bounds the number of non-empty Voronoi regions.

Theorem 6.9. *Let S be a set of n sites with minimum weight 1, total weight W , and random positions on the d -dimensional torus equipped with a p -norm, for constant d . For every fixed k , the expected number of regions of the weighted order- k Voronoi diagram of S is in $O(W)$. The same holds for random sites in a hypercube.*

Proof. For every subset $A \subseteq S$ with $|A| = k$, let X_A be the indicator random variable that has value 1 if and only if A has non-empty order- k Voronoi region. Moreover, let X be the sum of these random variables. Note that $\mathbb{E}[X]$ is exactly the quantity, we are interested in. Using linearity of expectation, we obtain

$$\mathbb{E}[\text{number of regions}] = \mathbb{E}[X] = \sum_{\substack{A \subseteq S \\ |A|=k}} \mathbb{E}[X_A].$$

Due to Lemma 6.5, a subset A with non-empty Voronoi region is also relevant. Thus, $\mathbb{E}[X_A] \leq \Pr[A \text{ is relevant}]$ and Lemma 6.8 yields

$$\sum_{\substack{A \subseteq S \\ |A|=k}} \mathbb{E}[X_A] \leq \sum_{\substack{A \subseteq S \\ |A|=k}} c \frac{\prod_{s_i \in A} w_i}{\left(\sum_{s_i \notin A} w_i\right)^{k-1}}. \quad (12)$$

For technical reasons, we assume c to be the maximum of 1 and the constant from Lemma 6.8. We continue by proving the following claim:

$$\sum_{\substack{A \subseteq S \\ |A|=k}} \mathbb{E}[X_A] \leq 4^{k^2} c W. \quad (13)$$

In addition to implying the theorem, this claim specifies a constant that comes on top of c , which is crucial for the rest of the proof.

We first prove the claim for the situation, in which W is not dominated by the highest k weights. Afterwards, we deal with the other somewhat special case. More formally, let the weights w_1, \dots, w_n be sorted increasingly and consider the case that $\sum_{i=1}^{n-k} w_i \geq 4^{-k} W$, that is, if we leave out the k largest weights, we still have a significant portion of the total weight. We can use this to estimate the denominator in Equation (12):

$$\begin{aligned} \sum_{\substack{A \subseteq S \\ |A|=k}} c \frac{\prod_{s_i \in A} w_i}{\left(\sum_{s_i \notin A} w_i \right)^{k-1}} &\leq \sum_{\substack{A \subseteq S \\ |A|=k}} c \frac{\prod_{s_i \in A} w_i}{(4^{-k} W)^{k-1}} \\ &= 4^{k(k-1)} c \cdot \frac{\sum_{A \subseteq S, |A|=k} \prod_{s_i \in A} w_i}{W^{k-1}}. \end{aligned}$$

To bound the fraction by W , observe that the binomial theorem yields

$$W^k = \left(\sum_{i=1}^n w_i \right)^k \geq \sum_{\substack{A \subseteq S \\ |A|=k}} \prod_{s_i \in A} w_i,$$

as each summand on the on the right-hand side also appears on the left-hand side. This proves the claim in Equation (13) for the case $\sum_{i=1}^{n-k} w_i \geq 4^{-k} W$.

For $\sum_{i=1}^{n-k} w_i < 4^{-k} W$, assume for contradiction that the claim in Equation (13) does not hold for every set of n weights. Then there exists a minimum counterexample, that is, a smallest number of n weights such that the expected number of non-empty regions exceeds $4^{k^2} c W$. We show that, based on this assumption, we can construct an even smaller counterexample; a contradiction. First note that $n > 2k$ for every counterexample, as there are fewer than $4^{k^2} c W$ subsets otherwise (recall that $c \geq 1$).

Now let w_1, \dots, w_n be the minimum counterexample and again assume that the weights are ordered increasingly. Moreover, fix the coordinates of the sites s_1, \dots, s_n and consider two order- k Voronoi diagrams: one on the set of all sites $S = \{s_1, \dots, s_n\}$, and the one on all but the k heaviest sites $S' = \{s_1, \dots, s_{n-k}\}$ (note that this is well defined as $n > 2k$). In the following, we call the former Voronoi diagram \mathcal{V} and the latter \mathcal{V}' . We define a mapping from the non-empty regions of \mathcal{V} to non-empty regions of \mathcal{V}' . Let $A \subseteq \{s_1, \dots, s_n\}$ be a subset of size k with non-empty region in \mathcal{V} and let p be an arbitrary point in this region. Moreover, let A' be the set of sites corresponding to the region of \mathcal{V}' containing p . Then we map the region of A to the region of A' . Note that A and A' share all sites that have not been deleted: $A \cap A' = A \cap S'$. Thus, any site A that is mapped to A' must satisfy $A \subseteq A' \cup (S \setminus S')$. This limits the number of different regions in \mathcal{V} that are mapped to the same region of \mathcal{V}' to at most 4^k . Thus, the number of regions in \mathcal{V}' is at least 4^{-k} times the number of regions in \mathcal{V} . As this holds for arbitrary coordinates, this also holds for the expected number of non-empty regions when choosing random coordinates.

As we assumed w_1, \dots, w_n to be a counterexample for Equation (13), the expected number of regions with these weights is more than $4^{k^2} c W$. Thus, by the above argument, the expected number of regions for the weights w_1, \dots, w_{n-k} is at least $4^{-k} \cdot 4^{k^2} c W$. As we consider the case $\sum_{i=1}^{n-k} w_i < 4^{-k} W$, we can substitute W to obtain that the weights w_1, \dots, w_{n-k} lead to at least $4^{k^2} c \sum_{i=1}^{n-k} w_i$ non-empty regions in expectation. Thus, the weights w_1, \dots, w_{n-k} also form a counterexample for the claim in Equation (13), which is a contradiction to the assumption that w_1, \dots, w_n is the minimum counterexample and thus to the assumption that there is a counterexample at all. ■

7 | GEOMETRIC SAT WITH NON-ZERO TEMPERATURE

In the case with temperature $T = 0$, we used the fact that every clause contains the k variables with smallest weighted distance; recall Section 4.2. This is no longer true for higher temperatures: for $T > 0$, a clause can, in principle, contain any variable. However, the probability to contain a variable that is far away is rather small. In the remainder of this section, we show that a constant fraction of clauses actually behave just like in the $T = 0$ case, that is, they contain the k closest variables. With this, we can then apply the argument outlined in Section 4.2.

7.1 | Expected number of nice clauses

Recall that a clause c is generated by drawing k variables without repetition with probabilities proportional to the connection weights. We call c *nice* if it contains the k variables with the highest connection weight and, additionally, the i th variable drawn for c has the i th highest connection weight with c . That is, c does not only contain the k variables with highest connection weight but, also, they are drawn in descending order. This is a slightly stronger property than just requiring c to contain the k variables with lowest weighted distance.

Let \bar{x} be the connection weight of a variable v that has rather high connection weight with c . To show that the probability for $v \in c$ is reasonably high, we prove that \bar{x} is large compared to the sum of connection weights over all variables with smaller weight. The following lemma bounds this sum for a given \bar{x} . We use the Iverson bracket to exclude the variables with weight larger than \bar{x} from the sum, that is, $[X(c, v) \leq \bar{x}]$ evaluates to 1 if $X(c, v) \leq \bar{x}$ and to 0 otherwise.

Lemma 7.1. *Let c be a clause at any position and let V be a set of n weighted variables with random positions in \mathbb{T}^d . For $T < 1$ and $\bar{x} \in \Omega(W^{1/T})$, the expected sum of connection weights smaller than \bar{x} is in $O(\bar{x})$, that is,*

$$\mathbb{E} \left[\sum_{v \in V} X(c, v) \cdot [X(c, v) < \bar{x}] \right] \in O(\bar{x}).$$

Proof. Using linearity of expectation, the term in the lemma's statement equals to the sum over the expectations $\mathbb{E}[X(c, v) \cdot [X(c, v) \leq \bar{x}]]$. To bound this expectation, we consider the three events $X(c, v) \leq (2^d w_v)^{1/T}$, $(2^d w_v)^{1/T} < X(c, v) < \bar{x}$, and $\bar{x} \leq X(c, v)$. Note that $[X(c, v) < \bar{x}]$ is 0 in the last event and 1 in the former two. Thus, we obtain

$$\begin{aligned} \mathbb{E}[X(c, v) \cdot [X(c, v) < \bar{x}]] \\ = \Pr[X(c, v) \leq (2^d w_v)^{1/T}] \cdot \mathbb{E}[X(c, v) | X(c, v) \leq (2^d w_v)^{1/T}] \end{aligned} \quad (14)$$

$$+ \Pr[(2^d w_v)^{1/T} < X(c, v) < \bar{x}] \cdot \mathbb{E}[X(c, v) | (2^d w_v)^{1/T} < X(c, v) < \bar{x}]. \quad (15)$$

We bound the first term from above by assuming $X(c, v) = (2^d w_v)^{1/T}$ whenever $X(c, v) \leq (2^d w_v)^{1/T}$. Moreover, using the CDF for $X(c, v)$ (2) yields

$$\begin{aligned} (14) &\leq \Pr[X(c, v) \leq (2^d w_v)^{1/T}] \cdot (2^d w_v)^{1/T} \\ &= (1 - \Pi_{d, \mathbf{p}} w_v (2^d w_v)^{-1}) \cdot (2^d w_v)^{1/T} \\ &= (1 - \Pi_{d, \mathbf{p}} 2^{-d}) \cdot (2^d w_v)^{1/T} \in \Theta(w_v^{1/T}). \end{aligned}$$

