

Settling the Query Complexity of Non-Adaptive Junta Testing*

Xi Chen¹, Rocco A. Servedio², Li-Yang Tan³, Erik Waingarten⁴,
and Jinyu Xie⁵

- 1 Columbia University, New York, NY, USA
xichen@cs.columbia.edu
- 2 Columbia University, New York, NY, USA
rocco@cs.columbia.edu
- 3 Toyota Technological Institute, Chicago, IL, USA
liyang@cs.columbia.edu
- 4 Columbia University, New York, NY, USA
eaw@cs.columbia.edu
- 5 Columbia University, New York, NY, USA
jinyu@cs.columbia.edu

Abstract

We prove that any non-adaptive algorithm that tests whether an unknown Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is a k -junta or ϵ -far from every k -junta must make $\tilde{\Omega}(k^{3/2}/\epsilon)$ many queries for a wide range of parameters k and ϵ . Our result dramatically improves previous lower bounds from [12, 38], and is essentially optimal given Blais's non-adaptive junta tester from [7], which makes $\tilde{O}(k^{3/2})/\epsilon$ queries. Combined with the adaptive tester of [8] which makes $O(k \log k + k/\epsilon)$ queries, our result shows that adaptivity enables polynomial savings in query complexity for junta testing.

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity

Keywords and phrases property testing, juntas, query complexity

Digital Object Identifier 10.4230/LIPIcs.CCC.2017.26

1 Introduction

This paper is concerned with the power of adaptivity in property testing, specifically property testing of Boolean functions. At a high level, a property tester for Boolean functions is a randomized algorithm which, given black-box query access to an unknown and arbitrary Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, aims to distinguish between the case that f has some particular property of interest versus the case that f is far in Hamming distance from every Boolean function satisfying the property. The main goals in the study of property testing algorithms are to develop testers that make *as few queries* as possible, and to establish lower bounds matching these query-efficient algorithms. Property testing has by now been studied for many different types of Boolean functions, including linear functions and low-degree polynomials over $GF(2)$ [11, 2, 6], literals, conjunctions, s -term monotone and non-monotone

* X.C. and J.X. were supported by NSF awards CCF-1149257 and CCF-1423100.

R.A.S. was supported by NSF awards CCF-1420349 and CCF-1563155.

L.-Y.T. was supported by NSF award CCF-1563122.

E.W. was supported by the NSF Graduate Research Fellowship under Grant No. DGE-16-44869.



© Xi Chen, Rocco A. Servedio, Li-Yang Tan, Erik Waingarten,
and Jinyu Xie;
licensed under Creative Commons License CC-BY

32nd Computational Complexity Conference (CCC 2017).

Editor: Ryan O'Donnell; Article No. 26; pp. 26:1–26:19



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



DNFs [32, 18], monotone and unate functions [23, 20, 13, 16, 15, 27, 4, 28, 14, 3], various types of linear threshold functions [30, 31, 9], size- s decision trees and s -sparse $GF(2)$ polynomials and parities [18, 9, 10], functions with sparse or low-degree Fourier spectrum [24], and much more. See e.g. [33, 34, 22] for some fairly recent broad overviews of property testing research.

In this work we consider the property of being a k -junta, which is one of the earliest and most intensively studied properties in the Boolean function property testing literature. Recall that f is a k -junta if it has at most k relevant variables, i.e., there exist k distinct indices i_1, \dots, i_k and a k -variable function $g: \{0, 1\}^k \rightarrow \{0, 1\}$ such that $f(x) = g(x_{i_1}, \dots, x_{i_k})$ for all $x \in \{0, 1\}^n$. Given $k = k(n): \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon = \epsilon(n): \mathbb{N} \rightarrow \mathbb{R}_{>0}$, we say an algorithm which has black-box access to an unknown and arbitrary $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is an ϵ -tester or ϵ -testing algorithm for k -juntas if it accepts with probability at least $5/6$ when f is a $k(n)$ -junta and rejects with probability at least $5/6$ when f is $\epsilon(n)$ -far from all $k(n)$ -juntas (meaning that f disagrees with any $k(n)$ -junta g on at least $\epsilon(n) \cdot 2^n$ many inputs).

Property testers come in two flavors, adaptive and non-adaptive. An adaptive tester receives the value of f on its i -th query string before deciding on its $(i + 1)$ -st query string, while a non-adaptive tester selects all of its query strings before receiving the value of f on any of them. Note that non-adaptive testers can evaluate all of their queries in one parallel stage of execution, while this is in general not possible for adaptive testers. This means that if evaluating a query is very time-consuming, non-adaptive algorithms may sometimes be preferable to adaptive algorithms even if they require more queries. For this and other reasons, it is of interest to understand when, and to what extent, adaptive algorithms can use fewer queries than non-adaptive algorithms (see [36, 35] for examples of property testing problems where indeed adaptive algorithms are provably more query-efficient than non-adaptive ones).

The query complexity of adaptive junta testing algorithms is at this point well understood. In [17] Chockler and Gutfreund showed that even adaptive testers require $\Omega(k)$ queries to distinguish k -juntas from random functions on $k + 1$ variables, which are easily seen to be constant-far from k -juntas. Blais [8] gave an adaptive junta testing algorithm that uses only $O(k \log k + k/\epsilon)$ queries, which is optimal (for constant ϵ) up to a multiplicative factor of $O(\log k)$.

