

Mildly Exponential Lower Bounds on Tolerant Testers for Monotonicity, Unateness, and Juntas

Xi Chen ^{*} Anindya De [†] Yuhao Li [‡] Shivam Nadimpalli [§] Rocco A. Servedio [¶]

Abstract

We give the first super-polynomial (in fact, mildly exponential) lower bounds for tolerant testing (equivalently, distance estimation) of monotonicity, unateness, and juntas with a *constant* separation between the “yes” and “no” cases. Specifically, we give

- A $2^{\Omega(n^{1/4}/\sqrt{\varepsilon})}$ -query lower bound for non-adaptive, two-sided tolerant monotonicity testers and unateness testers when the “gap” parameter $\varepsilon_2 - \varepsilon_1$ is equal to ε , for any $\varepsilon \geq 1/\sqrt{n}$;
- A $2^{\Omega(k^{1/2})}$ -query lower bound for non-adaptive, two-sided tolerant junta testers when the gap parameter is an absolute constant.

In the constant-gap regime no non-trivial prior lower bound was known for monotonicity, the best prior lower bound known for unateness was $\tilde{\Omega}(n^{3/2})$ queries, and the best prior lower bound known for juntas was $\text{poly}(k)$ queries.

1 Introduction

A *monotone* Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$ is one for which $f(x) \leq f(y)$ whenever x is coordinate-wise less than or equal to y , and a *k-junta* $f : \{0,1\}^n \rightarrow \{0,1\}$ is a function which depends on at most k of its n input coordinates. Monotonicity testing, the closely related problem of unateness¹ testing, and junta testing are among the most fundamental and intensively studied problems in the field of property testing of Boolean functions; indeed, many of the earliest and most influential works in this area study these problems [GGL⁺00, DGL⁺99, FLN⁺02, PRS02, FKR⁺04].

Many different variants of these testing problems have been analyzed by now, including functions with non-Boolean domains and ranges (see e.g. [DGL⁺99, FLN⁺02, HK07, BCGM12, CS13b, BCS18, BCS20, PRW22, BCS23, BKKM23, FKR⁺04, Bla09]) and distribution-free testing (see e.g. [HK07, LCS⁺19, Bsh19]). In this paper we study the original setting for each of these testing problems, i.e. we consider Boolean functions $f : \{0,1\}^n \rightarrow \{0,1\}$ and we measure distance between functions with respect to the uniform distribution over $\{0,1\}^n$.

After two decades of intensive research, the original problems of testing whether an unknown $f : \{0,1\}^n \rightarrow \{0,1\}$ is monotone/unate/ k -junta versus ε -far from being monotone/unate/ k -junta are now quite well understood. We briefly recall the current state of the art:

- **Monotonicity:** Building on [GGL⁺00, CS13a, CST14], Khot et al. [KMS18] gave an $\tilde{O}(\sqrt{n}/\varepsilon^2)$ -query non-adaptive monotonicity testing algorithm with one-sided error.² An $\tilde{\Omega}(\sqrt{n})$ -query lower bound for non-adaptive two-sided testers was given by Chen et al. [CWX17a] (strengthening earlier non-adaptive lower bounds in [FLN⁺02, CST14, CDST15]); [CWX17a] also gave an $\tilde{\Omega}(n^{1/3})$ -query lower bound for adaptive algorithms with two-sided error. Thus far, no known algorithms for monotonicity testing use adaptivity.

^{*}Columbia University. Email: xichen@cs.columbia.edu.

[†]University of Pennsylvania. Email: anindyad@seas.upenn.edu.

[‡]Columbia University. Email: yuhaoli@cs.columbia.edu.

[§]Columbia University. Email: sn2855@columbia.edu.

[¶]Columbia University. Email: ras2105@columbia.edu.

¹A Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$ is *unate* if $g(x) = f(x \oplus r)$ is monotone for some $r \in \{0,1\}^n$, where $\oplus$ denotes coordinate-wise XOR; equivalently, f is either nondecreasing or nonincreasing in each coordinate.

²A tester is *non-adaptive* if the choice of its i -th query point does not depend on the responses received to queries $1, \dots, i-1$. A *one-sided* tester for a class of functions is one which must accept every function in the class with probability 1.

- **Unateness:** In [CS16] Chakrabarti and Seshadhri gave an $\tilde{O}(n/\varepsilon)$ -query non-adaptive unateness tester with one-sided error (see also [GGL⁺00, BCP⁺20]), and Chen et al. [CWX17a] gave a near-matching $\tilde{\Omega}(n)$ -query lower bound for this setting. Chen and Waingarten [CW19] gave an $\tilde{O}(n^{2/3}/\varepsilon^2)$ -query adaptive algorithm with one-sided error, strengthening prior adaptive algorithms [KS16, BCP⁺20, CWX17c]. The [CW19] adaptive algorithm is near-optimal, since Chen et al. [CWX17b] gave an $\tilde{\Omega}(n^{2/3})$ -query lower bound for two-sided adaptive testers.
- **Juntas:** State-of-the-art adaptive algorithms for testing k -juntas due to [Bla09, Bsh19] use $\tilde{O}(k/\varepsilon)$ queries, and near-matching $\tilde{\Omega}(k/\varepsilon)$ -query lower bounds are known for adaptive algorithms [Sağ18] (see also [PRS02, CG04, FKR⁺04, BGMdW13, STW15] for a long line of earlier works). For non-adaptive testers, in [Bla08] Blais gave an $\tilde{O}(k^{3/2})/\varepsilon$ -query algorithm, and Chen et al. gave an essentially matching $\tilde{\Omega}(k^{3/2}/\varepsilon)$ lower bound in [CST⁺18].

To summarize, for all three of the properties we consider — being monotone, being unate, and being a k -junta — matching or near-matching polynomial upper and lower bounds are known, for both adaptive and non-adaptive algorithms, for testing whether a function *perfectly* satisfies the property or is ε -far from having the property.

Tolerant Testing. The requirement that functions in the “yes-case” of standard property testing must *exactly* satisfy the property means that algorithms under this framework may be undesirably brittle; thus it is natural to consider a noise-tolerant variant of standard property testing. With this motivation, in [PRR06] Parnas et al. introduced a natural generalization of standard property testing, which they called *tolerant* property testing. For parameters $0 \leq \varepsilon_1 < \varepsilon_2$, an $(\varepsilon_1, \varepsilon_2)$ -tolerant tester for a class of functions is a query algorithm which must accept with high probability if the input is ε_1 -close to some function in the class and reject with high probability if the input is ε_2 -far from every function in the class (so “standard” testing corresponds to $(0, \varepsilon)$ -tolerant testing).

It is well known [PRR06] and not difficult to see that tolerant property testing is essentially equivalent to the problem of estimating the distance to the nearest function that has the property. Thus, results on tolerant testing can alternately be phrased as results on distance estimation, but for simplicity throughout this paper we will describe our results in terms of tolerant testing.

Tolerant property testers clearly enjoy an attractive level of robustness that typically is not shared by standard property testers. For this and other reasons, much property testing research in recent years has attempted to give algorithms and lower bounds for tolerant testing of various properties. However, while it is widely believed that tolerant testing is generally a hard algorithmic problem, it has proved to be quite challenging to establish strong upper or lower bounds on the query complexity of tolerant testing; much less is known here than in the standard (non-tolerant) setting. This lack of understanding is particularly acute for the three problems of monotonicity, unateness, and juntas which are our focus. We recall the state of the art prior to our results:

- **Tolerant Monotonicity and Unateness Testing:** Little is known about non-trivial algorithms for tolerant monotonicity or unateness testing. It is folklore that the known $2^{\tilde{O}(\sqrt{n/\varepsilon})}$ -query agnostic learning algorithms for monotone (respectively, unriaible Boolean functions [BT96, KKMS08, FKV17] imply the existence of $(\varepsilon_1, \varepsilon_2)$ -tolerant non-adaptive testing algorithms with query complexity $2^{\tilde{O}(\sqrt{n/(\varepsilon_2 - \varepsilon_1)})}$. Strengthening earlier work of [FR10], Pallavoor et al. [PRW22] gave an efficient algorithm for the case when there is a large multiplicative gap between ε_1 and ε_2 ; their algorithm uses $\text{poly}(n, 1/\varepsilon)$ queries to non-adaptively $(\varepsilon, \tilde{O}(\sqrt{n}) \cdot \varepsilon)$ -test monotonicity. Turning to lower bounds, Pallavoor et al. [PRW22] showed that for $0 < \kappa < 1/2$, any non-adaptive $(1/n^{1-\kappa}, 1/\sqrt{n})$ -tolerant tester for monotonicity or unateness must make at least 2^{n^κ} queries. We note that while this can give a strong lower bound in certain regimes when all of $\varepsilon_1, \varepsilon_2, \varepsilon_2 - \varepsilon_1$ are inverse-polynomially small, it does not give any lower bound when any of these values is lower bounded by a constant. Prior to the current work, the only known lower bound in the constant-gap regime was due to Levi and Waingarten [LW19], who showed that for constants $0 < \varepsilon_1 < \varepsilon_2$, tolerant unateness testing requires $\tilde{\Omega}(n)$ (possibly adaptive) queries and requires $\tilde{\Omega}(n^{3/2})$ non-adaptive queries.
- **Tolerant Junta Testing:** De et al. [DMN19] obtained a $2^k \cdot \text{poly}(k, 1/(\varepsilon_2 - \varepsilon_1))$ -query non-adaptive algorithm for testing k -juntas, and subsequently Iyer et al. [ITW21] gave an adaptive algorithm which makes $2^{\tilde{O}(\sqrt{k/(\varepsilon_2 - \varepsilon_1)})}$ queries. In terms of lower bounds, the above-mentioned result of Pallavoor et al. [PRW22] for the “small gap” regime also holds for $n/2$ -junta testing, which implies that for $0 < \kappa < 1/2$, any

	Upper Bound	Prior Work	This Work
Monotonicity	$\exp(\tilde{O}(\sqrt{n}))$ (Folklore)	$\tilde{\Omega}(\sqrt{n})$ [CWX17a]	$\exp(\Omega(n^{1/4}))$ (Theorem 1.1)
Unateness	$\exp(\tilde{O}(\sqrt{n}))$ (Folklore)	$\tilde{\Omega}(n^{3/2})$ [LW19]	$\exp(\Omega(n^{1/4}))$ (Theorem 1.1)
Juntas	$\exp(O(k))$ [DMN19]	$k^{\Omega(1)}$ [LW19, CP23]	$\exp(\Omega(\sqrt{k}))$ (Theorem 1.2)

Table 1: A summary of tolerant property testing upper and lower bounds for non-adaptive algorithms in the constant gap regime (i.e. $\varepsilon_2 - \varepsilon_1 = \Theta(1)$).

non-adaptive $(1/k^{1-\kappa}, 1/\sqrt{k})$ -tolerant tester for k -juntas must make at least 2^{k^κ} many queries. In the constant-gap regime, Levi and Waingarten [LW19] showed that for constants $0 < \varepsilon_1 < \varepsilon_2$, any non-adaptive $(\varepsilon_1, \varepsilon_2)$ -tolerant k -junta tester must make $\Omega(k^2)$ queries. Very recently Chen and Patel [CP23] have given an $\Omega(k^{-\Omega(\log(\varepsilon_2 - \varepsilon_1))})$ -query lower bound for adaptive k -junta testers, which is a $\text{poly}(k)$ -query lower bound when ε_2 and ε_1 differ by an additive constant.

