Nearly All k-SAT Functions Are Unate

József Balogh
University of Illinois at Urbana-Champaign
Champaign, IL, USA
jobal@illinois.edu

Dingding Dong
Harvard University
Cambridge, MA, USA
ddong@math.harvard.edu

Bernard Lidický
Iowa State University
Ames, IA, USA
lidicky@iastate.edu

Nitya Mani
Massachusetts Institute of Technology
Cambridge, MA, USA
nmani@mit.edu

Yufei Zhao
Massachusetts Institute of Technology
Cambridge, MA, USA
yufeiz@mit.edu

ABSTRACT

We prove that $1-o(1)$ fraction of all k-SAT functions on n Boolean variables are unate (i.e., monotone after first negating some variables), for any fixed positive integer k and as $n \to \infty$. This resolves a conjecture by Bollobás, Brightwell, and Leader from 2003.

This paper is the second half of a two-part work solving the problem. The first part, by Dong, Mani, and Zhao, reduces the conjecture to a Turán problem on partially directed hypergraphs. In this paper we solve this Turán problem.

CCS CONCEPTS

• **Mathematics of computing** → **Combinatoric problems**; **Hypergraphs**.

KEYWORDS

k-SAT function, sum of squares, hypergraph container method, Turán problems.

ACM Reference Format:
József Balogh, Dingding Dong, Bernard Lidický, Nitya Mani, and Yufei Zhao. 2023. Nearly All k-SAT Functions Are Unate. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC '23), June 20–23, 2023, Orlando, FL, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3564246.3585123

1 INTRODUCTION

A *k-SAT function* is a Boolean function $f : \{0,1\}^n \to \{0,1\}$ on n *variables* $x_1, \ldots, x_n$ that has at least one representation as a k-SAT (specifically k-CNF) formula. Such a function is called *unate* if it has a formula where, for each i, x_i and its negation $\overline{x_i}$ do not both appear in the formula. In other words, a function/formula is unate if it is monotone after first negating some variables (here *monotone* means that only positive literals x_i appear and no negated literals $\overline{x_i}$ appear).

Our main result below proves a conjecture by Bollobás, Brightwell, and Leader [5].

STOC '23, June 20–23, 2023, Orlando, FL, USA

ACM ISBN 978-1-4503-9913-5/23/06.
https://doi.org/10.1145/3564246.3585123

Theorem 1.1. *Let k be a fixed positive integer. As $n \to \infty$, the number of k-SAT functions on n Boolean variables is $(1+o(1))2^{n+\binom{n}{k}}$, and $1-o(1)$ proportion of all k-SAT functions on n Boolean variables are unate.*

An easy argument shows that the number of unate k-SAT functions on n Boolean variables is at least $(1-o(1))2^{n+\binom{n}{k}}$. Basically, for each i, we can choose whether to negate x_i, and then for each k-element subset of the n available literals, whether to include it as a clause. There is a small overcount due to unused variables, but it is easy to analyze (see [5]).

Bollobás, Brightwell, and Leader [5] proved a weaker version of their conjecture for the $k=2$ case, namely that the number of 2-SAT functions on n Boolean variables is $2^{(1+o(1))\binom{n}{2}}$. Some special cases of Theorem 1.1 were previously known: $k=2$ by Allen [1] (also see [8] for an alternate proof) and $k=3$ by Ilinca and Kahn [9].

We call a k-SAT formula *minimal* (the term *non-redundant* was used in [8]) if deleting any clause changes the function. Every k-SAT function admits some minimal k-SAT formula representation. Some functions, such as unate functions [1], admit a unique minimal representation, while others may admit one or multiple such representations. The fact that unate functions have a unique minimal representation follows from the following observation: suppose F_1, F_2 are formulae where F_1 is unate and F_2 contains a clause C whose variables do not form a clause in F_1; F_1 and F_2 cannot represent the same function, since there exists an assignment where C evaluates to F (so that F_2 is not satisfied) but all clauses in F_1 are satisfied.

We establish the following result, which implies Theorem 1.1.

