Hardness of Clique Approximation for Monotone Circuits

Jarosław Błasiok
ETH Zurich, Switzerland

Linus Meierhöfer
ETH Zurich, Switzerland

Abstract

We consider a problem of approximating the size of the largest clique in a graph, using a monotone circuit. Concretely, we focus on distinguishing a random Erdős–Rényi graph $\mathcal{G}_{n,p}$, with $p = n^{-\frac{2}{\alpha-1}}$ chosen st. with high probability it does not even contain an α-clique, from a random clique on β vertices (where $\alpha \le \beta$). Using the approximation method of Razborov, Alon and Boppana showed in their influential work in 1987 that as long as $\sqrt{\alpha}\beta < n^{1-\delta}/\log n$, this problem requires a monotone circuit of size $n^{\Omega(\delta\sqrt{\alpha})}$, implying a lower bound of $2^{\tilde{\Omega}(n^{1/3})}$ for the exact version of the problem $\textsc{Clique}_k$ when $k \approx n^{2/3}$. Recently, Cavalar, Kumar, and Rossman improved their result by showing a tight lower bound $n^{\Omega(k)}$, in a limited range $k \le n^{1/3}$, implying a comparable $2^{\tilde{\Omega}(n^{1/3})}$ lower bound after choosing the largest admissible k.

We combine the ideas of Cavalar, Kumar and Rossman with recent breakthrough results on sunflower conjecture by Alweiss, Lovett, Wu, and Zhang to show that as long as $\alpha\beta < n^{1-\delta}/\log n$, any monotone circuit rejecting $\mathcal{G}_{n,p}$ graph while accepting a β-clique needs to have size at least $n^{\Omega(\delta^2\alpha)}$; this implies a stronger $2^{\tilde{\Omega}(\sqrt{n})}$ lower bound for the unrestricted version of the problem.

We complement this result with a construction of an explicit monotone circuit of size $O(n^{\delta^2\alpha/2})$ which rejects $\mathcal{G}_{n,p}$, and accepts any graph containing β-clique whenever $\beta > n^{1-\delta}$. In particular, those two theorems give a precise characterization of the smallest β-clique that can be distinguished from $\mathcal{G}_{n,1/2}$: when $\beta > n/2^{C\sqrt{\log n}}$, there is a polynomial-size circuit that solves it, while for $\beta < n/2^{\omega(\sqrt{\log n})}$ every circuit needs size $n^{\omega(1)}$.

2012 ACM Subject Classification Theory of computation → Circuit complexity

Keywords and phrases circuit lower bounds, monotone circuits, sunflower conjecture

Digital Object Identifier 10.4230/LIPIcs.CCC.2025.4

Related Version *Full Version*: https://arxiv.org/abs/2501.09545

1 Introduction

Monotone circuits form a restricted computation model, where computation is performed by a directed acyclic graph, with input nodes (of in-degree 0) labeled by the input variables and internal nodes (*gates*) labeled $\wedge$ or $\vee$ (computing logical *and*, or logical *or* respectively, of input wires). It is not difficult to see that every monotone function on n binary variables $f : \{0,1\}^n \to \{0,1\}$ (and only monotone functions), can be computed by such a circuit, and by a simple counting argument one can show that a random monotone function requires a monotone circuit of size $2^{\Omega(n)}$.

Remarkably, in contrast with more general Boolean circuits (where the negation gate $\neg$ is also allowed), since the breakthrough work of Razborov [19], there are known unconditional super-linear lower bounds for the size of monotone circuit computing explicit monotone functions. Concretely, Razborov showed that for any $k \le \log(n)$, the function $\textsc{Clique}_k : \{0,1\}^{\binom{n}{2}} \to \{0,1\}$ – interpreting its input as an adjecency matrix of a graph G, and outputting 1 if and only if said graph contains a clique on k vertices – requires a monotone circuit of size

$n^{\Omega(k)}$. Putting $k = \log n$, this gives a quasipolynomial lower bound $2^{\Omega(\log^2(n))}$; an analogous result for boolean circuits including negation would imply $P \neq NP$ and seems to be far beyond the reach of current techniques, almost 40 years later.

To prove his lower bound, Razborov used the sunflower lemma of Erdős and Rado [10], bringing it to the attention of the theoretical computer science community. In the highly influential follow-up work, Alon and Boppana [1] further utilized the approximation method introduced by Razborov. They showed a $n^{\Omega(\sqrt{k})}$ lower bound for the $\textsc{Clique}_k$ problem when $k \leq n^{2/3}/\log n$ – implying a $2^{\tilde{\Omega}(n^{1/3})}$ lower bound at the optimal value of $k = n^{2/3}/\log n$ – a result which has since become a landmark of circuit complexity. Interestingly and less commonly known, in the same paper, they also showed a stronger inapproximability result: Any monotone circuit that rejects every graph without even an α-clique and accepts every graph that has a β-clique (where $\alpha \leq \beta$) needs to have a size at least $n^{\Omega(\delta\sqrt{\alpha})}$ as long as $\sqrt{\alpha}\beta < n^{1-\delta}/\log n$. In fact, their result holds in the average case – they show that it is hard to reject a random $(\alpha - 1)$-partite graph while accepting a clique on β random vertices.

Since then, several other techniques for showing monotone circuit lower bounds for specific problems have been introduced over almost four decades (e.g., [3, 4, 15, 11, 12, 7]), but the result of Alon and Boppana stood as essentially the best-known lower bound for the clique problem.

Very recently, the breakthrough result of Alweiss, Lovett, Wu, and Zhang [2] made a major step towards resolving the sunflower conjecture – conjecture about the "right" quantitative dependency in the sunflower lemma – and was followed in short succession by a sequence of further improvements [5, 17, 23, 14, 18]. Particularly noteworthy is the exposition by Rao [18], presenting not only the improved sunflower bounds but also the resolution of the Kahn-Kalai conjecture in just a few pages.

Following up on these results, Cavalar, Kumar, and Rossman [8], showed how to utilize the new sunflower bounds to prove the first $2^{\Omega(\sqrt{n})}$ lower bound for any explicit monotone function (improving the earlier $2^{\tilde{\Omega}(n^{1/3})}$ lower bound), specifically for the Harnik-Raz function [13]. In the same paper, they also gave the first substantial improvement over the Alon-Boppana bound for the clique problem: they showed that as long as $k \leq n^{1/3-\delta}$ the $\textsc{Clique}_k$ problem requires $n^{\Omega(k)}$ circuit size – a tight lower bound, but in a further restricted range of parameters; choosing the optimal $k \approx n^{1/3}$, this leads to the same $2^{\tilde{\Omega}(n^{1/3})}$ lower bound for Clique. As it turns out, their result for the clique did not use the new sunflower bounds, and, in fact, it was unclear how to combine those techniques.

1.1 Our results

We prove a lower bound on the size of monotone circuits solving the promise problem $\mathrm{GAP}_{\alpha,\beta}(n)$.

▶ **Definition 1.** *For integers $n, \alpha, \beta \in \mathbb{N}$ with $\alpha < \beta$, define the promise problem $\mathrm{GAP}_{\alpha,\beta}(n)$, where the two promise subsets $\mathcal{A}, \mathcal{B}$ are defined as*

$$\mathcal{A} := \{G \in \{0,1\}^{\binom{n}{2}} \mid G \text{ does not contain an } \alpha\text{-clique}\},$$

and

$$\mathcal{B} := \{G \in \{0,1\}^{\binom{n}{2}} \mid G \text{ contains an } \beta\text{-clique}\}.$$

To prove a lower bound on $\mathrm{GAP}_{\alpha,\beta}(n)$, we introduce an associated strengthened distributional problem $\mathcal{D}(\mathcal{T}_\alpha^-, \mathcal{T}_\beta^+)$ to distinguish two graph distributions $\mathcal{T}_\alpha^-$ and $\mathcal{T}_\beta^+$.

▶ **Definition 2.** *For an integer n, α, β with $\alpha \le \beta$, let $\mathcal{T}^+_\beta$ be a uniformly random distribution of isolated β-cliques on n vertices. Let $\mathcal{T}^-_\alpha \sim \mathcal{G}_{n,p}$ the Erdős–Rényi random graph distribution on n vertices with probability parameter $p = n^{-2/(\alpha-1)}$.*

The problem $\mathcal{D}(\mathcal{T}^-_\alpha, \mathcal{T}^+_\beta)$ is to distinguish those two distributions – i.e. given on input a graph G drawn either from $\mathcal{T}^-_\alpha$ or $\mathcal{T}^+_\beta$ we want to reject with probability at least $2/3$ in the former case and accept with probability at least $2/3$ in the latter.

Note that α is chosen so that $\mathcal{D}(\mathcal{T}^-_\alpha, \mathcal{T}^+_\beta)$ is associated with $\mathrm{GAP}_{\alpha,\beta}(n)$ as the probability that a graph drawn from $\mathcal{T}^-_\alpha$ contains a clique of size α is small – a simple union bound can be used to argue that $\Pr_{G\in\mathcal{T}^-_\alpha}[G \text{ contains } \alpha\text{-clique}] < 1/4$ for $\alpha \ge 4$. As such, any circuit solving the problem $\mathrm{GAP}_{\alpha,\beta}$ can distinguish $\mathcal{T}^-_\alpha$ from $\mathcal{T}^+_\beta$, and conversely, a lower bound for $\mathcal{D}(\mathcal{T}^-_\alpha, \mathcal{T}^+_\beta)$ implies a lower bound for $\mathrm{GAP}_{\alpha,\beta}$.