Prior to the current work, the picture was significantly less clear for non-adaptive junta testing. In the first work on junta testing, Fischer et al. [19] gave a non-adaptive tester that makes $O(k^2(\log k)^2/\epsilon)$ queries. This was improved by Blais [7] with a non-adaptive tester that uses only $O(k^{3/2}(\log k)^3/\epsilon)$ queries. On the lower bounds side, [7] also showed that for all $\epsilon \geq k/2^k$, any non-adaptive algorithm for ϵ -testing k -juntas must make $\Omega(k/(\epsilon \log(k/\epsilon)))$ queries. Buhrman et al. [12] gave an $\Omega(k \log k)$ lower bound (for constant ϵ) for non-adaptively testing whether a function f is a size- k parity; their argument also yields an $\Omega(k \log k)$ lower bound (for constant ϵ) for non-adaptively ϵ -testing k -juntas. More recently, [38] obtained a new lower bound for non-adaptive junta testing that is incomparable to both the [7] and the [12] lower bounds. They showed that for all $\epsilon: k^{-o_k(1)} \leq \epsilon \leq o_k(1)$, any non-adaptive ϵ -tester for k -juntas must make

$$\Omega\left(\frac{k \log k}{\epsilon^c \log(\log(k)/\epsilon^c)}\right)$$

many queries, where c is any absolute constant less than 1. For certain restricted values of ϵ such as $\epsilon = 1/\log k$, this lower bound is larger than the $O(k/\epsilon + k \log k)$ upper bound for [8]'s adaptive algorithm, so the [38] lower bound shows that in some restricted settings, adaptive junta testers can outperform non-adaptive ones. However, the difference in performance is quite small, at most a $o(\log k)$ factor. We further note that all of the lower bounds [7, 12, 38]

are of the form $\tilde{\Omega}(k)$ for constant ϵ , and hence rather far from the $\tilde{O}(k^{3/2})/\epsilon$ upper bound of [7].

1.1 Our results

The main result of the paper is the following theorem:

► **Theorem 1.** *Let $\alpha \in (0.5, 1)$ be an absolute constant. Let $k = k(n): \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon = \epsilon(n): \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be two functions that satisfy $k(n) \leq \alpha n$ and $2^{-n} \leq \epsilon(n) \leq 1/6$ for all sufficiently large n . Then any non-adaptive ϵ -tester for k -juntas must make $\tilde{\Omega}(k^{3/2}/\epsilon)$ many queries.*

Together with the $\tilde{O}(k^{3/2})/\epsilon$ non-adaptive upper bound from [7], Theorem 1 settles the query complexity of non-adaptive junta testing up to poly-logarithmic factors.

1.2 High-level overview of our approach

Our lower bound approach differs significantly from previous work. Buhrman et al. [12] leveraged the connection between communication complexity lower bounds and property testing lower bounds that was established in the work of [9] and applied an $\Omega(k \log k)$ lower bound on the one-way communication complexity of k -disjointness to establish their lower bound. Both [7] and [38] are based on edge-isoperimetry results for the Boolean hypercube (the edge-isoperimetric inequality of Harper [25], Bernstein [5], Lindsey [29], and Hart [26] in the case of [7], and a slight extension of a result of Frankl [21] in [38]). In contrast, our lower bound argument takes a very different approach; it consists of a sequence of careful reductions, and employs an *upper* bound on the total variation distance between two Binomial distributions (see Claim 15).

Below we provide a high level overview of the proof of the lower bound given by Theorem 1. First, it is not difficult to show that Theorem 1 is a consequence of the following more specific lower bound for the case where $k = \alpha n$:

► **Theorem 2.** *Let $\alpha \in (0.5, 1)$ be an absolute constant. Let $k = k(n): \mathbb{N} \rightarrow \mathbb{N}$ and $\epsilon = \epsilon(n): \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be two functions that satisfy $k(n) = \alpha n$ and $2^{-(2\alpha-1)n/2} \leq \epsilon(n) \leq 1/6$ for sufficiently large n . Then any non-adaptive ϵ -tester for k -juntas must make $\tilde{\Omega}(n^{3/2}/\epsilon)$ many queries.*

See Appendix A for the proof that Theorem 2 implies Theorem 1.

We now provide a sketch of how Theorem 2 is proved. It may be convenient for the reader, on the first reading, to consider $\alpha = 3/4$ and to think of ϵ as being a small constant such as 0.01.

Fix a sufficiently large n . Let $k = \alpha n$ and $\epsilon = \epsilon(n)$ with ϵ satisfying the condition in Theorem 2. We proceed by Yao's principle and prove lower bounds for deterministic non-adaptive algorithms which receive inputs drawn from one of two probability distributions, \mathcal{D}_{yes} and \mathcal{D}_{no} , over n -variable Boolean functions. The distributions \mathcal{D}_{yes} and \mathcal{D}_{no} are designed so that a Boolean function $\mathbf{f} \leftarrow \mathcal{D}_{\text{yes}}$ is a k -junta with probability $1 - o(1)$ and $\mathbf{f} \leftarrow \mathcal{D}_{\text{no}}$ is ϵ -far from every k -junta with probability $1 - o(1)$. In Section 2 we define \mathcal{D}_{yes} and \mathcal{D}_{no} , and establish the above properties. By Yao's principle, it then suffices to show that any q -query non-adaptive deterministic algorithm (i.e., any set of q queries) that succeeds in distinguishing them must have $q = \tilde{\Omega}(n^{3/2}/\epsilon)$.

This lower bound proof consists of two components:

1. A reduction from a simple algorithmic task called Set-Size-Set-Queries (SSSQ for short), which we discuss informally later in this subsection and we define formally in Section 3. This reduction implies that the non-adaptive deterministic query complexity of distinguishing \mathcal{D}_{yes} and \mathcal{D}_{no} is at least as large as that of SSSQ.
2. A lower bound of $\tilde{\Omega}(n^{3/2}/\epsilon)$ for the query complexity of SSSQ.

Having outlined the formal structure of our proof, let us give some intuition which may hopefully be helpful in motivating our construction and reduction. Our yes-functions and no-functions have very similar structure to each other, but are constructed with slightly different parameter settings. The first step in drawing a random function from \mathcal{D}_{yes} is choosing a uniform random subset \mathbf{M} of $\Theta(n)$ “addressing” variables from x_1, \dots, x_n . A random subset \mathbf{A} of the complementary variables $\bar{\mathbf{M}}$ is also selected, and for each assignment to the variables in \mathbf{M} (let us denote such an assignment by i), there is an independent random function h_i over a randomly selected subset \mathbf{S}_i of the variables in \mathbf{A} . A random function from \mathcal{D}_{no} is constructed in the same way, except that now the random subset \mathbf{A} is chosen to be slightly larger than in the yes-case. This disparity in the size of \mathbf{A} between the two cases causes random functions from \mathcal{D}_{yes} to almost always be k -juntas and random functions from \mathcal{D}_{no} to almost always be far from k -juntas.

