

Randomized versus Deterministic Decision Tree Size

Arkadev Chattopadhyay* Yogesh Dahiya† Nikhil S. Mande‡
 Jaikumar Radhakrishnan§ Swagato Sanyal¶

Abstract

A classic result of Nisan [SICOMP '91] states that the deterministic decision tree *depth* complexity of every total Boolean function is at most the cube of its randomized decision tree *depth* complexity. The question whether randomness helps in significantly reducing the *size* of decision trees appears not to have been addressed. We show that the logarithm of the deterministic decision tree size complexity of every total Boolean function on n input variables is at most the fourth power of the logarithm of its bounded-error randomized decision tree size complexity, ignoring a polylogarithmic factor in the input size. Our result has the following consequences:

- The deterministic AND-OR query complexity of a total Boolean function is at most the fourth power of its randomized AND-OR query complexity, ignoring a polylog n factor.
- The deterministic AND (OR) query complexity of a total Boolean function is at most the cube of its randomized AND (OR) query complexity, ignoring a polylog n factor. This answers a recent open question posed by Knop, Lovett, McGuire and Yuan [SIGACT News '21].
- The notion of *rank* of a Boolean function was defined in a classic work of Ehrenfeucht and Haussler [Information and Computation'89] in the context of learning theory, and is characterized by the logarithm of decision tree size up to a logarithmic factor in the input size. Our results confirm a recent conjecture (ignoring a polylog n factor) of Cornelissen, Mande and Patro [FSTTCS '22], that asserted the equivalence of randomized and deterministic analogs of rank, upto polynomial factors, for all total Boolean functions.
- Combined with the above-mentioned work of Ehrenfeucht and Haussler, our result implies that the class of functions computable by randomized decision trees of polynomial size, is PAC-learnable in quasi-polynomial time.

To obtain our main result on decision tree size, we introduce the notion of *block number* of a Boolean function, which can be thought of as a counting analog of *block sensitivity* of a Boolean function that played a central role in Nisan's result mentioned above.

*Tata Institute of Fundamental Research, Mumbai. Partially supported by the MATRICS grant MTR/2019/001633 of the Science & Engineering Research Board of the DST, India arkadev.c@tifr.res.in

†The Institute of Mathematical Sciences, Chennai yogeshdahiya@imsc.res.in

‡QuSoft and CWI, Amsterdam. Supported by the Dutch Research Council (NWO), as part of the Quantum Software Consortium programme (project number 024.003.037) nikhil.s.mande@gmail.com

§School of Technology and Computer Science, TIFR, Mumbai, and International Centre for Theoretical Sciences, TIFR, Bengaluru jaikumar@tifr.res.in

¶Indian Institute of Technology, Kharagpur. Supported by an ISIRD grant by SRIC, IIT Kharagpur swagato@cse.iitkgp.ac.in

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Consequences | 1 |
| 1.1.1 | (AND, OR)-decision trees | 2 |
| 1.1.2 | AND-decision trees | 2 |
| 1.1.3 | Decision tree rank | 3 |
| 1.2 | Our Techniques | 4 |
| 2 | Preliminaries | 5 |
| 2.1 | Combinatorial measures of Boolean functions | 5 |
| 2.2 | Decision trees for Boolean functions | 6 |
| 2.3 | Required results | 7 |
| 3 | Randomized versus deterministic size complexity | 8 |
| 3.1 | A better bound for monotone functions | 10 |
| 4 | (AND,OR) query complexity and AND query complexity | 10 |
| 4.1 | Randomized versus deterministic (AND, OR) query complexity | 10 |
| 4.2 | Randomized versus deterministic AND query complexity | 11 |
| 5 | Conclusions | 13 |
| A | Lower bounds for randomized decision tree size | 16 |
| A.1 | Communication complexity | 16 |
| A.2 | Lower bounds via communication complexity | 17 |
| A.3 | A template for proving randomized decision tree size lower bounds | 17 |
| A.3.1 | Monotone functions | 19 |
| A.3.2 | Symmetric functions | 20 |
| A.3.3 | EQUALITY | 20 |
| B | Proofs for (AND,OR) decision trees | 21 |
| C | An alternative proof of Lemma 4.6 | 22 |

1 Introduction

Understanding the power of randomness is of central interest in computer science. There are two main resources for algorithms: time and space. The role randomness plays in reducing requirements of time and space is under intensive investigation, but a resolution seems quite out of reach of current techniques. Decision trees are one of the most basic models of computation that has been studied for a long time. How much does randomness help algorithms in this model?

The two natural complexity measures of a decision tree are its depth, also called height, and its size. While depth has most probably been looked at more intensively, size also arises often, especially in the context of learning theory (see for example [16, 8, 7]). Small depth directly implies small size, i.e., a depth t decision tree has size at most 2^t . But the converse is well known to be false, as witnessed by the AND or OR function that can be computed with size $(n+1)$ but need depth also to be n . A classic work of Nisan [33], first published in the late eighties, showed that deterministic decision tree depth complexity is at most a cube of its randomized counterpart, for every total function.¹ The question about the extent of possible savings from the use of randomness in terms of the size of decision trees is natural. It remained unaddressed, somewhat surprisingly, since Nisan's work. It can be inferred from long-existing results that some savings is indeed possible. For instance, consider the function f known as the AND-OR tree of size n . This function is represented by a monotone formula of depth $\log n$, with binary AND and OR gates alternating in the layers. Saks and Wigderson [35] showed that the randomized depth complexity of f is $\Theta(n^{0.753\dots})$ which implies that the randomized size complexity of f is $2^{O(n^{0.753\dots})}$. Snir's [36] early and elegant work already implies that the deterministic decision tree size of f is $2^{\Theta(n)}$. This was later rediscovered by Jukna, Razborov, Savický and Wegener [27] via a spectral technique. While this gap may appear significant, in the log-scale the gap is still polynomial, i.e.,

$$\log(\text{DSize}^{\text{dt}}(f)) \geq (\log(\text{RSize}^{\text{dt}}(f)))^{1.32\dots},$$

where $\text{DSize}^{\text{dt}}(f)$ and $\text{RSize}^{\text{dt}}(f)$ denote deterministic and bounded-error randomized decision tree complexity of f , respectively. In the log-scale, as far as we know, this is the largest known gap between randomized and deterministic size.² Analogous to depth complexity, one may ask if the gap remains polynomial for every total function. Our main result answers this in the affirmative, ignoring polylog(n) factors.

Theorem 1.1. *For every total Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$,*

$$\log \text{DSize}^{\text{dt}}(f) = O((\log \text{RSize}^{\text{dt}}(f))^4 \log^3(n)).$$

1.1 Consequences

Theorem 1.1 complements Nisan's derandomization of decision tree depth. In this section, we describe three consequences of our main result, the first two of which are derandomization of depth in stronger models of decision trees. In such models, internal nodes of a tree are allowed to evaluate any function of the input bits from a given class of functions. The class of allowed functions defines the model. One such model is called Parity decision tree (PDT), where the class of allowed functions is just the set of all possible parities of the n input bits³ PDT's have gained significant interest, see for example [38, 23, 18, 12]. Interestingly, it is well known that randomness can provide largest possible savings in depth (and size) for PDT's. A simple adversary-based argument shows that deterministic PDTs need depth n to compute the OR of n bits. However, randomized PDTs compute OR with bounded error in just $O(1)$ depth.

In a recent interesting survey, Knop, Lovett, McGuire and Yuan [28] advocate looking at other natural classes of allowed functions, like AND and OR, besides PARITY. They mainly argue that such models act as natural and insightful intermediate models between ordinary decision trees and Yao's 2-party communication model. The latter is of great independent interest, and can also be naturally thought

¹Whether this cubic gap can be narrowed down further, is an area of active research with relatively recent breakthroughs [32, 1].

²We observe in Corollary A.4 that the AND-OR tree does not witness a stronger separation than this.

³Note that this is a generalization of ordinary decision trees as they can be simulated by merely evaluating the parity of just the relevant bit.

of as a decision tree, known as a protocol tree, with each node evaluating an arbitrary function that depends on just Alice’s or Bob’s input bits. In particular, if there exists a PDT or AND/OR decision tree evaluating a function f in small depth, then there exists a protocol tree of at most twice the depth for computing the same function, irrespective of the partition of the input bits between the players.

The study of these models of decision trees also has well-known connections to combinatorial group testing. The natural quantum analog of AND decision trees has been studied in this context [3] with surprising applications in proving classical lower bounds [5]. In another direction, PDTs (and more generally parity decision DAGs) are closely related to proof systems that allow resolution over equations in \mathbb{F}_2 , known as Res-Lin that is a generalization of ordinary resolution of clauses (see [34]). AND (OR) decision trees are a special case of linear threshold decision trees that are closely connected to the cutting planes proof system (see [24]). Both these proof systems are poorly understood [25, 17] and it is believed that a better understanding of PDT’s and threshold decision trees, especially w.r.t. search problems, should lead to progress here.

1.1.1 (AND, OR)-decision trees

The two cardinal functions that are hard for decision tree depth are AND and OR, which are two of the most basic Boolean functions. It is thus natural to look at decision trees where query nodes can evaluate AND’s or OR’s of arbitrary subsets of input bits. Note that the model is equivalent up to a factor of 2 in query cost, to the model where AND (OR) alone is allowed, but over *literals* (i.e., possibly negated variables) and not just positive variables. Such decision trees have been studied for long, for example in the early work of Ben-Asher and Newman [4] and more recently in the context of graph problems, see for example [6, 9].

Observing that (AND, OR)-decision tree depth is equivalent to ordinary decision tree size after ignoring polylogarithmic factors in the input size (see Lemma 4.3 and Lemma 4.4), we conclude from Theorem 1.1 that the randomized and deterministic depth in this model, denoted respectively by $R^{(\wedge, \vee)\text{-dt}}(\cdot)$ and $D^{(\wedge, \vee)\text{-dt}}(\cdot)$, are polynomially related as stated below formally.

Theorem 1.2. *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$,*

$$D^{(\wedge, \vee)\text{-dt}}(f) = O(R^{(\wedge, \vee)\text{-dt}}(f)^4 \log^7(n)).$$

1.1.2 AND-decision trees

A recent technique that has generated a lot of insight and solved several longstanding problems in two-party communication complexity is that of *lifting* decision tree depth complexity of a function/problem f to communication complexity of a composed problem. In this setting, each input bit of f is distributed among Alice and Bob via a gadget $g : \{0, 1\}^b \times \{0, 1\}^b \rightarrow \{0, 1\}$. When the gadget g has nice (obfuscating) property and its size b is appropriately large, lifting theorems assert that the randomized (deterministic) communication complexity of the composed problem is asymptotically $R^{\text{dt}}(f) \cdot R^{\text{cc}}(g)$ ($D^{\text{dt}}(f) \cdot D^{\text{cc}}(g)$). Observe that this corresponds roughly to the communication cost of the naive protocol that simulates the optimal decision tree algorithm for f , solving the i -th instance of g when the decision tree algorithm stipulates querying its i th input bit. These theorems are attractive as they reduce the task of understanding communication complexity to that of decision tree depth complexity, the latter usually being a lot simpler. Several such theorems have been developed, see for example [20, 13, 21, 11, 23, 28], exploiting convenient properties of the gadget g . A well-known challenge in this area is to reduce the size of the gadget to a constant (see for example [31]), independent of the input length of f , and still prove such lifting theorems. There are essentially two one-bit gadgets, XOR and AND (OR).

It is not difficult to see that when XOR is the gadget, then a communication protocol can simulate a PDT for f , a more powerful model than ordinary decision trees. This explains why the EQ function, a composition of AND and XOR, has small randomized communication complexity: AND has an $O(1)$ -depth randomized PDT. For the XOR gadget, a deterministic lifting theorem⁴ was relatively recently developed by Hatami, Hosseini and Lovett [23]. No randomized lifting theorem is known for the XOR gadget.

⁴Ideally, lifting theorems should work even when f is not only a total function but just a partial function or a relation. Many applications, particularly in circuit and proof complexity, require this generality of a lifting theorem. There are no such general lifting theorems known with a constant-size gadget.

For the AND gadget, similarly, a communication protocol can simulate an AND-decision tree (ADT). Very recently, Knop, Lovett, McGuire and Yuan [29] lifted deterministic ADT depth complexity of f to deterministic communication complexity of $f \circ \text{AND}$. They left open the problem of proving a randomized lifting theorem. In a follow-up survey [28] it was argued that this inability is linked to our lack of understanding of basic questions about randomized ADT complexity. They conjectured, however, that the randomized and deterministic ADT depth complexity of total functions are polynomially related. Using a key inequality we establish in the proof of our main result, Theorem 1.1, we are able to confirm this conjecture as stated below.