For the second term, we have to integrate over the probability density function (PDF) $f_X(x)$ of the connection weights $X(c, v)$, which is the derivative of $F_X(x)$ (2). Thus, $f_X(x) = T\Pi_{d,p}w_v x^{-T-1}$ for $x \geq (2^d w_v)^{1/T}$, and we obtain

$$\begin{aligned} (15) &= \Pr \left[(2^d w_v)^{1/T} < X(c, v) < \bar{x} \right] \cdot \int_{(2^d w_v)^{1/T}}^{\bar{x}} \frac{x \cdot f_X(x)}{\Pr \left[(2^d w_v)^{1/T} < X(c, v) < \bar{x} \right]} dx \\ &= T\Pi_{d,p}w_v \cdot \int_{(2^d w_v)^{1/T}}^{\bar{x}} x^{-T} dx. \end{aligned}$$

For $T < 1$, this evaluates to

$$\begin{aligned} &= T\Pi_{d,p}w_v \cdot \left[\frac{x^{1-T}}{1-T} \right]_{(2^d w_v)^{1/T}}^{\bar{x}} \\ &= \frac{T\Pi_{d,p}w_v}{1-T} \cdot [\bar{x}^{1-T} - (2^d w_v)^{1/T-1}] \\ &\leq \frac{T\Pi_{d,p}w_v}{1-T} \cdot \bar{x}^{1-T} \\ &\in \Theta(w_v \bar{x}^{1-T}). \end{aligned}$$

Putting these bounds together yields

$$\begin{aligned} \mathbb{E} \left[\sum_v X(c, v) \cdot [X(c, v) \leq \bar{x}] \right] &= \sum_v ((14) + (15)) \\ &\in O \left(\sum_v w_v^{1/T} + \sum_v w_v \bar{x}^{1-T} \right) \\ &\subseteq O \left(\left(\sum_v w_v \right)^{1/T} + \bar{x}^{1-T} \cdot \sum_v w_v \right) \\ &\in O(W^{1/T} + \bar{x}^{1-T}W). \end{aligned}$$

As, $\bar{x} \in \Omega(W^{1/T})$, we have $W^{1/T} \in O(\bar{x})$, which handles the first term. The second term is also in $O(\bar{x})$, as $\bar{x} \in \Omega(W^{1/T})$ implies $W \in O(\bar{x}^T)$. Thus, this yields the claimed bound of $O(\bar{x})$. ■

This lets us show that each clause is nice with constant probability. The only assumption we need for this is the fact that no single weight is too large, that is, every weight w_i has to be asymptotically smaller than the total weight W .

Theorem 7.2. *Let Φ be a random formula drawn from the weighted geometric model with ground space \mathbb{T}^d equipped with a p -norm, with temperature $T < 1$, and with $w_v/W \in o(1)$ for $v \in V$. Let c be a clause of Φ . Then c is nice with probability $\Omega(1)$.*

Proof. We prove two things. First, we show that, with probability $\Omega(1)$, there are at least k variables sufficiently close to c that they have connection weight $\Omega(W^{1/T})$. Second, we use Lemma 7.1 to show that the k variables with highest connection weight are chosen for c with constant probability (in descending order).

For the first part, we show that there is a constant a such that, with constant probability, at least k variables have connection weight at least $aW^{1/T}$. For a fixed variable v , we can use the CDF of $X(c, v)$ (Equation 2) to obtain

$$\begin{aligned} \Pr[X(c, v) \geq aW^{1/T}] &= \Pi_{d, \mathbf{p}} w_v (aW^{1/T})^{-T} \\ &= \frac{\Pi_{d, \mathbf{p}} w_v}{a^T W} \\ &= 2k \frac{w_v}{W}, \text{ for } a = \left(\frac{\Pi_{d, \mathbf{p}}}{2k} \right)^{1/T}. \end{aligned}$$

Note that this is a valid probability, as $w_v/W \in o(1)$ implies that it is below 1. For the above choice of a , we obtain that the expected number of variables with connection weight at least $aW^{1/T}$ is $2k$. As the connection weights for the different variables are independent, we can apply the Chernoff–Hoeffding bound in Theorem A.7 to obtain that at least k variables have connection weight $aW^{1/T}$ with constant probability.

For the second part of the proof, let \bar{x} be the connection weight of the k th closest variable. With the argument above, we can assume $\bar{x} \in \Omega(W^{1/T})$ with constant probability, which lets us apply Lemma 7.1. To do so, consider the experiment of drawing the first variable for our clause c . Let v be the variable that maximizes the connection weight $X(c, v)$. The probability of drawing v equals $X(c, v)$ divided by the sum of all connection weights. By Lemma 7.1, the sum of all connection weights smaller than \bar{x} is in $O(\bar{x})$. Thus, the sum of all connection weights is in $O(X(c, v))$, which implies that v is chosen with constant probability. As we draw variables without repetition, the exact same argument applies for the second closest variable and so on. Thus, the probability that c contains the k closest variables drawn in order of descending connection weights is at least a constant, if there are k sufficiently close variables. As the latter holds with constant probability, c is nice with constant probability. ■

By the linearity of expectation, this immediately yields the following bound on the expected number of nice clauses.

Corollary 7.3. *Let Φ be a random formula with m clauses drawn from the weighted geometric model with ground space \mathbb{T}^d equipped with a \mathbf{p} -norm, with temperature $T < 1$, and with $w_v/W \in o(1)$ for $v \in V$. The expected number of nice clauses in Φ is $\Theta(m)$.*

7.2 | Concentration of nice clauses

We show that the number of nice clauses is concentrated around its expectation, that is, with high probability, a constant fraction of clauses is nice. Our main tool for this will be the method of typical bounded differences [61]; see Section A.6. To this end, we consider several random variables, for example, the coordinates of clauses and variables, that together determine the whole process of generating a random formula. The number of nice clauses is then a function f of these random variables and its expectation is $\Theta(m)$, due to Corollary 7.3. Roughly speaking, the method of bounded differences then states that the probability that f deviates too much from its expectation is low if changing a single random variable only slightly changes f .

7.2.1 | The random variables

So far, we viewed the generation of a random formula as a two-step process: first, sample coordinates for the variables and clauses; second, sample the variables contained in each clause based on their distances. The first step can be easily expressed via random variables. Let V_1, \dots, V_n and C_1, \dots, C_m be the coordinates⁸ of the n variables and m clauses, respectively. Though the second step heavily depends on the distances determined by the first, we can determine all random choices in advance. For all $i \in [m]$ and $j \in [k]$, let X_i^j be a random variable uniformly distributed in $[0, 1)$. The variable X_i^j determines the j th variable of the i th clause c_i in the following way. We partition the interval $[0, 1)$ such that each variable v not already chosen for c_i corresponds to a subinterval of length proportional to the connection weight $X(c_i, v)$. We order these subintervals by length such that the largest interval comes first. The j th variable of c_i is then the variable whose interval contains X_i^j . Note that this samples k different variables for each clause, with probabilities proportional to the connection weights $X(c_i, v)$. Note further that the whole generation process of a random formula is determined by evaluating the independent random variables $V_1, \dots, V_n, C_1, \dots, C_m, X_1^1, \dots, X_m^k$.

To formalize the concept of nice clauses in this context, we require some more notation. For $i \in [m]$, let \mathcal{V}_i be the sequence of all variables ordered decreasingly by connection weight with the clause c_i . Moreover, let $\mathcal{V}_i[a, b]$ denote the subsequence from the a th to the b th variable in this sequence, including the boundaries. To simplify notation, we abbreviate the unique element in $\mathcal{V}_i[a, a]$ with $\mathcal{V}_i[a]$. Recall that clause c_i is nice if, for each of k steps, we choose the variable with highest connection weight that has not been chosen before. With respect to the random variables, this happens if, for each $j \in [k]$, X_i^j is smaller than the connection weight of $\mathcal{V}_i[j]$ divided by the sum of all connection weights of the remaining variables $\mathcal{V}_i[j, n]$. We thus define the indicator variable

$$N_i = \begin{cases} 1, & \text{if } \forall j \in [k] : X_i^j < \frac{X(c_i, \mathcal{V}_i[j])}{\sum_{v \in \mathcal{V}_i[j, n]} X(c_i, v)}, \\ 0 & \text{otherwise,} \end{cases} \quad (16)$$

which is 1 if and only if the i th clause is nice. With this, we can define the number of nice clauses as $f(V_1, \dots, V_n, C_1, \dots, C_m, X_1^1, \dots, X_m^k) = \sum_{i \in [m]} N_i$.

7.2.2 | Bounding the effect on the number of nice clauses

To apply the method of bounded differences (Theorem A.9 or the more specific Corollary A.10), we have to bound the effect of changing the value of only one of these random variables on f . For the variables C_1, \dots, C_m , this is easy: Changing C_i moves the position of the clause c_i , which only makes a difference for c_i . Thus, the number of nice clauses changes by at most 1. Similarly, changing X_i^j only impacts the clause c_i , which implies that it changes the number of nice clauses by at most 1.

For the variables V_1, \dots, V_n , one can actually construct situations in which changing only a single position drops f from m to 0. There are basically two situations in which this can happen. First, if a single variable is close to many clauses, changing its position potentially impacts many clauses. Second, if many inequalities in Equation (16) are rather tight, then moving a single variable slightly closer to many clauses can increase the denominator on the right hand side by enough to change N_i for many clauses. We exclude both situations by defining unlikely bad events. By assuming these bad

⁸Technically, these are multivariate random variables, as they represent d -dimensional points in \mathbb{T}^d .

events do not happen, we can bound the effect of moving a single variable v by

$$\delta_v = w_v^{\frac{1}{1+T}} n^{\frac{T}{1+T}} \log^{\frac{2}{1+T}} n. \quad (17)$$

The following bound gives a simpler estimate for δ_v that will be useful later.