Our main results are lower and upper bounds on the monotone complexity of $\mathcal{D}(\mathcal{T}^-_\alpha, \mathcal{T}^+_\beta)$, captured in the following two theorems.

▶ **Theorem 3.** *Let $\alpha\beta < n^{1-\delta}/\log(n)$. Then there exists no monotone circuit solving $\mathcal{D}(\mathcal{T}^-_\alpha, \mathcal{T}^+_\beta)$ with size less than $n^{\Omega(\delta^2\alpha)}$.*

▶ **Theorem 4.** *Let $\beta \ge n^{1-\delta}$. Then there exists a monotone circuit C on $\binom{n}{2}$ inputs, solving $\mathcal{D}(\mathcal{T}^-_\alpha, \mathcal{T}^+_\beta)$ with size$(C) = \mathcal{O}(n^{\delta^2\alpha/2})$.*

We want to emphasize a special case of $\alpha = 2\log_2 n + 1$ (or, equivalently, the negative distribution $\mathcal{T}^-$ being an Erdős–Rényi graph $\mathcal{G}_{n,p}$ with $p = 1/2$). In this case, according to Theorem 4, we can choose $\delta = \sqrt{1/\log n}$ to get a polynomial-size circuit, rejecting $\mathcal{G}_{n,p}$ yet accepting a clique of size $\beta = n/2^{C\sqrt{\log n}}$. On the other hand, Theorem 3 states that if we wanted to reject $\mathcal{G}_{n,1/2}$ and accept a clique of size β, where $\beta = n/2^{\omega(\sqrt{\log n})}$, we would need a circuit of size $n^{\omega(1)}$.

As such, our two theorems provide a near-tight characterization of the power of polynomial-size monotone circuits to distinguish the $\mathcal{G}_{n,1/2}$ graph from a large clique.

Another interesting corollary of Theorem 3, is obtained by setting $\beta = \alpha$, and taking α as large as possible while satisfying the restriction $\alpha^2 < n^{1-\delta}/\log n$, in order to obtain a strongest possible lower bound for the clique problem.

▶ **Corollary 5.** *For any $k \le n^{1/2-\delta}/\log^2 n$, the monotone complexity of the $\textsc{Clique}_k$ is $n^{\Omega(\delta^2 k)}$. In particular monotone complexity of the $\textsc{Clique}$ problem is $2^{\tilde{\Omega}(\sqrt{n})}$.*

Note that no truly exponential lower bounds for the monotone circuit size are known for *any* explicit monotone problem. A lower bound of $2^{\tilde{\Omega}(\sqrt{n})}$ (where n is the size of the input) was shown for the arguably less-natural Harnik-Raz function [13] in [8], breaking the long-standing barrier of $2^{\tilde{\Omega}(n^{1/3})}$ – and using the same breakthroughs in sunflower lemmas that are crucial in our improvement.

In our case, the input is an adjacency matrix of a graph on n vertices hence the input has $m = \Theta(n^2)$ bits; our lower bound, in terms of the input size, is then $2^{\Omega(m^{1/4})}$ – the exponent is still by a factor of two worse than in the strongest known lower bound for any explicit monotone function.

Concurrent and independent work

Independently to our results, Suzanna de Rezende and Marc Vinyals proved a comparable monotone circuit lower bound $2^{\Omega(\sqrt{n})}$ for the clique problem [9]. Specifically, they show a lower bound of $n^{\Omega(k)}$for the problem of distinguishing a $k+1$-clique from a k-colorable graph

as long as $k \leq n^{1/2-\varepsilon}$. Their does not use the approximation method, instead leveraging a combinatorial reduction together with a new query-to-communication complexity lifting theorem, and known relationship between specific communication complexity lower bounds and monotone circuit lower bounds. Interestingly, the key step in their improved lifting theorem also leverged the [2] breakthrough in sunflower bounds.

1.2 Approximation method

Let us briefly recall the idea behind the approximation method for showing monotone circuit lower bounds, introduced by Razborov [19], and then utilized by [1, 8].

In order to show a lower bound for a size of a monotone circuit distinguishing two specific distributions $\mathcal{T}^+$ and $\mathcal{T}^-$, we proceed by inductively (gate-by-gate) approximating any given circuit C by a simpler circuit $\hat{C}$ (from some class of simple circuits).

If we can show that

1. Every simple circuit fails to distinguish between $\mathcal{T}^+$ and $\mathcal{T}^-$ – that is, for all simple $\hat{C}$, we have $\mathbb{E}_{x\sim\mathcal{T}^+}\hat{C}(x) - \mathbb{E}_{x\sim\mathcal{T}^-}\hat{C}(x) \leq o(1)$.
2. In each gate, by applying the approximation, we introduce error at most δ on both positive and negative distribution.

That implies a lower bound $\Omega(\delta^{-1})$ for the size of the smallest monotone circuit C distinguishing those two distributions.

For the clique problem (say, distinguishing random $\mathcal{G}_{n,p}$ graph which likely does not even have a clique of size α, from a uniformly random clique of size $\beta \geq \alpha$), following [19, 1, 8], the "simple" circuits we use to approximate a given circuit are just small DNF formulas, specifically conjunctions of clique indicators – that is circuits of form

$$\bigvee_i \mathcal{K}_{A_i},$$

where

$$\mathcal{K}_{A_i} := \bigwedge_{\{u,v\}\in A_i} x_{uv}.$$

Moreover, a circuit is simple if the size of each clique indicator is bounded by c (we will eventually choose $c := \delta\alpha$), and the number of clique indicators of each size $\ell \leq c$ is appropriately bounded.

At a given gate, we wish to approximate a conjunction or a disjunction of two such simple circuits again by a simple circuit. By applying de Morgan laws in the conjunction case, we can transform the circuit again into DNF (without introducing any error).

Then, we repeat the following three steps, transforming the DNF obtained this way into a simple one.

1. We replace all conjunctions $\mathcal{K}_{A_i} \wedge \mathcal{K}_{A_j}$ by clique indicators $\mathcal{K}_{A_i\cup A_j}$.
2. As long as there is a family of cliques on sets $A_{i_1}, \ldots A_{i_k}$ having a specific combinatorial structure, we replace all of them by a clique on the set C, where C is the intersection of A_{i_j}.
3. We remove all clique indicators $\mathcal{K}_{A_i}$ for A_i larger than threshold c.

As it turns out, step 1 does not introduce any error on both the positive and negative distributions.

Step 3 can introduce error only on the positive distribution, which can be bounded by union bound – any given large clique indicator is unlikely to be satisfied by a large random clique. The number of those large clique indicators we discard is bounded (otherwise, we would have been able to apply step 2).

Finally, step 2 is the crux of the argument – we need to show that if the number of indicators of a given size is too large, then there always is a subset of those indicators that can be replaced by its intersection C (the *core*), without introducing too much error on the negative distribution (clearly we will not introduce any error on the positive distribution in this case).

In the Razborov's proof and some presentations of the Alon-Boppana proof (see for example [16, Chapter 9]), one can focus on finding a sunflower consisting of many sets among the set system $\{A_i\}$ (a sunflower is a family of sets for which all pairwise intersections are the same, see Section 1.3), and show that replacing a large sunflower by its core introduces only small error on the negative distribution. This, together with the Erdős-Rado result that any large enough family of bounded sets contains a large sunflower, gives an upper bound on the number of clique indicators of size ℓ for a "simple circuit" (one on which the step 2 can no longer be applied), and hence yields a complete proof of a monotone circuit lower bound for the clique problem. Specifically, going carefully through the calculations, one could show this way a lower bound $n^{\Omega(\sqrt{k})}$ for the $\textsc{Clique}_k$ problem whenever $k \le n^{1/2-\delta}$, translating to a $2^{\tilde{\Omega}(n^{1/4})}$ lower bound for the $\textsc{Clique}$ problem after taking optimal k. [1]

The work of Cavalar, Kumar, and Rossman [8], introduced a notion of *robust clique-sunflower*, which abstracts exactly the property needed to bound the error introduced on the negative distribution $\mathcal{G}_{n,p}$ by replacing the robust clique sunflower by its core. They showed better quantitative bounds on the size of the set system needed to contain such a robust clique sunflower and were able to deduce a lower bound $n^{\tilde{\Omega}(k)}$ for $k \le n^{1/3-\delta}$ – matching the same $2^{\tilde{\Omega}(n^{1/3})}$ for the clique problem when picking largest admissible k.

1.3 Sunflowers, robust sunflowers and robust clique sunflowers

We define the robust clique sunflower (after [8]) – a notion abstracting the main property of combinatorial sunflowers used to obtain a monotone lower bound for the $\textsc{Clique}_k$ problem, where the negative distribution is $\mathcal{G}_{n,p}$. This property ensures that for a DNF formula

$$\bigvee_i \mathcal{K}_{A_i}$$

if some of those sets A_i form a robust clique sunflower, we can replace them by a single clique indicator on the common intersection C while introducing only a small error on the negative distribution (and no error on the positive distribution).