Theorem 1.3. *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, $D^{\wedge \text{dt}}(f) = O(R^{\wedge \text{dt}}(f)^3 \log^4(n))$.*

En route to proving the above, we establish a characterization, up to polynomial gap, of the randomized ADT depth complexity in terms of natural combinatorial quantities associated with a Boolean function. At an input $x \in \{0, 1\}^n$, a block $B \subseteq [n]$, is called a sensitive 0-block if $x_i = 0$ for every $i \in B$, and $f(x) \neq f(x \oplus 1_B)$, where $x \oplus 1_B$ is obtained from x by switching values of the bits in B to 1. A set of indices $H \subseteq [n]$ is called a 0-hitting set for f at x if it intersects every sensitive 0-block of x . The 0-hitting set complexity of f , denoted by $\text{HSC}_0(f)$, is the smallest number r , such that for every input $x \in \{0, 1\}^n$, there exists a 0-hitting set of size at most r . A natural fractional relaxation leads to a quantity called the 0-fractional hitting set complexity, denoted by $\text{FHSC}_0(f)$. We also consider the zero (one) cover number of f , denoted by $N_0(f)$ ($N_1(f)$), which is the minimum number of monochromatic subcubes needed to cover $f^{-1}(0)$ ($f^{-1}(1)$). The cover number of f , denoted by $N(f)$, is defined to be $N_0(f) + N_1(f)$. We show that $\log N(f)$ is bounded from above by $\tilde{O}(R^{\wedge \text{dt}}(f)^2)$,⁵ and that $\text{FHSC}_0(f) = O(R^{\wedge \text{dt}}(f))$ (see Proof of Theorem 4.5 for proofs of these inequalities). Finally we also show that $D^{\wedge \text{dt}}(f) = \tilde{O}(\text{FHSC}_0(f) \log N(f))$ (Lemma 4.7). Combining these together, we obtain the following combinatorial characterization of deterministic and randomized ADT depth complexity.

Theorem 1.4. *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$,*

$$\max \left\{ \Omega(\text{FHSC}_0(f)), \tilde{\Omega}(\log N(f)^{1/2}) \right\} = R^{\wedge \text{dt}}(f) \leq D^{\wedge \text{dt}}(f) = O(\text{FHSC}_0(f) \cdot (\log N_0(f) + \log N_1(f))).$$

Theorem 1.4 says that a function is hard for randomized ADTs if and only if its cover number is large or it has a large 0-fractional hitting set complexity. We note below two simple functions, one of which has large cover number and small 0-fractional hitting set complexity, and the other has small cover number but large 0-fractional hitting set complexity.

- $f = \text{OR}_n$: $\log N(f) = O(\log n)$, $\text{FHSC}_0(f) = n$, $R^{\wedge \text{dt}}(f) = \Omega(n)$.
- $f = \text{AND}_n \circ \text{OR}_2$: $\log N(f) = \Omega(n)$, $\text{FHSC}_0(f) = 2$, $R^{\wedge \text{dt}}(f) = \tilde{\Omega}(n)$.

On the other hand, the AND function has small cover number, as well as small 0-fractional hitting set complexity. This is not surprising as $D^{\wedge \text{dt}}(\text{AND}) = 1$.

1.1.3 Decision tree rank

Another measure of decision trees that has attracted interest recently is that of decision tree *rank* [15, 14]. This notion was introduced by Ehrenfeucht and Haussler [16] in the context of learning theory. The rank of a Boolean function f is defined as the minimum rank of a decision tree computing f . One can show that rank characterizes the logarithm of decision tree size up to a logarithmic factor (see [15, Proposition 2.5], for example). An analogous statement is easily seen to hold true for randomized decision tree size and the natural randomized variant of rank. Cornelissen, Mande and Patro [14] conjectured that rank and randomized rank are polynomially related for all total Boolean functions. By the equivalence of rank with the logarithm of decision tree size mentioned above, Theorem 1.1 resolves this conjecture in the affirmative up to a polylog n factor. Ehrenfeucht and Haussler showed that the class of functions on n variables with decision tree size polynomial in n is learnable from random examples in quasi-polynomial time. An immediate consequence of Theorem 1.1 is that even the class of functions on n variables with *randomized* decision tree size polynomial in n is learnable from random examples in quasi-polynomial time.

⁵Throughout this paper we use notations $\tilde{O}(\cdot)$, $\tilde{\Theta}(\cdot)$, $\tilde{\Omega}(\cdot)$ to hide polylogarithmic factors in the input size.

1.2 Our Techniques

As far as we are aware, the only previous work achieving full derandomization for a model of decision trees, is that of Nisan’s [33]. The key notion Nisan introduced for his beautiful argument is that of *block sensitivity* of a Boolean function f . The block sensitivity of f at an input x is just the maximum number of disjoint sensitive blocks of f at x . Another key notion required in Nisan’s argument is that of *certificate complexity*. The certificate complexity of f at x is the minimum number of bits of x one needs to reveal to force f to a constant. The certificate complexity of f is the maximum certificate complexity over all inputs. Nisan’s argument is completed essentially⁶ by two facts: first, the certificate complexity of f at any input x is bounded from above by the square of the block sensitivity of f . Using the older known fact, re-discovered by many authors in different contexts [22, 10, 37], showing that the square of the certificate complexity is an upper bound for deterministic depth, yields the fourth power of block sensitivity as an upper bound on the deterministic depth of f . Second, Nisan proved that block sensitivity is a valid lower bound on the randomized depth complexity of f .⁷ This is simply because f embeds at x a promised OR of arity as much as its block sensitivity at x .

Block sensitivity is a very influential notion, but is not helpful for us in proving randomized decision tree size lower bounds. For example, OR, AND have the largest possible block sensitivity, and even sensitivity, of n and yet have deterministic decision tree size merely $O(n)$. More generally, any measure certified by specifying the evaluation of f at s points of the Boolean cube can yield at most a lower bound of s on decision tree size. This is because there always exists a deterministic decision tree of size s that computes f restricted to those s points, as we describe now. We will construct a tree such that each leaf corresponds to a unique point amongst the s given points. Start with a single node, and as long as there exists a leaf ℓ which is reached by two distinct points $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$, query a variable i such that $x_i \neq y_i$ at ℓ . This way, ℓ is replaced by two other leaves, and x and y are separated. Continuing this way, we eventually end up with a tree with exactly s leaves, one for each point. We may then label the leaves with the function values of the corresponding points.

Consider the MAJORITY function, at the all zero point. It’s block sensitivity there is just 2. However, there are exponentially many, i.e., $\binom{n}{n/2}$, minimally sensitive blocks. These minimal blocks form an anti-chain with respect to inclusion?. Further, their minimality implies⁸ that each point in the cube obtained by flipping the bits of such a block yields a point where every (flipped) bit is sensitive. We first show (Lemma 3.2) that any large set S of such points, by itself, ensures that the randomized decision tree size is $\tilde{\Omega}(|S|)$. This holds more generally. Hence, a randomized decision tree computing a function f that has a large number of minimally sensitive blocks at some input x will be forced to have a large size.

To formalize this, we define the *block number* of a function f as the largest integer r for which there exists a point x where f has r minimally sensitive blocks. Thus, we are left to deal with functions that have small block number. These are functions f , which at every point x , have *few* minimally sensitive blocks. An example of such an f with large randomized size complexity is PARITY. In the second part of the argument (Lemma 3.3), we prove, using a simple boosting argument that for such an f , one would be able to extract a cover for both the ones and zeroes of f by monochromatic subcubes, using about as many subcubes as the size of the randomized decision tree. This is sufficient to construct a deterministic decision tree for f of not too large size, thanks to a classical result of Ehrenfeucht and Haussler [16] from the late eighties. This result states that there exists a deterministic decision tree of size at most $n^{\log^2 N(f)}$, where $N(f)$ is the minimal size of a cover of f by monochromatic subcubes. In particular, combining the above two parts yields the following interesting relationship between the cover number and randomized size complexity, as stated in Theorem 3.1, part (a):

$$\log N(f) = \tilde{O}\left(\log^2 \text{RSize}^{\text{dt}}(f)\right). \quad (1)$$

The derandomization of depth for (AND, OR)-decision trees follows directly, once we observe that such depth complexity of f is equivalent, ignoring a factor of $\log n$, to the logarithm of its ordinary decision tree size complexity. This equivalence is established using a simple tree-balancing argument that is possible as we have both AND and OR queries available.

⁶Here, for simplicity, we sketch the ideas achieving a fourth power gap, and not the cubic gap that Nisan proved.

⁷It was later shown [2] that the square root of block sensitivity is also a lower bound on the *quantum* analog of depth complexity (which is better known as quantum query complexity), showing that deterministic depth complexity and quantum query complexity are also polynomially related for all total Boolean functions.

⁸This property of minimally sensitive blocks is a well-known fact that Nisan’s argument also exploited.

It requires more work to deal with decision trees that have just AND (OR) queries. The main reason is that AND queries on unnegated variables are not sufficient to equate depth with the logarithm of ordinary size. This is illustrated by the fact, already mentioned before in Section 1.1.2, that the OR function has ordinary decision tree size only n and, yet, the ADT depth complexity of OR is full, i.e., n . However, as observed by Loff and Mukhopadhyay [30], the ADT depth complexity is equivalent, upto $\text{polylog}(n)$ factors, to the zero-depth complexity of ordinary decision trees. The zero-depth complexity of f is the smallest number r for which there exists a decision tree computing f in which every path from the root to a leaf encounters no more than r zeroes. Our main contribution here is an adaptation of Ehrenfeucht and Haussler’s argument to prove a convenient deterministic zero-depth upper bound in terms of the fractional zero-hitting set complexity of f and its cover number, as stated in Theorem 1.4. The second important observation is that the fractional zero-hitting set complexity is a lower bound on the randomized zero-depth complexity of f . Noting the fact that a decision tree of zero-depth d is also an ordinary decision tree of $O(n^d)$ size, we complete the derandomization of ADT depth using Eq. (1).

We remark here that the notion of fractional zero-hitting set complexity was introduced in [29]. However, the approach taken by the authors in that paper was algebraic, in the sense that they focused on the sparsity of the unique polynomial over the AND basis representing f . In fact, in their later survey [28], they suggested looking at the sparsity needed by polynomials to *approximate* f to tackle randomized ADT depth complexity. Our proof of their conjecture (our Theorem 1.3), however, deviates from this approach completely and involves purely combinatorial arguments.

2 Preliminaries

This section formally introduces the required notation, different query models and various complexity measures used in our proofs.

All logarithms in this paper are taken base 2. We use notations $\tilde{O}(\cdot)$, $\tilde{\Theta}(\cdot)$, $\tilde{\Omega}(\cdot)$ to hide polylogarithmic factors in the input size (and not just polylogarithmic factors in the argument). For a bit $b \in \{0, 1\}$, \bar{b} denotes the bit $1 - b$. For a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let $\bar{f} : \{0, 1\}^n \rightarrow \{0, 1\}$ be the Boolean function defined as $\bar{f}(x) = 1 - f(x)$ for all $x \in \{0, 1\}^n$. For a positive integer n , let $[n]$ denote the set $\{1, 2, \dots, n\}$. For $x, y \in \{0, 1\}^n$ we say $x \leq y$ if $x_i \leq y_i$ for all $i \in [n]$. A Boolean function f is *monotone* if $x \leq y$ implies $f(x) \leq f(y)$. A maxterm of a monotone Boolean function, $f : \{0, 1\}^n \rightarrow \{0, 1\}$, is a minimal set of variables $S \subseteq [n]$ such that setting the variables in S to 0 forces f to 0. A minterm of f is a minimal set of variables which, when set to 1 forces f to 1.

2.1 Combinatorial measures of Boolean functions

Definition 2.1 (Subcube). *A subcube of dimension d (equivalently, codimension $n - d$) is a subset of $\{0, 1\}^n$ obtained by fixing $n - d$ of the variables. In other words, a subcube is a set of all inputs consistent with a partial assignment of n bits. For $b \in \{0, 1\}$, a subcube A is said to be b -monochromatic with respect to f if $f(x) = b$ for all x in A . The codimension of a subcube A is denoted by $\text{co-dim}(A)$.*

Definition 2.2 (Cover number). *For a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $b \in \{0, 1\}$, we define its b -cover number, denoted by $N_b(f)$, as*

$$\min \{k : \exists b\text{-monochromatic subcubes wrt } f, S_1, \dots, S_k, \text{ such that } \bigcup_{i \in [k]} S_i = f^{-1}(b)\}.$$

The cover number of f , denoted by $N(f)$, is defined to be $N_0(f) + N_1(f)$.