Summarizing the above results, major gaps remain in our understanding of the complexity of $(\varepsilon_1, \varepsilon_2)$ -tolerant testing for all three properties of monotonicity, unateness, and being a k -junta. Given the presumed difficulty of tolerant testing, this lack of knowledge seems to be most acute on the lower bounds side, and particularly in the most interesting and natural regime in which $\varepsilon_1 < \varepsilon_2$ are constants independent of n . Indeed, in this setting, even for non-adaptive testers the best known lower bound prior to the current work was $\tilde{\Omega}(n^{3/2})$ for unateness [LW19] and $\text{poly}(k)$ for k -juntas [LW19, CP23], and no lower bound seems to have been known for monotonicity other than the $\tilde{\Omega}(\sqrt{n})$ lower bound for standard (non-tolerant) non-adaptive monotonicity testing [CWX17a].

1.1 Our Results In this paper we give the first super-polynomial (in fact, mildly exponential) lower bounds for non-adaptive tolerant testing of monotonicity, unateness and juntas when the separation $\varepsilon_2 - \varepsilon_1$ between the “yes” and “no” cases is a constant independent of n . (Equivalently, via the standard connection to distance estimation mentioned earlier, our results imply that a mildly-exponential number of queries are required to perform distance estimation even to constant additive accuracy.)

In more detail, we prove the following main results:

THEOREM 1.1. (LOWER BOUNDS ON TOLERANT TESTERS FOR MONOTONICITY AND UNATENESS) *For any $\varepsilon \in (0, 1)$ with $\varepsilon \geq 1/\sqrt{n}$, there exist $\varepsilon_1, \varepsilon_2 > 0$ with $\varepsilon_1 = \Theta(\varepsilon)$ and $\varepsilon_2 - \varepsilon_1 = \Theta(\varepsilon)$ such that any non-adaptive $(\varepsilon_1, \varepsilon_2)$ -tolerant tester for monotonicity (or unateness) must make $2^{\Omega(n^{1/4}/\sqrt{\varepsilon})}$ queries.*

THEOREM 1.2. (LOWER BOUNDS ON TOLERANT TESTERS FOR k -JUNTAS) *Any non-adaptive $(0.1, 0.2)$ -tolerant tester for the property of being a k -junta must make $2^{\Omega(k^{1/2})}$ queries.*

These results dramatically narrow the gap between lower and upper bounds when $\varepsilon_2 - \varepsilon_1$ is a constant. Indeed, since as noted earlier there are known $2^{\tilde{O}(n^{1/2})}$ -query non-adaptive monotonicity / unateness testers (based on learning [BT96, KKMS08, FKV17]) and there is a known $2^{O(k)}$ -query non-adaptive k -junta tester [DMN19] in this setting, our lower bounds are off from the best possible algorithms by at most a quadratic factor in the exponent. It is an intriguing goal for future work to locate the optimal non-adaptive query complexity of tolerant monotonicity / unateness testing in the range between $2^{n^{1/4}}$ and $2^{\sqrt{n}}$, and likewise for junta testing in the range between $2^{\sqrt{k}}$ and 2^k .

1.2 Techniques The central ingredient of our improved lower bounds for monotonicity and unateness testing is a refinement and strengthening of a construction from [PRW22]. [PRW22] consider distributions $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$ over n -variable “yes”- and “no”-functions, each of which involves a random partition of the n input variables into a set of $n/2$ “control variables” and a complementary set of $n/2$ “action variables.” Intuitively, the interesting

n -bit inputs for both “yes”- and “no”-functions are ones for which the $n/2$ control bits have $n/4$ coordinates set to 0 and $n/4$ coordinates set to 1; on inputs of this sort, the setting of the action variables determines the value of the function (a useful way to think of this is that each distinct input of this sort results in a different “action subcube” determining the value of the function). The values of the function on the action subcubes are carefully defined in such a way as to make it impossible for an algorithm to distinguish “yes”-functions from “no”-functions unless two inputs x, x' are queried which lie in the same action subcube (so the setting of the $n/2$ control bits is the same between the two inputs x and x'), but differ in many of the $n/2$ action variables. Very roughly speaking, the non-adaptive lower bounds of [PRW22] follow from the fact that since the partition of $[n]$ into control variables and action variables is random and unknown, it is very unlikely for any two input strings which are far apart (which they must be in order for the action variables to differ in many locations as sketched above) to differ only in the $n/2$ action variables while being completely identical in the $n/2$ control variables.

It turns out that the analysis of [PRW22] is not able to handle “no”-functions that are more than $\Theta(1/\sqrt{n})$ -far from monotone/unate/junta essentially because the action bits are only consulted in their construction if the $n/2$ control bits are set in a perfectly balanced way (intuitively, this is because the middle layer of the $n/2$ -dimensional hypercube has a $\Theta(1/\sqrt{n})$ -fraction of all points). Our key insight is to change the [PRW22] construction by using (a small extension of) a construction of random monotone DNF formulas due to Talagrand [Tal96]. The purpose of the extension is to let us handle general values of $\varepsilon_2 \geq 1/\sqrt{n}$, but for simplicity we restrict the following discussion to the case that ε_2 is a constant, in which case our construction coincides with Talagrand’s.

Talagrand’s random DNF is a $2^{\sqrt{m}}$ -term monotone DNF formula over $\{0, 1\}^m$ (for us m will be the number of control variables) which has the property that a *constant* fraction of points in $\{0, 1\}^m$ satisfy a *unique* term. Like the points in the middle layer of the control subcube, these uniquely-satisfied points have the property that any pair of points x, x' which respectively satisfy two distinct unique terms $T_i, T_{i'}$ must be incomparable to each other, i.e. they neither satisfy $x \leq x'$ nor $x' \leq x$; this turns out to be essential for us due to the way that the “yes” and “no” functions are defined within the action subcubes. Intuitively, we use these “uniquely-satisfied” points in $\{0, 1\}^m$ in place of the “middle layer” of the control subcube, and rather than having each middle-layer point map to a distinct action subcube we have all points which uniquely satisfy a given term T_i map to the *same* action subcube.

The fact that a constant fraction of points in the control subcube satisfy a unique term in Talagrand’s random DNF is ultimately why we are able to handle constant values of $\varepsilon_1, \varepsilon_2$. Our construction differs in some other regards from the [PRW22] construction as well; to optimally balance parameters it turns out to be best for us to have many fewer action variables than control variables, and our $2^{\Omega(n^{1/4}/\sqrt{\varepsilon})}$ query lower bounds ultimately come from trading off constraints which arise from our use of the Talagrand DNF rather than the “middle layer” as in [PRW22].

Our improved lower bound for junta testing follows a similar high-level approach, but the technical details are simpler; it turns out that in this case, since we are not concerned with monotonicity there is no need to use Talagrand DNF, and instead we use a straightforward indexing scheme in which each assignment to the control bits indexes a different action subcube. (Intuitively, avoiding the use of the Talagrand DNF and the resulting tradeoff which it necessitates is why we are able to achieve an exponent of $1/2$ for juntas rather than the $1/4$ that we achieve for monotonicity and unateness.)

Finally, it is natural to wonder whether using some other function in place of the Talagrand DNF might lead to an improved tradeoff, and hence a quantitatively better lower bound, for monotonicity or unateness testing. In Section 6, we argue that the Talagrand function is in fact optimal for our approach, so improving on our lower bounds for monotonicity or unateness will require a different construction.

2 Preliminaries

For an integer n , we use $[n]$ to denote the set $\{1, \dots, n\}$. For two integers $n_1 \leq n_2$, we use $[n_1 : n_2]$ to denote $\{n_1, \dots, n_2\}$.

We will denote the 0/1-indicator of an event A by $\mathbf{1}\{A\}$. All probabilities and expectations will be with respect to the uniform distribution over the relevant domain unless stated otherwise. We use boldfaced letters such as $\mathbf{x}, \mathbf{f}$, and $\mathbf{A}$ to denote random variables (which may be real-valued, vector-valued, function-valued, or set-valued; the intended type will be clear from the context). We write $\mathbf{x} \sim \mathcal{D}$ to indicate that the random variable $\mathbf{x}$ is distributed according to distribution $\mathcal{D}$.

NOTATION 1. Given a string $x \in \{0,1\}^n$ and a set $A \subseteq [n]$, we write $x_A \in \{0,1\}^A$ to denote the $|A|$ -bit string obtained by restricting x to coordinates in A , i.e. $x_A := (x_i)_{i \in A}$.

Given two strings $x, y \in \{0,1\}^n$, we write $x \leq y$ to indicate that $x_i \leq y_i$ for all i ; if moreover $x \neq y$ then we may write $x < y$.