Theorem 1.2. *Let k be a fixed positive integer. As $n \to \infty$, the number of minimal k-SAT formulae on n Boolean variables is $(1+o(1))2^{n+\binom{n}{k}}$, and a $1-o(1)$ proportion of all minimal k-SAT formulae are unate.*

As suggested by Bollobás, Brightwell, and Leader [5], these results open doors to a *theory of random k-SAT functions*. For example, the theorems imply that a typical k-SAT function admits a unique minimal k-SAT formula, and furthermore the formula has $(1/2+o(1))\binom{n}{k}$ clauses. Note that our model is very different from that of *random k-SAT formulae* where clauses are added at random (e.g., the recent breakthrough on the satisfiability conjecture [6]). Rather, Theorem 1.2 concerns a random k-SAT formula *conditioned* on minimality. In this light, our results are analogous to the theory

of Erdős–Rényi random graphs $G(n,p)$ in the *dense setting* with $p = 1/2$, which corresponds to counting. It would be interesting to further study the behavior of sparser random k-SAT formulae where each clause is introduced with some probability p_n decaying with n, and conditioned on the minimality of the formula. This leads to a rich source of problems on thresholds and typical structures and potentially interesting phenomena.

In a different direction, we can also ask what happens when the width k of k-SAT is allowed to grow with n. Since our proof techniques are quantitative, they allow k to grow rather slowly with n (we have not worked out the precise dependence, which is a function of the specific choice of hypergraph container lemma applied in our argumentation; our results likely continue to hold for $k = O(\log\log n)$). A bold conjecture by Bollobás and Brightwell [4] states that for all $k \le (1/2-\epsilon)n$ with any fixed $\epsilon > 0$, the number of k-SAT functions on n Boolean variables is $2^{(1+o(1))\binom{n}{k}}$. This conjecture appears to be outside the scope of current methods. We leave the following related question as an open problem: how quickly can $k = k(n)$ grow with n so that a typical k-SAT function on n Boolean variables is unate? Curiously, a completely different behavior emerges for $k > n/2$, as observed in [4], and this regime also has a lot of interesting open problems.

This paper is the second half of a two-part work establishing the above results. The first part [7] (by Dong, Mani, and Zhao) reduces the problem, for each k, to a Turán problem on partially directed hypergraphs.[1] In this paper we solve this Turán problem completely for every k.

The reduction in [7] applies the hypergraph container method [2, 3, 11], a major recent development in combinatorics. Our solution to the Turán problem is partly inspired by the method of flag algebras, introduced by Razborov [10], which reduces graph density inequalities to sums of squares.

2 PARTIALLY DIRECTED HYPERGRAPHS

A *partially directed k-graph* (abbreviated as *k-PDG*) is formed by starting with a k-uniform hypergraph (i.e., whose edges are k-element subsets of vertices), and then for each edge, either (i) leaving it as an *undirected edge* or (ii) converting it to a *directed edge* by choosing a special vertex in the edge (we say that the edge *points to* or is *directed toward* this special vertex). We notate a directed edge by putting a $\vee$ on top of the special vertex. An example of a 2-PDG (that we call $\vec{T}_2$) is illustrated below:

$\vec{T}_2 :=$ 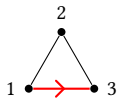 edges $= \{12, 1\check{3}, 23\}$.

Given a pair of k-PDGs, $\vec{H}$ and $\vec{G}$, we say that $\vec{H}$ is a *subgraph* of $\vec{G}$ if one can obtain $\vec{H}$ from $\vec{G}$ by a combination of (1) removing vertices, (2) removing edges, and (3) removing the orientation of some edges. As examples, the left 2-PDG below contains $\vec{T}_2$ as a subgraph, and the right does not contain $\vec{T}_2$ as a subgraph.

We say that a k-PDG $\vec{G}$ is *$\vec{H}$-free* if $\vec{G}$ does not contain $\vec{H}$ as a subgraph.