For $x \in \{0, 1\}^n$ and $B \subseteq [n]$, let 1_B denote the n -bit string that is 1 on bits in B and 0 otherwise.

Definition 2.3 (Sensitive block, block sensitivity). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$. A set $B \subseteq [n]$ is a sensitive block of f at input x if $f(x \oplus 1_B) \neq f(x)$; a sensitive block at x is minimal if no proper subset of it is a sensitive block at x . The block sensitivity of f at x , denoted by $\text{bs}(f, x)$, is the maximum r for which there exist disjoint blocks, B_1, B_2, \dots, B_r , each of which is a sensitive block of f at x . The block sensitivity of f , denoted by $\text{bs}(f)$, is defined as*

$$\text{bs}(f) = \max_{x \in \{0, 1\}^n} \text{bs}(f, x).$$

An index $i \in [n]$ is sensitive for x with respect to f if $\{i\}$ is a sensitive block of f at x . For a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $x \in \{0, 1\}^n$, let $\mathcal{W}(f, x)$ denote the set of all minimal sensitive blocks of f . A set $B \subseteq [n]$ is a *sensitive 0-block* of f at x if $f(x \oplus 1_B) \neq f(x)$ and $x_i = 0$ for each $i \in B$. Let $\mathcal{W}_0(f, x)$ denote the set of all minimal sensitive 0-blocks of f at x .

Next we define hitting set complexity and a new measure which we call *block number*. For sets $S, F \subseteq [n]$, we say that S *hits* F if $S \cap F \neq \emptyset$.

Definition 2.4 (Block number, hitting set complexity). *For a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $x \in \{0, 1\}^n$, we define the block number of f at x , denoted by $\#\text{bs}(f, x)$, to be the number of elements in $\mathcal{W}(f, x)$; the block number of f is*

$$\#\text{bs}(f) = \max_{x \in \{0, 1\}^n} \#\text{bs}(f, x).$$

The hitting set complexity of f at x , denoted by $\text{HSC}(f, x)$, is the minimum size of a set $S \subseteq [n]$ that hits each $F \in \mathcal{W}(f, x)$. The hitting set complexity of f , denoted by $\text{HSC}(f)$, is defined as

$$\text{HSC}(f) = \max_{x \in \{0, 1\}^n} \text{HSC}(f, x).$$

Next we define 0-hitting set complexity.

Definition 2.5 (0-hitting set complexity). *For $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and input $x \in \{0, 1\}^n$, the 0-hitting set complexity of f at x , denoted $\text{HSC}_0(f, x)$, is the minimum size of a set $S \subseteq [n]$ that hits each $F \in \mathcal{W}_0(f, x)$. The 0-hitting set complexity of f , denoted by $\text{HSC}_0(f)$, is defined as*

$$\text{HSC}_0(f) = \max_{x \in \{0, 1\}^n} \text{HSC}_0(f, x).$$

To study AND decision trees (see Section 2.2), Knop, Lovett, McGuire and Yuan [29] considered a fractional version of hitting set complexity, which we define below.

Definition 2.6 (0-fractional hitting set complexity). *For $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and input $x \in \{0, 1\}^n$, the 0-fractional hitting set complexity of f at x , denoted $\text{FHSC}_0(f, x)$, is $1/p$, where $p > 0$ is the smallest number for which there exists a distribution \mathcal{D} on indices $i \in [n]$ with the property that $\Pr_{i \sim \mathcal{D}}[i \in F] \geq p$ for each $F \in \mathcal{W}_0(f, x)$. Define $\text{FHSC}_0(f) = \max_{x \in \{0, 1\}^n} \text{FHSC}_0(f, x)$.*

We remark here that while we use different notation, Definition 2.5 and Definition 2.6 are the same as in [29, Definition 2.10] and [29, Definition 2.14], respectively.

2.2 Decision trees for Boolean functions

A deterministic decision tree (DT) on n variables is a binary tree whose internal nodes are labeled by variables and leaves are labeled $\{0, 1\}$. Every internal node has a left child and a right child; the edge leading to the left child is labeled 0 and the edge leading to the right child is labeled 1. On an input $x \in \{0, 1\}^n$, the tree's computation proceeds from the root down to a leaf as follows. When the computation reaches a certain node, the variable associated with that node is *queried*: if the value obtained is 0, the computation moves to the left child, otherwise it moves to the right child. The output of a decision tree T on input x , denoted by $T(x)$, is the label of leaf node reached by this computation. We say that a decision tree T computes the function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ (or T is a decision tree for f), if $T(x) = f(x)$ for all $x \in \{0, 1\}^n$. The deterministic decision tree depth complexity of f , is

$$\text{D}^{\text{dt}}(f) := \min_{T: T \text{ computes } f} \text{Depth}(T).$$

Analogously the deterministic decision tree size complexity of f is

$$\text{DSize}^{\text{dt}}(f) := \min_{T: T \text{ computes } f} \text{Size}(T),$$

where $\text{Size}(T)$ is the number of leaves in T .

A randomized query algorithm/decision tree \mathcal{A} is a distribution $\mathcal{D}_{\mathcal{A}}$ over deterministic decision trees. The computation of \mathcal{A} on input x proceeds by sampling a deterministic decision tree T according to $\mathcal{D}_{\mathcal{A}}$,

and outputs the label of the leaf reached by T on x . \mathcal{A} computes f with error at most ε if for every input x , the probability that $A(x) = f(x)$ is at least $1 - \varepsilon$. The cost of the randomized algorithm is measured by the number of worst-case queries made by \mathcal{A} on any input x , i.e., the maximum depth over all decision trees in the support of the distribution. The ε -error randomized decision tree complexity of f , denoted $R_\varepsilon^{\text{dt}}(f)$, is the minimum cost of such a randomized decision tree. The randomized decision tree size of f , denoted by $\text{RSize}_\varepsilon^{\text{dt}}(f)$, is defined similarly. That is,

$$\text{RSize}_\varepsilon^{\text{dt}}(f) = \min_{\substack{\mathcal{A} \text{ computing } f \\ \text{with error } \leq \varepsilon}} \max_{T: \mathcal{D}_\mathcal{A}(T) > 0} \text{Size}(T).$$

Each internal node of a standard decision tree corresponds to a variable being queried. Generalizing the class of permitted queries gives rise to many variants of decision trees that have been considered in different contexts. In each variant, the relevant deterministic and randomized decision tree complexities are defined just as above.

In our work, we are interested in three such classes.

- When the permitted queries are AND of (non-negated) variables, we refer to such trees as *AND-decision trees (ADT)*. We let $D^{\wedge \text{dt}}(\cdot)$ and $R_\varepsilon^{\wedge \text{dt}}(\cdot)$ denote the deterministic and ε -error randomized decision tree complexities, respectively, in this model.
- When the permitted queries are AND of a subset of variables (AND query) or an OR of a subset of variables (OR query), we refer to such trees as *(AND, OR)-decision trees ((AND, OR)-DT)*. We let $D^{(\wedge, \vee)\text{-dt}}(\cdot)$ and $R_\varepsilon^{(\wedge, \vee)\text{-dt}}(\cdot)$ denote the deterministic and ε -error randomized decision tree complexities, respectively, in this model.
- When the permitted queries are just variables, but the depth of a path is measured by the number of queries answered as 0 on the path, we refer to deterministic and ε -error randomized decision tree complexities as $D^{0\text{-dt}}(\cdot)$ and $R_\varepsilon^{0\text{-dt}}(\cdot)$, respectively. We use the term ‘0-depth’ of a decision tree to refer to the longest root-to-leaf path where only 0-labelled edges contribute to the length.

Throughout this paper, when we drop ε from the subscript of a randomized decision tree (possibly with generalized queries as above) measure, we assume $\varepsilon = 1/3$.

Definition 2.7 (Query tree, query set, transcript). *A deterministic query tree on n variables is a rooted binary tree where every node is labelled by a variable. Unlike a decision tree, a query tree does not return the value of any function at the leaves. (It represents a strategy for querying variables with the goal of finding a certificate.) By dropping its leaves, we regard a decision tree for a function as a query tree. A randomized query tree is modelled as a distribution over deterministic query trees. For a query tree T and an input $x \in \{0, 1\}^n$, let $Q_T(x)$ be the set of indices of variables queried by T on input x ; the transcript of T on input x is the sequence $\sigma_T(x) = ((i_t, x[i_t]) : t = 1, 2, \dots)$, where i_t is the index of the input bit of x probed by T in the t -th step. For a randomized query tree T , let $T^{\otimes \ell}$ be the randomized query tree corresponding to performing ℓ independent computations of the query tree T . We use $\text{Size}(T)$ to denote the number of distinct transcripts of T over all inputs $x \in \{0, 1\}^n$.*

2.3 Required results

This section contains some required results that give bounds on some combinatorial measures and on certain decision tree measures of Boolean functions defined in the earlier two subsections.

Cover number gives a lower bound on deterministic decision tree size. Each leaf of a decision tree corresponds to a monochromatic subcube generated by the partial assignments defined by the root to leaf path. Thus a deterministic decision tree computing a function f induces a monochromatic subcube partition of f , giving the following.

Proposition 2.8. *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, $\text{DSize}^{\text{dt}}(f) \geq N(f)$.*

In the other direction, Ehrenfeucht and Haussler [16, Lemma 1, Lemma 6] gave the following upper bound on the logarithm of the decision tree size of a function f . The version we state below can be found in [26, Theorem 14.32], for example.

Theorem 2.9 ([26, Theorem 14.32]). *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$,*

$$\log \text{DSize}^{\text{dt}}(f) = O(\log^2 N(f) \cdot \log n).$$

Lemma 2.10. *Suppose $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $y \in \{0, 1\}^n$. Suppose T is an ε -error randomized decision tree for f , $B \in \mathcal{W}(f, y)$ is a sensitive block of f at y , and $W \subseteq \mathcal{W}(f, y)$ is a set of sensitive blocks of f at y . Then,*

- (a) $\Pr[Q_T(y) \cap B = \emptyset] \leq 2\varepsilon$.
- (b) $\Pr[Q_{T^{\otimes \ell}}(y) \cap B = \emptyset] \leq (2\varepsilon)^\ell$.
- (c) $\Pr[Q_{T^{\otimes \ell}}(y) \text{ is a hitting set for } W] \geq 1 - (2\varepsilon)^\ell |W|$.

Proof. For part (a), we have

$$\begin{aligned} 1 &= \Pr[f(y) \neq f(y \oplus 1_B)] \\ &\leq \Pr[f(y) \neq T(y)] + \Pr[T(y) \neq T(y \oplus 1_B)] + \Pr[T(y \oplus 1_B) \neq f(y \oplus 1_B)] \\ &\leq 2\varepsilon + \Pr[T(y) \neq T(y \oplus 1_B)]. \end{aligned}$$

Thus, $\Pr[T(y) \neq T(y \oplus 1_B)] \geq 1 - 2\varepsilon$. Then,

$$\Pr[Q_T(y) \cap B = \emptyset] \leq \Pr[T(y) = T(y \oplus 1_B)] = 1 - \Pr[T(y) \neq T(y \oplus 1_B)] \leq 2\varepsilon.$$

Part (b) follows from part (a) because $T^{\otimes \ell}$ corresponds to ℓ independent computations performed on T ; in particular, $Q_{T^{\otimes \ell}}(y) = Q_1 \cup Q_2 \cup \dots \cup Q_\ell$, where the Q_i are independent and each has the same distribution as $Q_T(y)$. By a union bound over all $B \in W$, we have

$$\Pr[\exists B \in W : Q_{T^{\otimes \ell}}(y) \cap B = \emptyset] \leq (2\varepsilon)^\ell |W|.$$

Part (c) immediately follows. □

Proposition 2.11. *Q contains a certificate for f at x iff Q intersects every block in $\mathcal{W}(f, x)$.*

Proof. Suppose Q contains a certificate for f at x . If Q does not intersect a block B in $\mathcal{W}(f, x)$, then $y = x \oplus 1_B$ and x agree on all the positions in Q but $f(x) \neq f(y)$ —a contradiction; thus Q intersects every block in $\mathcal{W}(f, x)$. Next suppose Q intersects every block in $\mathcal{W}(f, x)$. Suppose for some y , we have $f(x) \neq f(y)$; we will show that x and y differ in some position in Q . Now, $x \oplus y$ is a characteristic vector of a sensitive block of f at x , and must include in it a block in $\mathcal{W}(f, x)$; Q intersects this block, so x and y differ in a position in Q . □

AND query complexity is tightly related to 0-depth query complexity, which was introduced by Loff and Mukhopadhyay [30].