Lemma 7.4. *Let $0 < T < 1$ and $w_v \in O(n^{1-\varepsilon})$ for an arbitrary $\varepsilon > 0$. Then $\delta_v \in O\left(\frac{\sqrt{w_v n}}{\log n}\right)$.*

Proof. We ignore logarithmic factors and show that $\delta_v/\sqrt{w_v n}$ converges polynomially to 0 for $n \rightarrow \infty$. As logarithmic factors grow slower than any polynomial, this proves the claim. We get

$$\frac{\delta_v}{\sqrt{w_v n}} = w_v^{\frac{1}{1+T} - \frac{1}{2}} n^{\frac{T}{1+T} - \frac{1}{2}}.$$

Rearranging the exponents yields

$$\frac{1}{1+T} - \frac{1}{2} = \frac{2 - (1+T)}{2(1+T)} = \frac{1-T}{2(1+T)}, \quad \text{and} \quad \frac{T}{1+T} - \frac{1}{2} = \frac{2T - (1+T)}{2(1+T)} = -\frac{1-T}{2(1+T)}.$$

Thus $\delta_v/\sqrt{w_v n} = (w_v/n)^c$ for a positive constant c . As $w_v \in O(n^{1-\varepsilon})$, this yields the claim. ■

The following lemma states that, with overwhelming probability, no point (and therefore no variable) is too close to too many clauses. This eliminates the first problematic situation (and will also help with the second). Note that this statement only assumes random clause positions and holds for arbitrary variable positions, that is, when moving a variable, we can assume that it holds before and after the movement.

Lemma 7.5. *Let $m \in O(n)$, $0 < T < 1$, and $w_v \in O(n^{1-\varepsilon})$ for every $v \in V$ and arbitrary constant $\varepsilon > 0$. Let $r = (w_v \log^2(n)/n)^{\frac{1}{d(1+T)}}$. With overwhelming probability, for every point \mathbf{p} , the ball $B_{\mathbf{p}}(r)$ around \mathbf{p} with radius r contains only $O(\delta_v)$ clauses.*

Proof. As there are uncountably many points \mathbf{p} , it is hard to argue about them directly. Thus, we first reduce the statement to one about finitely many positions, namely the positions of the clauses. Then it remains to show the statement for these positions.

Consider a fixed point \mathbf{p} . As $B_{\mathbf{p}}(r)$ has diameter $2r$, the pair-wise distance between clauses in $B_{\mathbf{p}}(r)$ is at most $2r$. Thus, if there exists a point \mathbf{p} such that $B_{\mathbf{p}}(r)$ contains too many clauses, then there exists a clause that has too many other clauses at distance at most $2r$. Thus, it suffices to show that for every clause $c \in C$, the number of clauses of distance at most $2r$ to c is in $O(\delta_v)$.

Let c_0 be a fixed clause (we later apply the union bound over all clauses). We want to bound the probability for another clause c to be closer than $2r$ to c_0 . For this, we use the CDF of the distance in Equation (1). Note that the restriction of Equation (1) to the interval $[0, 0.5]$ is not an issue here, as $w_v \in O(n^{1-\varepsilon})$ implies $r \in o(1)$ and thus $2r \leq 0.5$. Thus, we obtain

$$\begin{aligned}
\Pr[\|c_0 - c\| \leq 2r] &= \Pi_{d,p} 2^d r^d \\
&= \Pi_{d,p} 2^d \left(\frac{w_v \log^2 n}{n} \right)^{\frac{1}{1+T}} \\
&= \Pi_{d,p} 2^d \left(\frac{w_v}{n} \right)^{\frac{1}{1+T}} \log^{\frac{2}{1+T}} n.
\end{aligned}$$

As there are $m \in O(n)$ clauses, the expected number of clauses with distance at most $2r$ to c_0 is

$$m \Pi_{d,p} 2^d \left(\frac{w_v}{n} \right)^{\frac{1}{1+T}} \log^{\frac{2}{1+T}} n \in O \left(w_v^{\frac{1}{1+T}} n^{1-\frac{1}{1+T}} \log^{\frac{2}{1+T}} n \right) = O \left(w_v^{\frac{1}{1+T}} n^{\frac{T}{1+T}} \log^{\frac{2}{1+T}} n \right),$$

which is already the claimed bound of $O(\delta_v)$. As $0 < T < 1$, this upper bound grows polynomially in n . Thus, by the Chernoff–Hoeffding bound in Corollary A.8, it holds asymptotically with overwhelming probability. Applying the union bound over all $O(n)$ clauses yields the claim. ■

The above lemma is stated in terms of the distances. In the following it will be useful to think of it in terms of connection weights instead. The following lemma translates the radius r in Lemma 7.5 to the corresponding connection weight between a clause and a variable at distance r .

Lemma 7.6. *Let $v \in V$ be a variable and let $c \in C$ be a clause with distance $\|c - v\| = (w_v \log^2(n)/n)^{\frac{1}{d(1+T)}}$. They have connection weight $X(c, v) = w_v^{\frac{1}{1+T}} n^{\frac{1}{T(1+T)}} \log^{-\frac{2}{T(1+T)}} n$.*

Proof. Using the definition of the connection weight and inserting the above distance, we obtain

$$\begin{aligned}
X(c, v) &= \left(\frac{w_v}{\|c - v\|^d} \right)^{\frac{1}{T}} \\
&= \left(w_v \cdot \left(\frac{n}{w_v \log^2 n} \right)^{\frac{1}{1+T}} \right)^{\frac{1}{T}} \\
&= w_v^{\frac{1}{T} - \frac{1}{T(1+T)}} n^{\frac{1}{T(1+T)}} \log^{-\frac{2}{T(1+T)}} n \\
&= w_v^{\frac{1}{1+T}} n^{\frac{1}{T(1+T)}} \log^{-\frac{2}{T(1+T)}} n.
\end{aligned}$$

■

Combining Lemmas 7.5 and 7.6, we obtain that, for arbitrary variable positions (and random clause positions), no variable has a high connection weight to too many clauses, as summarized by the following corollary.

Corollary 7.7. *Let $m \in O(n)$, $0 < T < 1$, and $w_v \in O(n^{1-\varepsilon})$ for every $v \in V$ and arbitrary constant $\varepsilon > 0$. With overwhelming probability, for every variable v and every possible position of v , the number of clauses with connection weight at least $w_v^{\frac{1}{1+T}} n^{\frac{1}{T(1+T)}} \log^{-\frac{2}{T(1+T)}} n$ is in $O(\delta_v)$.*

For the second problematic situation mentioned above, consider for a clause c_i the k inequalities in Equation (16). We call c_i v -critical if for one of these inequalities the difference between the left and

right hand side is at most δ_v/n . In the following lemma, we first bound the number of critical clauses. Afterwards, we show that the concept of critical clauses works as intended in the sense that moving the variable v does only change the niceness status of v -critical clauses.

Lemma 7.8. *Let $m \in O(n)$, $0 < T < 1$ and let v be a variable. With overwhelming probability, there are only $O(\delta_v)$ v -critical clauses.*

Proof. A clause c_i can only be v -critical if one of the random variables X_i^j for $j \in [k]$ differs by at most δ_v/n to the right hand side of the inequality in Equation (16). The probability for this to happen for a single X_i^j is $2\delta_v/n$. As k is constant, c_i is v -critical with probability $O(\delta_v/n)$. Thus, as $m \in O(n)$, the expected number of v -critical clauses is in $O(\delta_v)$. As the event of being v -critical is independent for the different clauses, and as this bound is polynomial in n for $T > 0$ (see Equation (17)), the Chernoff–Hoeffding bound in Corollary A.8 yields the claim. ■

To prove that the movement of a single variable does not change the niceness status of too many clauses, we argue along the following lines. Let v be the variable we move and consider a clause c . If, before or after the movement, v is so close to c that we get a very high connection weight $X(c, v)$, we basically give up on c and assume that c changes its status (from being nice to not being nice or the other way round). By Corollary 7.7 this only happens for at most $O(\delta_v)$ clauses. Similarly, if c is v -critical, we also give up on c , which happens for at most $O(\delta_v)$ clauses by Lemma 7.8. Then it remains to show that in all other cases (i.e., when $X(c, v)$ is low before and after the movement and c is not v -critical), the status of c remains unchanged.

This is done as follows. As c is not v -critical, the difference between the right and left hand side of the inequality in Equation (16) is somewhat high. Thus, if moving v does not change the right hand side by too much, then c keeps its niceness status. To show this, we can use the fact that $X(c, v)$ is low before and after the movement and thus it cannot change by too much. This change of $X(c, v)$ has to be considered relative to the other connection weights, that is, changing $X(c, v)$ has less impact if there are other variables with higher connection weight. The following lemma establishes that these other variables with higher connection weight indeed exist.

Lemma 7.9. *Let $w_v \in O(n^{1-\varepsilon})$ for every $v \in V$ and arbitrary constant $\varepsilon > 0$. With overwhelming probability every clause has k variables with connection weight at least $W^{\frac{1}{T}} \log^{-\frac{2}{T}} n$.*

Proof. Let $x_0 = W^{\frac{1}{T}} \log^{-\frac{2}{T}} n$ be the above connection weight and let c be a clause with fixed position. For every variable v , the probability for $X(c, v) \geq x_0$ is $\Pi_{d,p} w_v x_0^{-T} = \Pi_{d,p} w_v W^{-1} \log^2 n$ by Equation (2). Note that we can apply Equation (2) as $x_0 \geq (2^d w_v)^{1/T}$ due to the condition $w_v \in O(n^{1-\varepsilon})$ and the fact that $W \geq n$. Summing this over all variables yields that the expected number of variables with connection weight at least x_0 is $\Pi_{d,p} \log^2 n$. By Corollary A.8, c has $\Omega(\log^2 n)$ variables with connection weight at least x_0 with overwhelming probability. Applying the union bound over all clauses and the fact that k is constant while $\log^2 n$ grows with n yields the claim. ■

Now, we are ready to bound the effect of moving just a single variable on the number of nice clauses.

Lemma 7.10. *Let $m \in O(n)$, $0 < T < 1$, and $w_v \in O(n^{1-\varepsilon})$ for every $v \in V$ and arbitrary constant $\varepsilon > 0$. With overwhelming probability, moving a single variable to an arbitrary position changes the number of nice clauses by only $O(\delta_v)$.*

Proof. We show the result for a fixed variable v . It then follows for all variables using the union bound.