An important k-PDG for us is denoted by $\vec{T}_k$, formed (for each $k \ge 3$) by starting with $\vec{T}_2$ and then adding $k-2$ common vertices to all three edges, e.g.,

$\vec{T}_3 =$ 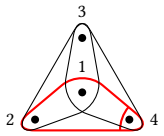 edges $= \{123, 12\check{4}, 134\}$,

and

$\vec{T}_4 =$ 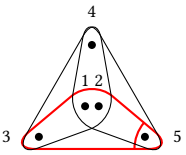 edges $= \{1234, 123\check{5}, 1245\}$.

We prove the following statement, which, by the reduction in [7], implies Theorems 1.1 and 1.2.

Theorem 2.1. *Let $k \ge 2$ be a positive integer. There exists some $\theta > \log_2 3$ such that every n-vertex $\vec{T}_k$-free k-PDG with $\alpha\binom{n}{k}$ undirected edges and $\beta\binom{n}{k}$ directed edges satisfies*

$$\alpha + \theta\beta \le 1 + o_{n\to\infty}(1).$$

Remark 2.2 (A sketch of the reduction). Let us sketch this reduction and defer to [7] for details. The reader is welcomed to skip this remark.

We wish to prove Theorem 1.2 and count minimal k-SAT formulae. We relax the minimality condition and actually upper bound the number of k-SAT formulae that avoid some short certificate of non-minimality. There is a useful analogy between the latter problem and counting triangle-free graphs on n vertices. The hypergraph container method allows us to reduce such an asymptotic enumeration problem to a Turán problem (along with supersaturation, which is automatic in the dense setting).

For counting minimal k-SAT formulae, containers are themselves k-SAT formulae (but not necessarily minimal). We say that a formula is *simple* if every k-element subset of variables supports at most one clause (e.g., $x_1x_2x_3$ and $x_1x_2\overline{x_3}$ do not both appear as clauses). We say that a formula is *semisimple* if every k-element subset of variables supports either at most one clause, or exactly two clauses differing by negation at exactly one variable (e.g., $\{x_1\overline{x_2}x_3, x_1\overline{x_2x_3}\}$, but not $\{x_1x_2x_3, x_1\overline{x_2x_3}\}$).

We wish to upper bound the number of simple minimal k-SAT formulae (it turns out there are negligibly many non-simple ones). The container theorem produces a small collection of container formulae, such that each simple minimal k-SAT formula is contained in some container. These containers satisfy additional properties. For example, each container is, up to removing $o(n^k)$ clauses, a semisimple formula. We can convert a semisimple k-SAT formula (arising from a container) to a k-PDG by converting each clause to an edge, and whenever two clauses are supported on the same set of

[1] The first paper [7] also solves the Turán problem for $k = 4$ using a different method from this work. A separate brute-force computer search by Mani and Yu yielded the $k = 5$ case, as documented in the appendix of the arXiv version 3 of [7] at arXiv:2107.09233v3.

variables, we direct the edge towards the unique differing variable (e.g., $\{x_1\overline{x_2}x_3, x_1\overline{x_2x_3}\}$ becomes $12\breve{3}$). The $\vec{T}_k$-free condition for k-PDGs corresponds to a certain short certificate of non-minimality in the semisimple formula (we actually need to consider a 2-blow-up in this step, but let us skip that discussion here).

Consider a container G, so that up to removing $o(n^k)$ clauses from G, we obtain a semisimple formula G'. Among all k-element subsets of variables (out of $\binom{n}{k}$ total), suppose $\alpha\binom{n}{k}$ of them support exactly one clause in G', and $\beta\binom{n}{k}$ of them support exactly two clauses in G'. To choose a simple subformula of G, we have two choices for each of the first type (i.e., whether to include or not), three choices for each of the second type (i.e., keep one of the two clauses, or discard both), and $2^{o(n^k)}$ choices for whether to include each clause in $G \setminus G'$. Then the number of simple subformulae of G is $2^{\alpha\binom{n}{k}}3^{\beta\binom{n}{k}}2^{o(n^k)}$. The container theorem guarantees us that the number of containers is $2^{o(n^k)}$. Thus, we find that the number of simple minimal formulae is at most $2^{\alpha\binom{n}{k}}3^{\beta\binom{n}{k}}2^{o(n^k)}$, and we wish to obtain an upper bound of $2^{(1+o(1))\binom{n}{k}}$, i.e., we want $\alpha + (\log_2 3)\beta \le 1 + o(1)$. This is the reason for the appearance of $\log_2 3$ in Theorem 2.1. This sketches the proof that the number of minimal k-SAT formulae is $2^{(1+o(1))\binom{n}{k}}$. To get the more precise count of $(1+o(1))2^{n+\binom{n}{k}}$, one needs an additional stability argument to handle the case when α is close to 1. The stability argument was introduced in [9] for 3-SAT, and extended in [7] to general k-SAT.