Claim 2.12 ([30], [29, Claim 4.4]). *For every Boolean function f , we have $D^{0\text{-dt}}(f) \leq D^{\wedge\text{dt}}(f) \leq D^{0\text{-dt}}(f) \log n$ and $R^{0\text{-dt}}(f) \leq R^{\wedge\text{dt}}(f) \leq R^{0\text{-dt}}(f) \log n$.*

While Loff and Mukhopadhyay [30], and Knop, Lovett, McGuire and Yuan [29] only deal with deterministic trees, the analogous result for randomized trees can easily be seen to hold true by applying the relevant deterministic result to all trees in the support of the relevant randomized tree.

3 Randomized versus deterministic size complexity

In this section, we present our main result; we establish an upper bound for $N(f)$ and $\text{DSize}^{\text{dt}}(f)$ in terms of $\text{RSize}^{\text{dt}}(f)$. In the next section, the first upper bound will provide a crucial ingredient in our proof of a similar result for decision trees with (AND,OR) queries and AND queries.

Theorem 3.1 (Main result). *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, we have the following.*

- (a) $\log N(f) = O((\log \text{RSize}^{\text{dt}}(f))^2 \log(n))$.

$$(b) \log \text{DSize}^{\text{dt}}(f) = O((\log \text{RSize}^{\text{dt}}(f))^4 \log^3(n)).$$

As we stated in the introduction, the notion of block number of a Boolean function plays a central role in our argument. In the following two lemmas the block number appears in two different roles, first as a lower bound and then as an upper bound.

Lemma 3.2. *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$. Then, for all $x \in \{0, 1\}^n$, we have*

$$\text{RSize}_{1/3}^{\text{dt}}(f)^{2 \log n} \geq \left(\frac{1}{2}\right) \#bs(f, x).$$

Lemma 3.3. *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$. Then,*

$$N(f) \leq n \cdot \text{RSize}_{1/3}^{\text{dt}}(f)^{2 \log \#bs(f)}.$$

We first verify that our main result follows from these lemmas; we then justify these lemmas formally.

Proof of Theorem 3.1. We assume $\text{RSize}_{1/3}^{\text{dt}}(f) \geq 2$. This is because if $\text{RSize}_{1/3}^{\text{dt}}(f) = 1$, then f must be a constant function, and the theorem follows trivially. From Lemma 3.3 and Lemma 3.2 we obtain

$$\begin{aligned} \log N(f) &\leq \log n + \log \text{RSize}_{1/3}^{\text{dt}}(f) \cdot 2 \log \#bs(f); \\ &\leq \log n + 2 \log \text{RSize}_{1/3}^{\text{dt}}(f) (\log \text{RSize}_{1/3}^{\text{dt}}(f) \cdot 2 \log n + 1) \\ &= O((\log \text{RSize}_{1/3}^{\text{dt}}(f))^2 \cdot \log n). \end{aligned}$$

Part (b) then follows from Theorem 2.9. □

We now prove the lemmas.

Proof of Lemma 3.2. Let T be an $(1/3)$ -error randomized decision tree for f . Fix $x \in \{0, 1\}^n$. For $B \in \mathcal{W}(f, x)$, let $y_B = x \oplus 1_B$. Since B is minimal, we have $f(y_B) \neq f(y_B \oplus 1_{\{i\}})$ for all $i \in B$; that is, $\{i\} \in \mathcal{W}(f, y_B)$. With $\ell = \lfloor 2 \log n \rfloor$ (assume n is large) and $W = \{\{i\} : i \in B\}$, we conclude from Lemma 2.10 (c) that

$$\Pr[Q_{T^{\otimes \ell}}(y_B) \supseteq B] \geq 1 - \left(\frac{2}{3}\right)^\ell |B| \geq 1 - \left(\frac{1}{2n}\right) |B| \geq \frac{1}{2}.$$

Thus, using linearity of expectation, we may fix a deterministic tree T^* in the support of $T^{\otimes \ell}$ such that for at least half of the $\#bs(f, x)$ choices for B , we have $Q_{T^*}(y_B) \supseteq B$; let W' be the set of these B . For each $B \in W'$, let $\sigma_B = \sigma_{T^*}(y_B)$ be the transcript of T^* on input y_B . We make two remarks.

- (a) Since the blocks in $\mathcal{W}(f, y)$ are minimal, for distinct $B, B' \in \mathcal{W}(f, x)$, there is an $i \in B \setminus B'$; for this i , $y_B[i] = x[i] + 1 \neq x_i = y_{B'}[i] \pmod{2}$. (There is a similar $i' \in B' \setminus B$, but we won't need it.) Thus, for distinct $B, B' \in W'$, we have $\sigma_B \neq \sigma_{B'}$.
- (b) Since T^* is obtained by picking ℓ trees in the support of T and running them one after the other, the total number of possibilities for σ_B is at most $\text{Size}(T)^\ell$.

Since T was an arbitrary $(\frac{1}{3})$ -error decision tree for f , we conclude from the above remarks that $\text{RSize}_{1/3}^{\text{dt}}(f)^{2 \log n} \geq |W'| \geq \frac{1}{2} \#bs(f, x)$. □

Proof of Lemma 3.3. Let T be a $(\frac{1}{3})$ -error randomized decision tree for f . Fix $x \in \{0, 1\}^n$, and let $W = \mathcal{W}(f, x)$. Recall from Proposition 2.11 that a set Q contains a certificate for f at x iff Q intersects every set in $\mathcal{W}(f, x)$. Now, set $\ell = 2 \log \#bs(f)$ and conclude from Lemma 2.10 (c) that

$$\begin{aligned} \Pr[Q_{T^{\otimes \ell}}(x) \text{ contains a certificate for } f \text{ at } x] &= \Pr[Q_{T^{\otimes \ell}}(x) \text{ is a hitting set for } \mathcal{W}(f, x)] \\ &\geq 1 - \left(\frac{2}{3}\right)^\ell \#bs(f) > \frac{1}{2}. \end{aligned}$$

Let T_1, T_2, \dots, T_n be n trees each picked with the same distribution as $T^{\otimes \ell}(x)$. Then, using a union bound over all $x \in \{0, 1\}^n$, we conclude that with non-zero probability, for all $x \in \{0, 1\}^n$, there is an i

such that $Q_{T_i}(x)$ contains a certificate of x . Fix a choice $T_1^*, T_2^*, \dots, T_n^*$ of deterministic query trees in the support of $T^{\otimes \ell}(x)$ such that for all x , there is an i such that $Q_{T_i^*}(x)$ includes a certificate of f at x . Let $\sigma_i(x) = \sigma_{T_i^*}(x)$ be the transcript of T_i^* on input x ; let $C_{i,x} = \{y : \sigma_i(x) = \sigma_i(y)\}$. Note that each $C_{i,x}$ is a subcube; furthermore, if $Q_{T_i^*}(x)$ contains a certificate for x , then $C_{i,x}$ is monochromatic. For every i , the number of possibilities for $\sigma_i(x)$ and hence $C_{i,x}$ is at most $\text{Size}(T)^\ell$. Then,

$$\{C_{i,x} : i \in [n], x \in \{0,1\}^n \text{ and } Q_{T_i^*}(x) \text{ contains a certificate for } x\}$$

is a cover for f with at most $n \cdot \text{Size}(T)^\ell$ subcubes. Since T was an arbitrary $(\frac{1}{3})$ -error protocol, the lemma follows. \square

3.1 A better bound for monotone functions

In this section, we observe that for monotone functions we can obtain better bounds than those provided by Theorem 3.1.

Theorem 3.4 (Improvement for monotone functions). *Suppose $f : \{0,1\}^n \rightarrow \{0,1\}$ is monotone. Then, we have the following.*

- (a) $\log N(f) = O(\log \text{RSize}^{\text{dt}}(f) \log(n))$;
- (b) $\log \text{DSize}^{\text{dt}}(f) = O((\log \text{RSize}^{\text{dt}}(f))^2 \log^3(n))$.

Proof. Note that part (b) follows from part (a) and Theorem 2.9. Part (a) clearly holds if f is a constant function, so we restrict attention to functions f that are not constant. Note that $\#\text{bs}(f, 0^n) = N_1(f)$, for we have a distinct minimal sensitive block from each minterm; similarly, $\#\text{bs}(f, 1^n) = N_0(f)$, for we have one minimal sensitive block from each maxterm. It follows that $N(f) \leq 2 \cdot \#\text{bs}(f)$. Then,

$$\log N(f) \leq \log(\#\text{bs}(f)) + 1 = O(\log \text{RSize}^{\text{dt}}(f) \log n),$$

where the last inequality follows from Lemma 3.2. \square

4 (AND,OR) query complexity and AND query complexity

In the previous section, we established that for every total function, there is a deterministic decision tree whose size is bounded by a quasi-polynomial of the *size* of its smallest randomized decision tree (ignoring some polynomial factors involving $\log n$ in the exponent). Though this result referred to decision tree *size*, we now use it to derive consequences for the randomized and deterministic *depth* in the (AND, OR) decision tree model, and in the AND (OR) decision tree model.

4.1 Randomized versus deterministic (AND, OR) query complexity

The following is our main theorem in this subsection, that shows that randomness does not offer much advantage in the (AND,OR) decision tree model.

Theorem 4.1. *For every Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$,*

$$D^{(\wedge, \vee)\text{-dt}}(f) = O(R^{(\wedge, \vee)\text{-dt}}(f)^4 \log^7(n)).$$

The translation is achieved by showing that for all functions, the logarithm of the *size* of the best ordinary decision tree and the depth of an (AND,OR) decision tree are closely related.

Lemma 4.2. *For every Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$, we have*

$$\log \text{DSize}^{\text{dt}}(f) / (2 \log n) \leq D^{(\wedge, \vee)\text{-dt}}(f) \leq 4 \log \text{DSize}^{\text{dt}}(f), \quad (2)$$

$$\log \text{RSize}^{\text{dt}}(f) / (2 \log n) \leq R^{(\wedge, \vee)\text{-dt}}(f) \leq 4 \log \text{RSize}^{\text{dt}}(f). \quad (3)$$

Proof of Theorem 4.1. We use the second inequality of Eq. (2) and the first inequality of Eq. (3). We then have

$$\begin{aligned} D^{(\wedge, \vee)\text{-dt}}(f) &\leq 2 \log \text{DSize}^{\text{dt}}(f) && \text{(by the second inequality in Eq. (2))} \\ &= O((\log \text{RSize}^{\text{dt}}(f))^4 (\log n)^3) && \text{(by Theorem 3.1 (b))} \\ &= O(\text{R}^{(\wedge, \vee)\text{-dt}}(f)^4 \log^7(n)) && \text{(by the first inequality in Eq. (3)).} \end{aligned}$$

□

It remains to establish Lemma 4.2. We need two observations, which we state in the following lemmas. We first observe that if (AND,OR) queries are allowed, then an arbitrary decision tree can be *balanced* so that the depth of the resulting (AND,OR) decision tree is logarithmic in the size of the original tree.

Lemma 4.3. *Suppose T is an ordinary decision tree for a function f . Then, there is an (AND, OR) decision tree C for f such that.*

$$\text{Depth}^{(\wedge, \vee)\text{-dt}}(C) \leq 4 \log \text{DSize}^{\text{dt}}(T).$$

This lemma follows using a tree-balancing argument that has been used in various settings (Boolean formulas, communication protocols, etc.) previously. A formal justification appears in Appendix B. Note that Lemma 4.3 immediately implies the second inequality in Eq. (2) and also Eq. (3). To justify the first inequality, we show that (AND,OR) decision trees can be converted to ordinary decision trees whose size is bounded by an exponential in the depth of the original tree. Our justification of Lemma 4.2 will be complete once we show the following.

Lemma 4.4. *For every (AND, OR) decision tree C of depth d computing f , there is an ordinary decision tree T computing f such that*

$$\log \text{Size}(T) \leq d \log \left(\frac{ne}{d} \right).$$

Again we provide a formal justification in Appendix B. An informal argument goes as follows. Replace each AND query in the tree by a skewed ordinary decision tree for computing AND. Label the edges corresponding to the value 0 as RED and the other edges as BLUE. Similarly, replace each OR query by a skewed ordinary decision tree, but this time label the edge with label 1 as RED and the edges with label 0 as BLUE. In the resulting tree, we may eliminate nodes where a variable queried earlier is queried again. Thus we obtain an ordinary decision tree with depth at most n and at most d RED edges on any root-to-leaf path. The number of paths to a leaf in this tree can be bounded by the number of sequences in $\{\text{RED}, \text{BLUE}\}^n$ with at most d RED edges; any such sequence can be completed to a sequence with exactly $n + d$ elements, with exactly d REDS; thus, the number of leaves is at most $\binom{n+d}{d} \leq \left(\frac{2en}{d}\right)^d$.