An intuitive explanation of why this construction is amenable to a lower bound for non-adaptive algorithms is as follows. Intuitively, for an algorithm to determine that it is interacting with (say) a random no-function rather than a random yes-function, it must determine that the subset \mathbf{A} is larger than it should be in the yes-case. Since the set \mathbf{M} of $\Theta(n)$ many “addressing” variables is selected randomly, if a non-adaptive algorithm uses two query strings x, x' that differ in more than a few coordinates, it is very likely that they will correspond to two different random functions $h_i, h_{i'}$. Hence every pair of query strings x, x' that correspond to the same h_i can differ only in a few coordinates in $\bar{\mathbf{M}}$, with high probability, which significantly limits the power of a non-adaptive algorithm in distinguishing \mathcal{D}_{yes} and \mathcal{D}_{no} no matter which set of query strings it picks. This makes it possible for us to reduce from the SSSQ problem to the problem of distinguishing \mathcal{D}_{yes} and \mathcal{D}_{no} at the price of only a small quantitative cost in query complexity, see Section 4.

At a high level, the SSSQ task involves distinguishing whether or not a hidden set (corresponding to \mathbf{A}) is “large.” An algorithm for this task can only access certain random bits, whose biases are determined by the hidden set and whose exact distribution is inspired by the exact definition of the random functions h_i over the random subsets \mathbf{S}_i . Although SSSQ is an artificial problem, it is much easier to work with compared to the original problem of distinguishing \mathcal{D}_{yes} and \mathcal{D}_{no} . In particular, we give a reduction from an even simpler algorithmic task called Set-Size-Element-Queries (SSEQ for short) to SSSQ (see Section 5.1) and the query complexity lower bound for SSSQ follows directly from the lower bound for SSEQ presented in Section 5.2.

Let us give a high-level description of the SSEQ task to provide some intuition for how we prove a query lower bound on it. Roughly speaking, in this task an oracle holds an unknown and random subset \mathbf{A} of $[m]$ (here $m = \Theta(n)$) which is either “small” (size roughly $m/2$) or “large” (size roughly $m/2 + \Theta(\sqrt{n} \cdot \log n)$), and the task is to determine whether \mathbf{A} is small or large. The algorithm may repeatedly query the oracle by providing it, at the j -th query, with an element $i_j \in [m]$; if $i_j \notin \mathbf{A}$ then the oracle responds “0” with probability 1, and if $i_j \in \mathbf{A}$ then the oracle responds “1” with probability ϵ/\sqrt{n} and “0” otherwise. Intuitively, the only way for an algorithm to determine that the unknown set \mathbf{A} is (say) large, is to determine that the fraction of elements of $[m]$ that belong to \mathbf{A} is $1/2 + \Theta(\log n/\sqrt{n})$ rather than $1/2$; this in turn intuitively requires sampling $\Omega(n/\log^2 n)$ many random elements of

$[m]$ and for each one ascertaining with high confidence whether or not it belongs to \mathbf{A} . But the nature of the oracle access described above for SSEQ is such that for any given $i \in [m]$, at least $\Omega(\sqrt{n}/\epsilon)$ many repeated queries to the oracle on input i are required in order to reach even a modest level of confidence as to whether or not $i \in \mathbf{A}$. As alluded to earlier, the formal argument establishing our lower bound on the query complexity of SSEQ relies on an upper bound on the total variation distance between two Binomial distributions.

1.3 Organization and Notation

We start with the definitions of \mathcal{D}_{yes} and \mathcal{D}_{no} as well as proofs of their properties in Section 2. We then introduce SSSQ in Section 3, and give a reduction from SSSQ to the problem of distinguishing \mathcal{D}_{yes} and \mathcal{D}_{no} in Section 4. More formally, we show that any non-adaptive deterministic algorithm that distinguishes \mathcal{D}_{yes} and \mathcal{D}_{no} can be used to solve SSSQ with only an $O(\log n)$ factor loss in the query complexity. Finally, we prove in Section 5 a lower bound for the query complexity of SSSQ. Theorem 2 then follows by combining this lower bound with the reduction in Section 4.

We use boldfaced letters such as $\mathbf{f}, \mathbf{A}, \mathbf{S}$ to denote random variables. Given a string $x \in \{0, 1\}^n$ and $\ell \in [n]$, we write $x^{(\ell)}$ to denote the string obtained from x by flipping the ℓ -th coordinate. An edge along the ℓ th direction in $\{0, 1\}^n$ is a pair (x, y) of strings with $y = x^{(\ell)}$. We say an edge (x, y) is *bichromatic with respect to a function f* (or simply *f -bichromatic*) if $f(x) \neq f(y)$. Given $x \in \{0, 1\}^n$ and $S \subseteq [n]$, we use $x|_S \in \{0, 1\}^S$ to denote the projection of x on S .

2 The \mathcal{D}_{yes} and \mathcal{D}_{no} distributions

Let $\alpha \in (0.5, 1)$ be an absolute constant. Let n be a sufficiently large integer, with $k = \alpha n$, and let ϵ be the distance parameter that satisfies

$$2^{-(2\alpha-1)n/2} \leq \epsilon \leq 1/6. \quad (1)$$

In this section we describe a pair of probability distributions \mathcal{D}_{yes} and \mathcal{D}_{no} supported over Boolean functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$. We then show that $\mathbf{f} \leftarrow \mathcal{D}_{\text{yes}}$ is a k -junta with probability $1 - o(1)$, and that $\mathbf{f} \leftarrow \mathcal{D}_{\text{no}}$ is ϵ -far from being a k -junta with probability $1 - o(1)$.


We start with some parameters settings.