3 SOLUTION OF THE TURÁN PROBLEM

We only consider $k \ge 4$ from now on, as [7] already gives simple proofs of Theorem 2.1 for $k = 2$ and $k = 3$. Here is the key lemma.

Lemma 3.1. *Let $k \ge 4$ be a positive integer. There exist $\theta > \log_2 3$ and $a, b \in \mathbb{R}$ (depending on k) such that the following holds. Suppose $\vec{F}$ is a $\vec{T}_k$-free k-PDG on $k+1$ vertices. Let $x_1, \dots, x_{k-1}, y, z$ be a permutation of vertices of $\vec{F}$ chosen uniformly at random.*

Define $x := x_1 \dots x_{k-1}$, let xy denote the event that $x_1 \dots x_{k-1}y$ forms an undirected edge in $\vec{F}$, let $x\breve{y}$ denote the event that $x_1 \dots x_{k-1}\breve{y}$ is a directed edge in $\vec{F}$, and let $\underline{xy}$ denote the event that there is no edge with vertices $x_1, \dots, x_{k-1}, y$ in $\vec{F}$. Then, we have that

$$\mathbb{P}(xy) + k\theta\mathbb{P}(x\breve{y}) + a^2\mathbb{P}(\underline{xy} \wedge \underline{xz}) - 2ab\mathbb{P}(\underline{xy} \wedge x\breve{z}) + b^2\mathbb{P}(x\breve{y} \wedge x\breve{z}) \le 1.$$

Remark 3.2. The proof gives $\theta = 1 + \frac{1}{\sqrt{2}} \ge 1.707 > 1.585 > \log_2 3$.

Proof of Theorem 2.1 for $k \ge 4$ using Lemma 3.1. Let $\vec{H}$ be an n-vertex $\vec{T}_k$-free k-PDG with $\alpha\binom{n}{k}$ undirected edges and $\beta\binom{n}{k}$ directed edges.

Let $x_1, \dots, x_{k-1}, y, z$ be vertices of $\vec{H}$ chosen without replacement uniformly at random. Applying Lemma 3.1 (by first conditioning on the set of $k+1$ randomly selected vertices) and using the notation of the lemma,

$$\mathbb{P}(xy) + k\theta\mathbb{P}(x\breve{y}) + a^2\mathbb{P}(\underline{xy} \wedge \underline{xz}) - 2ab\mathbb{P}(\underline{xy} \wedge x\breve{z}) + b^2\mathbb{P}(x\breve{y} \wedge x\breve{z}) \le 1.$$

Note that $\mathbb{P}(xy) = \alpha$ and $k\mathbb{P}(x\breve{y}) = \beta$. It remains to show that

$$a^2\mathbb{P}(\underline{xy} \wedge \underline{xz}) - 2ab\mathbb{P}(\underline{xy} \wedge x\breve{z}) + b^2\mathbb{P}(x\breve{y} \wedge x\breve{z}) \ge -o(1).$$