Consider how the niceness status of clauses changes when moving v . Due to Corollary 7.7, there are only $O(\delta_v)$ clauses with connection weight at least $w_v^{\frac{1}{1+T}} n^{\frac{1}{T(1+T)}} \log^{-\frac{2}{T(1+T)}} n$ before or after the movement. Moreover, due to Lemma 7.8 there are only $O(\delta_v)$ v -critical clauses. Thus, even if all these clauses change the status from being nice to not being nice or vice versa, the number of nice clauses changes by only $O(\delta_v)$.

Every remaining clause c is not v -critical and we have $X(c, v) \leq w_v^{\frac{1}{1+T}} n^{\frac{1}{T(1+T)}} \log^{-\frac{2}{T(1+T)}} n$ before and after the movement. In the following, we show that a clause c with these two properties is nice after the movement if and only if it is nice before the movement.

We first observe that v does not belong to the k variables closest to c due to Lemma 7.9: with overwhelming probability, there are k variables with connection weight at least $W^{\frac{1}{T}} \log^{-\frac{2}{T}} n$, which is asymptotically bigger than $X(c, v)$ as $w_v \in O(n^{1-\varepsilon})$.

Thus, in the right hand side of the inequality in Equation (16), the connection weight $X(c, v)$ only appears in the denominator. To show that the right hand side does not change by too much, let x be the numerator, let y be the denominator before the movement, and let y' be the denominator after the movement. Note that $|y' - y|$ is exactly the change in $X(c, v)$ caused by the movement of v . With this, the right hand side of the inequality in Equation (16) changes by

$$\left| \frac{x}{y} - \frac{x}{y'} \right| = \left| \frac{xy' - xy}{yy'} \right| = \frac{x}{y'} \cdot \frac{|y' - y|}{y}.$$

Note that x (the numerator) is the connection weight of one variable whose connection weight also appears in the sum of the denominator (after and before the movement). Thus, $\frac{x}{y'} \leq 1$ and the above change is upper bounded by $\frac{|y' - y|}{y}$. Note that the upper bound on $X(c, v)$ holds before and after the movement and thus $X(c, v)$ can only change by less than this upper bound, that is, $|y' - y| < w_v^{\frac{1}{1+T}} n^{\frac{1}{T(1+T)}} \log^{-\frac{2}{T(1+T)}} n$. Moreover, y is the sum of multiple connection weights including the weight of one of the k closest variables. Thus, by Lemma 7.9 and the fact that $W \geq n$ we can assume that $y \geq n^{\frac{1}{T}} \log^{-\frac{2}{T}} n$. Putting this together yields

$$\left| \frac{x}{y} - \frac{x}{y'} \right| \leq \frac{|y' - y|}{y} < \frac{w_v^{\frac{1}{1+T}} n^{\frac{1}{T(1+T)}} \log^{-\frac{2}{T(1+T)}} n}{n^{\frac{1}{T}} \log^{-\frac{2}{T}} n} = \left(\frac{w_v}{n} \right)^{\frac{1}{1+T}} \log^{\frac{2}{1+T}} n = \frac{\delta_v}{n}.$$

As c is not v -critical, the difference between the left and right side of the inequality in Equation (16) is at least $\frac{\delta_v}{n}$ before the movement. Thus, as the movement can change the right hand side by only less than $\frac{\delta_v}{n}$, the clause c is nice after the movement if and only if it was nice before. ■

With this we are ready to prove concentration using the method of typical bounded differences.

Theorem 7.11. *Let Φ be a random formula with n variables and $m \in \Theta(n)$ clauses drawn from the weighted geometric model with ground space \mathbb{T}^d equipped with a \mathfrak{p} -norm, with temperature $0 < T < 1$, with $W \in O(n)$, and with $w_v \in O(n^{1-\varepsilon})$ for every $v \in V$ and arbitrary constant $\varepsilon > 0$. With high probability, $\Theta(m)$ clauses are nice.*

Proof. We want to apply Corollary A.10. As defined in Section 7.2.1, the random variables are the variable positions V_1, \dots, V_n , the clause positions C_1, \dots, C_m , and the coin flips X_1^1, \dots, X_m^k , and the function f is the number of nice clauses. For $N = n + m + km$ note that $|f(X)| \leq m \leq N$. For the nice event Γ we assume that the statement from Lemma 7.10 holds. Due to Lemma 7.10, the probability for this is $\Pr[\Gamma] \geq 1 - N^{-c}$ for any constant c and sufficiently large N . Thus, when choosing $c \geq 3$, we satisfy the condition $|f(X)| \leq N^{c-2}$ of Corollary A.10.

Now we have to bound the change of f when changing only one of the random variables, assuming we start with an event in Γ , that is, we have to determine the Δ_i from Corollary A.10. As mentioned before, changing a clause position C_i or one of the X_i^j impacts only one clause and thus changes f by at most 1. Moreover, as we start with a configuration satisfying Lemma 7.10, f changes by only $O(\delta_v)$ for variable v . Thus, for the sum in Corollary A.10 we obtain

$$\sum_{i \in [N]} \Delta_i^2 \in O\left(m + km + \sum_{v \in V} \delta_v^2\right).$$

Due to Lemma 7.4, we have $\delta_v \in O(\sqrt{w_v n} / \log n)$. Thus, the above sum can be bounded by

$$\sum_{v \in V} \delta_v^2 \in O\left(\sum_{v \in V} \left(\frac{\sqrt{w_v n}}{\log n}\right)^2\right) = O\left(\frac{n}{\log^2 n} \sum_{v \in V} w_v\right) = O\left(\frac{n^2}{\log^2 n}\right).$$

As $\mathbb{E}[f] \in \Theta(m) = \Theta(n)$, this is exactly the bound required by Corollary A.10 and thus the number of nice clauses is in $\Theta(m)$ with high probability. ■

7.2.3 | Putting things together

Now, we are ready to prove our main theorem for the geometric model.

Theorem 7.12. *Let Φ be a formula with n variables and $m \in \Theta(n)$ clauses drawn from the weighted geometric model with ground space \mathbb{T}^d equipped with a \mathfrak{p} -norm, temperature $T < 1$, $W \in O(n)$, and $w_v \in O(n^{1-\varepsilon})$ for every $v \in V$ and any constant $\varepsilon > 0$. Then, Φ contains a. a. s. an unsatisfiable subformula of constant size, which can be found in $O(n \log n)$ time.*

Proof. Let m' be the number of clauses in Φ that consist of the k variables with minimum weighted distance. By Theorem 7.11 we have $m' \in \Theta(m) = \Theta(n)$. In the following, we consider only these clauses.

Consider the weighted order- k Voronoi diagram of the n variables and let n' be the number of non-empty regions. By Theorem 6.9 and due to $W \in O(n)$, we have

$\mathbb{E}[n'] \in O(n)$. Moreover, it follows from Markov's inequality that $n' \leq n \log n$ holds asymptotically almost surely:

$$\Pr[n' \geq n \log n] \leq \frac{\mathbb{E}[n']}{n \log n} \in O\left(\frac{1}{\log n}\right).$$

Now, determining the k variables of a clause c is equivalent to observing which region of the order- k Voronoi diagram contains c , or more precisely, which k variables define this region. Thus, choosing random positions for the clauses is like throwing m' balls into n' (non-uniform) bins. Thus, if $m' \in \Omega(n'/\text{polylog } n')$, we can apply Corollary A.5. With the above bounds, which hold asymptotically almost surely, it is not hard to see that this condition in fact holds: If $n' \leq n$, it clearly holds as $m' \in \Omega(n)$. Otherwise, we have $n' \leq n \log n \leq n \log n'$, which implies $n \geq n'/\log n'$, and thus $m \in \Omega(n'/\log n')$.

Applying Corollary A.5 tells us that, asymptotically almost surely, there is a bin with a superconstant number of balls. In other words, there is a superconstant number of clauses that share the same set of k variables. For sufficiently large n , this is bigger than 2^k , which implies an unsatisfiable subformula consisting of only 2^k clauses. Clearly, it can be found in $O(n \log n)$ time by sorting the clauses lexicographically with respect to the contained variables. ■

ACKNOWLEDGMENTS

The authors would like to thank Thomas Sauerwald for the fruitful discussions on random SAT models and bipartite expansion. This work is partially funded by the project *Scale-Free Satisfiability* (project no. 416061626) of the German Research Foundation (DFG). Open Access funding enabled and organized by Projekt DEAL.

REFERENCES

1. C. Ansótegui, M. L. Bonet, and J. Levy, "On the structure of industrial SAT instances," Principles and practice of constraint programming (CP), I. P. Gent (ed.), Springer, Berlin, 2009, pp. 127–141. https://doi.org/10.1007/978-3-642-04244-7_13.
2. C. Ansótegui, J. Giráldez-Cru, and J. Levy, "The community structure of SAT formulas," Theory and applications of satisfiability testing (SAT), A. Cimatti and R. Sebastiani (eds.), Springer, Berlin, 2012, pp. 410–423. https://doi.org/10.1007/978-3-642-31612-8_31.
3. C. Ansótegui, M. L. Bonet, and J. Levy, *Scale-free random SAT instances*, Algorithms **15** (2022), no. 6, 219. <https://doi.org/10.3390/a15060219>.
4. C. Ansótegui, M. L. Bonet, and J. Levy, "Towards industrial-like random SAT instances," Proc. Int. Jt. Conf. Artif. Intell. (IJCAI), IJCAI, Melbourne, 2009, pp. 387–392. <https://doi.org/10.5555/1661445.1661507>.
5. F. Aurenhammer and H. Edelsbrunner, *An optimal algorithm for constructing the weighted Voronoi diagram in the plane*, Pattern Recognit. **17** (1984), no. 2, 251–257. [https://doi.org/10.1016/0031-3203\(84\)90064-5](https://doi.org/10.1016/0031-3203(84)90064-5).
6. F. Aurenhammer, R. Klein, and D.-T. Lee, *Voronoi diagrams and Delaunay triangulations*, World Scientific, Singapore, 2013. <https://doi.org/10.1142/8685>.
7. P. Beame and A. Sabharwal, "Non-restarting SAT solvers with simple preprocessing can efficiently simulate resolution," Proc. Conf. Artif. Intell. (AAAI), AAAI Press, Washington, DC, 2014, pp. 2608–2615. <https://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/view/8397>.
8. E. Ben-Sasson and N. Galesi, *Space complexity of random formulae in resolution*, Random Struct. Algorithms **23** (2003), no. 1, 92–109. <https://doi.org/10.1002/rsa.10089>.
9. E. Ben-Sasson and A. Wigderson, *Short proofs are narrow—resolution made simple*, J. ACM **48** (2001), no. 2, 149–169. <https://doi.org/10.1145/375827.375835>.