4.2 Randomized versus deterministic AND query complexity

In the previous subsection, we saw an application of our main result to (AND,OR) decision tree depth. That application crucially depended on our ability to balance decision trees by using (AND,OR) queries. With just AND queries, we cannot use the same argument. Nevertheless, in this section, we show that randomness does not offer much advantage even in the AND decision tree model.

Theorem 4.5. *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, $D^{\wedge\text{dt}}(f) = O(\text{R}^{\wedge\text{dt}}(f)^3 \log^4(n))$.*

Recall from Claim 2.12 that AND decision tree depth complexity is bounded from below and above (up to an $O(\log n)$ factor) by 0-depth complexity, both in the deterministic and randomized settings. In this section, we will work mainly with 0-depth complexity and use this connection to derive our conclusion on AND decision tree depth complexity.

Once again, our proof relies on two observations. First, a randomized decision tree of small 0-depth provides us a way of efficiently hitting the 0-blocks of f at any input. We formalize this using the notion of 0-fractional hitting set complexity.

Lemma 4.6. $\text{FHSC}_0(f) = O(\text{R}^{0\text{-dt}}(f))$.

Our next observation allows us to build 0-depth decision trees from 0-fractional hitting sets. This argument, which can be regarded as the analog of the result of Ehrenfeucht and Haussler [16] (see Theorem 2.9), is the main technical argument of this section.

Lemma 4.7. $D^{0\text{-dt}}(f) = O(\text{FHSC}_0(f)(\log N_0(f) + \log N_1(f)))$.

Before we formally justify these observations, let us see how they imply Theorem 4.5.

Proof of Theorem 4.5. We have

$$\begin{aligned}
D^{\wedge\text{dt}}(f) &\leq D^{0\text{-dt}}(f) \log n && \text{(by Claim 2.12)} \\
&= O(\text{FHSC}_0(f) \cdot (\log N_0(f) + \log N_1(f))) \log n && \text{(by Lemma 4.7)} \\
&= O(R^{0\text{-dt}}(f) \cdot (\log N_0(f) + \log N_1(f)) \log n) && \text{(by Lemma 4.6)} \\
&= O(R^{0\text{-dt}}(f) \cdot \log N(f) \cdot \log n) \\
&= O(R^{0\text{-dt}}(f) \cdot (\log \text{RSize}^{\text{dt}}(f))^2 \cdot \log^2 n) && \text{(by Theorem 3.1 (a))} \\
&\stackrel{(*)}{=} O(R^{0\text{-dt}}(f)^3 \cdot \log^4 n) && \text{(by Eq. (4), see below)} \\
&= O(R^{\wedge\text{dt}}(f)^3 \cdot \log^4 n). && \text{(Claim 2.12)}
\end{aligned}$$

The inequality marked (*) holds because a decision tree T on n variables (that queries no variable twice on the same path) with 0-depth d has at most $\binom{n+d}{d}$ leaves (see Lemma 4.4); thus

$$\log \text{RSize}^{\text{dt}}(f) = O(R^{0\text{-dt}}(f) \log n). \quad (4)$$

□

Proof of Lemma 4.7. We first present an informal overview. Suppose $\text{FHSC}_0(f) = 1/p$. In particular, for every input z and a (multi-)set W of sensitive 0-blocks of f at z , there is a variable, by an averaging argument, that hits at least a fraction p of the blocks in W . Our strategy for building the decision tree for f is as follows. After having queried some variables, we restrict attention to the unset variables, and the function induced on these variables. If the restricted function is a constant function, then we announce the value. Otherwise, we present a specific *reference input* (consistent with the answers received to the queries so far) and consider a (multi-)set W of sensitive 0-blocks of this input. The property above ensures that there is a variable that lies in a fraction p of these blocks. We query this variable; the choice of the variable will ensure that if this bit is 0 (in the actual input, not in the reference input), then after querying this bit, the size of either the 0-cover or the 1-cover for the restricted function will be at most a factor $(1-p)$ of what it was before the current step. Thus, on any input, by the time $t^* = \lceil (1/p) \ln(N_0(f) \cdot N_1(f)) \rceil$ bits read turn out to be zero, either the 0-cover number or the 1-cover number for the restricted function will have become zero, thereby making the restricted function constant.

To implement this strategy, we need to specify how the following are determined at each step. Suppose we wish to compute $f(x)$.

Reference input: Our reference input z will initially be $(0, 0, \dots, 0)$. Suppose the input bit $x[i]$ is queried in a certain step; we obtain the next reference input by changing that bit of the current reference input from 0 to $x[i]$. Thus, at every stage, the reference input is consistent with the part of x discovered by previous queries; the variables not yet queried continue to be 0.

The set of 0-blocks: Let $S \subseteq [n]$ be the indices of variables queried so far and $\rho : S \rightarrow \{0, 1\}$ be their evaluations. Restrict attention to the variables with indices not in S . Assume that this restricted function is not constant, that is, the optimal size of 0-cover and 1-cover are at least 1. Note that the reference input z assigns zeros to all the unset variables. We have two cases.

- Suppose $f(z) = 0$: For each subcube C in the 1-cover consistent with the queried variables (for $i \in S$, $x[i]$ is unset or set to $\rho(i)$ in C), identify it with a conjunction of the form

$$\bigwedge_{j \in B_C} x[j] \wedge \bigwedge_{j \in B'_C} \overline{x[j]}, \quad (5)$$

where $B_C = \{i : i \notin S \text{ and } x[i] \text{ is set to 1 in } C\}$ and $B'_C = \{i : i \notin S \text{ and } x[i] \text{ is set to 0 in } C\}$. Since $f(z) = 0$, each B_C must be non-empty; furthermore, each B_C is a sensitive 0-block for f at z . We let W be the multiset $\{B_C : C \text{ is a subcube in the 1-cover consistent with the queried variables}\}$.

- Suppose $f(z) = 1$: We consider subcubes in the optimal 0-cover consistent with the queried variables, and identify each of them with a conjunction of the form (5). We obtain a multi-set W of sensitive 0-blocks for the reference input z as before.

The bit to be queried next: By construction, each block in W is a sensitive 0-block for the reference input z . So by the property stated at the outset, there is a variable $x[i]$ that belongs to at least a fraction p of the blocks in W (counted with multiplicity). Note that if this variable is queried and the input bit $x[i]$ is revealed to be zero, then the subcubes C in the cover corresponding to blocks in W where $i \in B_C$ do not contribute to the cover for the new restricted function.

At each step, if the value of f has not yet been determined, we determine the next variable to be queried as described above. This completes the description of our strategy. It remains to bound the 0-depth of the decision tree corresponding to this strategy. Let N_0^t and N_1^t denote the number of subcubes surviving (consistent with the queried variables) after t queries. Note that surviving subcubes form a cover for the restricted function, although it need not be optimal. Consider the quantity $M_t = N_0^t \cdot N_1^t$. Initially, $M_0 = N_0(f) \cdot N_1(f)$; and in general $M_t \leq M_0(1-p)^t$. Set $t^* = \lceil (1/p) \ln(N_0(f) \cdot N_1(f)) \rceil$. Suppose we read at least $t^* + 1$ zeros; then $M_{t^*} < 1$, that is, $M_{t^*} = 0$ (one of the covers has become empty); but then we would have stopped when we read t^* zeros—a contradiction. It follows that the 0-depth of the decision tree underlying this strategy is $O((1/p) \log(N_0(f) \cdot N_1(f)))$. \square

Proof of Lemma 4.6. Let T be an ε -error randomized decision tree for f with 0-depth d . We will show that for each $x \in \{0, 1\}^n$, we have

$$\text{FHSC}_0(f, x) \leq \frac{d}{1-2\varepsilon}. \quad (6)$$

Our claim will follow from this, because T was fixed arbitrarily. By Lemma 2.10, for each block $B \in \mathcal{W}_0(f, x)$, we have $\Pr[Q_T(x) \cap B \neq \emptyset] \geq 1-2\varepsilon$. Let $Q_T^0(x) = \{i \in Q_T(x) : x[i] = 0\}$; note that $|Q_T^0(x)| \leq d$. Since B is a zero block of f at x , we have that $x[i] = 0$ for each $i \in B$; thus, $\Pr[Q_T^0(x) \cap B \neq \emptyset] \geq 1-2\varepsilon$. Let the index I be chosen uniformly from $Q_T^0(x)$. Then, for each $B \in \mathcal{W}_0(f, x)$, we have $\Pr[I \in B] \geq (1-2\varepsilon)/d$. This establishes Eq. (6): we take the distribution of I as the distribution D required in the definition of $\text{FHSC}_0(f, x)$ (see Definition 2.6). \square

5 Conclusions

Nisan's result from the late eighties [33] can be viewed as the analog of derandomizing time for decision trees. The complexity measure of logarithm of decision tree size may naturally be thought of as the analog of space for decision trees. In this light, we would like to view our main theorem as providing a derandomization of space in the decision tree model. Nisan's result upper bounds the deterministic depth complexity of a function f by the cube of its randomized depth. It remains an outstanding problem to determine if this gap is achieved by some total function. Relatively recently, a new total function was defined by Göös, Pitassi and Watson [20] to separate deterministic partition number and communication complexity by a quadratic factor, solving a longstanding problem. A modification of this function, which we refer to henceforth as the modified GPW function, was shown in [1] to achieve a quadratic gap between the bounded-error randomized and deterministic decision tree depth. In particular, the modified GPW function has deterministic depth $\Omega(n)$ but randomized depth $\Theta(\sqrt{n})$. No better separation is known for depth for any function.

It becomes natural to wonder if our Theorem 1.1 obtains the optimal relationship between the logarithms of randomized and deterministic decision tree size. Interestingly, even though the modified GPW function provided a breakthrough in increasing the known gap between randomized and deterministic depth, it turns out that both its deterministic and randomized size is $2^{\tilde{\Theta}(\sqrt{n})}$. On the other hand, the AND-OR tree function, which achieved a smaller gap for depth than the modified GPW function, yields the largest known gap for log size as we stated in Section 1 and further elaborate upon in Appendix A.

This suggests that depth and size behave rather differently and efforts to better understand the gap between deterministic and randomized size could be challenging and rewarding.

For ADT's, again it can be shown that the modified GPW function fails to generate any advantage for randomized algorithms w.r.t. depth. The best known gap is again provided by the AND-OR tree function. It would be interesting to find either functions yielding larger gaps or to narrow down the gap from cubic given by our Theorem 1.3.