We will show that this inequality holds for every fixed choice of x. Conditioned on x, we see that y and z are uniformly chosen vertices without replacement in $V(\vec{F}) \setminus \{x_1, \dots, x_k\}$. When n is large, this is not much different than with replacement, in which case y and z would be conditionally independent given x. In particular, $\mathbb{P}(\underline{xy} \wedge \underline{xz}|x) = \mathbb{P}(\underline{xy}|x)^2 + o(1)$, and $\mathbb{P}(\underline{xy} \wedge x\breve{z}|x) = \mathbb{P}(\underline{xy}|x)\mathbb{P}(x\breve{y}|x) + o(1)$, and $\mathbb{P}(x\breve{y} \wedge x\breve{z}|x) = \mathbb{P}(x\breve{y}|x)^2 + o(1)$. Then we can prove the above displayed inequality, conditioned on any x, by observing that

$$\begin{aligned} &a^2\mathbb{P}(\underline{xy} \wedge \underline{xz}|x) - 2ab\mathbb{P}(\underline{xy} \wedge x\breve{z}|x) + b^2\mathbb{P}(x\breve{y}, x\breve{z}|x) \\ &\quad = a^2\mathbb{P}(\underline{xy}|x)^2 - 2ab\mathbb{P}(\underline{xy}|x)\mathbb{P}(x\breve{y}|x) + b^2\mathbb{P}(x\breve{y}|x)^2 - o(1) \\ &\quad = \left(a\mathbb{P}(\underline{xy}|x) - b\mathbb{P}(x\breve{y}|x)\right)^2 - o(1) \ge -o(1). \end{aligned}$$

□

Finally, it remains to prove Lemma 3.1.

Proof of Lemma 3.1. Let $\vec{H}$ be a $\vec{T}_k$-free k-PDG on $k+1$ vertices. Construct the following associated digraph D on the same vertex set as $\vec{H}$:

(1) For every directed edge in $\vec{H}$ missing vertex i and directed towards vertex j, add the directed edge $i \to j$ in D.
(2) For every undirected edge in $\vec{H}$ missing vertex i, add the loop $i \to i$ in D.

Notice that every vertex in D has out-degree at most 1. Moreover, since $\vec{H}$ does not contain $\vec{T}_k$ as a subgraph, D is free of the following forbidden pattern:

- *Forbidden pattern:* $i_1 \to i_2 \to *$ and $i_3 \to *$ for three distinct vertices i_1, i_2, i_3, where $*$ can be any vertex (the two $*$'s do not have to be the same).

Indeed, if D contains the above pattern, then $\vec{H}$ has an edge missing i_2, an edge missing i_3, and an edge containing both i_2 and i_3 that is directed at i_2. These three edges contain $\vec{T}_k$ as a subgraph.

It is easy to deduce the following exhaustive classification of all digraphs D on $k+1$ vertices with out-degree at most 1 at every vertex and without the above forbidden pattern.

(A) $i \to j \to i$ for distinct i, j. No other edges.
(B) $i \to j \to j$ for distinct i, j. No other edges.
(C) $i_1 \to i_2 \to i_3$ for distinct i_1, i_2, i_3. No other edges.
(D) No vertex has both positive in-degree and positive out-degree except isolated self-loops $i \to i$.
(E) A number of loops $i_1 \to i_1, \dots, i_m \to i_m$.

Write $i \to \varnothing$ to denote that i has out-degree 0 in D. The inequality in Lemma 3.1 then translates to

$$\begin{aligned} &\mathbb{P}(z \to z) + k\theta \cdot \mathbb{P}(z \to y) + a^2\mathbb{P}(y \to \varnothing \wedge z \to \varnothing) \\ &\quad - 2ab\mathbb{P}(y \to z \wedge z \to \varnothing) + b^2\mathbb{P}(y \to z \wedge z \to y) \le 1. \end{aligned} \tag{3.1}$$