10. M. Bienkowski, V. Damerow, F. M. der Heide, and C. Sohler, "Average case complexity of Voronoi diagrams of n sites from the unit cube," Proc. Eur. Workshop Comput. Geom. (EuroCG), 2005, pp. 167–170. <http://www.win.tue.nl/EWCG2005/Proceedings/43.pdf>.
11. T. Bläsius, C. Freiberger, T. Friedrich, M. Katzmann, F. Montenegro-Retana, and M. Thieffry, *Efficient shortest paths in scale-free networks with underlying hyperbolic geometry*, ACM Trans. Algorithms **18** (2022), no. 2, 19:1–19:32. <https://doi.org/10.1145/3516483>.
12. T. Bläsius, P. Fischbeck, T. Friedrich, and M. Katzmann, *Solving vertex cover in polynomial time on hyperbolic random graphs*, Theory Comput. Syst. **67** (2023), no. 1, 28–51. <https://doi.org/10.1007/s00224-021-10062-9>.
13. T. Bläsius, T. Friedrich, and A. M. Sutton, "On the empirical time complexity of scale-free 3-SAT at the phase transition," Tools and algorithms for the construction and analysis of systems (TACAS), T. Vojnar and L. Zhang (eds.), Springer, Cham, 2019, pp. 117–134. https://doi.org/10.1007/978-3-030-17462-0_7.
14. M. Bode, N. Fountoulakis, and T. Müller, "On the giant component of random hyperbolic graphs," Eur. Conf. Comb. Graph Theory Appl., J. Nešetřil and M. Pellegrini (eds.), Edizioni della Normale, Pisa, 2013, pp. 425–429. https://doi.org/10.1007/978-88-7642-475-5_68.
15. C. Bohler, P. Cheilaris, R. Klein, C.-H. Liu, E. Papadopoulou, and M. Zavershynskiy, *On the complexity of higher order abstract Voronoi diagrams*, Comput. Geom. **48** (2015), no. 8, 539–551. <https://doi.org/10.1016/j.comgeo.2015.04.008>.
16. J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec, *Voronoi diagrams in higher dimensions under certain polyhedral distance functions*, Discret. Comput. Geom. **19** (1998), no. 4, 485–519. <https://doi.org/10.1007/PL00009366>.
17. B. N. Boots, *Weighting Thiessen polygons*, Econ. Geogr. **56** (1980), no. 3, 248–259. <https://doi.org/10.2307/142716>.
18. K. Bringmann, R. Keusch, and J. Lengler, *Geometric inhomogeneous random graphs*, Theor. Comput. Sci. **760** (2019), 35–54. <https://doi.org/10.1016/j.tcs.2018.08.014>.
19. F. Chung and L. Linyuan, *The average distances in random graphs with given expected degrees*, Proc. Natl. Acad. Sci. USA **99** (2002), no. 25, 15879–15882. <https://doi.org/10.1073/pnas.252631999>.
20. F. Chung and L. Linyuan, *Connected components in random graphs with given expected degree sequences*, Ann. Comb. **6** (2002), no. 2, 125–145. <https://doi.org/10.1007/PL00012580>.
21. V. Chvátal and E. Szemerédi, *Many hard examples for resolution*, J. ACM **35** (1988), no. 4, 759–768. <https://doi.org/10.1145/48014.48016>.
22. C. Cooper, A. Frieze, and G. Sorkin, *Random 2SAT with prescribed literal degrees*, Algorithmica **48** (2007), 249–265.
23. B. Cundefinedrbunar, A. Grama, J. Vitek, and O. Cundefinedrbunar, *Redundancy and coverage detection in sensor networks*, ACM Trans. Sens. Netw. **2** (2006), no. 1, 94–128. <https://doi.org/10.1145/1138127.1138131>.
24. M. Davis and H. Putnam, *A computing procedure for quantification theory*, J. ACM **7** (1960), no. 3, 201–215. <https://doi.org/10.1145/321033.321034>.
25. E. Deza and M.-M. Deza, *Voronoi diagram distances*, Elsevier, Amsterdam, 2006, 253–261. <https://doi.org/10.1016/B978-044452087-6/50020-2>.
26. A. Driemel, S. Har-Peled, and B. Raichel, *On the expected complexity of Voronoi diagrams on terrains*, ACM Trans. Algorithms **12** (2016), no. 3, 37. <https://doi.org/10.1145/2846099>.
27. D. P. Dubhashi and A. Panconesi, *Concentration of measure for the analysis of randomized algorithms*, Cambridge University Press, Cambridge, 2012. <https://doi.org/10.1017/CBO9780511581274>.
28. R. A. Dwyer, *Higher-dimensional Voronoi diagrams in linear expected time*, Discret. Comput. Geom. **6** (1991), no. 3, 343–367. <https://doi.org/10.1007/BF02574694>.
29. J. Erickson, "Nice point sets can have nasty Delaunay triangulations," Proc. Annu. Symp. Comput. Geom. (SoCG), ACM, New York, 2001, pp. 96–105. <https://doi.org/10.1145/378583.378636>.
30. J. Erickson, *Dense point sets have sparse Delaunay triangulations: Or "... but not too nasty"*, Discret. Comput. Geom. **33** (2005), 83–115. <https://doi.org/10.1007/s00454-004-1089-3>.
31. J. L. Esteban and J. Torán, *Space bounds for resolution*, Inf. Comput. **171** (2001), no. 1, 84–97. <https://doi.org/10.1006/inco.2001.2921>.
32. C. Fan and B. Raichel, "Linear expected complexity for directional and multiplicative Voronoi diagrams," Annu. Eur. Symp. Algorithms (ESA), Dagstuhl Publishing, Wadern, Germany, 2020, pp. 45:1–45:18. <https://doi.org/10.4230/LIPIcs.ESA.2020.45>.
33. T. Friedrich and R. Rothenberger, "The satisfiability threshold for non-uniform random 2-SAT," Int. Colloq. Autom. Lang. Program. (ICALP), Vol **132**, Dagstuhl Publishing, Wadern, Germany, 2019, pp. 61:1–61:14. <https://doi.org/10.4230/LIPIcs.ICALP.2019.61>.
34. T. Friedrich, A. Krohmer, R. Rothenberger, T. Sauerwald, and A. M. Sutton, "Bounds on the satisfiability threshold for power law distributed random SAT," Annu. Eur. Symp. Algorithms (ESA), K. Pruhs and C. Sohler (eds.), Dagstuhl Publishing, Wadern, Germany, 2017, pp. 37:1–37:15. <https://doi.org/10.4230/LIPIcs.ESA.2017.37>.
35. T. Friedrich, A. Krohmer, R. Rothenberger, and A. M. Sutton, "Phase transitions for scale-free SAT formulas," Proc. Conf. Artif. Intell. (AAAI), AAAI Press, Washington, DC, 2017, pp. 3893–3899. <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14755>.