Given two Boolean functions $f, g : \{0,1\}^n \rightarrow \{0,1\}$, we define the *distance* between f and g (denoted by $\text{dist}(f, g)$) to be the normalized Hamming distance between f and g , i.e.

$$\text{dist}(f, g) := \Pr_{\mathbf{x} \sim \{0,1\}^n} [f(\mathbf{x}) \neq g(\mathbf{x})].$$

A *property* $\mathcal{P}$ is a collection of Boolean functions; we say that a function $f : \{0,1\}^n \rightarrow \{0,1\}$ is ε -far from the property $\mathcal{P}$ if f satisfies

$$\text{dist}(f, \mathcal{P}) := \min_{g \in \mathcal{P}} \text{dist}(f, g) \geq \varepsilon.$$

We say f is ε -close from $\mathcal{P}$ if $\text{dist}(f, \mathcal{P}) \leq \varepsilon$.

We recall the following basic fact about the middle layers of the hypercube $\{0,1\}^n$:

FACT 2.1. We have

$$\frac{1}{4\sqrt{n}} \leq \frac{\binom{n}{n/2+\ell}}{2^n} \leq \frac{1}{\sqrt{n}}, \quad \text{for every integer } \ell \text{ with } |\ell| \leq 0.1\sqrt{n}.$$

2.1 Lower Bounds for Testing Algorithms Our query-complexity lower bounds for tolerant testing algorithms are obtained via Yao's minimax principle [Yao77], which we recall below. (We remind the reader that an algorithm for the problem of $(\varepsilon_1, \varepsilon_2)$ -tolerant testing is correct on an input function f provided that it outputs “yes” if f is ε_1 -close to the property and outputs “no” if f is ε_2 -far from the property; if the distance to the property is between ε_1 and ε_2 then the algorithm is correct regardless of what it outputs.)

THEOREM 2.1. (YAO'S PRINCIPLE) *To prove a q -query lower bound on the worst-case query complexity of any non-adaptive randomized testing algorithm, it suffices to give a distribution $\mathcal{D}$ on instances such that for any q -query non-adaptive deterministic algorithm $\mathcal{A}$, we have*

$$\Pr_{\mathbf{f} \sim \mathcal{D}} [\mathcal{A} \text{ is correct on } \mathbf{f}] \leq c,$$

where $0 \leq c < 1$ is a universal constant.

3 Talagrand's Random DNF

We define a useful distribution over Boolean functions that will play a central role in the proofs of our tolerant lower bound for monotonicity and unateness. The construction is a slight generalization of a distribution over DNF (disjunctive normal form) formulas that was constructed by Talagrand [Tal96]. (The generalization is that we allow a parameter ε to control the size of each term and the number of terms; the original construction corresponds to $\varepsilon = 1$.)

DEFINITION 2. (TALAGRAN'S RANDOM DNF) *Let $\varepsilon \in (0,1)$ and let $L := 0.1 \cdot 2^{\sqrt{n}/\varepsilon}$. Let $\text{Talagrand}(n, \varepsilon)$ be the following distribution on ordered tuples of L monotone terms: for each $i = 1, \dots, L$, the i -th term is obtained by independently drawing a set $\mathbf{T}_i \subseteq [n]$ where each set $\mathbf{T}_i$ is obtained by drawing $\sqrt{n}/\varepsilon$ elements of $[n]$ independently and with replacement. We use $\mathbf{T}$ to denote the ordered tuple $\mathbf{T} = (\mathbf{T}_1, \dots, \mathbf{T}_L)$ which is a draw from $\text{Talagrand}(n, \varepsilon)$. Then a “Talagrand DNF” is given by*

$$\mathbf{f}(x) = \bigvee_{\ell=1}^L \left(\bigwedge_{j \in \mathbf{T}_\ell} x_j \right).$$

It is clear that any Talagrand DNF obtained by a draw from $\text{Talagrand}(n, \varepsilon)$ is a monotone function.

We will frequently view $T_i \subseteq [n]$ as the term $\bigwedge_{j \in T_i} x_j$, where we say $T_i(x) = 1$ if and only if $x_j = 1$ for all $j \in T_i$. We may also write $T = (T_1, \dots, T_k)$ to represent a DNF, which is defined by the disjunction of the terms T_i . We will often be interested in the probability of a random input $\mathbf{x} \sim \{0,1\}^n$ satisfying a unique term $\mathbf{T}_i$ in a Talagrand DNF; towards this, we introduce the following notation:

NOTATION 3. Given a DNF $T = (T_1, \dots, T_k)$ where each T_i is a term, we define the collection of terms of T satisfied by x , written $S_T(x)$, as

$$S_T(x) := \{\ell \in [k] : T_\ell(x) = 1\}.$$

The following claim shows that on average over the draw of $\mathbf{T} \sim \text{Talagrand}(n, \varepsilon)$, an $\Omega(\varepsilon)$ fraction of strings from $\{0, 1\}^n$ satisfy a *unique* term in the Talagrand DNF (i.e. $|S_{\mathbf{T}}(x)| = 1$ for $\Omega(\varepsilon)$ -fraction of $x \in \{0, 1\}^n$). We note that an elegant argument of Kane [Kan13] gives this for $\varepsilon = \Theta(1)$, but this argument does not extend to the setting of small ε which we require. The proof below is based on an argument due to O'Donnell and Wimmer (cf. Theorem 2.1 of [OW07]).

PROPOSITION 3.1. For $\varepsilon \in (0, 1)$, let $\mathbf{T} \sim \text{Talagrand}(n, \varepsilon)$ be as in Definition 2. Then

$$\Pr_{\mathbf{T}, x} [|S_{\mathbf{T}}(x)| = 1] = \Omega(\max\{\varepsilon, 1/\sqrt{n}\}).$$

Proof. Note that Proposition 3.1 is immediate if the following holds: For every string $x \in \{0, 1\}^n$ with $|x| \in [n/2, n/2 + 0.05\varepsilon\sqrt{n}]$,³ we have

$$(3.1) \quad \Pr_{\mathbf{T}} [|S_{\mathbf{T}}(x)| = 1] = \Omega(1).$$

This is because a straightforward application of the Chernoff bound, and the well-known middle binomial coefficient bound $\binom{n}{n/2}/2^n = \Theta(1/\sqrt{n})$, together imply that

$$\Pr_x [|x| \in [n/2, n/2 + 0.05\varepsilon\sqrt{n}]] = \Omega(\max\{\varepsilon, 1/\sqrt{n}\}).$$

We prove Equation (3.1) in the rest of the proof. Fix such an $x \in \{0, 1\}^n$ and let $\mathbf{T}_i$ be one of the $0.1 \cdot 2^{\sqrt{n}/\varepsilon}$ terms of $\mathbf{T} \sim \text{Talagrand}(n, \varepsilon)$. Recalling that $\mathbf{T}_i$ consists of $\sqrt{n}/\varepsilon$ many variables (with repetition), we have

$$\Pr_{\mathbf{T}_i} [\mathbf{T}_i(x) = 1] \leq \left(\frac{1}{2} + \frac{0.05\varepsilon}{\sqrt{n}}\right)^{\sqrt{n}/\varepsilon} = \left(\frac{1}{2} + \frac{0.1\varepsilon}{2\sqrt{n}}\right)^{\sqrt{n}/\varepsilon} \leq 2^{-\sqrt{n}/\varepsilon} \exp(0.1).$$

where in the second inequality we used the fact that $1 + x \leq e^x$ for all $x \in \mathbb{R}$. It follows by the linearity of expectation that for any x as above, we have

$$\mathbf{E}_{\mathbf{T}} [|S_{\mathbf{T}}(x)|] = \sum_{i=1}^{0.1 \cdot 2^{\sqrt{n}/\varepsilon}} \Pr_{\mathbf{T}_i} [\mathbf{T}_i(x) = 1] \leq 0.1 \cdot \exp(0.1) < 0.12.$$

Markov's inequality then implies that

$$(3.2) \quad \Pr_{\mathbf{T}} [|S_{\mathbf{T}}(x)| \geq 2] \leq \frac{1}{2} \cdot \mathbf{E}_{\mathbf{T}} [|S_{\mathbf{T}}(x)|] \leq 0.06.$$

On the other hand, since $|x| \geq n/2$, we have

$$(3.3) \quad \Pr_{\mathbf{T}} [|S_{\mathbf{T}}(x)| = 0] \leq \left(1 - \left(\frac{1}{2}\right)^{\sqrt{n}/\varepsilon}\right)^{0.1 \cdot 2^{\sqrt{n}/\varepsilon}} \leq \exp(-0.1) < 0.91,$$

where the second inequality used again $1 + x \leq e^x$. Combining Equations (3.2) and (3.3), we get that

$$\Pr_{\mathbf{T}} [|S_{\mathbf{T}}(x)| = 1] > 0.03,$$

thus establishing Equation (3.1), which in turn completes the proof. $\square$

³Note that when $0.05\varepsilon\sqrt{n} < 1$, only strings $x \in \{0, 1\}^n$ with $|x| = n/2$ are considered.

4 Lower Bounds on Tolerant Testers for Monotonicity and Unateness

We start with some objects that we need in the construction of the two distributions $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$.

Let $\varepsilon \in (0, 1)$ be a parameter with $\varepsilon \geq c_0/\sqrt{n}$ for some sufficiently large constant c_0 . We partition the variables $x_1, \dots, x_n$ into *control* variables and *action* variables as follows: Let $a := \sqrt{n}/\varepsilon$ and $A \subseteq [n]$ be a fixed subset of $[n]$ of size a . Let $C := [n] \setminus A$. We refer to the variables x_i for $i \in C$ as *control variables* and the variables x_i for $i \in A$ as *action variables*. We first define two pairs of functions over $\{0, 1\}^A$ on the action variables as follows (we will use these functions later in the definition of $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$): Let $h^{(+,0)}, h^{(+,1)}, h^{(-,0)}$ and $h^{(-,1)}$ be Boolean functions over $\{0, 1\}^A$ defined as follows:

$$h^{(+,0)}(x_A) = \begin{cases} 0 & |x_A| > \frac{a}{2} + c_1\sqrt{a}; \\ 0 & |x_A| \in [\frac{a}{2} - c_1\sqrt{a}, \frac{a}{2} + c_1\sqrt{a}]; \\ 0 & |x_A| < \frac{a}{2} - c_1\sqrt{a}. \end{cases} \quad h^{(+,1)}(x_A) = \begin{cases} 1 & |x_A| > \frac{a}{2} + c_1\sqrt{a}; \\ 0 & |x_A| \in [\frac{a}{2} - c_1\sqrt{a}, \frac{a}{2} + c_1\sqrt{a}]; \\ 1 & |x_A| < \frac{a}{2} - c_1\sqrt{a}. \end{cases}$$

and

$$h^{(-,0)}(x_A) = \begin{cases} 1 & |x_A| > \frac{a}{2} + c_1\sqrt{a}; \\ 0 & |x_A| \in [\frac{a}{2} - c_1\sqrt{a}, \frac{a}{2} + c_1\sqrt{a}]; \\ 0 & |x_A| < \frac{a}{2} - c_1\sqrt{a}. \end{cases} \quad h^{(-,1)}(x_A) = \begin{cases} 0 & |x_A| > \frac{a}{2} + c_1\sqrt{a}; \\ 0 & |x_A| \in [\frac{a}{2} - c_1\sqrt{a}, \frac{a}{2} + c_1\sqrt{a}]; \\ 1 & |x_A| < \frac{a}{2} - c_1\sqrt{a}. \end{cases}$$

for some sufficiently small constant c_1 to be specified later.