Let u and d denote the number of undirected and directed edges in $\vec{H}$, respectively. We calculate the terms in inequality (3.1):

$$\mathbb{P}(z \to z) = \frac{u}{k+1}, \quad \mathbb{P}(z \to y) = \frac{d}{k(k+1)},$$
$$\mathbb{P}(y \to \varnothing \wedge z \to \varnothing) = \frac{(k+1-u-d)(k-u-d)}{(k+1)k},$$
$$\mathbb{P}(y \to z \wedge z \to \varnothing) = \begin{cases} 0 & \text{in cases (A), (B) and (E),} \\ \frac{1}{(k+1)k} & \text{in case (C),} \\ \frac{d}{(k+1)k} & \text{in case (D),} \end{cases}$$
$$\mathbb{P}(y \to z \wedge z \to y) = \begin{cases} \frac{2}{(k+1)k} & \text{in case (A),} \\ 0 & \text{in cases (B), (C), (D) and (E).} \end{cases}$$

Then (3.1) reduces to simpler inequalities in each of the four cases:

(A) When our digraph is of the form $i \to j \to i$ for distinct i, j with no other edges, (3.1) simplifies to

$$\frac{2\theta}{k+1} + \frac{(k-1)(k-2)}{(k+1)k}a^2 + \frac{2}{(k+1)k}b^2 \le 1. \tag{3.2}$$

(B) Since $\theta \ge 1$, the associated inequality in this case is implied by (A); we note the inequality below, but do not need to consider this case separately:

$$\frac{1}{k+1} + \frac{\theta}{k+1} + \frac{(k-1)(k-2)}{(k+1)k}a^2 \le 1.$$

(C) The inequality in this case is also implied by (A) (we will end up choosing $a, b > 0$) so we also need not consider it separately:

$$\frac{2\theta}{k+1} + \frac{(k-1)(k-2)}{k(k+1)}a^2 - \frac{2}{k(k+1)}ab \le 1.$$

(D) In this case, if $d = 0$ then $u \le k+1$, and if $d \ge 1$ then $u+d \le k$. The inequality reduces to

$$\frac{u}{k+1} + \frac{\theta d}{k+1} + \frac{(k+1-u-d)(k-u-d)}{(k+1)k}a^2 - \frac{2d}{(k+1)k}ab \le 1. \tag{3.3}$$

(E) In this case, since $\vec{H}$ has no directed edge, (3.1) simplifies to

$$\frac{u}{k+1} + \frac{(k+1-u)(k-u)}{(k+1)k}a^2 \le 1, \tag{3.4}$$

which holds as long as $0 \le a \le 1$.

Thus it remains to show that we can choose parameters $\theta > \log_2 3$ and $a, b > 0$ to make the inequalities implied by (A), (D) and (E) true. We choose

$$\theta = 1 + \frac{1}{\sqrt{2}} > 1.707, \text{ satisfying } 2\theta^2 - 4\theta + 1 = 0$$

$$a = \frac{1}{\sqrt{2}}, \quad b = \frac{k(\theta-1)-1}{\sqrt{2}}.$$

(3.4) immediately follows from the fact that $0 \le a \le 1$.

We can verify (3.2), establishing the desired inequality for case (A) (and thus cases (B) and (C)) via a direct substitution:

$$\begin{aligned} &\frac{2\theta}{k+1} + \frac{(k-1)(k-2)}{(k+1)k}a^2 + \frac{2}{(k+1)k}b^2 - 1 \\ &\quad = \frac{2\theta}{k+1} + \frac{(k-1)(k-2)}{2(k+1)k} + \frac{(k(\theta-1)-1)^2}{(k+1)k} - 1 \\ &\quad = \frac{4k\theta + (k-1)(k-2) + 2(k(\theta-1)-1)^2 - 2(k+1)k}{2(k+1)k} \\ &\quad = \frac{4-k+k^2(2\theta^2-4\theta+1)}{2(k+1)k} = \frac{4-k}{2(k+1)k} \le 0. \end{aligned}$$

To verify (3.3) for (D), observe that the left-hand side of (3.3) is non-decreasing in u. Indeed, replacing u by $u+1$ increases the expression by at least

$$\frac{1}{k+1} - \frac{(k+1)k - k(k-1)}{(k+1)k}a^2 = \frac{1}{k+1} - \frac{2}{k+1}a^2 = 0.$$