▶ **Definition 6** (Robust clique sunflower). *A family of sets $\mathcal{S} \subset 2^{[n]}$ is an (p,ε)-robust clique sunflower (with core $C = \bigcap_{S\in\mathcal{S}} S$), if*

$$\Pr_G(\exists S \in \mathcal{S}, K_S \subset G \cup K_C) \ge 1-\varepsilon,$$

where G is a random Erdős–Rényi graph $\mathcal{G}_{n,p}$ sampled by including each edge independently with probability p, and $K_S \subset \binom{[n]}{2}$ is a clique on vertices S, i.e. $K_S := \{\{u,v\} : u,v \in S, u \ne v\}$.

[1] In the actual Alon-Boppana paper, they used a slightly more general combinatorial structure than sunflowers: a sequence of distinct sets $A_{i_1}, \dots A_{i_k}$, together with a set $C \subset A_t$ (not necessarily distinct from $A_{i_1}, \dots A_{i_k}$), s.t. all pairwise intersections $A_{i_s} \cap A_{i_r} \subset C$ – and they replaced $A_{i_1}, \dots A_{i_k}$ by C in step 2 here. Clearly any sunflower $A_{i_1} \dots A_{i_k}$ with the core $C = \bigcap A_{i_j}$ satifies this property. By using this more general structure, they were able to show a lower bound $n^{\sqrt{k}}$ for $k \le n^{2/3}$ – matching the result one would get insisting on using a sunflower here if the sunflower conjecture was true; yet without having to prove this conjecture.

This notion is intimately related to a similar notion of a robust sunflower, originally introduced in [20].

▶ **Definition 7** (Robust sunflower). *A family of sets $\mathcal{S} \subset 2^{[n]}$ is an (p, ε)-robust sunflower (with core $C = \bigcap_{S \in \mathcal{S}} S$), if*

$$\Pr_W(\exists S \in \mathcal{S}, S \subset W \cup C) \geq 1 - \varepsilon,$$

where W is a random subset of $[n]$ chosen by including each element independently with probability p.

As it turns out, any large enough family of sets of size ℓ contains a robust sunflower or a robust clique sunflower. We introduce a notation for the dependence between the size of the family and the parameters of this sunflower.

▶ **Definition 8.** *We define $RB(\ell, p, \varepsilon)$ to be the smallest number such that any ℓ-uniform[2] set system of size at least $RB(\ell, p, \varepsilon)$ contains a (p, ε)-robust sunflower.*

Similarly, we define $RCB(\ell, p, \varepsilon)$ to be the smallest number such that any ℓ-uniform set system of size at least $RCB(\ell, p, \varepsilon)$ contains a (p, ε)-robust clique-sunflower.

Note that a family $\mathcal{S} \subset 2^{[n]}$ is a robust clique sunflower (as in the Definition 6), if and only if the family of edge-sets $\{K_S : S \in \mathcal{S}\}$ is a robust sunflower. This observation implies a simple (yet unsatisfactory) bound

$$RCB(\ell, p, \varepsilon) \leq RB\left(\binom{\ell}{2}, p, \varepsilon\right). \tag{1}$$

In Section 1.4 (for instance Theorem 14) we provide a quantitative statement for how the upper bounds on RCB can be translated into lower bounds for the Clique problem.

For completeness, let us discuss a way to reinterpret a result similar to the Alon-Boppana (with the negative distribution being $\mathcal{G}_{n,p}$ instead of random $(\alpha - 1)$-partite graph) in this framework, by deducing a bound on the robust clique sunflowers from the sunflower bound.

▶ **Definition 9** (k-Sunflower). *An family $\mathcal{S} \subset 2^{[n]}$ of k sets is a k-sunflower (with a core $C = \cap_{S \in \mathcal{S}} S$), if all pairwise intersections of sets from $\mathcal{S}$ are C, i.e. for all $S_1 \neq S_2 \in \mathcal{S}$, we have $S_1 \cap S_2 = C$.*

What we call a k-sunflower is sometimes called in the literature *a sunflower with k-petals*. We chose a slightly more concise terminology.

A sunflower lemma [10] now says that for every ℓ and k, there is a finite number $S(\ell, k) \leq \ell! \, (k-1)^\ell$ such that every ℓ-uniform set system of size at least $S(\ell, k)$ has a k-sunflower. Subsequent results, including the breakthrough by [2] and improvements [5, 17, 23, 14, 18] provide successively smaller upper bounds on $S(\ell, k)$, concluding with the best currently known bound

$$S(\ell, k) \leq O(k \log \ell)^\ell. \tag{2}$$

while the famous sunflower conjecture ([10]) stipulates that $S(\ell, k) \leq O(k)^\ell$.

The following observation, similar in form to the core of the Razborov and Alon-Boppana results, is a simple way to connect sunflowers to robust clique sunflowers.

[2] ℓ-uniform set system is just a family of subsets of a universe, each subset of size exactly ℓ.

▶ **Lemma 10.** *Every ℓ-uniform k-sunflower is an $(p, \exp(-kp^{\binom{\ell}{2}}))$-robust clique sunflower. In particular $RCB(\ell, p, \varepsilon) \le S(\ell, \log(1/\varepsilon)p^{-\binom{\ell}{2}})$.*

Proof. Consider a k-sunflower $S_1, \dots S_k$ with core C. By the definition of sunflower the sets of edges $K_{S_i} \setminus K_C$ are disjoint, hence the indicator random variables $R_i := \mathbf{1}[K_{S_i} \subset G \cup K_C]$ are independent. Since $\mathbb{E}\, R_i \ge p^{\binom{\ell}{2}}$, the probability that all R_i are zero is at most $(1 - p^{\binom{\ell}{2}})^k \le \exp(-kp^{\binom{\ell}{2}})$. ◀

This, combined only with the classical upper bound $S(\ell, k) \le O(\ell k)^\ell$ by Erdős-Rado yields an upper bound

$$RCB(\ell, p, \varepsilon) \le O(1/p)^{\ell^3} (\ell \log(1/\varepsilon))^\ell,$$

which can be used to recover $n^{\Omega(\sqrt{\alpha})}$ lower bound via the approximation method outlined in the previous section (see for instance Theorem 14).

Crucially for us, a bound on robust sunflower numbers is an essential part of the new, improved series of results on the bounds for sunflower numbers. Specifically, the following theorem was used to deduce those new bounds and is important for our application.

▶ **Theorem 11** ([2, 5, 18])**.** *The robust sunflower numbers are bounded as*

$$RB(\ell, p, \varepsilon) \le O(p^{-1}(\log(1/\varepsilon) + \log \ell))^\ell.$$

This theorem fairly quickly implies the breakthrough upper bound on sunflower numbers (2). For our purposes, though, the sunflower consequence of Theorem 11 is less relevant, and we will leverage the existence of robust sunflowers directly.

In Section 3, we prove a much stronger reduction than (1), showing how robust sunflower bounds can be used to give bounds on robust clique-sunflowers, which leads to our main improvement.

▶ **Theorem 12.** *For any $\ell \ge 1$, $p \in (0, 1)$ and $\varepsilon \in (0, 1)$ we have*

$$RCB(\ell, p, \varepsilon) \le RB(\ell, p^\ell, \varepsilon/\ell^2),$$

where $RCB(\ell, p, \varepsilon)$ and $RB(\ell, p, \varepsilon)$ are defined as in Definition 8.

Combining the upper bound for robust sunflowers in Theorem 11 with our comparison theorem (Theorem 12), we get

▶ **Corollary 13.** *The clique sunflower numbers are bounded as*

$$RCB(\ell, p, \varepsilon) \le O(p^{-\ell}(\log(1/\varepsilon) + \log \ell))^\ell.$$

This result is an improvement over a bound from [8, Lemma 3.2] of the same quantity. They directly showed an upper bound

$$RCB(\ell, p, \varepsilon) \le p^{-\binom{\ell}{2}} O(\ell \log(1/\varepsilon))^\ell. \tag{3}$$

As we will see later, reducing the $O(\ell)^\ell$ factor down to $O(\log \ell)^\ell$ allowed us to show a lower bound $2^{\tilde{\Omega}(\sqrt{n})}$ for the clique problem as opposed to $2^{\tilde{\Omega}(n^{1/3})}$ from [1, 8].

In order to prove Theorem 12 we show that any $(p^\ell, \varepsilon/\ell^2)$-robust sunflower is an (p, ε)-robust clique sunflower. This statement is easier to interpret for sunflowers with empty core: given a ℓ uniform family of sets $\mathcal{F} \subset 2^{[n]}$ (with empty common intersection), if we are very

likely to cover one of those sets while sampling vertices of $[n]$ independently at random with probability p^ℓ, we are also very likely to cover one of those with a clique, when sampling edges of $\binom{[n]}{2}$ independently with probability p.

As it turns out, one can set up specific stochastic processes $\{Y_S\}_{S\in\mathcal{F}}$, and $\{Y'_S\}_{S\in\mathcal{F}}$ associated with both of those experiments – the expected supremum of these processes will be directly connected (respectively) with the probability of covering one of the sets of $S \in \mathcal{F}$ by a random set (with the inclusion probability p^ℓ), or the probability of covering one of the cliques on the vertices of $S \in \mathcal{F}$ by a random graph (with edge probability p). The central technical part of the proof is the comparison lemma stating that $\mathbb{E}\sup_{S\in\mathcal{F}} Y_S \leq \mathbb{E}\sup_{S\in\mathcal{F}} Y'_S$ – theorems of that form appeared in the literature on the theory of stochastic processes (most well known is the Slepian lemma, or Gaussian comparison principle [22, Corollary 2.10.12]). Even though we could not apply any of the known comparison lemmas black-box, we adapt the ideas that were used in [21] to show a comparison lemma for coordinate-wise contractions of canonical Bernoulli processes, and we show the desired inequality for our two specific stochastic processes in question.