References

- [1] Andris Ambainis, Kaspars Balodis, Aleksandrs Belovs, Troy Lee, Miklos Santha, and Juris Smotrovs. Separations in query complexity based on pointer functions. *J. ACM*, 64(5):32:1–32:24, 2017. Earlier version in STOC 2016.
- [2] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001. Earlier version in FOCS 1998.
- [3] Aleksandrs Belovs. Quantum algorithms for learning symmetric juntas via the adversary bound. *Comput. Complex.*, 24(2):255–293, 2015. Earlier version in CCC 2014.
- [4] Yosi Ben-Asher and Ilan Newman. Decision trees with Boolean threshold queries. *Journal of Computer and System Sciences*, 51(3):495–502, 1995.
- [5] Shalev Ben-David, Adam Bouland, Ankit Garg, and Robin Kothari. Classical lower bounds from quantum upper bounds. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 339–349. IEEE Computer Society, 2018.
- [6] Gal Beniamini and Noam Nisan. Bipartite perfect matching as a real polynomial. In *53rd Annual ACM SIGACT Symposium on Theory of Computing STOC*, pages 1118–1131. ACM, 2021.
- [7] Guy Blanc, Jane Lange, Mingda Qiao, and Li-Yang Tan. Properly learning decision trees in almost polynomial time. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 920–929. IEEE, 2021.
- [8] Guy Blanc, Jane Lange, and Li-Yang Tan. Top-down induction of decision trees: Rigorous guarantees and inherent limitations. In *11th Innovations in Theoretical Computer Science Conference, ITCS*, volume 151 of *LIPIcs*, pages 44:1–44:44. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [9] Joakim Blikstad, Jan van den Brand, Yuval Efron, Sagnik Mukhopadhyay, and Danupon Nanongkai. Nearly optimal communication and query complexity of bipartite matching. *CoRR*, abs/2208.02526, 2022. In FOCS'22.
- [10] Manuel Blum and Russell Impagliazzo. Generic oracles and oracle classes. In *28th Annual Symposium on Foundations of Computer Science (sfcs 1987)*, pages 118–126. IEEE, 1987.
- [11] Arkadev Chattopadhyay, Yuval Filmus, Sajin Koroth, Or Meir, and Toniann Pitassi. Query-to-communication lifting using low-discrepancy gadgets. *SIAM J. Comput.*, 50(1):171–210, 2021. Preliminary version in ICALP, 2019.
- [12] Arkadev Chattopadhyay, Ankit Garg, and Suhail Sherif. Towards stronger counterexamples to the log-approximate-rank conjecture. In *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS*, volume 213 of *LIPIcs*, pages 13:1–13:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [13] Arkadev Chattopadhyay, Michal Koucký, Bruno Loff, and Sagnik Mukhopadhyay. Simulation theorems via pseudo-random properties. *Comput. Complex.*, 28(4):617–659, 2019.
- [14] Arjan Cornelissen, Nikhil S. Mande, and Subhasree Patro. Improved quantum query upper bounds based on classical decision trees. In *42nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS*, volume 250 of *LIPIcs*, pages 15:1–15:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

- [15] Yogesh Dahiya and Meena Mahajan. On (simple) decision tree rank. In *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS*, volume 213 of *LIPIcs*, pages 15:1–15:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [16] Andrzej Ehrenfeucht and David Haussler. Learning decision trees from random examples. *Information and Computation*, 82(3):231 – 246, 1989. Earlier version in COLT’88.
- [17] Ankit Garg, Mika Göös, Pritish Kamath, and Dmitry Sokolov. Monotone circuit lower bounds from resolution. *Theory Comput.*, 16:1–30, 2020. Earlier version in STOC’18.
- [18] Uma Girish, Avishay Tal, and Kewen Wu. Fourier growth of parity decision trees. In *36th Computational Complexity Conference, CCC*, volume 200 of *LIPIcs*, pages 39:1–39:36. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [19] Mika Göös and T. S. Jayram. A composition theorem for conical juntas. In *31st Conference on Computational Complexity, CCC*, volume 50 of *LIPIcs*, pages 5:1–5:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [20] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. *SIAM J. Comput.*, 47(6):2435–2450, 2018. Earlier version in FOCS 2015.
- [21] Mika Göös, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for BPP. *SIAM J. Comput.*, 49(4), 2020. Preliminary version in FOCS, 2017.
- [22] Juris Hartmanis and Lane A Hemachandra. One-way functions, robustness, and the non-isomorphism of NP-complete sets. Technical report, Cornell University, 1986.
- [23] Hamed Hatami, Kaave Hosseini, and Shachar Lovett. Structure of protocols for XOR functions. *SIAM J. Comput.*, 47(1):208–217, 2018. Earlier version in FOCS 2016.
- [24] Russell Impagliazzo, Toniann Pitassi, and Alasdair Urquhart. Upper and lower bounds for tree-like cutting planes proofs. In *Proceedings of the Ninth Annual Symposium on Logic in Computer Science, LICS*, pages 220–228. IEEE Computer Society, 1994.
- [25] Dmitry Itsykson and Dmitry Sokolov. Resolution over linear equations modulo two. *Ann. Pure Appl. Log.*, 171(1), 2020.
- [26] Stasys Jukna. *Boolean Function Complexity - Advances and Frontiers*, volume 27 of *Algorithms and Combinatorics*. Springer, 2012.
- [27] Stasys Jukna, Alexander A. Razborov, Petr Savický, and Ingo Wegener. On P versus $NP \cap co-NP$ for decision trees and read-once branching programs. *Comput. Complex.*, 8(4):357–370, 1999.
- [28] Alexander Knop, Shachar Lovett, Sam McGuire, and Weiqiang Yuan. Guest column: Models of computation between decision trees and communication. *SIGACT News*, 52(2):46–70, 2021.
- [29] Alexander Knop, Shachar Lovett, Sam McGuire, and Weiqiang Yuan. Log-rank and lifting for AND-functions. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC*, pages 197–208, 2021.
- [30] Bruno Loff and Sagnik Mukhopadhyay. Lifting theorems for equality. In *36th International Symposium on Theoretical Aspects of Computer Science, STACS*, volume 126 of *LIPIcs*, pages 50:1–50:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [31] Shachar Lovett, Raghu Meka, Ian Mertz, Toniann Pitassi, and Jiapeng Zhang. Lifting with sunflowers. In *13th Innovations in Theoretical Computer Science Conference, ITCS*, volume 215 of *LIPIcs*, pages 104:1–104:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [32] Sagnik Mukhopadhyay, Jaikumar Radhakrishnan, and Swagato Sanyal. Separation between deterministic and randomized query complexity. *SIAM J. Comput.*, 47(4):1644–1666, 2018.
- [33] Noam Nisan. CREW prams and decision trees. *SIAM J. Comput.*, 20(6):999–1007, 1991. Earlier version in STOC’89.

- [34] Ran Raz and Iddo Tzameret. Resolution over linear equations and multilinear proofs. *Ann. Pure Appl. Log.*, 155(3):194–224, 2008.
- [35] Michael E. Saks and Avi Wigderson. Probabilistic boolean decision trees and the complexity of evaluating game trees. In *27th Annual Symposium on Foundations of Computer Science, FOCS*, pages 29–38. IEEE Computer Society, 1986.
- [36] Marc Snir. Lower bounds on probabilistic linear decision trees. *Theor. Comput. Sci.*, 38:69–82, 1985.
- [37] Gábor Tardos. Query complexity, or why is it difficult to separate $\text{NP}^A \text{ cap } \text{coNP}^A$ from P^A by random oracles A ? *Comb.*, 9(4):385–392, 1989.
- [38] Hing Yin Tsang, Chung Hoi Wong, Ning Xie, and Shengyu Zhang. Fourier sparsity, spectral norm, and the log-rank conjecture. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 658–667. IEEE Computer Society, 2013.
- [39] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the 11th Annual ACM Symposium on Theory of Computing, STOC*, pages 209–213. ACM, 1979.
- [40] Andrew Chi-Chih Yao. Lower bounds by probabilistic arguments (extended abstract). In *24th Annual Symposium on Foundations of Computer Science, FOCS*, pages 420–428. IEEE Computer Society, 1983.

A Lower bounds for randomized decision tree size

In this section we first observe that randomized communication complexity gives a lower bound on randomized decision tree size. As a consequence, using a communication lower bound due to Göös and Jayram [19], we conclude a lower bound on the randomized decision tree size of the AND-OR tree on n variables, nearly matching the upper bound implied by the following result.

Theorem A.1 ([35]). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be the AND-OR tree on n variables. Then,*

$$R_{1/3}^{\text{dt}}(f) = \Theta(n^{0.753\dots}).$$

This immediately implies

$$\text{RSize}_{1/3}^{\text{dt}}(f) = 2^{O(n^{0.753\dots})}. \tag{7}$$

After this, we give a general template to prove randomized decision tree size lower bounds and demonstrate its use by showing lower bounds for specific classes of Boolean functions.

A.1 Communication complexity

The model of interest to us is the two-party classical *bounded-error* model of communication introduced by Yao [39]. In this model, two parties, Alice and Bob, receive inputs $x \in \{0, 1\}^n, y \in \{0, 1\}^n$, respectively. They wish to jointly compute a function $F : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ on the input (x, y) using a communication protocol Π . They have unbounded computational power individually and wish to minimize the number of bits communicated. Alice and Bob have access to an arbitrarily long string of public random bits. At any point in the protocol, the party that sends the next message is decided by the previous messages communicated. The message communicated must be a function of the party’s input, the previous messages communicated and the random string. We say that a protocol Π computes F to error ε if $\Pr[\Pi(x, y) = F(x, y)] \geq 1 - \varepsilon$ for each $x \in \{0, 1\}^n, y \in \{0, 1\}^n$. The cost of the protocol is the largest number of bits communicated in the worst case over all possible inputs and outcomes of the random coin tosses. The ε -error randomized communication complexity of F , denoted $R_\varepsilon^{\text{cc}}(F)$, is the minimum cost of a protocol that computes F to error ε .

We require the following theorem due to Göös and Jayram [19, Theorem 2.3].

Theorem A.2. Let $F : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ be the AND-OR tree on $2n$ inputs, where Alice gets the first bit of every bottom gate and Bob gets the rest. Then,

$$R_\varepsilon^{\text{cc}}(F) = \tilde{\Omega}(n^{0.753\dots}).$$

A.2 Lower bounds via communication complexity

We observe the following.

Lemma A.3. Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be a Boolean function such that $R_\varepsilon^{\text{cc}}(f) \geq c$ under some partition of the inputs. Then,

$$\text{RSize}_\varepsilon^{\text{dt}}(f) = 2^{\Omega(c)}.$$

Proof. Consider an arbitrary partition of inputs between Alice and Bob. Consider a randomized decision tree of size s computing f to error ε . One obtains a communication protocol for f in the following way: Alice and Bob sample a tree from the underlying distribution of the randomized decision tree using public randomness. They then evaluate this tree from the root down using a communication protocol; if a node has the label of a variable held by Alice (Bob), she (he) sends its value to Bob (Alice) and they continue in the relevant subtree. When they reach a leaf they output the value at that leaf. Note that the randomized protocol thus obtained has exactly the same structure as the randomized decision tree. It is well known that communication protocols can be balanced (using the same ideas as in the proof of Lemma 4.3). We thus obtain a randomized protocol of cost $O(\log s)$ that computes f to error ε . This proves the lemma. \square

Along with Theorem A.2, this implies the following, yielding near-tightness of the bound in Eq. (7).

Corollary A.4. Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be the AND-OR tree on n inputs. Then,

$$\text{RSize}_{1/3}^{\text{dt}}(f) = 2^{\tilde{\Omega}(n^{0.753\dots})}.$$

In view of the above, one might ask if it is necessary for the randomized communication complexity of a function to be large under some partition of inputs in order for its randomized decision tree size to be large. Specifically, can one show that the randomized decision tree size of the well-studied Equality function is large, even though its randomized communication complexity is a constant? In Section A.3 we show that the randomized decision tree size of Equality is indeed large, as an application of a more general template that we introduce to prove lower bounds against randomized decision tree size.

A.3 A template for proving randomized decision tree size lower bounds

In this section, we give a set of sufficient conditions which imply randomized decision tree size lower bounds. In particular, we can use these to recover Theorem 3.1 for the particular case of monotone functions and we can also use these in other cases to prove better randomized size lower bounds. We start by giving some definitions. For a subcube A , its 0-codimension is the number of variables fixed to 0 by A and its 1-codimension is the number of variables fixed to 1 by A . We define a notion of spreadness for a distribution on $\{0,1\}^n$. For a set $S \subseteq \{0,1\}^n$ and a distribution μ on $\{0,1\}^n$, let $\mu(S) = \Pr_{x \sim \mu}[x \in S]$.

Definition A.5 (Spreadness). A probability distribution μ on $\{0,1\}^n$ is

- (ℓ, κ) -spread if for any subcube S of codimension ℓ , $\mu(S) \leq \kappa$.
- $(\ell, \kappa, 1)$ -spread if for any subcube S of 1-codimension ℓ , $\mu(S) \leq \kappa$.
- $(\ell, \kappa, 0)$ -spread if for any subcube S of 0-codimension ℓ , $\mu(S) \leq \kappa$.

Next, we define a notion of sensitivity of a distribution. For $b \in \{0,1\}$, a set $B \subseteq [n]$ is a *sensitive b -block* of f at input x if $f(x \oplus 1_B) \neq f(x)$ and $x_i = b$ for each $i \in B$. The *b -block sensitivity* of f at x , denoted by $\text{bs}_b(f, x)$, is the maximum integer r for which there exist r disjoint sensitive b -blocks of f at x .

Definition A.6 (Sensitivity of a distribution). For a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $b \in \{0, 1\}$, a distribution μ on $\{0, 1\}^n$ is

- (ℓ, b) -sensitive w.r.t. f if for each input x in the support of μ , we have $f(x) = b$ and $\text{bs}(f, x) \geq \ell$.
- $(\ell, b, 0)$ -sensitive w.r.t. f if for each input x in the support of μ , we have $f(x) = b$ and $\text{bs}_0(f, x) \geq \ell$.
- $(\ell, b, 1)$ -sensitive w.r.t. f if for each input x in the support of μ , we have $f(x) = b$ and $\text{bs}_1(f, x) \geq \ell$.