Define

$$\begin{aligned} \delta &\stackrel{\text{def}}{=} 1 - \alpha \in (0, 0.5), & p &\stackrel{\text{def}}{=} \frac{1}{2}, & q &\stackrel{\text{def}}{=} \frac{1}{2} + \frac{\log n}{\sqrt{n}}, \\ m &\stackrel{\text{def}}{=} 2\delta n + \delta\sqrt{n} \log n, & t &\stackrel{\text{def}}{=} n - m = (2\alpha - 1)n - \delta\sqrt{n} \log n, & N &\stackrel{\text{def}}{=} 2^t. \end{aligned}$$

A function $\mathbf{f} \leftarrow \mathcal{D}_{\text{yes}}$ is drawn according to the following randomized procedure:

1. Sample a random subset $\mathbf{M} \subset [n]$ of size t . Let $\mathbf{\Gamma} = \Gamma_{\mathbf{M}}: \{0, 1\}^n \rightarrow [N]$ be the function that maps $x \in \{0, 1\}^n$ to the integer encoded by $x|_{\mathbf{M}}$ in binary plus one. Note that $|\overline{\mathbf{M}}| = n - t = m$.
2. Sample an $\mathbf{A} \subseteq \overline{\mathbf{M}}$ by including each element of $\overline{\mathbf{M}}$ in \mathbf{A} independently with probability p .
3. Sample independently a sequence of N random subsets $\mathbf{S} = (\mathbf{S}_i: i \in [N])$ of \mathbf{A} as follows: for each $i \in [N]$, each element of \mathbf{A} is included in \mathbf{S}_i independently with probability ϵ/\sqrt{n} . Next we sample a sequence of N functions $\mathbf{H} = (\mathbf{h}_i: i \in [N])$, by letting $\mathbf{h}_i: \{0, 1\}^n \rightarrow \{0, 1\}$ be a random function over the coordinates in \mathbf{S}_i , i.e., we sample an unbiased bit $z_i(b)$ for each string $b \in \{0, 1\}^{\mathbf{S}_i}$ independently and set $\mathbf{h}_i(x) = z_i(x|_{\mathbf{S}_i})$.



■ **Figure 1** An example of how an input $x \in \{0, 1\}^n$ is evaluated by $f \sim \mathcal{D}_{\text{yes}}$ (or \mathcal{D}_{no}). The relevant variables of x are shaded gray. All variables in \mathbf{M} index x into h_2 , which is a random function over the variables \mathbf{S}_2 , which are sampled from \mathbf{A} by including each with probability ϵ/\sqrt{n} .

4. Finally, $f = f_{\mathbf{M}, \mathbf{A}, \mathbf{H}}: \{0, 1\}^n \rightarrow \{0, 1\}$ is defined using \mathbf{M} , \mathbf{A} and \mathbf{H} as follows:

$$f(x) = h_{\Gamma_{\mathbf{M}}(x)}(x), \quad \text{for each } x \in \{0, 1\}^n.$$

In words, an input x is assigned the value $f(x)$ as follows: according to the coordinates of x in the set \mathbf{M} (which intuitively should be thought of as unknown), one of the N functions h_i (each of which is, intuitively, a random function over an unknown subset \mathbf{S}_i of coordinates) is selected and evaluated on x 's coordinates in \mathbf{S}_i . For intuition, we note that both \mathbf{M} and $\overline{\mathbf{M}}$ will always be of size $\Theta(n)$, the size of \mathbf{A} will almost always be $\Theta(n)$, and for a given $i \in [N]$ the expected size of \mathbf{S}_i will typically be $\Theta(\epsilon\sqrt{n})$ (though the size of \mathbf{S}_i may not be as highly concentrated as the other sets when ϵ is tiny).

A function $f \leftarrow \mathcal{D}_{\text{no}}$ is generated using the same procedure except that \mathbf{A} is a random subset of $\overline{\mathbf{M}}$ drawn by including each element of $\overline{\mathbf{M}}$ in \mathbf{A} independently with probability q (instead of p). See Figure 1 for an example of how an input $x \in \{0, 1\}^n$ is evaluated by $f \sim \mathcal{D}_{\text{yes}}$ or \mathcal{D}_{no} .

2.1 Most functions drawn from \mathcal{D}_{yes} are k -juntas

We first prove that $f \leftarrow \mathcal{D}_{\text{yes}}$ is a k -junta with probability $1 - o(1)$.

► **Lemma 3.** *A function $f \leftarrow \mathcal{D}_{\text{yes}}$ is a k -junta with probability $1 - o(1)$.*

Proof. By the definition of \mathcal{D}_{yes} , all the relevant variables of $f \sim \mathcal{D}_{\text{yes}}$ belong to $\mathbf{M} \cup \mathbf{A}$. Note that $|\mathbf{M}| = t$. On the other hand, the expected size of \mathbf{A} is $\delta n + \delta\sqrt{n} \log n/2$. By a Chernoff bound,

$$|\mathbf{A}| \leq \delta n + \frac{\delta\sqrt{n} \log n}{2} + \frac{\delta\sqrt{n} \log n}{4} < \delta n + \delta\sqrt{n} \log n$$

with probability $1 - o(1)$. When this happens we have $|\mathbf{M} \cup \mathbf{A}| < \alpha n = k$. ◀

2.2 Most functions drawn from \mathcal{D}_{no} are ϵ -far from k -juntas

Next we prove that $f \leftarrow \mathcal{D}_{\text{no}}$ is ϵ -far from any k -junta with probability $1 - o(1)$. The details of the argument are somewhat technical so we start by giving some high-level intuition, which is relatively simple. Since $q = p + \log(n)/\sqrt{n}$, a typical outcome of \mathbf{A} drawn from \mathcal{D}_{no} is slightly larger than a typical outcome drawn from \mathcal{D}_{yes} , and this difference causes almost every outcome of $|\mathbf{M} \cup \mathbf{A}|$ in \mathcal{D}_{no} (with $\mathbf{M} \cup \mathbf{A}$ being the set of relevant variables

for $\mathbf{f} \leftarrow \mathcal{D}_{no}$) to be larger than k by at least $9\sqrt{n}$. As a result, the relevant variables of any k -junta must miss either (a) at least one variable from \mathbf{M} , or (b) at least $9\sqrt{n}$ variables from \mathbf{A} . Missing even a single variable from \mathbf{M} causes the k -junta to be far from \mathbf{f} (this is made precise in Claim 6 below). On the other hand, missing $9\sqrt{n}$ variables from \mathbf{A} means that with probability at least $\Omega(\epsilon)$, at least one variable is missing from a typical \mathbf{S}_i (recall that these are random (ϵ/\sqrt{n}) -dense subsets of \mathbf{A}). Because \mathbf{h}_i is a random function over the variables in \mathbf{S}_i , missing even a single variable would lead to a constant fraction of error when \mathbf{h}_i is the function determining the output of \mathbf{f} .