Now we are ready to define the distributions $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$ over $f : \{0, 1\}^n \rightarrow \{0, 1\}$. We follow the convention that random variables are in boldface and fixed quantities are in the standard typeface.

Let $m = n - a$. A function $\mathbf{f}_{\text{yes}} \sim \mathcal{D}_{\text{yes}}$ is drawn as follows. We start by sampling a subset $\mathbf{A} \subseteq [n]$ of size a uniformly at random and let $\mathbf{C} := [n] \setminus \mathbf{A}$. Note that there are in total $n - a$ control variables. We let $L := 0.1 \cdot 2^{\sqrt{n-a}/\varepsilon}$ and draw an L -term monotone Talagrand DNF $\mathbf{T} \sim \text{Talagrand}(n - a, \varepsilon)$ on $\mathbf{C}$ as described in Definition 2. Finally, we sample L random bits $\mathbf{b} \in \{0, 1\}^L$ uniformly at random. Given $\mathbf{A}, \mathbf{T}$ and $\mathbf{b}$, $\mathbf{f}_{\text{yes}}$ is defined by letting

$$\mathbf{f}_{\text{yes}}(x) = \begin{cases} 1 & |S_{\mathbf{T}}(x_{\mathbf{C}})| > 1 \text{ or } |x_{\mathbf{C}}| > m/2 + 0.05\varepsilon\sqrt{m}; \\ 0 & |S_{\mathbf{T}}(x_{\mathbf{C}})| = 0 \text{ or } |x_{\mathbf{C}}| < m/2; \\ h^{(+,0)}(x_{\mathbf{A}}) & S_{\mathbf{T}}(x_{\mathbf{C}}) = \{\ell\}, |x_{\mathbf{C}}| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}] \text{ and } \mathbf{b}_{\ell} = 0; \\ h^{(+,1)}(x_{\mathbf{A}}) & S_{\mathbf{T}}(x_{\mathbf{C}}) = \{\ell\}, |x_{\mathbf{C}}| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}] \text{ and } \mathbf{b}_{\ell} = 1. \end{cases}$$

To draw a function $\mathbf{f}_{\text{no}} \sim \mathcal{D}_{\text{no}}$, we sample $\mathbf{A}, \mathbf{T}$ and $\mathbf{b}$ exactly as in the definition of $\mathcal{D}_{\text{yes}}$ above, and we replace $h^{(+,0/1)}$ by $h^{(-,0/1)}$ in the construction described above. In more detail, $\mathbf{f}_{\text{no}}$ is defined by letting

$$\mathbf{f}_{\text{no}}(x) = \begin{cases} 1 & |S_{\mathbf{T}}(x_{\mathbf{C}})| > 1 \text{ or } |x_{\mathbf{C}}| > m/2 + 0.05\varepsilon\sqrt{m}; \\ 0 & |S_{\mathbf{T}}(x_{\mathbf{C}})| = 0 \text{ or } |x_{\mathbf{C}}| < m/2; \\ h^{(-,0)}(x_{\mathbf{A}}) & S_{\mathbf{T}}(x_{\mathbf{C}}) = \{\ell\}, |x_{\mathbf{C}}| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}] \text{ and } \mathbf{b}_{\ell} = 0; \\ h^{(-,1)}(x_{\mathbf{A}}) & S_{\mathbf{T}}(x_{\mathbf{C}}) = \{\ell\}, |x_{\mathbf{C}}| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}] \text{ and } \mathbf{b}_{\ell} = 1. \end{cases}$$

4.1 Distance to Monotonicity and Unateness It is easy to verify that every function from $\mathcal{D}_{\text{yes}}$ is close to monotone (and thus, unate):

LEMMA 4.1. *Every function in the support of $\mathcal{D}_{\text{yes}}$ is $(0.1c_1\varepsilon)$ -close to monotonicity.*

Proof. Let f be a function in the support of $\mathcal{D}_{\text{yes}}$ defined using $\mathbf{A}, \mathbf{T}$ and $\mathbf{b}$. Let f' be the partial function obtained from f by replacing $f(x)$ with nil for any $x \in \{0, 1\}^n$ that satisfies $S_{\mathbf{T}}(x_{\mathbf{C}}) = \{\ell\}$ for some ℓ and the following two conditions:

$$(4.4) \quad |x_{\mathbf{C}}| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}] \quad \text{and} \quad |x_{\mathbf{A}}| \in [a/2 - c_1\sqrt{a}, a/2 + c_1\sqrt{a}].$$




Figure 1: *

(a) A draw of $\mathbf{f}_{\text{yes}} \sim \mathcal{D}_{\text{yes}}$




Figure 2: *

(b) A draw of $\mathbf{f}_{\text{no}} \sim \mathcal{D}_{\text{no}}$

Figure 3: The left hand side depicts the control subcube $\{0,1\}^C$ with the terms of the Talagrand DNF, and the right hand side depicts an action subcube $\{0,1\}^A$. The cross-hatched region in the control subcube corresponds to outcomes of the control bits for which the action subcube determines the value of the function and the dashed lines indicate Hamming weight levels.

Note that the fraction of points erased, by Fact 2.1, is at most $0.05\varepsilon \cdot 2c_1 = 0.1c_1\varepsilon$. So it suffices to show that the partial function f' is monotone.

To prove this, it suffices to show that for any $x < y$, if $f(x) = 1$ and $f(y) \neq \text{nil}$, then $f(y) = 1$. The case when x satisfies $|S_T(x_C)| > 1$ or $|x_C| > m/2 + 0.05\varepsilon\sqrt{m}$ is trivial. Otherwise, x satisfies $S_T(x_C) = \{\ell\}$ for some ℓ and $|x_C| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}]$, and $b_\ell = 1$. Given that $x < y$, either $|S_T(y_C)| > 1$, in which case $f(y) = 1$ and we are done, or $|y_C| > m/2 + 0.05\varepsilon\sqrt{m}$, in which case $f(y) = 1$ and we are also done, or $S_T(y_C) = \{\ell\}$ and $|y_C| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}]$. For the latter, we have either y_A is in the middle layers and $f(y) = \text{nil}$ or y_A is not in the middle layers and thus, $f(y) = 1$. This finishes the proof of the lemma. $\square$

Next we show that with high probability a function drawn from $\mathcal{D}_{\text{no}}$ is far from unate (and thus, far from monotone as well). Before that, we first show a property of the random choices of $\mathbf{A}, \mathbf{T}$ and $\mathbf{b}$.

LEMMA 4.2. *Recall that $m = n - a$. With probability at least 0.01 over $\mathbf{A}, \mathbf{T}$ and $\mathbf{b}$, the number of $x \in \{0, 1\}^C$ that satisfies $S_T(x) = \{\ell\}$ for some ℓ , $|x_C| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}]$ and $b_\ell = 1$ is $\Omega(\varepsilon) \cdot 2^m$. Symmetrically, with probability at least 0.01 over $\mathbf{A}, \mathbf{T}$ and $\mathbf{b}$, the number of $x \in \{0, 1\}^C$ that satisfies the same conditions above except $b_\ell = 0$ is also $\Omega(\varepsilon) \cdot 2^m$.*

Proof. We first introduce the event $\text{GoodTalagrand}(T)$, which states that there exists an $\Omega(\varepsilon)$ -fraction of points $x \in \{0, 1\}^m$ such that $|S_T(x)| = 1$. Formally, let $t \cdot \varepsilon$ be the number of x in $\{0, 1\}^m$ such that $|x| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}]$ and let $\text{GoodTalagrand}(T)$ be the event that $\mathbb{E}_{|x| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}]} [\mathbf{1}\{|S_T(x)| = 1\}] \geq t \cdot \varepsilon/100$. We will show that

$$\Pr_{T \sim \text{Talagrand}(m, \varepsilon)} [\text{GoodTalagrand}(T)] \geq 0.02.$$

Let p denote $\Pr_{T \sim \text{Talagrand}(m, \varepsilon)} [\text{GoodTalagrand}(T)]$. Recall that in Proposition 3.1, we have shown that $\Pr_T [|S_T(x)| = 1] \geq 0.03$ for any $x \in \{0, 1\}^m$ with $|x|_1 \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}]$. So we have

$$p \cdot t\varepsilon + (1 - p) \cdot t\varepsilon/100 \geq t\varepsilon \Pr_T [|S_T(x)| = 1] \geq 0.03t\varepsilon,$$

which implies that $p \geq 2/99 > 0.02$.

Fix an arbitrary T in the support of $\text{Talagrand}(m, \varepsilon)$ such that $\text{GoodTalagrand}(T)$ happens. Note that for each x such that $S_T(x) = \{\ell\}$ for some $\ell \in [L]$, $\mathbb{E}[b_\ell] = 1/2$. So by linearity of expectation and Markov's inequality, we know that with probability at least 99%, there is an $\Omega(\varepsilon)$ -fraction of points x such that $S_T(x) = \{\ell\}$ and $b_\ell = 1$. The other symmetric statement follows from the same argument. $\square$

LEMMA 4.3. *With probability at least 0.01, $\mathbf{f} \sim \mathcal{D}_{\text{no}}$ is $\Omega(\varepsilon)$ -far from unate.*

Proof. We first include a folklore claim and its proof for completeness.