1.4 Lower bounds for monotone circuits depending on sunflower bounds

We carry over the high-level structure of the proof outlined in Section 1.2 in more detail in Section 2. This part of the proof is technically very similar to known results applying the approximation method (for example [19, 1, 8]). However, in contrast with many of the previous work, we made an effort to provide a statement of the lower bound theorem parameterized by the robust clique sunflower numbers $RCB(\ell, p, \varepsilon)$ – which, in turn, can be upper bounded in several ways by the robust sunflower numbers $RB(\ell, p, \varepsilon)$, or sunflower numbers $S(\ell, k)$ as discussed in Section 1.3.

Making these dependencies explicit instead of choosing optimal values of relevant parameters ahead of time based on currently best-known bounds on sunflower numbers makes the interplay between sunflower-like combinatorial statements and monotone lower bounds for the clique problem much clearer.

Moreover, our analysis allows us to show the inapproximability results (Theorem 3) – i.e., hardness of distinguishing a random clique of size β, from a random graph $\mathcal{G}_{n,p}$ which is unlikely to even have a clique of size α for $\alpha < \beta$.

The analysis of [19, 8] focused on the exact case, i.e., $\beta = \alpha$; whereas in [19, 1], the negative distribution was chosen to be a random complete $(\beta - 1)$-partite graph. However, this choice was not crucial in their reasoning – only simple modifications are needed to adapt their proofs to the $\mathcal{G}_{n,p}$ being the negative distribution.

As we have been made aware recently, an even more general statement can be found [6, Theorem 5.6.4] – the following theorem can be deduced as a corollary of their more general framework, by a concrete instantiation of their notion of the notion of *abstract sunflower* to denote the robust clique sunflower.

▶ **Theorem 14.** *Fix some $c \leq n$ and take $p := n^{-\frac{2}{\alpha-1}}$. If for every $\ell \leq 2c$, we have*

$$(\beta/n)\, RCB(\ell, p, \varepsilon)^{1/\ell} \leq \gamma \leq o(1)$$

then the size of any monotone circuit C distinguishing $\mathcal{T}^-_\alpha$ and $\mathcal{T}^+_\beta$ satisfies

$$size(C) \geq \Omega(\min\{\gamma^{-c},\ n^{-2c}/\varepsilon\}).$$

In particular, choosing $\varepsilon = n^{-4c}$, if for all $\ell \leq 2c$ we have

$$(\beta/n)\, RCB(\ell, p, n^{-4c})^{1/\ell} < n^{-\delta}$$

then the monotone complexity of distinguishing $\mathcal{T}^-_\alpha$ and $\mathcal{T}^+_\beta$ is $\Omega(n^{\delta c})$.

It is instructive to see how to essentially recover the Alon-Boppana bounds from Theorem 14, using a simple lemma stating that large enough sunflowers are robust clique sunflowers (Lemma 10) together with the new bounds on the sunflower numbers (2). In this case, we have an upper bound on the

$$RCB(\ell, p, \varepsilon)^{1/\ell} \leq p^{-\ell^2}(\log(1/\varepsilon) + \log \ell) \leq n^{O(\frac{c^2}{\alpha})} c \log n,$$

so taking $c := O(\sqrt{\delta\alpha})$ for a small constant δ, such that $RCB(\ell, p, \varepsilon)^{1/\ell} \leq n^{\delta}\sqrt{\alpha}\log n$, we end up with a complexity lower bound $n^{\Omega(\delta c)} = n^{\Omega(\sqrt{\alpha})}$, as long as

$$\sqrt{\alpha}\beta \log n \lesssim n^{1-2\delta},$$

recovering [1, Theorem 3.11]. Setting $\alpha = \beta$, gives a lower bound of form $n^{\Omega(\sqrt{\alpha})}$ as long as $\alpha \leq n^{2/3-\delta}$.

In a similar vein, plugging in the upper bound (3) for $RCB(\ell, p, \varepsilon)$ (originally shown in [8, Lemma 3.2]), yields the lower bound $n^{\Omega(\alpha)}$, as long as $\alpha^2\beta \leq n^{1-\delta}/\log n$, and again choosing $\alpha = \beta$ yields an $n^{\Omega(\alpha)}$ for $\alpha \leq n^{1/3-\delta}$, recovering their [8, Theorem 3.23].

Finally, Theorem 14 together with our new bounds for the robust clique numbers (Corollary 13) directly imply the lower bound claimed in Theorem 3.

2 Approximating cliques

In this chapter, we will give an extended version of the approximation method (showing Theorem 14) and take advantage of our improved robust clique sunflower bound to strengthen the lower bound for $\textsc{Clique}_k$ to $n^{\Omega(k)}$ if $k \leq n^{1/2-\delta}$.

2.1 Abstract approximation method

For the sake of abstraction, we describe a general inductive procedure of converting any monotone circuit C computing a distributional decision problem $\mathcal{D}(\mathcal{T}^-, \mathcal{T}^+)$, gate by gate, into an approximation circuit $\hat{C} \in \mathcal{A}$ from a set of approximation circuits $\mathcal{A}$. This procedure depends on a pair of "compression" functions $\mathcal{P}^\wedge, \mathcal{P}^\vee : \mathcal{A} \times \mathcal{A} \to \mathcal{A}$. The error introduced by the conversion, with respect to the distributions $\mathcal{T}^+$ and $\mathcal{T}^-$, will depend solely on $\mathcal{P}$.

▶ **Definition 15.** *For any monotone circuit C and a pair of compression functions $\mathcal{P} = (\mathcal{P}^\wedge, \mathcal{P}^\vee)$ with $\mathcal{P}^\wedge, \mathcal{P}^\vee : \mathcal{A} \times \mathcal{A} \to \mathcal{A}$ define $\hat{C} \in \mathcal{A}$ as the result of the following procedure.*

1. ***Input variables**: Let $C = x_i$ be an input variable. Define*

   $$\hat{C} = C = x_i.$$

2. $\wedge$ ***-gate**: Assume $C = \hat{C}_0 \wedge \hat{C}_1$. Define*

   $$\hat{C} = \mathcal{P}^\wedge(\hat{C}_0, \hat{C}_1).$$

3. $\vee$ ***-gate**: Assume $C = \hat{C}_0 \vee \hat{C}_1$. Define*

   $$\hat{C} = \mathcal{P}^\vee(\hat{C}_0, \hat{C}_1).$$

Depending on the choice of $\mathcal{P}$, $\hat{C}$ can be very different from C. For the sake of analysis, we are interested in the maximum error that a single transformation step can introduce on either distribution.

▶ **Definition 16.** *Let $\mathcal{I}(\mathcal{P}')$ be the image of $\mathcal{P}'$ and define the positive approximation error*

$$\zeta^+_{\mathcal{P}} = \max_{\odot\in\{\wedge,\vee\}} \max_{\hat{C}_0,\hat{C}_1\in\mathcal{I}(\mathcal{P}^\odot)} \{\Pr_{G\in\mathcal{T}^+}[(\hat{C}_0 \odot \hat{C}_1) = 1 \text{ and } \mathcal{P}^\odot(\hat{C}_0 \odot \hat{C}_1)] = 0\}$$

and the negative approximation error

$$\zeta^-_{\mathcal{P}} = \max_{\odot\in\{\wedge,\vee\}} \max_{\hat{C}_0,\hat{C}_1\in\mathcal{I}(\mathcal{P}^\odot)} \{\Pr_{G\in\mathcal{T}^-}[(\hat{C}_0 \odot \hat{C}_1) = 0 \text{ and } \mathcal{P}^\odot(\hat{C}_0 \odot \hat{C}_1)] = 1\}$$

We can also formalize the earlier intuition about proving circuit size lower bounds using the relative and absolute approximation errors.

▶ **Definition 17.** *We say that a circuit C distinguishes distributions $\mathcal{T}^+$ and $\mathcal{T}^-$ if $\mathbb{E}_{X\sim\mathcal{T}^+}[C(X)] - \mathbb{E}_{X\sim\mathcal{T}^-}[C(X)] \geq 2/3$.*

We say that a circuit is $\mathcal{P}$-simple if it is in the image of the compression function $\mathcal{P}$.

▶ **Lemma 18.** *If every $\mathcal{P}$-simple circuit $\hat{C}$ satisfies $\mathbb{E}_{X\sim\mathcal{T}^+}[\hat{C}(X)] - \mathbb{E}_{X\sim\mathcal{T}^-}[\hat{C}(X)] \leq 1/3$, then every circuit C distinguishing $\mathcal{T}^+$ and $\mathcal{T}^-$ has*

$$\text{size}(C) \geq \Omega(\min(1/\zeta^+_{\mathcal{P}}, 1/\zeta^-_{\mathcal{P}})).$$

Proof. The transformation of C into $\hat{C}$ introduces at most ζ^+ error on $\mathcal{T}^+$ and ζ^- error on $\mathcal{T}^-$ per gate. As C distinguishes the distributions, but $\mathbb{E}_{X\sim\mathcal{T}^+}[\hat{C}(X)] - \mathbb{E}_{X\sim\mathcal{T}^-}[\hat{C}(X)] \leq 1/3$, the error of the $\mathcal{P}$-transformation must match the absolute error of the $\mathcal{P}$-simple circuit $\hat{C}$, so the theorem follows. ◀

2.2 Proving clique lower bounds

In the following, we will specialize the definition of the approximation circuits $\mathcal{A}$ and the compression functions $(\mathcal{P}^\wedge, \mathcal{P}^\vee)$ to prove the lower bounds on $\mathcal{D}(\mathcal{T}^-_\alpha, \mathcal{T}^+_\beta)$.