► **Lemma 4.** *A function $\mathbf{f} \leftarrow \mathcal{D}_{no}$ is ϵ -far from being a k -junta with probability $1 - o(1)$.*

Proof. Fix any subset $M \subset [n]$ of size t , and we consider $\mathbf{f} = \mathbf{f}_{M,\mathbf{A},\mathbf{H}}$ where \mathbf{A} and \mathbf{H} are sampled according to the procedure for \mathcal{D}_{no} . With probability $1 - o(1)$ over the choice of \mathbf{A} , we have

$$|\mathbf{A}| \geq qm - \frac{\delta\sqrt{n} \log n}{2} \geq \delta n + 2\delta\sqrt{n} \log n \quad \text{and} \quad |\mathbf{M} \cup \mathbf{A}| \geq k + \delta\sqrt{n} \log n. \quad (2)$$

We assume this is the case for the rest of the proof and fix any such set $A \subset \overline{M}$. It suffices to show that $\mathbf{f} = \mathbf{f}_{M,A,\mathbf{H}}$ is ϵ -far from k -juntas with probability $1 - o(1)$, where \mathbf{H} is sampled according to the rest (steps 3 and 4) of the procedure for \mathcal{D}_{no} (by sampling \mathbf{S}_i from A and then \mathbf{h}_i over \mathbf{S}_i).

The plan for the rest of the proof is the following. For each $V \subset M \cup A$ of size $9\sqrt{n}$, we use \mathbf{E}_V to denote the size of the *maximum* set of vertex-disjoint, \mathbf{f} -bichromatic edges along directions in V only. We will prove the following claim:

► **Claim 5.** *For each $V \subset M \cup A$ of size $9\sqrt{n}$, we have $\mathbf{E}_V \geq \epsilon 2^n$ with probability $1 - \exp(-2^{\Omega(n)})$.*

Note that when $\mathbf{E}_V \geq \epsilon 2^n$, we have $\text{dist}(\mathbf{f}, g) \geq \epsilon$ for every function g that does not depend on any variable in V . This is because, for every \mathbf{f} -bichromatic edge $(x, x^{(\ell)})$ along a coordinate $\ell \in V$, we must have $\mathbf{f}(x) \neq \mathbf{f}(x^{(\ell)})$ since the edge is bichromatic but $g(x) = g(x^{(\ell)})$ as g does not depend on the ℓ th variable. As a result, \mathbf{f} must disagree with g on at least $\epsilon 2^n$ many points.

Assuming Claim 5 for now, we can apply a union bound over all

$$\binom{|M \cup A|}{9\sqrt{n}} \leq \binom{n}{9\sqrt{n}} \leq 2^{O(\sqrt{n} \log n)}$$

possible choices of $V \subset M \cup A$ to conclude that with probability $1 - o(1)$, $\mathbf{f} = \mathbf{f}_{M,A,\mathbf{H}}$ is ϵ -far from all functions that do not depend on at least $9\sqrt{n}$ variables in $M \cup A$. By (2), this set includes all k -juntas. This concludes the proof of the Lemma 4 modulo the proof of Claim 5. ◀

In the rest of the section, we prove Claim 5 for a fixed subset $V \subset M \cup A$ of size $9\sqrt{n}$. We start with the simpler case when $V \cap M$ is nonempty.

► **Claim 6.** *If $V \cap M \neq \emptyset$, then we have $\mathbf{E}_V \geq 2^n/5$ with probability $1 - \exp(-2^{\Omega(n)})$.*

Proof. Fix an $\ell \in V \cap M$; we will argue that with probability $1 - \exp(-2^{\Omega(n)})$ there are at least $2^n/5$ \mathbf{f} -bichromatic edges along direction ℓ . This suffices since such edges are clearly vertex-disjoint.

Observe that since $\ell \in M$, every $x \in \{0, 1\}^n$ has $\Gamma(x) \neq \Gamma(x^{(\ell)})$. For each $b \in \{0, 1\}^M$, let X_b be the set of $x \in \{0, 1\}^n$ with $x|_S = b$. We partition $\{0, 1\}^n$ into 2^{t-1} pairs X_b and

$X_{b^{(\ell)}}$, where b ranges over the 2^{t-1} strings in $\{0,1\}^M$ with $b_\ell = 0$. For each such pair, we use \mathbf{D}_b to denote the number of \mathbf{f} -bichromatic edges between X_b and $X_{b^{(\ell)}}$. We are interested in lower bounding $\sum_b \mathbf{D}_b$.

We will apply Hoeffding's inequality. For this purpose we note that the \mathbf{D}_b 's are independent (since they depend on distinct \mathbf{h}_i 's), always lie between 0 and 2^m , and each one has expectation 2^{m-1} . The latter is because each edge $(x, x^{(\ell)})$ has $\mathbf{f}(x)$ and $\mathbf{f}(x^{(\ell)})$ drawn as two independent random bits, which is the case since $\Gamma(x) \neq \Gamma(x^{(\ell)})$. Thus, the expectation of $\sum_b \mathbf{D}_b$ is 2^{n-2} . By Hoeffding's inequality, we have

$$\Pr \left[\left| \sum_b \mathbf{D}_b - 2^{n-2} \right| \geq \frac{2^n}{20} \right] \leq 2 \cdot \exp \left(-\frac{2(2^n/20)^2}{2^{t-1} \cdot 2^{2m}} \right) = \exp \left(-2^{\Omega(n)} \right)$$

since $t = \Omega(n)$. This finishes the proof of the claim. \blacktriangleleft

Now we may assume that $V \subset A$ (and $|V| = 9\sqrt{n}$). We use \mathbf{I} to denote the set of $i \in [N]$ such that $\mathbf{S}_i \cap V \neq \emptyset$. The following claim shows that \mathbf{I} is large with extremely high probability:

► Claim 7. *We have $|\mathbf{I}| \geq 4.4\epsilon N$ with probability at least $1 - \exp(-2^{\Omega(n)})$ over the choice of \mathbf{S} .*

Proof. For each $i \in [N]$ we have (using $1 - x \leq e^{-x}$ for all x and $1 - x/2 \geq e^{-x}$ for $x \in [0, 1.5]$):

$$\Pr [i \in \mathbf{I}] = 1 - \left(1 - \frac{\epsilon}{\sqrt{n}} \right)^{9\sqrt{n}} \geq 1 - e^{-9\epsilon} \geq 4.5\epsilon,$$

since ϵ/\sqrt{n} is the probability of each element of A being included in \mathbf{S}_i and $\epsilon \leq 1/6$ so $9\epsilon \leq 1.5$.