CLAIM 4. *For integer $0 \leq w \leq a$, let P_w denote the set of points in $\{0, 1\}^a$ with Hamming weight w , i.e. $P_w = \{x \in \{0, 1\}^a : |x|_1 = w\}$. Then for any $0 \leq w \leq a/2$, the bipartite graph (P_w, P_{a-w}) with the poset relations as edges has a perfect matching.*

Proof. The key point is that (P_w, P_{a-w}) is a k -regular bipartite graph, where $k = \binom{a-w}{w}$.

We apply Hall's theorem to show this claim. Consider any subset $S \subseteq P_{a-w}$. The number of edges associated with S is exactly $k|S|$. Let $\mathcal{N}(S)$ be the neighborhood of S . Then we know the number of edges associated with $\mathcal{N}(S)$ is exactly $k|\mathcal{N}(S)|$, and these edges include the $k|S|$ edges above. So we have $k|S| \leq k|\mathcal{N}(S)|$, which means $|S| \leq |\mathcal{N}(S)|$. $\square$

We will use the following claim about the two functions $h^{(-,0)}$ and $h^{(-,1)}$:

CLAIM 5. *Fix any set $A \subset [n]$ of size a . For any $r \in \{0, 1\}^A$, either*

$$h_0 := h^{(-,0)}(x \oplus r) \quad \text{or} \quad h_1 := h^{(-,1)}(x \oplus r)$$

is $\Omega(1)$ -far from monotone.

Proof. Fix a string $r \in \{0, 1\}^A$. Without loss of generality, we assume that r satisfies $|r| \leq a/2$ and show that h_1 is $\Omega(1)$ -far from monotone. For the case when $|r| < a/2$, we can take $r' = r \oplus 1$ with $|r'| \geq a/2$. Then what we prove below shows that $h^{(-,1)}(x \oplus r')$ is $\Omega(1)$ -far from monotone. On the other hand, we have $h^{(-,1)}(x) = h^{(-,0)}(x \oplus 1)$ and thus,

$$h^{(-,0)}(x \oplus r) = h^{(-,1)}(x \oplus r \oplus 1) = h^{(-,1)}(x \oplus r')$$

is $\Omega(1)$ -far from monotone.

In the rest of the proof we focus on the case when $|r| \leq a/2$ and show that h_1 is $\Omega(1)$ -far from monotone. By the symmetry of h_1 , we assume without loss of generality that $r_1 = \dots = r_k = 0$, where $k = a - |r| \geq a/2$. To prove that h_1 is $\Omega(1)$ -far from monotone, it suffices to give $\Omega(2^a)$ many disjoint pairs (x, y) such that $x < y$, $h_1(x) = 1$ and $h_1(y) = 0$.

We start by picking a string $z \in \{0, 1\}^{a-k}$ such that

$$|z| \in [(a-k)/2 - 0.01\sqrt{a}, (a-k)/2 + 0.01\sqrt{a}].$$

Note that the number of such z , by Fact 2.1, is $\Omega(2^{a-k})$.

Next, for any such string z , we build disjoint pairs (x, y) such that $x < y$,

$$x_{[k+1:a]} = y_{[k+1:a]} = z, \quad |x_{[k]}| \leq k/2 - 0.02\sqrt{a} \quad \text{and} \quad |y_{[k]}| \geq k/2 + 0.02\sqrt{a}.$$

On the one hand, it follows from $k \geq a/2$, Fact 2.1 and Definition 4 that the number of such disjoint pairs is $\Omega(2^k)$. On the other hand, we have $h_1(x) = 1$ and $h_1(y) = 0$ when c_1 is sufficiently small. For example, for $h_1(x)$, we have

$$|x \oplus r| = |x_{[k]}| + (a-k) - |z| \leq k/2 - 0.02\sqrt{a} + (a-k)/2 + 0.01\sqrt{a} = a/2 - 0.01\sqrt{a},$$

which is smaller than $a/2 - c_1\sqrt{a}$ when $c_1 < 0.01$.

As a result, the total number of disjoint pairs is $\Omega(2^a)$ and the claim follows. $\square$

To prove Lemma 4.3, we have from Lemma 4.2 that with probability at least 0.01 over $\mathbf{A}, \mathbf{T}$ and $\mathbf{b}$, the number of $y \in \{0, 1\}^C$ that satisfies $S_{\mathbf{T}}(y) = \{\ell\}$ for some ℓ , $|y_C| \in [m/2, m/2 + 0.05\epsilon\sqrt{m}]$ and $\mathbf{b}_\ell = 1$ is $\Omega(\epsilon) \cdot 2^{n-a}$ and (symmetrically) the number of $y \in \{0, 1\}^C$ that satisfies the same conditions with $\mathbf{b}_\ell = 0$ is also $\Omega(\epsilon) \cdot 2^{n-a}$. The lemma then follows from Definition 5. $\square$

To summarize, by setting the constant c_1 sufficiently small, there are ϵ_1 and ϵ_2 satisfying

$$\epsilon_1 = \Theta(\epsilon) \quad \text{and} \quad \epsilon_2 - \epsilon_1 = \Theta(\epsilon)$$

such that every function drawn from $\mathcal{D}_{\text{yes}}$ is ϵ_1 -close to monotone (and hence ϵ_1 -close to unate) and a function drawn from $\mathcal{D}_{\text{no}}$ is ϵ_2 -far from unate with probability at least 0.01.

4.2 Indistinguishability of $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$ To prove Theorem 1.1, we show that no non-adaptive deterministic algorithm $\mathcal{A}$ that makes $q = 2^{c_2 n^{1/4}/\sqrt{\epsilon}}$ queries, for some sufficiently small constant c_2 , can distinguish $\mathcal{D}_{\text{yes}}$ from $\mathcal{D}_{\text{no}}$. Specifically, for any nonadaptive deterministic algorithm $\mathcal{A}$ with query complexity $q = 2^{c_2 n^{1/4}/\sqrt{\epsilon}}$, we show that

$$(4.5) \quad \Pr_{\mathbf{f}_{\text{yes}} \sim \mathcal{D}_{\text{yes}}} [\mathcal{A} \text{ accepts } \mathbf{f}_{\text{yes}}] \leq \Pr_{\mathbf{f}_{\text{no}} \sim \mathcal{D}_{\text{no}}} [\mathcal{A} \text{ accepts } \mathbf{f}_{\text{no}}] + o_n(1).$$

To this end, we define **Bad** to be the event that there are two strings x and y queried by $\mathcal{A}$ that satisfy $S_{\mathbf{T}}(x_C) = S_{\mathbf{T}}(y_C) = \{\ell\}$ for some ℓ and $|x_C|, |y_C| \in [m/2, m/2 + 0.05\epsilon m]$ such that one is in the top region and the other is in the bottom region of the action cube, namely $|x_{\mathbf{A}}| > a/2 + c_1\sqrt{a}$ and $|y_{\mathbf{A}}| < a/2 - c_1\sqrt{a}$. We will first show in Lemma 4.4 that the algorithm can distinguish $\mathcal{A}$ only when **Bad** occurs. On the other hand, in Lemma 4.5, we show **Bad** occurs with probability $o_n(1)$ when the number of queries is $2^{c_2 n^{1/4}/\sqrt{\epsilon}}$ and c_2 is sufficiently small (compared to c_1).

The formal argument proceeds as follows. We write $\mathcal{A}(f)$ to denote the sequence of q answers to the queries made by $\mathcal{A}$ to f . We write $\text{view}_{\mathcal{A}}(\mathcal{D}_{\text{yes}})$ (respectively $\text{view}_{\mathcal{A}}(\mathcal{D}_{\text{no}})$) to be the distribution of $\mathcal{A}(f)$ for $f \sim \mathcal{D}_{\text{yes}}$ (respectively $f \sim \mathcal{D}_{\text{no}}$). The following claim asserts that conditioned on **Bad** not happening, the distributions $\text{view}_{\mathcal{A}}(\mathcal{D}_{\text{yes}}|\overline{\text{Bad}})$ and $\text{view}_{\mathcal{A}}(\mathcal{D}_{\text{no}}|\overline{\text{Bad}})$ are identical.

LEMMA 4.4. $\text{view}_{\mathcal{A}}(\mathcal{D}_{\text{yes}}|_{\overline{\text{Bad}}}) = \text{view}_{\mathcal{A}}(\mathcal{D}_{\text{no}}|_{\overline{\text{Bad}}})$.

Proof. Let $Q_{\mathcal{A}}$ be the set of points queried by $\mathcal{A}$.

Recall that the distributions of the partition of $[n]$ into control variables $\mathbf{C}$ and action variables $\mathbf{A}$ are identical for $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$. So fix an arbitrary partition C and A . As the distribution of Talagrand DNF $\mathbf{T} \sim \text{Talagrand}(m, \varepsilon)$ is also identical for $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$, we fix an arbitrary T . Let $\mathbf{f}_{\text{yes}} \sim \mathcal{D}_{\text{yes}}$ be a random function drawn from $\mathcal{D}_{\text{yes}}$.

Note that for any point $x \in \{0, 1\}^n$ such that $|S_T(x_C)| \neq 1$, $|x_C| \notin [m/2, m/2 + 0.05\varepsilon\sqrt{m}]$ or $|x_A| \in [a/2 - c_1\sqrt{a}, a/2 + c_1\sqrt{a}]$, by construction we have that $\mathbf{f}(x)$ can be determined directly in the same way for both $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$. So it suffices for us to consider the points x such that $|S_T(x_C)| = 1$, $|x_C| \in [m/2, m/2 + 0.05\varepsilon\sqrt{m}]$, and $|x_A| \notin [a/2 - c_1\sqrt{a}, a/2 + c_1\sqrt{a}]$. We call these points *important* points.

We divide these important points into disjoint groups according to $S_T(x_C)$. More precisely, for every $\ell \in [L]$, let $X_{\ell} = \{x \mid x \text{ is important, } S_T(x_C) = \{\ell\}\}$. Let $\mathbf{f}_{\ell}(x)$ denote the function $\mathbf{f}(x)$ restricted to X_{ℓ} . Note that for a fixed $\ell \in [L]$, the functions $\mathbf{f}_{\ell}(x)$ only depends on the random bit $\mathbf{b}_{\ell}$. As a result, the distributions of functions $\mathbf{f}_{\ell}(x)$ for different ℓ are independent.