Consequently, it remains to verify case (D) when u is as large as possible, meaning $u + d \in \{k, k+1\}$, which makes the third term in (3.3) zero. The left-hand side of (3.3) becomes

$$\frac{u}{k+1} + \frac{\theta d}{k+1} - \frac{2d}{(k+1)k}ab.$$

If $d = 0$, then $u = k+1$, and the inequality clearly holds. Otherwise, $d + u = k$. Since the above expression is linear in u (or equivalently, linear in d), the maximum is attained at one of endpoints $(u, d) = (k, 0)$ or $(0, k)$. The only nontrivial situation to check is $(u, d) = (0, k)$, in which case the above expression is

$$\frac{k\theta}{k+1} - \frac{2}{k+1}ab = \frac{k\theta}{k+1} - \frac{k(\theta-1)-1}{k+1} = 1.$$

Therefore, (3.1) always holds. □

ACKNOWLEDGMENTS

The collaboration that led to this paper started when Zhao presented [7] at an Oberwolfach workshop in 2022, with Balogh also in attendance. This work used the computing resources at the Center for Computational Mathematics, University of Colorado Denver, including the Alderaan cluster, supported by the National Science Foundation award OAC-2019089.

Balogh was supported in part by NSF grants DMS-1764123 and RTG DMS-1937241, FRG DMS-2152488, the Arnold O. Beckman Research Award (UIUC Campus Research Board RB 22000), the Langan Scholar Fund (UIUC). Lidický was supported in part by NSF grant FRG DMS-2152490 and Scott Hanna fellowship. Mani was supported by the NSF Graduate Research Fellowship Program and a Hertz Graduate Fellowship. Zhao was supported in part by NSF CAREER award DMS-2044606, a Sloan Research Fellowship, and the MIT Solomon Buchsbaum Fund.

REFERENCES

[1] Peter Allen. 2007. Almost every 2-SAT function is unate. *Israel J. Math.* 161 (2007), 311–346. https://doi.org/10.1007/s11856-007-0081-z

[2] József Balogh, Robert Morris, and Wojciech Samotij. 2015. Independent sets in hypergraphs. *J. Amer. Math. Soc.* 28, 3 (2015), 669–709. https://doi.org/10.1090/S0894-0347-2014-00816-X

[3] József Balogh, Robert Morris, and Wojciech Samotij. 2018. The method of hypergraph containers. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures.* World Sci. Publ., 3059–3092. https://doi.org/10.1142/9789813272880_0172

[4] Béla Bollobás and Graham R. Brightwell. 2003. The number of k-SAT functions. *Random Structures Algorithms* 22, 3 (2003), 227–247. https://doi.org/10.1002/rsa.10079
[5] Béla Bollobás, Graham R. Brightwell, and Imre Leader. 2003. The number of 2-SAT functions. *Israel J. Math.* 133 (2003), 45–60. https://doi.org/10.1007/BF02773060
[6] Jian Ding, Allan Sly, and Nike Sun. 2022. Proof of the satisfiability conjecture for large k. *Ann. of Math.* 196, 1 (2022), 1–388. https://doi.org/10.4007/annals.2022.196.1.1
[7] Dingding Dong, Nitya Mani, and Yufei Zhao. 2022. Enumerating k-SAT functions. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2141–2184. https://doi.org/10.1137/1.9781611977073.85
[8] L. Ilinca and J. Kahn. 2009. On the number of 2-SAT functions. *Combin. Probab. Comput.* 18, 5 (2009), 749–764. https://doi.org/10.1017/S096354830900995X
[9] L. Ilinca and J. Kahn. 2012. The number of 3-SAT functions. *Israel J. Math.* 192, 2 (2012), 869–919. https://doi.org/10.1007/s11856-012-0056-6
[10] Alexander A. Razborov. 2007. Flag algebras. *J. Symbolic Logic* 72, 4 (2007), 1239–1282. https://doi.org/10.2178/jsl/1203350785
[11] David Saxton and Andrew Thomason. 2015. Hypergraph containers. *Invent. Math.* 201, 3 (2015), 925–992. https://doi.org/10.1007/s00222-014-0562-8

Received 2022-11-07; accepted 2023-02-06