Using $\epsilon \geq 2^{-(2\alpha-1)n/2}$ from (1), we have $\mathbf{E}[|\mathbf{I}|] \geq 4.5\epsilon N = 2^{\Omega(n)}$. Since the \mathbf{S}_i 's are independent, a Chernoff bound implies that $|\mathbf{I}| \geq 4.4\epsilon N$ with probability $1 - \exp(-2^{\Omega(n)})$. \blacktriangleleft

By Claim 7, we fix S_1, \dots, S_N to be any sequence of subsets of A that satisfy $|I| \geq 4.4\epsilon N$ in the rest of the proof, and it suffices to show that over the random choices of $\mathbf{h}_1, \dots, \mathbf{h}_N$ (where each \mathbf{h}_i is chosen to be a random function over S_i), $\mathbf{E}_V \geq \epsilon 2^n$ with probability at least $1 - \exp(-2^{\Omega(n)})$.

To this end we use $\rho(i)$ for each $i \in I$ to denote the first coordinate of S_i in V , and Z_i to denote the set of $x \in \{0,1\}^n$ with $\Gamma(x) = i$. Note that the Z_i 's are disjoint. We further partition each Z_i into disjoint $Z_{i,b}$, $b \in \{0,1\}^{S_i}$, with $x \in Z_{i,b}$ iff $x \in Z_i$ and $x|_{S_i} = b$. For each $i \in I$ and $b \in \{0,1\}^{S_i}$ with $b_{\rho(i)} = 0$, we use $\mathbf{D}_{i,b}$ to denote the number of \mathbf{f} -bichromatic edges between $Z_{i,b}$ and $Z_{i,b^{(\rho(i))}}$ along the $\rho(i)$ th direction. It is clear that such edges, over all i and b , are vertex-disjoint and thus,

$$\mathbf{E}_V \geq \sum_{i \in I} \sum_{\substack{b \in \{0,1\}^{S_i} \\ b_{\rho(i)} = 0}} \mathbf{D}_{i,b}. \quad (3)$$

We will apply Hoeffding's inequality. Note that $\mathbf{D}_{i,b}$ is $2^{m-|S_i|}$ with probability $1/2$, and 0 with probability $1/2$. Thus, the expectation of the RHS of (3) is

$$\sum_{i \in I} 2^{|S_i|-1} \cdot 2^{m-|S_i|-1} = |I| \cdot 2^{m-2} \geq 1.1\epsilon 2^n,$$

using $|I| \geq 4.4\epsilon N$. Since all the $\mathbf{D}_{i,b}$'s are independent, by Hoeffding's inequality we have

$$\begin{aligned} \Pr \left[\left| \text{RHS of (3)} - |I| \cdot 2^{m-2} \right| \geq 0.01 |I| \cdot 2^{m-2} \right] &\leq 2 \cdot \exp \left(- \frac{2(0.01 |I| \cdot 2^{m-2})^2}{\sum_{i \in I} 2^{|S_i|-1} \cdot 2^{2(m-|S_i|)}} \right) \\ &\leq \exp \left(-2^{\Omega(n)} \right), \end{aligned}$$

since $|I| \geq \Omega(\epsilon N) = 2^{\Omega(n)}$. When this does not happen, we have $\mathbf{E}_V \geq 0.99 \cdot |I| \cdot 2^{m-2} > \epsilon 2^n$.

3 The Set-Size-Set-Queries (SSSQ) Problem

We first introduce the Set-Size-Set-Queries (SSSQ for short) problem, which is an artificial problem that we use as a bridge to prove Theorem 2. We use the same parameters p, q and m from the definition of \mathcal{D}_{yes} and \mathcal{D}_{no} , with n being sufficiently large (so $m = \Omega(n)$ is sufficiently large as well).

We start by defining \mathcal{A}_{yes} and \mathcal{A}_{no} , two distributions over subsets of $[m]$: $\mathbf{A} \sim \mathcal{A}_{\text{yes}}$ is drawn by independently including each element of $[m]$ with probability p and $\mathbf{A} \sim \mathcal{A}_{\text{no}}$ is drawn by independently including each element with probability q . In SSSQ, the algorithm needs to determine whether an unknown $A \subseteq [m]$ is drawn from \mathcal{A}_{yes} or \mathcal{A}_{no} . (For intuition, to see that this task is reasonable, we observe here that a straightforward Chernoff bound shows that almost every outcome of $\mathbf{A} \sim \mathcal{A}_{\text{yes}}$ is larger than almost every outcome of $\mathbf{A} \sim \mathcal{A}_{\text{no}}$ by $\Omega(\sqrt{n} \log n)$.)

Let A be a subset of $[m]$ which is hidden in an oracle. An algorithm accesses A (in order to tell whether it is drawn from \mathcal{A}_{yes} or \mathcal{A}_{no}) by interacting with the oracle in the following way: each time it calls the oracle, it does so by sending a subset of $[m]$ to the oracle. The oracle responds as follows: for each j in the subset, it returns a bit that is 0 if $j \notin A$, and is 1 with probability ϵ/\sqrt{n} and 0 with probability $1 - \epsilon/\sqrt{n}$ if $j \in A$. The cost of such an oracle call is the size of the subset provided to the oracle.

More formally, a deterministic and non-adaptive algorithm $\text{ALG} = (g, T)$ for SSSQ accesses the set A hidden in the oracle by submitting a list of queries $T = (T_1, \dots, T_d)$, for some $d \geq 1$, where each $T_i \subseteq [m]$ is a set. (Thus, we call each T_i a *set query*, as part of the name SSSQ.)