So it suffices for us to fix an arbitrary $\ell \in [L]$ and only consider points that are in X_{ℓ} . The condition that **Bad** does not happen implies that $|x_A| > a/2 + c_1\sqrt{a}$ for all $x \in Q_{\mathcal{A}} \cap X_{\ell}$ or $|x_A| < a/2 - c_1\sqrt{a}$ for all $x \in Q_{\mathcal{A}} \cap X_{\ell}$. In particular, this means $\mathbf{f}_{\ell}(x) = \mathbf{f}_{\ell}(y)$ for all $x, y \in Q_{\mathcal{A}} \cap X_{\ell}$, and this holds for both $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$.

Since $\mathbf{f}_{\ell}(x)$ are the same for all $x \in Q_{\mathcal{A}} \cap X_{\ell}$, the distribution of $\mathbf{f}_{\ell}$ is actually one random bit, which only depends on the uniform random bit $\mathbf{b}_{\ell}$. Indeed, $\mathbf{f}_{\ell}(x^i) = 0$ with probability $1/2$ and $\mathbf{f}_{\ell}(x^i) = 1$ with probability $1/2$, which holds for both $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$.

This finishes the proof. $\square$

Next, we show that the probability that **Bad** happens is small (recall that $q = 2^{c_2 n^{1/4}/\sqrt{\varepsilon}}$):

LEMMA 4.5. *For any set of points $Q_{\mathcal{A}} = \{x^1, \dots, x^q\} \subseteq \{0, 1\}^n$, $\Pr[\text{Bad}] = o_n(1)$.*

Proof. Fix any $x, y \in \{0, 1\}^n$. We will upper bound the probability that $S_T(x_C) = S_T(y_C) = \{\ell\}$ for some $\ell \in [L]$ and $|x_A| < \frac{a}{2} - c_1\sqrt{a}$ and $|y_A| > \frac{a}{2} + c_1\sqrt{a}$. Call this specific event **Bad** _{xy} .

Let I_{01} be the set of i with $x_i = 0$ and $y_i = 1$. On one hand, for **Bad** _{xy} to happen, we have:

$$(\diamond) \quad |I_{01} \cap \mathbf{A}| \geq 2c_1\sqrt{a}.$$

On the other hand, to have $S_T(x_C) = S_T(y_C) = \{\ell\}$, we must have:

$$(\star) \quad \text{There exists } \ell \in [L] \text{ such that } S_T(x) = S_T(y) = \{\ell\}.$$

It follows that

$$\Pr[\text{Bad}_{xy}] \leq \min(\Pr[\diamond], \Pr[\star]);$$

we will in fact show that

$$\min(\Pr[\diamond], \Pr[\star]) \leq 2^{-0.25c_1 n^{1/4}/\sqrt{\varepsilon}}.$$

Let $t = |I_{01}|$. By the random choice of the coordinates defining the action cube $\mathbf{A}$, we have

$$\begin{aligned} \Pr[\diamond] &\leq \Pr\left[\text{Bin}\left(a, \frac{t}{n-a}\right) \geq 2c_1\sqrt{a}\right] \leq \binom{a}{2c_1\sqrt{a}} \cdot \left(\frac{t}{n-a}\right)^{2c_1\sqrt{a}} \\ &\leq \left(\frac{ea}{2c_1\sqrt{a}}\right)^{2c_1\sqrt{a}} \cdot \left(\frac{t}{n-a}\right)^{2c_1\sqrt{a}} \leq \left(\frac{et\sqrt{a}}{2c_1(n-a)}\right)^{2c_1\sqrt{a}} \leq \left(\frac{et\sqrt{a}}{2c_1n(1-\frac{1}{c_0})}\right)^{2c_1\sqrt{a}}. \end{aligned}$$

To bound $\Pr[\star]$, we use

$$\begin{aligned} \Pr[\star] &= \Pr[S_T(x) = S_T(y) \text{ and there exists } \ell \in [L] \text{ such that } S_T(y) = \{\ell\}] \\ &\leq \Pr[S_T(x) = S_T(y) \mid \text{there exists } \ell \in [L] \text{ such that } S_T(y) = \{\ell\}] \\ &\leq \max_{\ell \in [L]} \Pr[S_T(x) = S_T(y) \mid S_T(y) = \{\ell\}] \\ &\leq \left(1 - \frac{t}{n-a}\right)^{\sqrt{n-a}/\varepsilon} \leq e^{-t/(\varepsilon\sqrt{n-a})} \leq e^{-t/(\varepsilon\sqrt{n})}, \end{aligned}$$

where the last line above is by the definition of the random process $\mathbf{T} \sim \text{Talagrand}(m, \varepsilon)$. Next, note that

$$\begin{aligned} \text{If } t \leq c_1 \frac{1}{4} n^{3/4} \cdot \sqrt{\varepsilon}, \text{ then } \Pr[\diamond] &\leq 2^{-c_1 n^{1/4}/\sqrt{\varepsilon}}; \text{ and} \\ \text{If } t > c_1 \frac{1}{4} n^{3/4} \cdot \sqrt{\varepsilon}, \text{ then } \Pr[\star] &< 2^{-0.25c_1 n^{1/4}/\sqrt{\varepsilon}}. \end{aligned}$$

Overall, we thus get that

$$\Pr[\text{Bad}_{xy}] \leq \min(\Pr[\diamond], \Pr[\star]) \leq 2^{-0.25c_1 n^{1/4}/\sqrt{\varepsilon}}.$$

By a union bound for all pairs of points of $Q_{\mathcal{A}}$, we know that

$$\Pr[\text{Bad}] \leq 2^{-0.25c_1 n^{1/4}/\sqrt{\varepsilon}} \cdot \left(2^{c_2 n^{1/4}/\sqrt{\varepsilon}}\right)^2 = o_n(1)$$

as long as c_2 is sufficiently small (compared to c_1). This completes the proof. $\square$

Now we are ready to prove Theorem 1.1.

Proof. [Proof of Theorem 1.1] Let $\mathcal{D} = \frac{1}{2}\{\mathcal{D}_{\text{yes}} + \mathcal{D}_{\text{no}}\}$. Then we have

$$\begin{aligned} \Pr_{\mathcal{D}}[\mathcal{A} \text{ is correct on } \mathbf{f}] &= \frac{1}{2} \left(\Pr_{\mathcal{D}_{\text{yes}}}[\mathcal{A} \text{ is correct on } \mathbf{f}_{\text{yes}}] + \Pr_{\mathcal{D}_{\text{no}}}[\mathcal{A} \text{ is correct on } \mathbf{f}_{\text{no}}] \right) \\ (4.6) \quad &= \frac{1}{2} \left(\Pr_{\mathcal{D}_{\text{yes}}}[\mathcal{A} \text{ accepts } \mathbf{f}_{\text{yes}}] + \Pr_{\mathcal{D}_{\text{no}}}[\mathcal{A} \text{ is correct on } \mathbf{f}_{\text{no}}] \right) \\ (4.7) \quad &\leq \frac{1}{2} \left(\Pr_{\mathcal{D}_{\text{yes}}}[\mathcal{A} \text{ accepts } \mathbf{f}_{\text{yes}}] + 0.99 + 0.01 \Pr_{\mathcal{D}_{\text{no}}}[\mathcal{A} \text{ rejects } \mathbf{f}_{\text{no}}] \right) \\ &= \frac{1}{2} \left(\Pr_{\mathcal{D}_{\text{yes}}}[\mathcal{A} \text{ accepts } \mathbf{f}_{\text{yes}}] + 1 - 0.01 \Pr_{\mathcal{D}_{\text{no}}}[\mathcal{A} \text{ accepts } \mathbf{f}_{\text{no}}] \right) \\ &\leq \frac{199}{200} + \frac{1}{200} \left(\Pr_{\mathcal{D}_{\text{yes}}}[\mathcal{A} \text{ accepts } \mathbf{f}_{\text{yes}}] - \Pr_{\mathcal{D}_{\text{no}}}[\mathcal{A} \text{ accepts } \mathbf{f}_{\text{no}}] \right) \\ (4.8) \quad &= \frac{199}{200} + \frac{\Pr[\text{Bad}]}{200} \left(\Pr_{\mathcal{D}_{\text{no}}|\text{Bad}}[\mathcal{A} \text{ accepts } \mathbf{f}_{\text{no}}] - \Pr_{\mathcal{D}_{\text{yes}}|\text{Bad}}[\mathcal{A} \text{ accepts } \mathbf{f}_{\text{yes}}] \right) \\ &\leq \frac{199}{200} + \frac{\Pr[\text{Bad}]}{200} \\ (4.9) \quad &\leq \frac{199}{200} + o_n(1), \end{aligned}$$

where Equation (4.6) is because of Lemma 4.1, Equation (4.7) is because $\mathbf{f}_{\text{no}}$ is not c_2 -far from unate with probability at most 0.99 thanks to Lemma 4.3, Equation (4.8) is from Lemma 4.4, and Equation (4.9) follows from Lemma 4.5. Theorem 1.1 now follows from Yao's minimax principle (Theorem 2.1). $\square$

5 Lower Bounds on Tolerant Testers for Juntas

We use a different, simpler pair of distributions $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$ for juntas. Let $a = n/2$ and $A \subset [n]$ be a set of size a . The four functions $h^{(+,0)}, h^{(+,1)}, h^{(-,0)}$ and $h^{(-,1)}$ over $\{0,1\}^A$ are defined in the same way as in Section 4 with the constant c_1 fixed to be 0.05.

To draw a function $\mathbf{f}_{\text{yes}} \sim \mathcal{D}_{\text{yes}}$, we first sample a set $\mathbf{A} \subset [n]$ of size a uniformly at random, set $\mathbf{C} = [n] \setminus \mathbf{A}$, and sample a Boolean function $\mathbf{b}$ over $\{0,1\}^{\mathbf{C}}$ uniformly at random. Then the Boolean function $\mathbf{f}_{\text{yes}}$ over $\{0,1\}^n$ is defined using $\mathbf{A}$ and $\mathbf{b}$ as follows:

$$\mathbf{f}_{\text{yes}}(x) = \begin{cases} h^{(+,0)}(x_{\mathbf{A}}) & \mathbf{b}(x_{\mathbf{C}}) = 0 \\ h^{(+,1)}(x_{\mathbf{A}}) & \mathbf{b}(x_{\mathbf{C}}) = 1 \end{cases}$$

To draw $\mathbf{f}_{\text{no}} \sim \mathcal{D}_{\text{no}}$, we first sample $\mathbf{A}$ and $\mathbf{b}$ in the same way as in $\mathcal{D}_{\text{yes}}$, and $\mathbf{f}_{\text{no}}$ is defined as

$$\mathbf{f}_{\text{no}}(x) = \begin{cases} h^{(-,0)}(x_{\mathbf{A}}) & \mathbf{b}(x_{\mathbf{C}}) = 0 \\ h^{(-,1)}(x_{\mathbf{A}}) & \mathbf{b}(x_{\mathbf{C}}) = 1 \end{cases}$$

We first show that every function in the support of $\mathcal{D}_{\text{yes}}$ is close to a $(n/2)$ -junta.