36. L. C. Galvão, A. G. N. Novaes, J. E. Souza de Cursi, and J. C. Souza, *A multiplicatively-weighted Voronoi diagram approach to logistics districting*, Comput. Oper. Res. **33** (2006), no. 1, 93–114. <https://doi.org/10.1016/j.cor.2004.07.001>.
37. A. Gemsa, D. T. Lee, C.-H. Liu, and D. Wagner, “Higher order city Voronoi diagrams,” Scandinavian workshop on algorithm theory (SWAT), F. V. Fomin and P. Kaski (eds.), Springer, Berlin, 2012, pp. 59–70. https://doi.org/10.1007/978-3-642-31155-0_6.
38. J. Giráldez-Cru and J. Levy, “A modularity-based random SAT instances generator,” Proc. Int. Jt. Conf. Artif. Intell. (IJCAI), IJCAI, Melbourne, 2015, pp. 1952–1958. <https://doi.org/10.5555/2832415.2832520>.
39. J. Giráldez-Cru and J. Levy, “Description of popularity-similarity SAT instances,” Proc. SAT Compet. 2017: Solver Benchmark Descr., T. Balyo, M. J. H. Heule, and M. Jarvisalo (eds.), University of Helsinki, Helsinki, 2017, pp. 49–50. <https://helda.helsinki.fi/handle/10138/224324>.
40. J. Giráldez-Cru and J. Levy, “Locality in random SAT instances,” Proc. Int. Jt. Conf. Artif. Intell. (IJCAI), IJCAI, Melbourne, 2017, pp. 638–644. <https://doi.org/10.24963/ijcai.2017/89>.
41. M. J. Golin and H.-S. Na, *On the average complexity of 3D-Voronoi diagrams of random points on convex polytopes*, Comput. Geom. **25** (2003), no. 3, 197–231. [https://doi.org/10.1016/S0925-7721\(02\)00123-2](https://doi.org/10.1016/S0925-7721(02)00123-2).
42. S. Har-Peled and B. Raichel, *On the complexity of randomly weighted multiplicative Voronoi diagrams*, Discret. Comput. Geom. **53** (2015), no. 3, 547–568. <https://doi.org/10.1007/s00454-015-9675-0>.
43. V. Klee, *On the complexity of d-dimensional Voronoi diagrams*, Arch. Math. **34** (1980), no. 1, 75–80. <https://doi.org/10.1007/BF01224932>.
44. D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguñá, *Hyperbolic geometry of complex networks*, Phys. Rev. E **82** (2010), no. 3, 036106. <https://doi.org/10.1103/PhysRevE.82.036106>.
45. N.-M. Lê, *On Voronoi diagrams in the L_p -metric in \mathbb{R}^d* , Discret. Comput. Geom. **16** (1996), no. 2, 177–196. <https://doi.org/10.1007/BF02716806>.
46. D.-T. Lee, *On k-nearest neighbor Voronoi diagrams in the plane*, IEEE Trans. Comput. **31** (1982), 478–487. <https://doi.org/10.1109/TC.1982.1676031>.
47. C.-H. Liu, E. Papadopoulos, and D. T. Lee, “An output-sensitive approach for the L_1/L_∞ k-nearest-neighbor Voronoi diagram,” Annu. Eur. Symp. Algorithms (ESA), C. Demetrescu and M. M. Halldórsson (eds.), Springer, Berlin, 2011, pp. 70–81. https://doi.org/10.1007/978-3-642-23719-5_7.
48. N. Mull, D. J. Fremont, and S. A. Seshia, “On the hardness of SAT with community structure,” Theory and applications of satisfiability testing (SAT), N. Creignou and D. Le Berre (eds.), Springer, Cham, 2016, pp. 141–159. https://doi.org/10.1007/978-3-319-40970-2_10.
49. K. Mulmuley, *On levels in arrangements and Voronoi diagrams*, Discret. Comput. Geom. **6** (1991), no. 3, 307–338. <https://doi.org/10.1007/BF02574692>.
50. T. Müller and M. Staps, *The diameter of KPKVB random graphs*, Adv. Appl. Probab. **51** (2019), no. 2, 358–377. <https://doi.org/10.1017/apr.2019.23>.
51. Z. Newsham, V. Ganesh, S. Fischmeister, G. Audemard, and L. Simon, “Impact of community structure on SAT solver performance,” Theory and applications of satisfiability testing (SAT), C. Sinz and U. Egly (eds.), Springer, Cham, 2014, pp. 252–268. https://doi.org/10.1007/978-3-319-09284-3_20.
52. O. Omelchenko and A. A. Bulatov, *Satisfiability threshold for power law random 2-SAT in configuration model*, Theor. Comput. Sci. **888** (2021), 70–94. <https://doi.org/10.1016/j.tcs.2021.07.028>.
53. F. Papadopoulos, M. Maksim Kitsak, Á. Serrano, M. Boguñá, and D. Krioukov, *Popularity versus similarity in growing networks*, Nature **489** (2012), 537–540. <https://doi.org/10.1038/nature11459>.
54. K. Pipatsrisawat and A. Darwiche, *On the power of clause-learning SAT solvers as resolution engines*, Artif. Intell. **175** (2011), no. 2, 512–525. <https://doi.org/10.1016/j.artint.2010.10.002>.
55. M. Raab and A. Steger, ““Balls into bins”—a simple and tight analysis,” Randomization and approximation techniques in computer science, M. Luby, J. D. P. Rolim, and M. Serna (eds.), Springer, Berlin, 1998, pp. 159–170. https://doi.org/10.1007/3-540-49543-6_13.
56. R. Seidel, “On the number of faces in higher-dimensional Voronoi diagrams,” Proc. Annu. Symp. Comput. Geom. (SoCG), ACM, New York, 1987, pp. 181–185. <https://doi.org/10.1145/41958.41977>.
57. M. I. Shamos and D. Hoey, “Closest-point problems,” Proc. Annu. Symp. Found. Comput. Sci., ACM, New York, 1975, pp. 151–162. <https://doi.org/10.1109/SFCS.1975.8>.
58. P. van Beek, “Backtracking search algorithms,” Foundations of artificial intelligence, Handbook of Constraint Programming, Vol 2, F. Rossi, P. van Beek, and T. Walsh (eds.), Elsevier, Amsterdam, 2006, pp. 85–134. [https://doi.org/10.1016/S1574-6526\(06\)80008-8](https://doi.org/10.1016/S1574-6526(06)80008-8).
59. M. Y. Vardi, *Boolean satisfiability: Theory and engineering*, Commun. ACM **57** (2014), no. 3, 5. <https://doi.org/10.1145/2578043>.
60. I. Voitalov, P. van der Hoorn, R. van der Hofstad, and D. Krioukov, *Scale-free networks well done*, Phys. Rev. Res. **1** (2019), 033034. <https://doi.org/10.1103/PhysRevResearch.1.033034>.
61. L. Warnke, *On the method of typical bounded differences*, Comb. Probab. Comput. **25** (2016), no. 2, 269–299. <https://doi.org/10.1017/S0963548315000103>.

62. E. Zulkoski, R. Martins, C. M. Wintersteiger, R. Robere, J. Liang, K. Czarnecki, and V. Ganesh. Relating complexity-theoretic parameters with SAT solver performance. *CoRR* abs/1706.08611, 2017. <http://arxiv.org/abs/1706.08611>.

How to cite this article: T. Bläsius, T. Friedrich, A. Göbel, J. Levy, and R. Rothenberger, *The impact of heterogeneity and geometry on the proof complexity of random satisfiability*, *Random Struct. Alg.* **63** (2023), 885–941. <https://doi.org/10.1002/rsa.21168>

APPENDIX A: BASIC TECHNICAL TOOLS

This section is a collection of tools we use throughout the article that were either known before or are straight-forward to prove but distract from the core arguments we make in the article.

A.1 | Discrete power-law weights

The following lemma summarizes some properties of the probability distribution given by the discrete power-law weights.

Lemma A.1. *Let $\beta > 2$ and*

$$p_i = \frac{i^{-1/(\beta-1)}}{\sum_{j=1}^n j^{-1/(\beta-1)}},$$

for $i \in [n]$. It holds that

$$\sum_{j=1}^n j^{-1/(\beta-1)} = (1 + o(1)) \frac{\beta-1}{\beta-2} \cdot n^{(\beta-2)/(\beta-1)},$$

$$F(i) := \sum_{j=1}^i p_j \in O\left(\left(\frac{i}{n}\right)^{(\beta-2)/(\beta-1)}\right),$$

and

$$\sum_{j=1}^n p_j^2 \in \begin{cases} \Theta\left(n^{-2\frac{\beta-2}{\beta-1}}\right), & \beta < 3; \\ \Theta(\ln n/n), & \beta = 3; \\ \Theta(n^{-1}), & \beta > 3. \end{cases} \quad (\text{A1})$$

Proof. Since $j^{-1/(\beta-1)}$ is monotonically decreasing, it holds that

$$\begin{aligned} \sum_{j=1}^n j^{-1/(\beta-1)} &\leq 1 + \int_{j=1}^n j^{-1/(\beta-1)} dj \\ &= 1 + \frac{\beta-1}{\beta-2} (n^{(\beta-2)/(\beta-1)} - 1) = \frac{\beta-1}{\beta-2} \cdot n^{(\beta-2)/(\beta-1)} - \frac{1}{\beta-2}, \end{aligned}$$

and

$$\begin{aligned}\sum_{j=1}^n j^{-1/(\beta-1)} &\geq n^{-1/(\beta-1)} + \int_{j=1}^n j^{-1/(\beta-1)} dj \\ &= \frac{\beta-1}{\beta-2} \cdot n^{(\beta-2)/(\beta-1)} - \frac{\beta-1}{\beta-2} + n^{-1/(\beta-1)}.\end{aligned}$$

Equivalently, we get

$$\begin{aligned}F(i) &= \sum_{j=1}^i p_j = \frac{\sum_{j=1}^i j^{-1/(\beta-1)}}{\sum_{j=1}^n j^{-1/(\beta-1)}} \\ &\leq \frac{1}{\sum_{j=1}^n j^{-1/(\beta-1)}} \left(1 + \int_{j=1}^i j^{-1/(\beta-1)} dj \right) \\ &\leq \frac{1}{\sum_{j=1}^n j^{-1/(\beta-1)}} \left(\frac{\beta-1}{\beta-2} \cdot i^{(\beta-2)/(\beta-1)} - \frac{1}{\beta-2} \right) \\ &\in O\left(\left(\frac{i}{n}\right)^{(\beta-2)/(\beta-1)}\right).\end{aligned}$$

Finally, we want to bound

$$\sum_{j=1}^n p_j^2 = \frac{\sum_{j=1}^n j^{-2/(\beta-1)}}{\left(\sum_{j=1}^n j^{-1/(\beta-1)}\right)^2}.$$

First, note that for $\beta = 3$ this equation yields

$$\sum_{j=1}^n p_j^2 = \frac{H_n}{\left(\sum_{j=1}^n j^{-1/(\beta-1)}\right)^2} \in \Theta(\ln n/n),$$

where H_n denotes the n th harmonic number. For $\beta \neq 3$ we can achieve

$$\begin{aligned}\sum_{j=1}^n p_j^2 &\leq \frac{1}{\left(\sum_{j=1}^n j^{-1/(\beta-1)}\right)^2} \left(1 + \int_{j=1}^n j^{-2/(\beta-1)} dj \right) \\ &= \frac{1}{\left(\sum_{j=1}^n j^{-1/(\beta-1)}\right)^2} \left(1 + \frac{\beta-1}{\beta-3} \cdot (n^{(\beta-3)/(\beta-1)} - 1) \right),\end{aligned}$$

and

$$\begin{aligned}\sum_{j=1}^n p_j^2 &\geq \frac{1}{\left(\sum_{j=1}^n j^{-1/(\beta-1)}\right)^2} \left(n^{-2/(\beta-1)} + \int_{j=1}^n j^{-2/(\beta-1)} dj \right) \\ &= \frac{1}{\left(\sum_{j=1}^n j^{-1/(\beta-1)}\right)^2} \left(n^{-2/(\beta-1)} + \frac{\beta-1}{\beta-3} \cdot (n^{(\beta-3)/(\beta-1)} - 1) \right).\end{aligned}$$

If $\beta < 3$, the expressions above yield

$$\sum_{j=1}^n p_j^2 \in \Theta \left(\frac{1}{\left(\sum_{j=1}^n j^{-1/(\beta-1)} \right)^2} \right) \subseteq \Theta \left(n^{-2(\beta-2)/(\beta-1)} \right).$$

For $\beta > 3$, they yield

$$\sum_{j=1}^n p_j^2 \in \Theta \left(\frac{n^{(\beta-3)/(\beta-1)}}{n^{2(\beta-2)/(\beta-1)}} \right) \subseteq \Theta \left(n^{-1} \right).$$

This proves all statements of the lemma. \blacksquare

A.2 | CDF of connection weights in the geometric model

The CDF $F_X(x)$ of the connection weights $X(c, v)$ in the geometric SAT model satisfies the following lemma.