We now show that the existence of a distribution that is spread and sensitive w.r.t. f implies large randomized decision tree size complexity.

Theorem A.7. For a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let $b \in \{0, 1\}$ and let μ_b be a distribution on $\{0, 1\}^n$ which satisfies any of the following properties:

1. μ_b is (ℓ, b) -sensitive w.r.t. f and $(\lceil \ell/2 \rceil, \kappa)$ -spread.
2. There exists $c \in \{0, 1\}$ such μ_0 is (ℓ, b, c) -sensitive w.r.t. f and $(\lceil \ell/2 \rceil, \kappa, c)$ -spread.

Then $\text{RSize}_{1/10}^{\text{dt}}(f) = \Omega(1/\kappa)$.

Proof. We prove (1) above, the proof of (2) follows along similar lines. Without loss of generality, assume $b = 0$. Thus, μ_0 is a distribution over inputs in $f^{-1}(0)$ that is $(\ell, 0)$ -sensitive. Define μ_1 , a distribution over $f^{-1}(1)$, as follows: Sample an input x according to μ_0 and let B_1, B_2, \dots, B_r be disjoint sensitive block of f at x where $r = \text{bs}(f, x)$. Sample a i from $[r]$ and return $x \oplus 1_{B_i}$. Define the distribution $\mu = (\mu_0 + \mu_1)/2$. We show below that $\text{DSize}_{\mu, 1/10}^{\text{dt}}(f) = \Omega(1/\kappa)$, thereby showing $\text{RSize}_{1/10}^{\text{dt}}(f) = \Omega(1/\kappa)$ by applying Proposition C.1. Let T be a deterministic decision tree that computes f correctly on at least a $(9/10)$ -mass of inputs sampled according to μ . Since μ can be viewed as sampling according to μ_0 with probability $1/2$ and according to μ_1 with probability $1/2$, T must be correct on at least a $(4/5)$ mass of μ_0 as well as a $(4/5)$ mass of μ_1 . Let L_0 and L_1 be set of all 0-leaves and 1-leaves of T , respectively. Define

$$\begin{aligned} \rho_0 &= \sum_{v \in L_0} \Pr_{x \sim \mu_0} [x \text{ reaches } v], \\ \rho_1 &= \sum_{v \in L_1} \Pr_{x \sim \mu_1} [x \text{ reaches } v]. \end{aligned}$$

That is, ρ_0 (ρ_1) denotes the μ_0 -mass (μ_1 -mass) captured by 0-leaves (1-leaves). By the argument above, both ρ_0 and ρ_1 must be at least $4/5$. We observe following about the 0-paths (paths leading to 0-leaves).

1. (Short Paths). Firstly, we show that the 0-paths of length less than $\lceil \ell/2 \rceil$ must capture less than $2/5$ mass of μ_0 , i.e.,

$$\sum_{\substack{v \in L_0 \\ |v| < \lceil \ell/2 \rceil}} \Pr_{x \sim \mu_0} [x \text{ reaches } v] \leq 2/5,$$

where $|v|$ denotes the length of the path leading to v . Towards a contradiction, assume

$$\sum_{\substack{v \in L_0 \\ |v| < \lceil \ell/2 \rceil}} \Pr_{x \sim \mu_0} [x \text{ reaches } v] > 2/5.$$

Let S be a subcube with codimension less than $\lceil \ell/2 \rceil$ corresponding to a 0-leaf. For each x supported by μ_0 which lies in S , our assumption on μ_0 implies that $\text{bs}(f, x) \geq \ell$. Thus, there must exist at least $\text{bs}(f, x)/2$ blocks such that none of their variables are read on the path to this leaf. Flipping any one of these blocks gives an input supported by μ_1 and that still reaches this leaf. Hence $\mu_1(S) \geq (1/2)\mu_0(S)$. Thus,

$$\sum_{\substack{v \in L_0 \\ |v| < \lceil \ell/2 \rceil}} \Pr_{x \sim \mu_1} [x \text{ reaches } v] \geq \frac{1}{2} \cdot \sum_{\substack{v \in L_0 \\ |v| < \lceil \ell/2 \rceil}} \Pr_{x \sim \mu_0} [x \text{ reaches } v] > 1/5.$$

This contradicts the fact that T is correct on at least a $4/5$ mass of μ_1 .

2. (Long Paths). Secondly, a 0-path of length at least $\lceil \ell/2 \rceil$ captures at most κ of μ_0 mass. This follows from the spreadness property of μ_0 .

Since ρ_0 must be at least $4/5$, from the first bullet above, we have that 0-paths in T of length at least $\lceil \ell/2 \rceil$ must capture at least $2/5$ of μ_0 mass. By the second bullet above, there must be at least $2/(5\kappa)$ 0-paths in T of length at least $\lceil \ell/2 \rceil$, proving an $\Omega(1/\kappa)$ lower bound on the number of leaves in T .

The analysis for (2) is along similar lines. The only difference will be in using sensitive c -blocks in place of sensitive blocks and defining short and long paths for the case analysis above based on c -codimension of the corresponding subcubes. \square

Below we give some examples to showcase the usage of the above template to prove randomized decision tree size lower bounds for specific classes of Boolean functions.

A.3.1 Monotone functions

In this section, we use our template to give an alternative proof of Theorem 3.4. We prove part (a) of Theorem 3.4 as part (b) follows from part (a) and Theorem 2.9.

Proof. (Part (a) Theorem 3.4). We show that f has a sub-function (f under some restriction) g that has large randomized decision tree size. The statement of theorem trivially holds when $N(f) \leq n^c$ for any fixed constant c , so we restrict our attention to the case when $N(f) \geq n^c$ for some appropriate c to be fixed later. Without loss of generality, let $N_1(f) \geq N(f)/2$, the case when $N_0(f) \geq N(f)/2$ is similar. For a monotone function f , the subcubes corresponding to minterms of f form the unique minimum 1-cover for f . For a minterm M , let x_M be the input such that $x_i = 1$ for $i \in M$ and 0 otherwise. Let $S = \{x_M | M \text{ is a minterm of } f\}$. Note that $|S| = N_1(f)$ as each input in S is a 1-input of f and each minterm forms a minimal sensitive block for 0^n . We have that for each $x \in S$, $\text{bs}_1(f, x) = |x|$, i.e., each 1-bit position of x is sensitive. For $i \in [n]$, let S_i denote inputs of Hamming weight i in S . By an averaging argument, there exists $j \in [n]$ such that $|S_j| \geq N_1(f)/n$. We give an iterative procedure to find our desired sub-function g .

1. Initialise $V = [n]$, $g = f$, $\ell = j$, $R = S_j$ and $k = \log(N_1(f)/n)/(2(\lceil \log n \rceil + 1))$. We use V to denote the input domain of g . Throughout the process, R will be a set of 1-inputs of g over $\{0, 1\}^V$ each of Hamming weight and 1-block sensitivity ℓ w.r.t. g .
2. Find a subcube P over $\{0, 1\}^V$ of 1-codimension $\lceil \ell/2 \rceil$ such that it contains more than 2^{-k} fraction of inputs in R . Let $Q \subseteq V$ be the set of variables fixed by P . If no such subcube exists, stop the procedure. Else, update V to $[n] \setminus Q$, ℓ to $\lfloor \ell/2 \rfloor$ and g to $g|_P$ (g restricted to P). For $V \subseteq T$ and $x \in \{0, 1\}^T$, let x_V in $\{0, 1\}^V$ denote the projection of x on to V . Update R to $\{x_V | x \in R \text{ and } x \text{ lies in } P\}$. Note that after this update, g is a function on $\{0, 1\}^V$ and $R \subseteq \{0, 1\}^V$. Also, each $x \in R$ is a 1-input of g , has Hamming weight ℓ , and each 1-bit position of x forms a sensitive 1-block for g i.e., $\text{bs}_1(g, x) = \ell$.
3. Repeat the above step.

We claim the above procedure can not run forever and must terminate within the second step's $(\lceil \log n \rceil + 1)$ 'th repetition. The procedure starts with each input in R having Hamming weight ℓ , and with every repetition of the second step, the updated R contains inputs with Hamming weight halved from the previous iteration. So if the second step repeats $(\lceil \log n \rceil + 1)$ times, R can have only one input, namely 0^n . On the other hand, the size of R only reduces by a factor of 2^{-k} with each repetition of the second step because of the way we choose the subcube P , and so by the end of $(\lceil \log n \rceil + 1)$ repetitions, the size of R must be

$$|R| \geq |S_j| 2^{-k(\lceil \log n \rceil + 1)} = (N_1(f)/n) 2^{-k(\lceil \log n \rceil + 1)} = (N_1(f)/n)^{1/2} \geq n^{\frac{c-1}{2}}.$$

So if we choose $c > 1$, the size of $|R|$ will be strictly larger than 1, contradicting our earlier conclusion, which says the size of R can be at most 1 after $(\lceil \log n \rceil + 1)$ repetitions of the second step. So, the above procedure does terminate and at the end of the process we get g , a subfunction of f on $\{0, 1\}^V$, and $R \subseteq \{0, 1\}^V$ with following properties

- $g(x) = 1$ for all $x \in R$.
- (Spreadness). Each $x \in R$ has Hamming weight ℓ and no subcube of 1-codimension $\lceil \ell/2 \rceil$ contains more than 2^{-k} fraction of inputs of R .
- (Sensitivity). Each $x \in R$ has $\text{bs}_1(g, x) = \ell$.

Let μ_1 be the distribution that randomly samples an input from R . From the properties of R stated above, we get μ_1 is a distribution over 1-inputs of g which is $(\lceil \ell/2 \rceil, 2^{-k}, 1)$ -spread and $(\ell, 1, 1)$ -sensitive w.r.t. g . Applying Theorem A.7, we get $\log \text{RSize}_{1/10}^{\text{dt}}(f) \geq \log \text{RSize}_{1/10}^{\text{dt}}(g) = \Omega(k) = \Omega(\log N_1(f)/\log n) = \Omega(\log N(f)/\log n)$. \square

A.3.2 Symmetric functions

In this section we show randomized decision tree size lower bounds for symmetric functions, as an application of our template.

Definition A.8 (Symmetric functions). *A function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is said to be symmetric if for all $\sigma \in S_n$ and for all $x \in \{0, 1\}^n$ we have $f(x) = f(\sigma(x))$.*

In other words, a Boolean function is symmetric if its value only depends on the Hamming weight of its input. For a symmetric Boolean function f on n variables, let $f_0, f_1, \dots, f_n \in \{0, 1\}$ denote the values of the function f on inputs of Hamming weight $0, 1, \dots, n$ respectively. We consider the case when $f_k \neq f_{k+1}$ for some $0 \leq k \leq n/2$. The case when $k > n/2$ is analogous and symmetric.

Theorem A.9. *Let f be a symmetric Boolean function and let $0 \leq k \leq n/2$ be such that $f_k \neq f_{k+1}$. Then, $\text{RSize}^{\text{dt}}(f) = 2^{\Omega(k)}$.*

Proof. Without loss of generality, let $f_k = 0$. Define $t = k/n$. We exhibit a spread and sensitive distribution with respect to f . Let μ_0 be a distribution which samples a uniformly random input string of Hamming weight exactly k . We claim that μ_0 has the necessary spreadness and sensitivity properties. It is clear that μ_0 is $(n - k, 0, 0)$ -sensitive. For spreadness, let S be a subcube of 0-codimension at least $\lceil \frac{n-k}{2} \rceil$, we have

$$\Pr_{x \sim \mu_0} [x \in S] = \frac{|\{x | x \in S \text{ and } |x| = k\}|}{\binom{n}{k}} \leq \frac{\binom{n - \lceil \frac{n-k}{2} \rceil}{k}}{\binom{n}{k}} = \prod_{i=0}^{k-1} \frac{m-i}{n-i}$$

where $m = n - \lceil \frac{n-k}{2} \rceil$. Since the factors inside the product are decreasing, we get an upper bound of $(m/n)^k \leq (\frac{2}{1+t})^{-tn}$. So we get μ_0 is $(\lceil \frac{n-k}{2} \rceil, (\frac{2}{1+t})^{-tn}, 0)$ -spread. Applying Theorem A.7, we get $\text{RSize}_{1/10}^{\text{dt}}(f) = \Omega((\frac{2}{1+t})^{tn}) = 2^{\Omega(k)}$. \square

An immediate corollary of the above result is an exponential randomized decision tree size lower bound for the MAJORITY function.

Corollary A.10. $\text{RSize}^{\text{dt}}(\text{MAJORITY}) = 2^{\Omega(n)}$.