LEMMA 5.1. *Every function in the support of $\mathcal{D}_{\text{yes}}$ is 0.1-close to a $(n/2)$ -junta.*

Proof. Given A, b and the function f they define in the support of $\mathcal{D}_{\text{yes}}$, we let g be the function such that $g(x) = b(x_{\mathbf{C}})$. It is clear that g is a $(n/2)$ -junta and the distance between f and g , using Fact 2.1, is at most 0.1. $\square$

Next we show that with high probability, $\mathbf{f}_{\text{no}} \sim \mathcal{D}_{\text{no}}$ is far from a $(n/2)$ -junta.

LEMMA 5.2. *With probability at least $1 - o_n(1)$, $\mathbf{f}_{\text{no}} \sim \mathcal{D}_{\text{no}}$ is 0.2-far from any $(n/2)$ -junta.*

Proof. First of all, it follows from Chernoff bound and a union bound that with probability at least $1 - o_n(1)$, $\mathbf{A}$ and $\mathbf{b}$ satisfy the following condition:

1. For every $i \in \mathbf{C}$, there are at least $0.24 \cdot 2^{n/2}$ many strings $x \in \{0, 1\}^{\mathbf{C}}$ with $x_i = 0$ such that $\mathbf{b}(x) \neq \mathbf{b}(x^{(i)})$.

We assume below that A and b satisfy the above condition, and show that the function f in the support of $\mathcal{D}_{\text{no}}$ defined using A and b is 0.2-far from $(n/2)$ -juntas.

Let g be any $(n/2)$ -junta and $I \subset [n]$ of size $n/2$ be its influential variables. Given the condition above, we have that the number of bichromatic edges of f along each direction $i \in \mathbf{C}$ is at least

$$0.24 \cdot 2^{n/2} \cdot (1 - 0.1) \cdot 2^{n/2} > 0.2 \cdot 2^n.$$

So if $I \neq \mathbf{C}$, then g must be at least 0.2-far from f . On the other hand, if $I = \mathbf{C}$, then from the construction of f , g is at least 0.45-far from f . This finishes the proof of the lemma. $\square$

Let $\mathcal{A}$ be a non-adaptive deterministic algorithm that makes $q = 2^{0.01\sqrt{n}}$ queries. Let $\mathbf{A}$ and $\mathbf{b}$ be drawn as in the definition of $\mathcal{D}_{\text{yes}}$ and $\mathcal{D}_{\text{no}}$. Let **Bad** be the following event: there are two points x and y queried by $\mathcal{A}$ that satisfy $x_{\mathbf{C}} = y_{\mathbf{C}}$, $|x_{\mathbf{A}}| > a/2 + 0.05\sqrt{a}$ and $|y_{\mathbf{A}}| < a/2 - 0.05\sqrt{a}$. Following the proof of Lemma 4.4, $\mathcal{A}$ can potentially distinguish $\mathcal{D}_{\text{yes}}$ from $\mathcal{D}_{\text{no}}$ only when **Bad** occurs. The next lemma shows that **Bad** occurs with probability $o_n(1)$, from which Theorem 1.2 follows:

LEMMA 5.3. *The probability of the event **Bad** is $o_n(1)$.*

Proof. Let x and y be two points queried by $\mathcal{A}$. For $|x_{\mathbf{A}}| > a/2 + 0.05\sqrt{a}$ and $|y_{\mathbf{A}}| < a/2 - 0.05\sqrt{a}$ to hold, it must be the case that x and y have Hamming distance at least $0.1\sqrt{a}$. However, for any x and y with Hamming distance at least $0.1\sqrt{a} \geq 0.07\sqrt{n}$, the probability of $x_{\mathbf{C}} = y_{\mathbf{C}}$ is at most $2^{-0.06\sqrt{n}}$. The lemma then follows from a union bound over all pairs of points queried by $\mathcal{A}$. $\square$

The proof of Theorem 1.2 follows from the same steps as that of Theorem 1.1.

6 Discussion

Our results suggest several intriguing directions and possibilities for future work.

The Role of Adaptivity in Tolerant Testing. Recall that the best known upper bound for tolerant monotonicity (resp. unateness) testing is a $2^{\tilde{O}(\sqrt{n})}$ -query non-adaptive algorithm that goes through agnostic learners for monotone (resp. unate) functions [BT96, KKMS08, FKV17]. No non-trivial *adaptive* lower bounds are known for tolerant monotonicity or unateness testing, and no tolerant testing algorithms are known which employ adaptivity. For tolerant k -junta testing in the constant gap setting (i.e. when $\varepsilon := \varepsilon_2 - \varepsilon_1 = \Theta(1)$), recall that De, Mossel, and Neeman [DMN19] give a $2^{O(k)}$ -query non-adaptive algorithm. Subsequent work by Iyer, Tal, and Whitmeyer [ITW21] gives a $2^{O(\sqrt{k})}$ -query *adaptive* algorithm for the same task; very recently, Chen and Patel [CP23] gave a $k^{\Omega(1)}$ -query lower bound against adaptive tolerant junta testers.

The prior discussion highlights a gap in our understanding: Does adaptivity help for tolerant monotonicity, unateness, or junta testing? Proving improved lower bounds against adaptive algorithms, or alternatively designing efficient testers that make adaptive queries for these problems, is a natural next step for future work.

A Barrier to Improving the Monotonicity & Unateness Lower Bound. One might hope to improve the $2^{\Omega(n^{1/4})}$ -query lower bound against non-adaptive monotonicity and unateness testers from Section 4 by using another Boolean function on the control subcube instead of the Talagrand DNF. Tracing through the proof of Theorem 1.1 in Section 4, it is not too difficult to establish the following:

PROPOSITION 6.1. *Let $\varepsilon \in (0, 1)$. For $0 \leq \tau \leq \frac{1}{2}$, suppose there exist L disjoint sets $S_1, S_2, \dots, S_L \subseteq \{0, 1\}^n$ with $L := 0.01 \cdot 2^{n^{1-\tau}/\varepsilon}$ such that*

1. $2^{-n}|\mathcal{S}| \geq \Omega\left(\max\left\{\varepsilon, \frac{1}{\sqrt{n}}\right\}\right)$ where $\mathcal{S} := S_1 \sqcup S_2 \sqcup \dots \sqcup S_L$;

2. *If $x, y \in \mathcal{S}$ and $x \leq y$, then $x, y \in S_i$ for some $i \in [L]$; and*

3. *Define the function $\text{sens}_{\mathcal{S}}^- : \mathcal{S} \rightarrow [n]$ as*

$$\text{sens}_{\mathcal{S}}^-(x) := |\{i \in [n] : x_i = 1, x^{\oplus i} \notin \mathcal{S}\}|$$

where $x^{\oplus i} := (x_1, \dots, 1 - x_i, \dots, x_n)$; then we have

$$\mathbf{E}_{\mathbf{x} \sim \mathcal{S}} [\text{sens}_{\mathcal{S}}^-(\mathbf{x})] = \Omega(n^{1-\tau}).$$

Then any non-adaptive algorithm for $(\varepsilon_1, \varepsilon_2)$ -tolerant monotonicity or unateness testing must make $2^{n^{\sqrt{(1-\tau)/\varepsilon}}}$ queries, where $\varepsilon_1 = \Theta(\varepsilon)$, $\varepsilon_2 - \varepsilon_1 = \Theta(\varepsilon)$ as in Theorem 1.1.

To align Proposition 6.1 with the Talagrand DNF construction in Sections 3 and 4, note that the L disjoint sets $S_1, \dots, S_L$ correspond to the uniquely-satisfying assignments of each of the L terms of the Talagrand DNF. Item 2 ensures that the “yes” (resp. “no”) functions are indeed close (resp. far) from being monotone or unate, and Item 3 ensures indistinguishability of the distributions (cf. Lemma 4.5). Note that the function $\text{sens}_{\mathcal{S}}^-$ can be viewed as a “directed” variant of the standard notion of *sensitivity* of a subset of the Boolean hypercube (cf. Chapter 2 of [O’D14]).

Proposition 6.1 shows that if it were possible to obtain quantitatively stronger parameters with a variant of the Talagrand DNF, our approach would yield an improved lower bound for tolerant monotonicity or unateness testing. However, it turns out that no such improvement is possible:

PROPOSITION 6.2. *Suppose $\varepsilon = \Theta(1)$ and $\tau < 1/2$. Then there is no $\mathcal{S}$ as in Proposition 6.1.*

Proof. Suppose, for the sake of contradiction, that such an $\mathcal{S}$ exists. Consider the monotone Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ obtained by taking the upward closure of $\mathcal{S}$, i.e.

$$f(x) = \begin{cases} 1 & \text{there exists } y \in \mathcal{S} \text{ such that } x \geq y \\ 0 & \text{otherwise} \end{cases}.$$

Writing $\text{sens}_f(x)$ for the *sensitivity* of the function f at x (cf. Chapter 2 of [O’D14]), we then have from Item 3 of Proposition 6.1 that

$$\mathbf{E}_{\mathbf{x} \sim \{0,1\}^n} [\text{sens}_f(\mathbf{x})] \geq \Omega(n^{1-\tau}).$$

However, it is well known that the average sensitivity of a monotone Boolean function is at most $O(\sqrt{n})$ (cf. Theorem 2.33 of [O’D14]), resulting in a contradiction for $\tau < \frac{1}{2}$. $\square$

This suggests that quantitatively improving on Theorem 1.1 may require a substantially new construction.