Lemma A.2. $F_X(x) = 1 - \Pi_{d,p} w_v x^{-T}$ for $x \geq (2^d w_v)^{1/T}$.

Proof. Inserting the definition of the connection weight and rearranging slightly yields

$$\begin{aligned} F_X(x) &= \Pr [X(c, v) \leq x] \\ &= \Pr \left[\left(\frac{w_v}{\|c - v\|^d} \right)^{1/T} \leq x \right] \\ &= \Pr \left[\|c - v\| \geq w_v^{1/d} x^{-T/d} \right] \\ &= 1 - \Pr \left[\|c - v\| < w_v^{1/d} x^{-T/d} \right]. \end{aligned}$$

As c and v are two random points, we can use the CDF for the distances between random points in Equation (1) to obtain

$$F_X(x) = 1 - \Pi_{d,p} w_v x^{-T} \quad \text{for } x \geq (2^d w_v)^{1/T},$$

which concludes the proof. \blacksquare

A.3 | Volume of balls in a hypercube

We are regularly concerned with the asymptotic behavior of a ball's volume depending on its radius. The following lemma helps us to deal with the edge case, where the ball stretches beyond the boundary of our ground space.

Lemma A.3. *Let H be a d -dimensional unit-hypercube in \mathbb{R}^d equipped with a p -norm. There exists a constant $c > 0$ such that, for every $p \in H$ and $r > 0$, the intersection of H with the ball $B_p(r)$ of radius r around p has volume at least $\min\{1, cr^d\}$.*

Proof. In the following, we assume $H = [-0.5, 0.5]^d$ (rather than $[0, 1]^d$), as it makes the proof more convenient. If r is sufficiently small, then $B_p(r)$ is completely contained in H . Thus, in this case, the claim follows from the fact that the volume of a ball with radius r in d -dimensional space is proportional to r^d . Thus, we have to prove that the parts of $B_p(r)$ outside of H are asymptotically not relevant.

Let p_1, \dots, p_d be the coordinates of \mathbf{p} and assume without loss of generality that \mathbf{p} lies in the all-negative orthant, that is, $p_i \leq 0$ for $i \in [d]$. We proof the claim by defining a box B with the following three properties. First, the box B has volume proportional to r^d . Second, B is a subset of the ball $B_p(r)$. Third, B is a subset of the hypercube H or H is a subset of $B_p(r)$. Note that the lemma's statement clearly holds if H is a subset of $B_p(r)$ as the intersection has volume 1 in this case. If H is not a subset of $B_p(r)$, the second and third property imply that B is a subset of the intersection of $B_p(r)$ and H . Thus, the volume of B given by the first property is a lower bound for the volume of the intersection, which proves the claim.

It remains to define B and prove the three properties. The box B has \mathbf{p} as corner and extends from there in the direction of the all-positive orthant. The side lengths are chosen proportional to the distance from the edge of H in this direction. Formally, the corners of B are $\{p_1, p_1 + r(0.5 - p_1)/\sqrt[p]{d}\} \times \dots \times \{p_d, p_d + r(0.5 - p_d)/\sqrt[p]{d}\}$.

To prove the first property, note that the side length of B in dimension i is $r(0.5 - p_i)/\sqrt[p]{d}$. As $p_i \leq 0$, this is at least $0.5r/\sqrt[p]{d}$, which implies that the volume of B is at least $(0.5r/\sqrt[p]{d})^d$. For the second property, note that the point of B with maximum distance from \mathbf{p} is the opposite corner, that is, the point with coordinates $(p_i + r(0.5 - p_i)/\sqrt[p]{d})$. The distance from \mathbf{p} is given by

$$\sqrt[p]{\sum_{i=1}^d \left(\frac{r(0.5 - p_i)}{\sqrt[p]{d}} \right)^p} \leq \sqrt[p]{\sum_{i=1}^d \frac{r^p}{d}} = r.$$

Finally, for the third property, assume $r = \sqrt[p]{d}$. Then the coordinates $p_i + r(0.5 - p_i)/\sqrt[p]{d}$ of the corners of B simplify to 0.5. Thus, all corners of B are still in the hypercube H if $r \leq \sqrt[p]{d}$. On the other hand, if $r \geq \sqrt[p]{d}$, then H is completely contained in $B_p(r)$, which concludes the proof of the last property. We note that it is easy to verify that all above arguments also hold for the limit $p = \infty$. ■

A.4 | Derivative of the incomplete gamma function

We need the following somewhat technical bound that is easy to verify.

Lemma A.4. *Let Γ be the gamma function. For any $\alpha, \beta, \gamma, d \in \mathbb{R}$ with $\beta, \gamma, d > 0$,*

$$\int_0^\gamma x^{\alpha d - 1} \exp(-\beta x^d) dx \leq \frac{\Gamma(\alpha)}{\beta^\alpha d}.$$

Proof. Let $\Gamma(\alpha, x)$ be the incomplete gamma function. Its derivative is

$$\frac{\partial \Gamma(\alpha, x)}{\partial x} = -x^{\alpha-1} \exp(-x).$$

Thus, it follows that

$$\frac{\partial}{\partial x} \left(-\frac{\Gamma(\alpha, \beta x^d)}{\beta^\alpha d} \right) = \beta d x^{d-1} \frac{(\beta x^d)^{\alpha-1} \exp(-\beta x^d)}{\beta^\alpha d} = x^{\alpha d-1} \exp(-\beta x^d).$$

Using this, the given integral evaluates to

$$\begin{aligned} \int_0^\gamma x^{\alpha d-1} \exp(-\beta x^d) dx &= \left[-\frac{\Gamma(\alpha, \beta x^d)}{\beta^\alpha d} \right]_0^\gamma \\ &= \frac{1}{\beta^\alpha d} (\Gamma(\alpha, 0) - \Gamma(\alpha, \beta \gamma^d)) \\ &\leq \frac{\Gamma(\alpha, 0)}{\beta^\alpha d}. \end{aligned}$$

■

A.5 | Balls into heterogeneous bins

Consider throwing m balls into n uniform bins, that is, for each ball we draw one of the n bins uniformly at random and place the ball into the drawn bin. The *maximum load* L is the random variable that describes the maximum number of balls that are together in the same bin. From the analysis by Raab and Steger [55, Theorem 1], we immediately get the following corollary.

Corollary A.5 ([55, Theorem 1]). *Throw m balls into n uniform bins and let L be the maximum load. If $m \in \Omega(\frac{n}{\text{polylog } n})$, then $L \in \Omega(\frac{\log n}{\log \log n})$ asymptotically almost surely.*

Now assume we have non-uniform bins, that is, the probability for each ball to end up in the i th bin is p_i with $\sum_i p_i = 1$. Intuitively, Corollary A.5 should still hold in this setting, as increasing the probability of some bins only makes it more likely that a bin gets many balls. Making this argument formal yields the following theorem.

Theorem A.6. *Corollary A.5 also holds for the non-uniform bins.*

Proof. Let $B = [n]$ be the set of all bins and let B' be the subset of bins with probability at least $1/(2n)$. These are the bins whose probability either increased, or decreased by a factor of at most 2. Without loss of generality, let $B' = [n']$. Note that the probability for a ball to land in a bin of B' is at least a constant, as every bin not in B' has probability at most $1/(2n)$. Thus, by the Chernoff–Hoeffding bound in Corollary A.8, a constant fraction of the balls end up in a bin of B' with high probability. We make a case distinction on how large n' is.

First, assume $n' \leq m/\log n$. Thus, with high probability, we end up with $\Theta(m)$ balls in at most $m/\log n$ bins, which means that at least one bin contains $\Omega(\log n)$ balls. Thus, clearly $L \in \Omega(\log n/\log \log n)$.

Second, assume $n' > m/\log n$. Recall that each bin in B' has probability at least $1/(2n)$. We consider the alternative experiment where, for every ball, each bin in B' has probability exactly $1/(2n)$ to get the ball. Balls not landing in B' are discarded. Let L' denote the maximum number of balls that share a bin in B' . Clearly, we can couple the two experiments such that $L \geq L'$ holds in every outcome. It remains to show that $L' \in \Omega(\log n/\log \log n)$. For this, let m' be the number of balls ending up in B' . Note that m' is a random variable.