▶ **Definition 19.** *Let $A \subseteq [n]$ be a subset of vertices. Let $\mathcal{K}_A$ be the clique indicator function on A, such that $\mathcal{K}_A(G) = 1 \Leftrightarrow K_A \subset G$. For any family of subsets of $[n]$, $\{A_1, \ldots, A_m\}$, define the associated approximation circuit as*

$$A = \bigvee_{i=1}^{m} \mathcal{K}_{A_i}$$

Let $\mathcal{A}$ be the set of all approximation circuits.

We will often freely switch between the formal definition of an approximator as a Boolean circuit $A = \bigvee_{i=1}^m \mathcal{K}_{A_i}$ and its representation as a set system over the universe of graph vertices $A = \{A_1, \ldots, A_m\} \subseteq 2^{2^{[n]}}$.

▶ **Definition 20.** *For an "inner-compression" function $\tau : \mathcal{A} \to \mathcal{A}$, define $\mathcal{P} = (\mathcal{P}^\wedge, \mathcal{P}^\vee)$ with*

$$\mathcal{P}^\wedge(\bigvee_i^u \mathcal{K}_{X_i}, \bigvee_j^v \mathcal{K}_{Y_i}) = \tau(\bigvee_i^u \bigvee_j^v \mathcal{K}_{X_i\cup Y_j})$$

$$\mathcal{P}^\vee(\bigvee_i^u \mathcal{K}_{X_i}, \bigvee_j^v \mathcal{K}_{Y_i}) = \tau(\bigvee_i^u \mathcal{K}_{X_i} \vee \bigvee_j^v \mathcal{K}_{Y_i})$$

We observe that for the identity $\tau = id$, $\mathcal{P}$ introduces no error on the graph distributions $\mathcal{T}^-_\alpha$ and $\mathcal{T}^+_\beta$. Indeed, the positive distribution is supported on the β cliques. A β-clique G contains both cliques K_{X_i} and K_{Y_j} if and only if it contains a clique $K_{X_i\cup Y_i}$. On the other hand, on the negative distribution, transforming a conjunction $K_{X_i} \wedge K_{X_j}$ into the clique indicator $K_{X_i\cup Y_j}$ can only reject more instances. Thus, the error only depends on the choice of τ. For constructing such τ, we will use the robust clique sunflowers.

2.3 Robust closures

For a DNF formula $A := \bigvee_{i=1}^{m} \mathcal{K}_{A_i}$ we define $cl_{p,\varepsilon}(A)$ to be a formula obtained from A by repeatedly replacing any (p,ε)-robust clique sunflower $A_{i_1}, \dots A_{i_k}$ by its core $C := \bigcap_j A_{i_j}$, as long as there is such a robust clique sunflower, interleaved with removing all sets A_j, s.t. $A_i \subset A_j$ for some A_i. The $\mathrm{trim}_c(A)$ will be a circuit obtained from A by removing all sets $|A_i| > c$.

Finally, we will choose the inner compression function (used in Definition 20) as

$$\tau(A) := \mathrm{trim}_c(cl_{p,\varepsilon}(A)).$$

We say that a circuit A is (p,ε)-closed, if there are no (p,ε)-robust sunflowrers among the sets $A_1, \dots A_m$. The bound on the number of sets A_j of a given size for an (p,ε)-closed circuit is a direct consequence of the definition of a closed circuit and the RCB numbers.

▶ **Fact 21.** *The number $\mathcal{M}_l(A)$ of clique indicators A_i of size l of in a (p,ϵ)-closed circuit A is bounded by*

$$\mathcal{M}_l(A) < RCB(\ell, p, \varepsilon).$$

2.4 The lower bound

A crucial property that we will exploit both in the analysis of the approximator as well as in the construction of an explicit algorithm solving $\mathcal{D}(\mathcal{T}_\alpha^-, \mathcal{T}_\beta^+)$ is that the positive test distribution $\mathcal{T}_\beta^+$ is "well-spread" in a sense that it is unlikely for a large clique-indicator to accept many instances of $\mathcal{T}_\beta^+$.

▶ **Lemma 22.** *For a clique K_B with $|B| = \ell$,*

$$\mathbb{E}_{G \sim \mathcal{T}_\beta^+}[\mathcal{K}_B(G)] < (\beta/n)^\ell.$$

Proof. A graph G, sampled from $G \sim \mathcal{T}_\beta^+$ is an isolated β-clique A on n vertices. Thus $K_B(G) = 1$ if and only if $B \leq A$ such that

$$Pr_{G \sim \mathcal{T}_\beta^+}[\mathcal{K}_B(G) = 1] = Pr_{K_A \sim \binom{n}{\beta}}[B \subseteq A] = \frac{\binom{n-\beta}{\beta-\ell}}{\binom{n}{\beta}} \leq (\beta/n)^\ell. \quad \blacktriangleleft$$

This lemma, together with a bound on the number of clique indicators of each size (Fact 21), allows us to easily show that any $\mathcal{P}$-simple circuit cannot distinguish $\mathcal{T}_\alpha^+$ from $\mathcal{T}_\beta^-$ – either a circuit like that trivially accepts every output (and hence makes a large error on the negative distribution), or it accepts only small fraction of inputs from the positive distribution.

▶ **Lemma 23.** *If $\hat{C} \neq 1$ is an approximator obtained by the compression function $\mathcal{P}$, then*

$$\mathbb{E}_{G \sim \mathcal{T}_\beta^+}[\hat{C}(G)] \leq \sum_{l=2}^{c} (\beta/n)^l \, RCB(l, \epsilon, p).$$

Proof. The proof is a simple application of a union bound. If $G \sim \mathcal{T}_\beta^+$ and $\hat{C}(G) = 1$, then there must exist some term $\mathcal{K}_{A_i}$ of $\hat{C}$ with $K_{A_i} \subset G$. By Fact 21 there are at most $\mathcal{M}_l(f)$ terms of size l (for all $l \leq c$) and by Lemma 22 for every term of size l the probability that it accepts a positive instance is at most $(\beta/n)^l$. Thus

$$Pr_{G \sim \mathcal{T}_\beta^+}[\hat{C}(G) = 1] \leq \sum_{l=2}^{c} (\beta/n)^l \mathcal{M}_l(f) \leq \sum_{l=2}^{c} (\beta/n)^l \, RCB(l, \epsilon, p). \quad \blacktriangleleft$$

Single-step Approximation Errors

We give bounds on the single-step approximation errors ζ^+ and ζ^-, induced by $\mathcal{P}$.

▶ **Lemma 24.** *The single-step error introduced on the positive distribution is bounded as follows*

$$\zeta^+_{\mathcal{P}} \leq \sum_{l=c}^{2c} (\beta/n)^l \, RCB(l, \epsilon, p).$$

Proof. If we have for some $G \sim \mathcal{T}^+_\beta$ that $C(G) = 1$ but $\tau(C)(G) = 0$, then there was some term $\mathcal{K}_{A_i}$ of $cl_{p,\varepsilon}(C)$ of size larger than c with $\mathcal{K}_A(G) = 1$, which was discarded during the trimming process.

We can union bound the probability of this happening, in a similar way as Lemma 23, using the upper bound $M_\ell(cl_{p,\varepsilon}(C)) \leq RCB(\ell, p, \varepsilon)$ (Fact 21) on the number of terms of a given size ℓ, and the upper bound on the probability that a given term of size ℓ accepts the positive distribution (Lemma 22). ◀

We also notice that for a fixed closure factor, the negative approximation error is bounded by the probability bound given by the robust clique-sunflower.

▶ **Lemma 25.** *The single-step error introduced on the negative distribution is bounded as*

$$\zeta^-_{\mathcal{P}} \leq \epsilon n^{2c}.$$

Proof. Note that while applying $cl_{p,\varepsilon}(A)$ operation, each set $S \subset [n]$ of size at most $2c$ can be added as a core of some robust clique sunflower at most once, as the core is always a strict subset of the sunflower indicators. Hence, we will repeat the process of replacing a robust sunflower by its core at most n^{2c} times, in each step introducing error at most ε on the negative distribution (by the definition of robust clique sunflower). ◀

Complexity

By the results of the last sections, we can proceed to prove a quantified condition on the existence of a monotone circuit lower complexity bound on $\mathcal{D}(\mathcal{T}^+_\beta, \mathcal{T}^-_\alpha)$, depending on the robust-clique sunflower bound $RCB(l, \epsilon, p)$. The proof will follow by a simple application of Lemma 18 together with bounds on the error introduced in each step of the process, established in the previous section. We recall the statement of the theorem from Section 1.4.

▶ **Theorem 14.** *Fix some $c \leq n$ and take $p := n^{-\frac{2}{\alpha-1}}$. If for every $\ell \leq 2c$, we have*

$$(\beta/n) \, RCB(\ell, p, \varepsilon)^{1/\ell} \leq \gamma \leq o(1)$$

then the size of any monotone circuit C distinguishing $\mathcal{T}^-_\alpha$ and $\mathcal{T}^+_\beta$ satisfies

$$size(C) \geq \Omega(\min\{\gamma^{-c}, \; n^{-2c}/\varepsilon\}).$$

In particular, choosing $\varepsilon = n^{-4c}$, if for all $\ell \leq 2c$ we have

$$(\beta/n) \, RCB(\ell, p, n^{-4c})^{1/\ell} < n^{-\delta}$$

then the monotone complexity of distinguishing $\mathcal{T}^-_\alpha$ and $\mathcal{T}^+_\beta$ is $\Omega(n^{\delta c})$.