Acknowledgements

X.C. is supported by NSF grants IIS-1838154, CCF-2106429, and CCF-2107187. A.D. is supported by NSF grants CCF-1910534 and CCF-2045128. Y.L. is supported by NSF grants IIS-1838154, CCF-2106429 and CCF-2107187. S.N. is supported by NSF grants IIS-1838154, CCF-2106429, CCF-2211238, CCF-1763970, and CCF-2107187. R.A.S. is supported by NSF grants IIS-1838154, CCF-2106429, and CCF-2211238. This work was partially completed while some of the authors were visiting the Simons Institute for the Theory of Computing at UC Berkeley.

References

- [BCGM12] Jop Briët, Sourav Chakraborty, David García-Soriano, and Arie Matsliah. Monotonicity testing and shortest-path routing on the cube. *Comb.*, 32(1):35–53, 2012.
- [BCP⁺20] Roksana Baleshazar, Deeparnab Chakrabarty, Ramesh Krishnan S. Pallavoor, Sofya Raskhodnikova, and C. Seshadhri. Optimal unateness testers for real-valued functions: Adaptivity helps. *Theory Comput.*, 16:1–36, 2020.
- [BCS18] Hadley Black, Deeparnab Chakrabarty, and C. Seshadhri. A $o(d) \cdot \text{polylog } n$ Monotonicity Tester for Boolean Functions over the Hypergrid $[n]^d$. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 2133–2151. SIAM, 2018.
- [BCS20] Hadley Black, Deeparnab Chakrabarty, and C. Seshadhri. Domain reduction for monotonicity testing: A $o(d)$ tester for boolean functions in d -dimensions. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1975–1994. SIAM, 2020.
- [BCS23] Hadley Black, Deeparnab Chakrabarty, and C. Seshadhri. Directed Isoperimetric Theorems for Boolean Functions on the Hypergrid and an $\tilde{O}(n\sqrt{d})$ Monotonicity Tester. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 233–241. ACM, 2023.
- [BGMdW13] Harry Buhrman, David García-Soriano, Arie Matsliah, and Ronald de Wolf. The non-adaptive query complexity of testing k -parities. *Chic. J. Theor. Comput. Sci.*, 2013, 2013.
- [BKKM23] Mark Braverman, Subhash Khot, Guy Kindler, and Dor Minzer. Improved monotonicity testers via hypercube embeddings. In *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPIcs*, pages 25:1–25:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [Bla08] Eric Blais. Improved bounds for testing juntas. In *Proc. RANDOM*, pages 317–330, 2008.
- [Bla09] Eric Blais. Testing juntas nearly optimally. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 151–158, 2009.
- [Bsh19] Nader H. Bshouty. Almost optimal distribution-free junta testing. In Amir Shpilka, editor, *34th Computational Complexity Conference, CCC 2019, July 18-20, 2019, New Brunswick, NJ, USA*, volume 137 of *LIPIcs*, pages 2:1–2:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [BT96] N. Bshouty and C. Tamon. On the Fourier spectrum of monotone functions. *Journal of the ACM*, 43(4):747–770, 1996.
- [CDST15] Xi Chen, Anindya De, Rocco A. Servedio, and Li-Yang Tan. Boolean Function Monotonicity Testing Requires (Almost) $n^{1/2}$ Non-adaptive Queries. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC 2015*, pages 519–528, 2015.
- [CG04] H. Chockler and D. Gutfreund. A lower bound for testing juntas. *Information Processing Letters*, 90(6):301–305, 2004.
- [CP23] Xi Chen and Shyamal Patel. New lower bounds for adaptive tolerant junta testing. *FOCS*, 2023. to appear.
- [CS13a] Deeparnab Chakrabarty and C. Seshadhri. A $o(n)$ monotonicity tester for boolean functions over the hypercube. In *Proceedings of the 45th ACM Symposium on Theory of Computing*, pages 411–418, 2013.
- [CS13b] Deeparnab Chakrabarty and C. Seshadhri. Optimal bounds for monotonicity and lipschitz testing over hypercubes and hypergrids. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 419–428. ACM, 2013.
- [CS16] Deeparnab Chakrabarty and C. Seshadhri. A $\tilde{O}(n)$ non-adaptive tester for unateness. *CoRR*, abs/1608.06980, 2016.
- [CST14] Xi Chen, Rocco A. Servedio, and Li-Yang Tan. New algorithms and lower bounds for testing monotonicity. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science*, pages 286–295, 2014.
- [CST⁺18] Xi Chen, Rocco A. Servedio, Li-Yang Tan, Erik Waingarten, and Jinyu Xie. Settling the query complexity of non-adaptive junta testing. *J. ACM*, 65(6):40:1–40:18, 2018.

- [CW19] Xi Chen and Erik Waingarten. Testing unateness nearly optimally. In *Proceedings of the 51th ACM Symposium on Theory of Computing*, pages 547–558, 2019.
- [CWX17a] Xi Chen, Erik Waingarten, and Jinyu Xie. Beyond Talagrand functions: new lower bounds for testing monotonicity and unateness. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 523–536, 2017.
- [CWX17b] Xi Chen, Erik Waingarten, and Jinyu Xie. Beyond talagrand functions: New lower bounds for testing monotonicity and unateness. *CoRR*, abs/1702.06997, 2017.
- [CWX17c] Xi Chen, Erik Waingarten, and Jinyu Xie. Boolean unateness testing with $\tilde{O}(n^{3/4})$ adaptive queries. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 868–879, 2017.
- [DGL⁺99] Yevgeniy Dodis, Oded Goldreich, Eric Lehman, Sofya Raskhodnikova, Dana Ron, and Alex Samorodnitsky. Improved testing algorithms for monotonicity. In *Proceedings of RANDOM*, pages 97–108, 1999.
- [DMN19] Anindya De, Elchanan Mossel, and Joe Neeman. Junta correlation is testable. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9–12, 2019*, pages 1549–1563. IEEE Computer Society, 2019.
- [FKR⁺04] E. Fischer, G. Kindler, D. Ron, S. Safra, and A. Samorodnitsky. Testing juntas. *J. Computer & System Sciences*, 68(4):753–787, 2004.
- [FKV17] Vitaly Feldman, Pravesh Kothari, and Jan Vondrak. Tight bounds on ℓ_1 approximation and learning of self-bounding functions. In Steve Hanneke and Lev Reyzin, editors, *Proceedings of the 28th International Conference on Algorithmic Learning Theory*, volume 76 of *Proceedings of Machine Learning Research*, pages 540–559. PMLR, 15–17 Oct 2017.
- [FLN⁺02] E. Fischer, E. Lehman, I. Newman, S. Raskhodnikova, R. Rubinfeld, and A. Samorodnitsky. Monotonicity testing over general poset domains. In *Proc. 34th Annual ACM Symposium on the Theory of Computing*, pages 474–483, 2002.
- [FR10] Shahar Fattal and Dana Ron. Approximating the distance to monotonicity in high dimensions. *ACM Trans. Algorithms*, 6(3):52:1–52:37, 2010.
- [GGL⁺00] Oded Goldreich, Shafi Goldwasser, Eric Lehman, Dana Ron, and Alex Samordinsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000.
- [HK07] S. Halevy and E. Kushilevitz. Distribution-Free Property Testing. *SIAM J. Comput.*, 37(4):1107–1138, 2007.
- [ITW21] Vishnu Iyer, Avishay Tal, and Michael Whitmeyer. Junta distance approximation with sub-exponential queries. In Valentine Kabanets, editor, *36th Computational Complexity Conference, CCC 2021, July 20–23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, volume 200 of *LIPIcs*, pages 24:1–24:38. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [Kan13] Daniel M. Kane. A monotone function given by a low-depth decision tree that is not an approximate junta. *Theory Comput.*, 9:587–592, 2013.
- [KKMS08] A. Kalai, A. Klivans, Y. Mansour, and R. Servedio. Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6):1777–1805, 2008.
- [KMS18] Subhash Khot, Dor Minzer, and Muli Safra. On monotonicity testing and boolean isoperimetric-type theorems. *SIAM J. Comput.*, 47(6):2238–2276, 2018.
- [KS16] Subhash Khot and Igor Shinkar. An $\tilde{o}(n)$ queries adaptive tester for unateness. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2016, September 7–9, 2016, Paris, France*, volume 60 of *LIPIcs*, pages 37:1–37:7. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [LCS⁺19] Zhengyang Liu, Xi Chen, Rocco A. Servedio, Ying Sheng, and Jinyu Xie. Distribution-free junta testing. *ACM Trans. Algorithms*, 15(1):1:1–1:23, 2019.
- [LW19] Amit Levi and Erik Waingarten. Lower bounds for tolerant junta and unateness testing via rejection sampling of graphs. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10–12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 52:1–52:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [O’D14] R. O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [OW07] Ryan O’Donnell and Karl Wimmer. Approximation by DNF: examples and counterexamples. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland, July 9–13, 2007, Proceedings*, volume 4596 of *Lecture Notes in Computer Science*, pages 195–206. Springer, 2007.
- [PRR06] Michal Parnas, Dana Ron, and Ronitt Rubinfeld. Tolerant property testing and distance approximation. *J. Comput. Syst. Sci.*, 72(6):1012–1042, 2006.
- [PRS02] M. Parnas, D. Ron, and A. Samorodnitsky. Testing Basic Boolean Formulae. *SIAM J. Disc. Math.*, 16:20–46, 2002.

- [PRW22] Ramesh Krishnan S. Pallavoor, Sofya Raskhodnikova, and Erik Waingarten. Approximating the distance to monotonicity of boolean functions. *Random Struct. Algorithms*, 60(2):233–260, 2022.
- [Sağ18] Mert Sağlam. Near log-convexity of measured heat in (discrete) time and consequences. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 967–978. IEEE Computer Society, 2018.
- [STW15] R.A. Servedio, L.-Y. Tan, and J. Wright. Adaptivity helps for testing juntas. In *Proceedings of the 30th IEEE Conference on Computational Complexity*, pages 264–279, 2015.
- [Tal96] M. Talagrand. How much are increasing sets positively correlated? *Combinatorica*, 16(2):243–258, 1996.
- [Yao77] A. Yao. Probabilistic computations: Towards a unified measure of complexity. In *Proc. Seventeenth Annual Symposium on Foundations of Computer Science (STOC)*, pages 222–227, 1977.