- Given T , the oracle returns a list of random vectors $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_d)$, where $\mathbf{v}_i \in \{0, 1\}^{T_i}$ and each bit $\mathbf{v}_{i,j}$ is independently distributed as follows: if $j \notin A$ then $\mathbf{v}_{i,j} = 0$, and if $j \in A$ then

$$\mathbf{v}_{i,j} = \begin{cases} 1 & \text{with probability } \epsilon/\sqrt{n} \\ 0 & \text{with probability } 1 - (\epsilon/\sqrt{n}). \end{cases} \quad (4)$$

Note that the random vectors in \mathbf{v} depend on both T and A .

- Given $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_d)$, ALG returns (deterministically) the value of $g(\mathbf{v}) \in \{\text{"yes"}, \text{"no"}\}$.

The performance of $\text{ALG} = (g, T)$ is measured by its *query complexity* and its *advantage*.

- The query complexity of ALG is defined as $\sum_{i=1}^d |T_i|$, the total size of all the set queries. On the other hand, the advantage of ALG is defined as

$$\Pr_{\mathbf{A} \sim \mathcal{A}_{\text{yes}}} [\text{ALG}(\mathbf{A}) = \text{"yes"}] - \Pr_{\mathbf{A} \sim \mathcal{A}_{\text{no}}} [\text{ALG}(\mathbf{A}) = \text{"yes"}].$$

► **Remark 8.** In the definition above, g is a deterministic map from all possible sequences of vectors returned by the oracle to “yes” or “no.” Considering only deterministic as opposed to randomized g is without loss of generality since given any query sequence T , the highest possible advantage can always be achieved by a deterministic map g .

We prove the following lower bound for any deterministic, non-adaptive ALG in Section 5.

► **Lemma 9.** *Any deterministic, non-adaptive ALG for SSSQ with advantage at least $2/3$ satisfies*

$$\sum_{i=1}^d |T_i| \geq \frac{n^{3/2}}{\epsilon \cdot \log^3 n \cdot \log^2(n/\epsilon)}.$$

4 Reducing from SSSQ to distinguishing \mathcal{D}_{yes} and \mathcal{D}_{no}

In this section we reduce from SSSQ to the problem of distinguishing the pair of distributions \mathcal{D}_{yes} and \mathcal{D}_{no} . More precisely, let $\text{ALG}^* = (h, X)$ denote a deterministic and nonadaptive algorithm that makes $q \leq (n/\epsilon)^2$ string queries¹ $X = (x_1, \dots, x_q)$ to a hidden function f drawn from either \mathcal{D}_{yes} or \mathcal{D}_{no} , applies the (deterministic) map h to return $h(f(x_1), \dots, f(x_q)) \in \{\text{“yes”}, \text{“no”}\}$, and satisfies

$$\Pr_{f \sim \mathcal{D}_{\text{yes}}} [\text{ALG}^*(f) = \text{“yes”}] - \Pr_{f \sim \mathcal{D}_{\text{no}}} [\text{ALG}^*(f) = \text{“yes”}] \geq 3/4. \quad (5)$$

We show how to define from $\text{ALG}^* = (h, X)$ an algorithm $\text{ALG} = (g, T)$ for the problem SSSQ with query complexity at most $\tau \cdot q$ and advantage $2/3$, where $\tau = c_\alpha \cdot 5 \log(n/\epsilon)$ and

$$c_\alpha = -\frac{1}{\log(1.5 - \alpha)} > 0 \quad \text{with} \quad (1.5 - \alpha)^{c_\alpha} = 1/2$$

is a constant that depends on α . Given this reduction it follows from Lemma 9 that $q \geq \tilde{\Omega}(n^{3/2}/\epsilon)$. This finishes the proof of Theorem 2.

We start with some notation. Recall that in both \mathcal{D}_{yes} and \mathcal{D}_{no} , \mathbf{M} is a subset of $[n]$ of size t drawn uniformly at random. For a fixed M of size t , we use $\mathcal{E}_{\text{yes}}(M)$ to denote the distribution of \mathbf{A} and \mathbf{H} sampled in the randomized procedure for \mathcal{D}_{yes} , conditioning on $\mathbf{M} = M$. We define $\mathcal{E}_{\text{no}}(M)$ similarly. Then conditioning on $\mathbf{M} = M$, $f \sim \mathcal{D}_{\text{yes}}$ is distributed as $f_{M, \mathbf{A}, \mathbf{H}}$ with $(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{yes}}(M)$ and $f \sim \mathcal{D}_{\text{no}}$ is distributed as $f_{M, \mathbf{A}, \mathbf{H}}$ with $(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{no}}(M)$. This allows us to rewrite (5) as

$$\frac{1}{\binom{n}{t}} \cdot \sum_{M: |M|=t} \left(\Pr_{(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{yes}}(M)} [\text{ALG}^*(f_{M, \mathbf{A}, \mathbf{H}}) = \text{“yes”}] - \Pr_{(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{no}}(M)} [\text{ALG}^*(f_{M, \mathbf{A}, \mathbf{H}}) = \text{“yes”}] \right) \geq \frac{3}{4}.$$

We say $M \subset [n]$ is *good* if any two queries x_i and x_j in X with Hamming distance $\|x_i - x_j\|_1 \geq \tau$ have different projections on M , i.e., $(x_i)_{|M} \neq (x_j)_{|M}$. We prove below that most M 's are good.

► **Claim 10.** $\Pr_{\mathbf{M}} [\mathbf{M} \text{ is not good}] = o(1)$.

¹ Any algorithm that makes more than this many queries already fits the $\tilde{\Omega}(n^{3/2}/\epsilon)$ lower bound we aim for.