Proof. This statement follows directly from Lemma 18. First, note that by Lemma 23 any $\mathcal{P}$-simple circuit either accepts all negative instances or accepts a positive instance with probability

$$\sum_{2\le\ell\le c} (\beta/n)^\ell \, RCB(\ell,p,\varepsilon) \le \sum_{2\le\ell\le c} \gamma^\ell \le O(\gamma^2) \le o(1).$$

hence a $\mathcal{P}$-simple circuit cannot distinguish it $\mathcal{T}^-_\alpha$ from $\mathcal{T}^+_\beta$.

The error $\zeta^-_{\mathcal{P}} \le \varepsilon n^{2c}$ was shown as Lemma 25. For the $\zeta^+_{\mathcal{P}}$ we can use Lemma 24, to obtain a bound

$$\zeta^+_{\mathcal{P}} \le \sum_{\ell=c+1}^{2c} (\beta/n)^\ell \, RCB(\ell,p,\varepsilon) \le \sum_{\ell=c+1}^{\infty} \gamma^\ell \le O(\gamma^c).$$

With those two bounds, we can now directly apply Lemma 18 to deduce the first part of the theorem. The second part follows by choosing $\varepsilon := n^{-4c}$, $\gamma := n^{-\delta}$, and simple algebraic manipulations. ◀

We can now deduce Theorem 3 directly from this more general statement and the new upper bound on the RCB (Corollary 13).

Proof of Theorem 3. By Corollary 13, taking $\varepsilon = n^{-4c}$, for $\ell \le 2c$ we have

$$RCB(\ell,p,\varepsilon)^{1/\ell} \lesssim p^{-\ell}(\log(1/\varepsilon) + \log \ell) \lesssim n^{\frac{4c}{\alpha-1}} c \log n,$$

hence if we pick $c := \delta(\alpha-1)/8$, we will have

$$\beta \, RCB(\ell,p,\varepsilon)^{1/\ell} \lesssim \alpha\beta n^{\delta/2} \log n.$$

Now, when $\alpha\beta \le n^{1-\delta}/\log n$, this yields

$$\beta \, RCB(\ell,p,\varepsilon)^{1/\ell} \lesssim n^{1-\delta/2},$$

and finally, applying Theorem 14, we get a lower bound of form

$$\Omega(n^{\delta c/2}) \ge \Omega(n^{\delta^2\alpha/16}) \ge n^{\Omega(\delta^2\alpha)}.$$ ◀

3 Robust clique numbers are upper bounded by robust sunflower numbers

In this section, we will show a reduction, proving that any robust sunflower bound implies a corresponding robust-clique-sunflower bound with slightly different parameters. For our convenience, let us recall the statement of the theorem.

▶ **Theorem 12.** *For any $\ell \ge 1$, $p \in (0,1)$ and $\varepsilon \in (0,1)$ we have*

$$RCB(\ell,p,\varepsilon) \le RB(\ell,p^\ell,\varepsilon/\ell^2),$$

where $RCB(\ell,p,\varepsilon)$ and $RB(\ell,p,\varepsilon)$ are defined as in Definition 8.

The main technical lemma towards Theorem 12 is a comparison principle for a concrete pair of stochastic processes. The proof of this result has been inspired by [21]. Before we state the comparison lemma, let us provide a necessary definition.

▶ **Definition 26.** *We say that* $\phi : \binom{[n]}{k} \to \binom{[n]\times[n]}{k\ell}$ *is a* proper lifting *if for every* $i \in S$*, we have* $|\phi(S) \cap (\{i\} \times [n])| = \ell$*.*

Note that since the image of any set under the proper lifting has exactly $k\ell$ elements, and it has exactly ℓ elements in each of the k rows corresponding to elements of S, we have for every $i \notin S$, that $\phi(S) \cap (\{i\} \times [n]) = \emptyset$. The canonical examples of proper liftings to have in mind are $\phi(S) := S \times S$ or $\phi(S) := S \times [\ell]$, although we will need a slightly more complicated one in the proof of Theorem 12.

The following comparison lemma shows that among the specific class of the stochastic processes associated with a lifting ϕ, the smallest process is associated with the lifting $\phi(S) = S \times [\ell]$.

▶ **Lemma 27.** *Let* $\mathcal{F}$ *be a* k*-uniform family of subsets of* $[n]$*, and random variables* $\{A_{i,j}\}_{i\in[n],j\in[\ell]}, \{\tilde{A}_{i,j}\}_{i,j\in[n]}$ *be independent and identically distributed.*

If ϕ *is any proper lifting (as in Definition 26), then*

$$\mathbb{E} \sup_{S\in\mathcal{F}} \sum_{(i,j)\in S\times[\ell]} A_{i,j} \leq \mathbb{E} \sup_{S\in\mathcal{F}} \sum_{(i,j)\in\phi(S)} \tilde{A}_{i,j}. \tag{4}$$

We will prove this lemma later in Section 3.1. Before we do, let us see how it implies Theorem 12. More concretely, we will show that every $(p^\ell, \varepsilon/\ell^2)$-robust sunflower is also an (p, ε)-robust clique sunflower, and Theorem 12 will be a direct corollary.

It is easier to gain the intuition of how to prove a statement of this form from Lemma 27 if we consider a special case of a robust sunflower $\mathcal{S}$ with an empty core C – in this case, we would apply Lemma 27 with proper lifting $\phi(S) := S \times S$, and $A_{i,j}$ being independent $\{0,1\}$ valued random variables with $\mathbb{E}\, A_{i,j} = p$. The i-th row of a matrix A has a sum equal to ℓ independently with probability p^ℓ. Therefore, by the robust sunflower definition, with high probability there exists a square of form $S \times [\ell]$ with sum ℓ^2 – this implies that the expected supremum over sums over $S \times [\ell]$ as in the left-hand side of (4) is very close to ℓ^2, and hence also expected supremum over all sums over $S \times S$ for $S \in \mathcal{F}$ (by the aforementioned lemma). We can deduce that with high probability there is a square $S \times S$ in the matrix $\tilde{A}$ with sum ℓ^2 – a statement corresponding to the fact that some clique K_S is covered by a random $\mathcal{G}_{n,p}$ graph.

The handling of the more general case, where the core C is not necessarily empty, is only slightly more technical.

▶ **Lemma 28.** *If* ℓ*-uniform family* $\mathcal{S} \subset 2^{[n]}$ *is a* $(p^\ell, \varepsilon/\ell^2)$*-robust sunflower, it is also a* (p, ε)*-robust clique sunflower.*

Proof. Consider a matrix $A \in \{0,1\}^{n\times\ell}$, and $\tilde{A} \in \{0,1\}^{n\times n}$, with independent entries, each equal to 1 with probability p. We can treat the part of the matrix $\tilde{A}$ below the diagonal (i.e., $\tilde{A}_{i,j}$ for $j < i$) as determining the adjacency matrix for a random graph G (sampled according to the Erdős–Rényi distribution $\mathcal{G}_{n,p}$).

Let us assume without loss of generality that $C = [\ell - k]$ for some k. We can look at a k-uniform family $\mathcal{S}' := \{S \setminus C : S \in \mathcal{S}\}$. Consider also the set W of all indices such that the corresponding row in the matrix A is all-one: $W := \{i : \sum_{j\leq\ell} A_{i,j} = \ell\}$. Note that each element in W is included independently with probability p^ℓ. Since $\mathcal{S}$ is a robust sunflower with the core C, directly by definition of robust sunflower with probability at least $1 - \varepsilon$, there is a set $S \in \mathcal{S}'$ covered by W, i.e.

$$\sum_{i,j\in S\times[\ell]} A_{ij} = \ell k.$$

This means that with probability $1 - \varepsilon/\ell^2 \geq 1 - \varepsilon/(\ell k)$, the value of

$$\sup_{S \in \mathcal{S}'} \sum_{i,j \in S \times [\ell]} A_{ij} = \ell k,$$

and therefore

$$\mathbb{E} \sup_{i,j \in S \times [\ell]} A_{ij} \geq (1 - \varepsilon/\ell k)\ell k = \ell k - \varepsilon.$$

We will now apply Lemma 27 to a family $\mathcal{S}'$ together with a lifting function $\phi(S) := S \times ([\ell - k] \cup S)$ to deduce

$$\mathbb{E} \sup_{S \in \mathcal{S}'} \sum_{i,j \in \phi(S)} \tilde{A}_{ij} \geq \ell k - \varepsilon. \tag{5}$$

On the other hand, denoting by $1 - p_0$ the probability that there exists $S \in \mathcal{S}'$, s.t. all entries of $\tilde{A}_{ij}$ over $(i,j) \in \phi(S)$ are 1 we have

$$\mathbb{E} \sup_{S \in \mathcal{S}'} \sum_{i,j \in \phi(S)} \tilde{A}_{ij} \leq (1 - p_0)\ell k + p_0(\ell k - 1) = \ell k - p_0, \tag{6}$$

and combining (5) and (6), we get $p_0 \leq \varepsilon$.