However, if we condition on m' , then we are back to the normal homogeneous balls into bins, except that we throw m' balls into n' bins. If we show that $m' \in \Omega(n'/\text{polylog } n')$, then Corollary A.5 tells us that $L' \in \Omega(\log n' / \log \log n')$. First note that this is sufficient for our purpose: as $n' > m / \log n$ and $m \in \Omega(n/\text{polylog } n)$, we get

$$\begin{aligned} \frac{\log n'}{\log \log n'} &> \frac{\log m - \log \log n}{\log(\log m - \log \log n)} \\ &\in \Omega\left(\frac{\log n - \log \text{polylog } n - \log \log n}{\log(\log n - \log \text{polylog } n - \log \log n)}\right) \\ &\subseteq \Omega\left(\frac{\log n}{\log \log n}\right). \end{aligned}$$

It remains to show that $m' \in \Omega(n'/\text{polylog } n')$ so that we can actually apply Corollary A.5. To do so, recall that B' has n' bins, each with probability $1/(2n)$. Thus, the probability that a single ball lands in B' is $n'/(2n)$, which shows that m' is $mn'/(2n)$ in expectation. As n' is almost m (up to logarithmic factors) and m is almost n , this expectation is almost linear in n . Thus, by the Chernoff–Hoeffding bound in Corollary A.8, we can assume that

$$m' \in \Theta\left(\frac{mn'}{n}\right)$$

holds with high probability. Using that $n' > m / \log n$ and $m \in \Omega(n/\text{polylog } n)$, we obtain

$$\frac{mn'}{n} > \frac{m^2}{n \log n} \in \Omega\left(\frac{n}{\text{polylog } n}\right).$$

As $n' \leq n$, it follows that $m \in \Omega(n'/\text{polylog } n')$, which concludes the proof. ■

A.6 | Concentration bounds

For a random experiment, we say that an event happens *with high probability* (w. h. p.) if the probability is at least $1 - O(1/n)$. The type of event we are usually interested in is that a random variable assumes a value close to its expectation, that is, that the random variable is concentrated. In the following, we state two well known techniques to prove concentration, namely a Chernoff–Hoeffding bound and the method of bounded differences. In both cases we derive asymptotic variants that suite our purpose better than the original exact bounds.

A.6 | Chernoff–Hoeffding

Theorem A.7 (Theorem 1.1 in [27]). *Let X_1, \dots, X_n be independent random variables with values in $\{0, 1\}$ and let $X = \sum_{i \in [n]} X_i$ be their sum. Then, for all $0 < \varepsilon < 1$,*

$$\begin{aligned} \Pr[X > (1 + \varepsilon) \cdot \mathbb{E}[X]] &\leq \exp\left(-\frac{\varepsilon^2}{3} \mathbb{E}[X]\right), \text{ and} \\ \Pr[X < (1 - \varepsilon) \cdot \mathbb{E}[X]] &\leq \exp\left(-\frac{\varepsilon^2}{2} \mathbb{E}[X]\right). \end{aligned}$$

We use this bound multiple times in a similar way, which is captured by the following direct corollary.

Corollary A.8. *Let X_1, \dots, X_n , and X be as in Theorem A.7. Let $f(n) \in \omega(\log n)$ be an upper or lower bound for $\mathbb{E}[X]$. With overwhelming probability, $X \in O(f(n))$ and $X \in \Omega(f(n))$, respectively.*

Proof. Assume $f(n)$ is a lower bound, that is, $f(n) \leq \mathbb{E}[X]$. We show $X \in \Omega(f(n))$ with the desired probability. By the second inequality of Theorem A.7, we have

$$\begin{aligned} \Pr[X < (1 - \varepsilon) \cdot f(n)] &\leq \Pr[X < (1 - \varepsilon) \cdot \mathbb{E}[X]] \\ &\leq \exp\left(-\frac{\varepsilon^2}{2} \mathbb{E}[X]\right) \\ &\leq \exp\left(-\frac{\varepsilon^2}{2} f(n)\right) = n^{-\omega(1)}, \end{aligned}$$

where the last equality is due to the fact that $f(n) \in \omega(\log n)$. Thus, for any constant c , this probability is below n^{-c} for sufficiently large n . Hence, for any constant $\varepsilon \in (0, 1)$, $X \geq (1 - \varepsilon) \cdot f(n)$ with probability $1 - n^{-c}$.

Assume $f(n)$ is an upper bound, that is, $\mathbb{E}[X] \leq f(n)$. Let X' be a random variable with $f(n) = \mathbb{E}[X']$ such that X' dominates X in the sense that $X \leq X'$ for every outcome. We show that $X' \in O(f(n))$ with probability $1 - n^{-c}$, which implies $X \in O(f(n))$ with at least the same probability. The first inequality of Theorem A.7 yields

$$\begin{aligned} \Pr[X' > (1 + \varepsilon) \cdot f(n)] &= \Pr[X' > (1 + \varepsilon) \cdot \mathbb{E}[X']] \\ &\leq \exp\left(-\frac{\varepsilon^2}{3} \mathbb{E}[X']\right) \\ &= \exp\left(-\frac{\varepsilon^2}{3} f(n)\right) = n^{-\omega(1)}. \end{aligned}$$

As before, the last inequality comes from the fact that $f(n) \in \omega(\log n)$. All remaining arguments are as in the case where $f(n)$ was a lower bound. ■

A.6 | Method of typical bounded differences

Theorem A.9 (Theorem 2 in [61]). *Let $X = (X_1, \dots, X_N)$ be a family of independent random variables with X_k taking values in Λ_k and let $\Lambda = \prod_{j \in [N]} \Lambda_j$. Let $\Gamma \subseteq \Lambda$ be an event and assume that the function $f : \Lambda \rightarrow \mathbb{R}$ satisfies the following typical Lipschitz condition.*

(TL) *There are numbers $(c_k)_{k \in [N]}$ and $(d_k)_{k \in [N]}$ with $c_k \leq d_k$ such that whenever $x, \tilde{x} \in \Lambda$ differ only in the k th coordinate, we have*

$$|f(x) - f(\tilde{x})| \leq \begin{cases} c_k & \text{if } x \in \Gamma, \\ d_k & \text{otherwise.} \end{cases}$$

For any numbers $(\gamma_k)_{k \in [N]}$ with $\gamma_k \in (0, 1]$, there is an event $\mathcal{B} = \mathcal{B}(\Gamma, (\gamma_k)_{k \in [N]})$ satisfying

$$\Pr[\mathcal{B}] \leq \sum_{k \in [N]} \gamma_k^{-1} \cdot \Pr[X \notin \Gamma] \text{ and } \neg \mathcal{B} \subseteq \Gamma,$$

such that for $\mu = \mathbb{E}[f]$, $e_k = \gamma_k(d_k - c_k)$ and any $t \geq 0$, we have

$$\Pr[f(X) \geq \mu + t \text{ and } \neg \mathcal{B}] \leq \exp\left(-\frac{t^2}{2 \sum_{k \in [N]} (c_k + e_k)^2}\right).$$

We derive the following corollary from this, which is more convenient for our purpose and uses a notation more compatible with the rest of the article.

Corollary A.10. Let $X = (X_1, \dots, X_N) \in \Lambda$ be a family of independent random variables and let $\Gamma \subseteq \Lambda$ be an event with $\Pr[\Gamma] \geq 1 - N^{-c}$. Moreover, let $f : \Lambda \rightarrow \mathbb{R}$ with $|f(X)| \leq N^{c-2}$ and let $(\Delta_i)_{i \in [N]} \in \Omega(1)$ be numbers such that for any two $x \in \Gamma$ and $\tilde{x} \in \Lambda$ that differ only in the i th coordinate, we have $|f(x) - f(\tilde{x})| \leq \Delta_i$. If $\sum_{i \in [N]} \Delta_i^2 \in O(\mathbb{E}[f]^2 / \log^2 N)$ then $f(X) \in \Theta(\mathbb{E}[f])$ holds with high probability.

Proof. We want to apply Theorem A.9. First note that $|f(X)| \leq N^{c-1}$ implies that f satisfies the typical Lipschitz condition when setting $c_i = \Delta_i$ and $d_i = 2N^{c-2}$ for every i . We set γ_i in Theorem A.9 to $\gamma_i = 1/d_i$ yielding $e_i \leq 1$. Thus, we get the event \mathcal{B} with

$$\begin{aligned} \Pr[\mathcal{B}] &\leq \sum_{i \in [N]} \gamma_i^{-1} \cdot \Pr[\neg \Gamma] \\ &= \sum_{i \in [N]} d_i \cdot \Pr[\neg \Gamma] \\ &\leq \sum_{i \in [N]} 2N^{c-2} \cdot N^{-c} \\ &= 2N^{-1} \in O(N^{-1}), \end{aligned}$$

such that

$$\Pr[f(X) \geq \mathbb{E}[f] + t \text{ and } \neg \mathcal{B}] \leq \exp\left(-\frac{t^2}{2 \sum_{i \in [N]} (\Delta_i + e_i)^2}\right).$$

As $\Delta_i \in \Omega(1)$ and $e_i \leq 1$, we get that the sum in the denominator is up to constants equal to $\sum_{i \in [N]} \Delta_i^2 \in O(\mathbb{E}[f]^2 / \log^2 N)$, that is, for sufficiently large N , there exists a positive constant a such that

$$\Pr[f(X) \geq \mathbb{E}[f] + t \text{ and } \neg \mathcal{B}] \leq \exp\left(-\frac{a \cdot t^2 \cdot \log^2 N}{\mathbb{E}[f]^2}\right).$$

Choosing $t = b \cdot \mathbb{E}[f]$ for any positive constant b yields an upper bound of $N^{-ab^2 \log N} \in O(N^{-1})$ for the probability. Thus, we obtain

$$\Pr[f(x) \geq (1+b)\mathbb{E}[f]] \leq \Pr[f(x) \geq (1+b)\mathbb{E}[f] \text{ and } \neg \mathcal{B}] + \Pr[\mathcal{B}] \in O(N^{-1}).$$

As already noted by Warnke [61], one can obtain the same bound as stated in Theorem A.9 for the opposite direction ($\Pr[f(x) \leq \mathbb{E}[f] - t]$) by using $-f$. The above argument works exactly the same for this case, yielding $\Pr[f(x) \leq (1-b)\mathbb{E}[f]] \in O(N^{-1})$. Note that for this direction it is crucial that we can choose b to be an arbitrary positive constant. This yields the claim that $f(x) \in \Theta(\mathbb{E}[f])$ with high probability. ■