A.3.3 EQUALITY

The Equality function, denoted by $\text{EQ}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, is defined by $\text{EQ}(x, y) = 1$ iff $x_i = y_i$ for all $i \in [n]$. Consider the set $F = \{(x, x) | x \in \{0, 1\}^n\}$; it is easy to observe that one 1-subcube cannot cover more than one input in F . Thus, we get that $N_1(\text{EQ}_n) = 2^n$. Applying Theorem 3.1, we get that $\text{RSize}^{\text{dt}}(\text{EQ}_n) = 2^{\Omega((\frac{n}{\log n})^{1/2})}$. In this section, we apply our template to prove a self-contained $2^{\Omega(n)}$ randomized decision tree size lower bound for EQ_n .

Theorem A.11. $\text{RSize}^{\text{dt}}(\text{EQ}_n) = \Omega(2^{n/2})$.

Proof. Let μ_1 be a probability distribution which samples a uniformly random string $x \in \{0, 1\}^n$ and outputs (x, x) . We show μ_1 has the necessary spreadness and sensitivity properties to apply Theorem A.7 to give a randomized decision tree size lower bound for EQ_n . It is clear from the definition that μ_1 is $(2n, 1)$ -sensitive w.r.t. EQ_n . We show that μ_1 is also $(n, 2^{-n/2})$ -spread. Let S be a subcube of codimension $\ell \geq n$. Let P and Q be the set of variables of x and y fixed by S , respectively. Let $\rho_x : \{0, 1\}^P \rightarrow \{0, 1\}$ and $\rho_y : \{0, 1\}^Q \rightarrow \{0, 1\}$ be the maps according to which the variables in P and Q are fixed by S . Note that $|P| + |Q| = \ell$. Let U denote the uniform distribution over $\{0, 1\}^n$. Note that for all inputs in the support of μ_1 we have $x = y$, and that the marginal distributions of both x and y are the uniform distribution over $\{0, 1\}^n$. We have

$$\begin{aligned} \Pr_{(x,y) \sim \mu_1} [(x, y) \in S] &= \Pr_{(x,y) \sim \mu_1} \left[\bigwedge_{i \in P} (x_i = \rho_x(i)) \text{ and } \bigwedge_{j \in Q} (y_j = \rho_y(j)) \right] \\ &= \prod_{i \in P \setminus Q} \Pr_{x \sim U} [x_i = \rho_x(i)] \prod_{j \in Q \setminus P} \Pr_{y \sim U} [y_j = \rho_y(j)] \prod_{i \in P \cap Q} \Pr_{(x,y) \sim \mu_1} [x_i = y_i = \rho_x(i) = \rho_y(i)] \\ &= (1/2)^{|P \cup Q|} \leq (1/2)^{\ell/2} \leq (1/2)^{n/2}. \end{aligned}$$

Applying Theorem A.7, we get $\text{RSize}_{1/10}^{\text{dt}}(\text{EQ}_n) = \Omega(2^{n/2})$. \square

B Proofs for (AND,OR) decision trees

Proof of Lemma 4.3. Let T be an ordinary decision tree computing f with L leaves. We argue existence of an (AND, OR) DT C computing f of size $2 \log_{3/2} L$ by induction on L . The base case $L = 1$ is trivially true. For the induction step, we obtain a node v in T such that the number of leaves in the subtree rooted at v , T_v , lies in the range $(L/3, 2L/3]$. Such a node always exists and can be found by starting at the root and traversing down to the child node with more leaves until the number of leaves in the subtree rooted at the current node satisfies the condition. Let S_0 and S_1 be the set of variables evaluated to 0 and 1 on the path from the root to v , respectively. The (AND, OR) DT C starts by querying $Q_0 = \bigvee_{i \in S_0} x_i$ and $Q_1 = \bigwedge_{i \in S_1} x_i$. If Q_0 evaluates to 0 and Q_1 evaluates to 1, proceed by recursively constructing the (AND, OR) decision tree for T_v . Otherwise, proceed by recursively constructing the (AND, OR) decision tree for T' where T' is the DT obtained by replacing node v in T with a leaf node with an arbitrary label. C computes f is clear from the construction. Moreover since both T_v and T' have at most $2L/3$ many leaves, using induction, we get

$$\text{Depth}(C) \leq 2 \log_{3/2} L \leq 4 \log L.$$

\square

Proof of Lemma 4.4. Let C be an (AND, OR) decision tree computing f of depth d . We argue by induction on the number of variables n and depth d of C . Let $L(n, d)$ be the size of an optimal ordinary decision tree computing a n variable Boolean function for which there exists an (AND, OR) decision tree of depth d . Let Q be the root query in C . Without loss of generality, let Q be an AND query; the case for OR query is similar. Consider an ordinary decision tree T , which starts by querying x_i , a variable in Q . If x_i evaluates to 0, proceed by recursively constructing ordinary DT for the left subtree of the root node in C . If x_i evaluates to 1, proceed by recursively constructing an ordinary DT for C' , where C' is a reduced (AND, OR) decision tree obtained by setting x_i to 1 in each query in C . This gives us the following recurrence on $L(n, d)$.

$$L(n, d) \leq L(n-1, d-1) + L(n-1, d) \text{ with } L(0, d) = L(n, 0) = 1.$$

By induction on n and d and using the identity $\binom{n}{k} = \binom{n-1}{k-1} + \binom{n-1}{k}$, we get

$$L(n, d) \leq \sum_{i=0}^d \binom{n}{i},$$

which is bounded by

$$\sum_{i=0}^d \binom{n}{i} \leq \sum_{i=0}^d \frac{n^i}{i!} = \sum_{i=0}^d \frac{d^i}{i!} \left(\frac{n}{d}\right)^i \leq \left(\frac{n}{d}\right)^d \sum_{i=0}^d \frac{d^i}{i!} \leq \left(\frac{ne}{d}\right)^d.$$

Thus, there exists an ordinary decision tree computing f whose size is bounded from above by $(\frac{ne}{d})^d$. \square

C An alternative proof of Lemma 4.6

In this section we give an alternative, dual-based proof of Lemma 4.6. We first require some definitions and preliminaries.

For a probability distribution \mathcal{D} over $\{0, 1\}^n$, the $(\mathcal{D}, \varepsilon)$ -distributional query complexity of f , denoted $D_{\mathcal{D}, \varepsilon}^{\text{dt}}(f)$, is the minimum depth of a deterministic decision tree that gives a correct answer on $1 - \varepsilon$ fraction of inputs weighted by \mathcal{D} . Distributional size complexity is defined similarly. That is, with $x \sim \mathcal{D}$ denoting that x is sampled according to \mathcal{D} ,

$$\text{DSize}_{\mathcal{D}, \varepsilon}^{\text{dt}} = \min \left\{ \text{Size}(T) \mid T \text{ is a deterministic decision tree; } \Pr_{x \sim \mathcal{D}}[T(x) = f(x)] \geq 1 - \varepsilon \right\}.$$

Distributional complexity provides a technique to prove randomized query and size lower bounds via Yao's minmax principle [40].

Proposition C.1 ([40]). *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $\varepsilon > 0$, $R_{\varepsilon}^{\text{dt}}(f) = \max_{\mathcal{D}} D_{\mathcal{D}, \varepsilon}^{\text{dt}}(f)$ and $\text{RSize}_{\varepsilon}^{\text{dt}}(f) = \max_{\mathcal{D}} \text{DSize}_{\mathcal{D}, \varepsilon}^{\text{dt}}(f)$*

We remark that the proposition above holds true even when the decision trees are generalized, as in Section 2.2, for example.

Monotone block sensitivity and fractional monotone block sensitivity are variants of block sensitivity and fractional block sensitivity where we restrict changing the input by only flipping 0's to 1's.

Definition C.2 (Monotone block sensitivity). *For $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and input $x \in \{0, 1\}^n$, let $S = \{i : x_i = 0\}$. The monotone block sensitivity of f at x , denoted by $\text{MBS}(f, x)$, is the maximum r for which there exist r disjoint 0-blocks in $\mathcal{W}_0(f, x)$. The monotone block sensitivity of f , denoted by $\text{MBS}(f)$, is defined as*

$$\text{MBS}(f) = \max_{x \in \{0, 1\}^n} \text{MBS}(f, x).$$

Definition C.3 (Fractional monotone block sensitivity). *A distribution \mathcal{D} over $2^{[n]}$ is said to be p -smooth if for any $i \in [n]$ it holds that $\Pr_{B \sim \mathcal{D}}[i \in B] \leq p$. For $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and input $x \in \{0, 1\}^n$, the fractional monotone block sensitivity of f at x , denoted $\text{FMBS}(f, x)$, is equal to $1/p$, where $p > 0$ is the smallest number for which there exists a p -smooth distribution \mathcal{D} supported on a subset of $\mathcal{W}_0(f, x)$. Define*

$$\text{FMBS}(f) = \max_{x \in \{0, 1\}^n} \text{FMBS}(f, x).$$

Monotone block sensitivity and 0-hitting set complexity can be phrased as packing and covering integer programs, respectively. Their fractional versions can be seen as fractional packing and covering linear programs. By definitions and using linear programming duality, we get the following simple relationship between them.

Claim C.4 ([29, Claim 2.16]). *For every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and an input $x \in \{0, 1\}^n$,*

$$\text{MBS}(f, x) \leq \text{FMBS}(f, x) = \text{FHSC}_0(f, x) \leq \text{HSC}_0(f, x).$$

An alternative proof of Lemma 4.6. First, by Claim C.4, $\text{FHSC}_0(f, x) = \text{FMBS}(f, x)$ for every $x \in \{0, 1\}^n$. Thus it suffices to prove that $R^{0\text{-dt}}(f) = \Omega(\text{FMBS}(f))$. Let $a \in \{0, 1\}^n$ be an input achieving the fractional monotone block sensitivity, i.e., $\text{FMBS}(f, a) = \text{FMBS}(f)$. By definition, there exists a $1/\text{FMBS}(f)$ -smooth distribution, say \mathcal{D}' , on the sensitive 0-blocks of a . Using \mathcal{D}' we construct a hard distribution \mathcal{D} on inputs in $\{0, 1\}^n$ such that $D^{0\text{-dt}}(f)_{\mathcal{D}, 1/3}(f)$ (which is the distributional zero-depth

complexity of f w.r.t. \mathcal{D}) is at least $\text{FMBS}(f)/3$, thereby showing $R^{0\text{-dt}}(f) \geq \text{FMBS}(f)/3$ by invoking Proposition C.1. Let \mathcal{S} be the set of all sensitive 0-blocks of a . The hard distribution is as follows.

$$\mathcal{D}(x) = \begin{cases} 1/2 & \text{if } x = a \\ (1/2)\mathcal{D}'(B) & \text{if } x = a \oplus 1_B \text{ where } B \in \mathcal{S} \\ 0 & \text{otherwise} \end{cases}$$

Let T be any deterministic decision tree that gives the correct answer for f on $2/3$ mass of inputs when sampled from \mathcal{D} . We argue that the zero-depth of T must be at least $\text{FMBS}(f)/3$. Consider the path P traversed on a by T . We argue that path P must have at least $\text{FMBS}(f)/3$ variables whose query outcomes were 0. Suppose not, and let I be the set of indices of variables whose query outcomes were 0 on path P . By the $(1/\text{FMBS}(f))$ -smoothness of \mathcal{D}' , we have

$$\begin{aligned} \Pr_{B \sim \mathcal{D}'}[I \cap B \neq \emptyset] &\leq \sum_{i \in I} \Pr_{B \sim \mathcal{D}'}[i \in B] && \text{union bound} \\ &\leq |I| \cdot 1/\text{FMBS}(f) && \mathcal{D}' \text{ is } (1/\text{FMBS}(f))\text{-smooth} \\ &< 1/3. && |I| < \text{FMBS}(f)/3 \end{aligned}$$

Let $\mathcal{X} = \{a \oplus 1_B \mid B \neq \emptyset, B \cap I = \emptyset\}$. Note that each $x \in \mathcal{X}$ has $f(x) \neq f(a)$ and is consistent with all of the answers on the path P (and hence reaches the same leaf as a). Moreover, by the analysis above $\Pr_{x \sim \mathcal{D}}[x \in \mathcal{X}] \geq 1/3$. Now if $f(a)$ equals the label at the leaf of P , then T errors on inputs in \mathcal{X} with probability mass more than $1/3$ and if $f(a)$ does not equal this label, then T errs on a with probability mass $1/2$. Either way, this contradicts the assumption that T answers correctly on $2/3$ probability mass when inputs are sampled according to \mathcal{D} . \square