Finally, as discussed above, if we consider a graph G on $[n]$, where the edge $\{i,j\} \in G$ for $j < i$ if $\tilde{A}_{i,j} = 1$, this graph has exactly the distribution of $\mathcal{G}_{n,p}$. Moreover, when there exists $S \in \mathcal{S}'$, s.t. all entries $\tilde{A}_{i,j} = 1$ for $i \in S, j \in S \cup C$, then, in particular, the clique $K_{S \cup C}$ is contained in the graph $G \cup K_C$. Indeed, all edges within C are already contained in K_C, and $\tilde{A}_{i,j} = 1$ for $i \in S, j \in S \cup C$ implies that all edges between S and $S \cup C$ are present in G (since C consists of the first $|C|$ elements of $[n]$).

The probability of this happening is $1 - p_0 \geq 1 - \varepsilon$, as desired, finishing the proof that $\mathcal{S}$ is indeed a robust-clique-sunflower. ◀

Proof of Theorem 12. This theorem follows as a direct corollary of Lemma 28 and relevant definitions.

Indeed, every ℓ-uniform family of size at least $RB(p^\ell, \varepsilon/\ell^2)$ contains a $(p^\ell, \varepsilon/\ell^2)$-robust sunflower, which by Lemma 28 is (p, ε)-robust clique sunflower. ◀

3.1 Proof of Lemma 27

We will prove this statement by induction. Specifically, let us define $B_t := [t] \times [n]$, and $\bar{B}_t := ([n] \setminus [t]) \times [n]$ (so that $B_t \cup \bar{B}_t$ forms a partition of $[n] \times [n]$ into the first t rows, and the remaining $n - t$ rows).

Take $\phi_t : \binom{[n]}{k} \to \binom{[n] \times [n]}{\ell k}$ defined as $\phi_t(S) := ((S \times [\ell]) \cap B_t) \cup (\phi(S) \cap \bar{B}_t)$ – i.e. on the first t rows ϕ_t agrees with $S \times [\ell]$, and on the remaining $n - t$ rows it agrees with $\phi(S)$. Note that $\phi_t(S)$ has always ℓk elements, and moreover $\phi_0(S) = S \times [l]$, and $\phi_n(S) = \phi(S)$. As such, it is enough to show that for any t we have

$$\mathbb{E} \sup_{S \in \mathcal{F}} \sum_{(i,j) \in \phi_{t+1}(S)} \tilde{A}_{i,j} \leq \mathbb{E} \sup_{S \in \mathcal{F}} \sum_{(i,j) \in \phi_t(S)} A_{i,j}, \tag{7}$$

where $\{A_{ij}\}$ and $\{\tilde{A}_{ij}\}$ are two collections of independent and identically distributed random variables (with the same distribution).

In order to prove this statement, we will consider a coupling where $A_{i,j} = \tilde{A}_{i,j}$ for all $i \neq t+1$. Let us now condition on all variables $A_{i,j}$ for $i \neq t+1$ – we will show that for a fixed values of all those other variables, the desired inequality between expectation still holds, where expectation is only taken over the variables $A_{t+1,j}, \tilde{A}_{t+1,j}$.

Concretely, for any given S, we can decompose:

$$\sum_{(i,j)\in\phi_t(S)} A_{i,j} = \sum_{(i,j)\in\phi_t(S), i\neq t+1} A_{i,j} + \sum_{j:(t+1,j)\in\phi_t(S)} A_{t+1,j},$$

and similarly for ϕ_t. Note that when $i \neq t+1$ the pair $(i,j) \in \phi_t(S)$ if and only if $(i,j) \in \phi_{t+1}(S)$, so if we denote

$$a(S) := \sum_{(i,j)\in\phi_t(S) i\neq t+1} A_{i,j},$$

and by $p_0(S) := \{j : (t+1,j) \in \phi_t(S)\}$, and similarly $p_1(S) := \{j : (t+1,j) \in \phi_{t+1}(S)\}$, we have

$$\sum_{(i,j)\in\phi_t(S)} A_{i,j} = a(S) + \sum_{j\in p_0(S)} A_{t+1,j},$$

and

$$\sum_{(i,j)\in\phi_{t+1}(S)} \tilde{A}_{i,j} = a(S) + \sum_{j\in p_1(S)} \tilde{A}_{t+1,j},$$

since $\tilde{A}_{i,j} = A_{i,j}$ for $i \neq t+1$.

Finally, notice that if $t \in S$, $|p_0(S)| = \ell$ and $p_1(S) = [\ell]$, whereas if $t \notin S$, we have $p_0(S) = p_1(S) = \emptyset$. To finish the induction, we only need to show the following claim.

▷ Claim 29. Consider a finite sequence of pairs $(S_1, a_1), (S_2, a_2), \ldots (S_m, a_m)$, where $S_i \subset [n]$, $a_i \in \mathbb{R}$ and $|S_i|$ is either ℓ or 0.

Let $\mathcal{D}$ be a distribution over $\mathbb{R}$, and let $A_1, \ldots A_n$ and $\tilde{A}_1, \ldots \tilde{A}_n$ be two sequences of independent random variables distributed according to $\mathcal{D}$. Moreover, let $b_0 = \max\{a_i : S_i = \emptyset\}$ and $b_1 = \max\{a_i : S_i \neq \emptyset\}$. Then

$$\mathbb{E} \sup_{i\leq m}(a_i + \sum_{j\in S_i} A_j) \geq \mathbb{E} \max(b_1 + \sum_{i\in[\ell]} \tilde{A}_i, b_0).$$

Proof. If all sets S_i are empty, the statement is trivial. Let i_0 be an index s.t. $S_{i_0} = \emptyset$ and $a_{i_0} = b_0$, and similarly let i_1 be an index such that $S_{i_1} \neq \emptyset$, and $a_{i_1} = b_1$. Consider any permutation $\pi : [n] \to [n]$, s.t. $\pi(S_{i_1}) = [\ell]$. We can extend the sequence $\tilde{A}_i$ to n variables $\{\tilde{A}_i\}_{i\leq n}$, and consider a coupling between A and $\tilde{A}$, given by $A_i := \tilde{A}_{\pi(i)}$. Clearly, under this coupling, both marginal distributions of $\{A_i\}_{i\leq n}$ and $\{\tilde{A}_i\}_{i\leq \ell}$ are the right ones, all we need to show is that for any given realization of the joint process we have

$$\sup_{i\leq m}(a_i + \sum_{j\in S_i} A_j) \geq \max(b_1 + \sum_{i\leq l} \tilde{A}_i, b_0).$$

Indeed, by construction, we have

$$\begin{aligned}\sup_{i\leq m}(a_i + \sum_{j\in S_i} A_j) &= \sup_{i\leq m}(a_i + \sum_{j\in S_i} \tilde{A}_{\pi(j)}) \\ &\geq \sup_{i\in\{i_0,i_1\}}(a_i + \sum_{j\in S_i} \tilde{A}_{\pi(j)}) = \max(b_1 + \sum_{i\leq l} \tilde{A}_i, b_0).\end{aligned}$$ ◁

This lemma completes the proof of the inequality (7), and we can conclude the proof of Lemma 27, by chaining the inequalities (7) across all t, since $\phi_0(S) = \phi(S)$ and $\phi_n(S) = S \times [\ell]$.

4 The upper bound

We now turn towards the upper monotone complexity bound for $\mathcal{D}(\mathcal{T}^-_\alpha, \mathcal{T}^+_\beta)$ where we will construct an explicit circuit $\mathcal{C}$ for distinguishing $\mathcal{T}^+_\beta$ and $\mathcal{T}^-_\alpha$.

A useful subroutine that can be implemented using monotone boolean circuits is the ability to "sort" the input in polynomial size, as a sorting network on a boolean input can be easily used to construct a monotone boolean sorting circuit. [3] This also implies that monotone circuits can implement the threshold function T_τ, which accepts the input if there are at least τ input variables set to 1, simply wiring the output to the τ-th output of the sorting network.

In order to distinguish T^-_α and T^+_β, we will take a random family of subsets of size l, and accept the graph if the number of covered clique indicators on those subsets exceeds some threshold τ – by wiring the outputs of clique indicators to the inputs of the threshold function T_τ.

Probabilities of including a random clique in $\mathcal{T}^+_\beta$ and $\mathcal{T}^-_\alpha$

$\mathcal{T}^+_\beta$ and $\mathcal{T}^-_\alpha$ have an important distinctive property that we will use for constructing $\mathcal{C}_n$.

▶ **Lemma 30.** *For a clique indicator $\mathcal{K}_A$ with $|A|= l$, the probability of clique inclusion is*

$$Pr_{G\sim\mathcal{T}^+_\beta}[\mathcal{K}_A(G) = 1] = (\beta/n)^l$$

and

$$Pr_{G\sim\mathcal{T}^-_\alpha}[\mathcal{K}_A(G) = 1] = p^{\binom{l}{2}} = (n^{-2/\alpha-1})^{\binom{l}{2}}$$

Proof. The first equation has already been proven in Lemma 22. The second equation follows trivially from the definition of $\mathcal{G}_{n,p}$ – each of the $\binom{l}{2}$ edges of A is present in G independently with probability p. ◀

The interesting property is that, for the right choice of parameters, the "clique-probability" K_A is noticeably larger on $\mathcal{T}^+_\beta$ than on $\mathcal{T}^-_\alpha$. Take, for instance $p := \Pr_{G\sim\mathcal{T}^-_\alpha}(\mathcal{K}_A(G) = 1)$ and $q := \Pr_{G\sim\mathcal{T}^+_\beta}(\mathcal{K}_A(G) = 1)$, if we can pick the size of the clique l such that $q > 10p$, we will be able to construct a circuit of size $\Omega(1/q)$ that accepts $\mathcal{T}^+_\beta$ and rejects $\mathcal{T}^-_\alpha$, by looking at randomly placed $m = \Omega(1/q)$ clique indicators, accepting the input if at least $9mp$ of those clique indicators are accepting on G.

An explicit circuit

The probabilistic [4] monotone circuit $\mathcal{C}_n$ for an n-boolean input is defined by the following construction. For parameters m, l, τ

1. Choose m l-element subsets $A_1, \ldots, A_m \subseteq [n]$ independent uniformly at random and connect the associated clique indicators $K_{A_1}, \ldots, K_{A_m}$ to the input.
2. Connect the output of the clique indicators to the threshold function T^m_τ.
3. Connect the output of T^m_τ to the output of $\mathcal{C}_n$.

[3] For this, note that a sorting network essentially is a sequence of binary operations on the input, swapping two inputs if they are out of order. A swap of two booleans x, y can be trivially implemented by an $AND = \min\{x, y\}$ and an $OR = \max\{x, y\}$ gate.

[4] We construct an explicit probabilistic circuit here. However, the easy direction in the Yaos Principle [24] also directly implies a deterministic upper bound for the distributional problem.

The size of the entire circuit like that is $O(m \log m)$. We would like to see how to chose m and l in order for the circuit to reject $\mathcal{T}^-_\alpha$ and accept $\mathcal{T}^+_\beta$ with probability 3/4.

We will use the following well-known fact to analyze the above circuit construction.

▶ **Fact 31.** *Let $X_1, \dots X_m$ be Bernoulli random variables with $\mathbb{E}\, X_i \le p$ (not necessarily independent). Then $\mathbb{E} \sum X_i \le pm$, and by Markov Inequality, $\Pr(\sum X_i > 4pm) \le 1/4$.*

On the other hand, if $X'_1, \dots X'_m$ are independent Bernoulli random variables with $\mathbb{E}\, X'_i \ge 5p$, then as soon as $m \gtrsim 1/p\delta$, the Chebyshev inequality implies $\Pr(\sum X'_i \le 4pm) \le O(\delta^2)$.

We will use X_i to be a clique indicator on random l vertices under the distribution $\mathcal{T}^-_\alpha$, and X'_i to be a clique indicator on random l vertices under $\mathcal{T}^+_\beta$.

In what follows, we will discuss how to choose l in a way such that indeed $\mathbb{E}\, X'_i \ge 5\, \mathbb{E}\, X_i$, and take δ as a small constant to obtain a circuit distinguishing $\mathcal{T}^-_\alpha$ from $\mathcal{T}^+_\beta$. For random variables X_i (defined as clique indicators on the negative distribution), we use only the Markov inequality, and we do not need to worry about the correlations between them.

Let us quickly discuss that indeed, on the positive distribution, the random variables X'_i and X'_j are indeed independent – that is the case since conditioned on any specific realization G of $\mathcal{T}^+_\beta$ – i.e. G being a clique on random β vertices, the random variables X'_i and X'_j for $i \ne j$ are independent (since they are clique indicators on subsets of size l chosen independently at random), and have $\mathbb{E}[X_i|G] = \mathbb{E}[X_i]$ – the position of the original β vertices in a clique does not affect the probability that a random l vertices are the subset of those β vertices.

This fact gives a bound on the number of clique indicators needed to differentiate the graph distributions when the individual clique probabilities are sufficiently far apart.

Circuit Analysis

All we need to do now is to choose the parameter l in such a way that $q \ge 5p$, where $p := \mathbb{E}_{G \sim \mathcal{T}^-\alpha} \mathcal{K}_A(G)$ and $q := \mathbb{E}_{G \sim \mathcal{T}^+\beta} \mathcal{K}_A(G)$. The resulting circuit size will be of order $O(1/q)$.

Proof of Theorem 4. We consider circuit C_n constructed as discussed above, with the threshold $\tau := 9(\beta/n)^l$.

According to Lemma 30 and Fact 31 in order to make sure that the circuit indeed rejects $\mathcal{T}^-_\alpha$ and accepts $\mathcal{T}^+_\beta$ for given $\alpha \le \beta$, we would like to pick the smallest l such that

$$(\beta/n)^l \ge 5(n^{-\frac{2}{\alpha-1}})^{l^2}$$

leading to a circuit of size $O((n^{\frac{2}{\alpha-1}})^{l^2})$. After a sequence of simple algebraic manipulations and taking

$$\gamma := \frac{\alpha-1}{2} \frac{\log(n/\beta)}{\log n},$$

this condition is implied by

$$l \ge \gamma + C/\gamma$$

for some universal constant C, since in this case we have $l^2 \ge \gamma^2 + 2C$. Taking $l := \gamma + C/\gamma$, and observing that by assumption we have

$$\frac{\log(n/\beta)}{\log n} \le \delta$$

we obtain the circuit size

$$\text{size}(C_n) \le O(n^{\frac{2l^2}{\alpha-1}}) \le O(n^{\frac{2\gamma^2}{\alpha-1}}) \le O(n^{\alpha\delta^2/2}).$$ ◀

References

1 N. Alon and R. B. Boppana. The monotone circuit complexity of boolean functions. *Combinatorica 7*, 1987.

2 Ryan Alweiss, Shachar Lovett, Kewen Wu, and Jiapeng Zhang. Improved bounds for the sunflower lemma. *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, 2020.

3 Aleksandr Egorovich Andreev. A method for obtaining lower bounds on the complexity of individual monotone functions. In *Doklady Akademii Nauk*, volume 282(5), pages 1033–1037. Russian Academy of Sciences, 1985.

4 Alexander E Andreev. A method for obtaining efficient lower bounds for monotone complexity. *Algebra and Logic*, 26(1):1–18, 1987.

5 Tolson Bell, Suchakree Chueluecha, and Lutz Warnke. Note on sunflowers. *Discrete Mathematics*, 344(7):112367, 2021. doi:10.1016/j.disc.2021.112367.

6 Bruno Cavalar. Sunflower theorems in monotone circuit complexity. In *Anais do XXXIV Concurso de Teses e Dissertações*, pages 73–78, Porto Alegre, RS, Brasil, 2021. SBC. doi:10.5753/ctd.2021.15761.

7 Bruno P. Cavalar and Igor C. Oliveira. Constant-depth circuits vs. monotone circuits. In *Proceedings of the Conference on Proceedings of the 38th Computational Complexity Conference*, CCC '23, Dagstuhl, DEU, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2023.29.

8 Bruno Pasqualotto Cavalar, Mrinal Kumar, and Benjamin Rossman. Monotone circuit lower bounds from robust sunflowers. *Algorithmica*, 84(12):3655–3685, 2022. doi:10.1007/S00453-022-01000-3.

9 Susanna de Rezende and Marc Vinyals. Lifting with colorful sunflowers. *Unpublished manuscript*, 2025.

10 Paul Erdős and Richard Rado. Intersection theorems for systems of sets. *Journal of the London Mathematical Society*, 1960.

11 Ankit Garg, Mika Göös, Pritish Kamath, and Dmitry Sokolov. Monotone circuit lower bounds from resolution. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 902–911, 2018. doi:10.1145/3188745.3188838.

12 Mika Göös, Pritish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in monotone complexity and TFNP. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*, pages 38:1–38:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.

13 Danny Harnik and Ran Raz. Higher lower bounds on monotone size. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, STOC '00, pages 378–387, New York, NY, USA, 2000. Association for Computing Machinery. doi:10.1145/335305.335349.

14 Lunjia Hu, 2021. URL: https://theorydish.blog/2021/05/19/entropy-estimation-via-two-chains-streamlining-the-proof-of-the-sunflower-lemma/.

15 Stasys Jukna. Combinatorics of monotone computations. *Combinatorica*, 19(1):65–85, 1999. doi:10.1007/S004930050046.

16 Stasys Jukna et al. *Boolean function complexity: advances and frontiers*, volume 27. Springer, 2012. doi:10.1007/978-3-642-24508-4.

17 Anup Rao. Coding for Sunflowers. *Discrete Analysis*, February 2020. doi:10.19086/da.11887.

18 Anup Rao. Sunflowers: from soil to oil. *Bulletin of the American Mathematical Society*, 60(1):29–38, 2023.

19 Alexander Razborov. Lower bounds on the monotone complexity of some boolean function. In *Soviet Math. Dokl.*, volume 31, pages 354–357, 1985.

20 Benjamin Rossman. The monotone complexity of k-clique on random graphs. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 193–201, 2010. doi:10.1109/FOCS.2010.26.

21 Michel Talagrand. Regularity of Infinitely Divisible Processes. *The Annals of Probability*, 21(1):362–432, 1993. doi:10.1214/aop/1176989409.

22 Michel Talagrand. *Upper and lower bounds for stochastic processes*, volume 60. Springer, 2014.

23 Terrence Tao, 2020. URL: https://terrytao.wordpress.com/2020/07/20/the-sunflower-lemma-via-shannon-entropy/.

24 Andrew Chi-Chin Yao. Probabilistic computations: Toward a unified measure of complexity. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 222–227. IEEE Computer Society, 1977.