Proof. For each pair of strings x_i and x_j in X with Hamming distance at least τ , the probability of them having the same projection on \mathbf{M} (drawn uniformly from all size- t subsets) is at most

$$\begin{aligned} \frac{\binom{n-\tau}{t}}{\binom{n}{t}} &= \frac{(n-\tau-t+1)\cdots(n-t)}{(n-\tau+1)\cdots n} \leq \left(1 - \frac{t}{n}\right)^\tau \leq (2(1-\alpha) + o(1))^\tau < (1.5-\alpha)^\tau \\ &\leq O\left(\frac{\epsilon}{n}\right)^5, \end{aligned}$$

by our choices of c_α and τ . The claim follows by a union bound over at most $q^2 \leq (n/\epsilon)^4$ pairs. \blacktriangleleft

We can split the sum (5) into two sums: the sum over good M and the sum over bad M . By Claim 10 the contribution from the bad M is at most $o(1)$, and thus we have that

$$\frac{1}{\binom{n}{t}} \cdot \sum_{\text{good } M} \left(\Pr_{(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{yes}}(M)} [\text{ALG}^*(\mathbf{f}_{M, \mathbf{A}, \mathbf{H}}) = \text{“yes”}] - \Pr_{(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{no}}(M)} [\text{ALG}^*(\mathbf{f}_{M, \mathbf{A}, \mathbf{H}}) = \text{“yes”}] \right)$$

is at least $3/4 - o(1)$. Thus, there must exist a good set $M \subset [n]$ of size t with

$$\Pr_{(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{yes}}(M)} [\text{ALG}^*(\mathbf{f}_{M, \mathbf{A}, \mathbf{H}}) = \text{“yes”}] - \Pr_{(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{no}}(M)} [\text{ALG}^*(\mathbf{f}_{M, \mathbf{A}, \mathbf{H}}) = \text{“yes”}] \geq 2/3. \quad (6)$$

Fix such a good M . We use $\text{ALG}^* = (h, X)$ and M to define an algorithm $\text{ALG} = (g, T)$ for SSSQ as follows (note that the algorithm ALG below actually works over the universe \overline{M} (of size m) instead of $[m]$ as in the original definition of SSSQ but this can be handled by picking any bijection between \overline{M} and $[m]$; accordingly $\mathbf{A} \sim \mathcal{A}_{\text{yes}}$ is drawn by including each element of \overline{M} with probability p and $\mathbf{A} \sim \mathcal{A}_{\text{no}}$ is drawn by including each element of \overline{M} with probability q). We start with T :

1. First we use M to define an equivalence relation \sim over the query set X , where $x_i \sim x_j$ if $(x_i)_M = (x_j)_M$. Let X_1, \dots, X_d , $d \geq 1$, denote the equivalence classes of X , and let us write $\rho(\ell)$ for each $\ell \in [d]$ to denote the value $\Gamma(x) \in [N]$ that is shared by all strings $x \in X_\ell$.
2. Next we define a sequence of subsets of \overline{M} , $T = (T_1, \dots, T_d)$, as the set queries of ALG , where

$$T_\ell = \{i \in \overline{M} : \exists x, y \in X_\ell \text{ such that } x_i \neq y_i\}. \quad (7)$$

To upper bound $|T_\ell|$, fixing an arbitrary string $x \in X_\ell$ and recalling that M is good, we have that

$$|T_\ell| \leq \sum_{y \in X_\ell} \|x - y\|_1 \leq \sum_{y \in X_\ell} \tau = \tau \cdot |X_\ell|.$$

As a result, the query complexity of ALG (using T as its set queries) is at most

$$\sum_{\ell=1}^d |T_\ell| \leq \tau \cdot \sum_{\ell=1}^d |X_\ell| \leq \tau \cdot q.$$

It remains to define h and then prove that the advantage of $\text{ALG} = (g, T)$ for SSSQ is at least $2/3$. Indeed the g that we define is a randomized map and we describe it as a randomized procedure below (by Remark 8 one can extract from g a deterministic map that achieves the same advantage):

26:12 Settling the Query Complexity of Non-Adaptive Junta Testing

1. Given $v_1, \dots, v_d, v_\ell \in \{0, 1\}^{T_\ell}$, as the strings returned by the oracle upon being given T , let

$$R_\ell = \{j \in T_\ell : v_{\ell,j} = 1\}. \quad (8)$$

For each $\ell \in [d]$, the procedure draws a random function $\mathbf{f}_\ell : \{0, 1\}^{R_\ell} \rightarrow \{0, 1\}$, by flipping $2^{|R_\ell|}$ many independent and unbiased random bits.

2. Next for each query $x \in X_\ell$, $\ell \in [d]$, we feed $\mathbf{f}_\ell(x|_{R_\ell})$ to h as the bit that the oracle returns upon the query x . Finally the procedure returns the result (“yes” or “no”) that h returns.

In the rest of the proof we show that the advantage of $\text{ALG} = (g, T)$ is exactly the same as the LHS of (6) and thus, is at least $2/3$.

For convenience, we use \mathcal{V}_{yes} to denote the distribution of responses $\mathbf{v} = (v_1, \dots, v_d)$ to T when $\mathbf{A} \sim \mathcal{A}_{\text{yes}}$, and \mathcal{V}_{no} to denote the distribution when $\mathbf{A} \sim \mathcal{A}_{\text{no}}$. Then the advantage of ALG is

$$\Pr_{\mathbf{v} \sim \mathcal{V}_{\text{yes}}} [g(\mathbf{v}) = \text{“yes”}] - \Pr_{\mathbf{v} \sim \mathcal{V}_{\text{no}}} [g(\mathbf{v}) = \text{“yes”}].$$

It suffices to show that

$$\Pr_{\mathbf{v} \sim \mathcal{V}_{\text{yes}}} [g(\mathbf{v}) = \text{“yes”}] = \Pr_{(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{yes}}(M)} [\text{ALG}^*(\mathbf{f}_{M, \mathbf{A}, \mathbf{H}}) = \text{“yes”}] \quad \text{and} \quad (9)$$

$$\Pr_{\mathbf{v} \sim \mathcal{V}_{\text{no}}} [g(\mathbf{v}) = \text{“yes”}] = \Pr_{(\mathbf{A}, \mathbf{H}) \sim \mathcal{E}_{\text{no}}(M)} [\text{ALG}^*(\mathbf{f}_{M, \mathbf{A}, \mathbf{H}}) = \text{“yes”}]. \quad (10)$$

We show (9); the proof of (10) is similar. From the definition of \mathcal{V}_{yes} and $\mathcal{E}_{\text{yes}}(M)$ the distribution of $(\mathbf{R}_\ell : \ell \in [d])$ derived from $\mathbf{v} \sim \mathcal{V}_{\text{yes}}$ using (8) is the same as the distribution of $(\mathbf{S}_{\rho(\ell)} \cap T_\ell : \ell \in [d])$: both are sampled by first drawing a random subset \mathbf{A} of \overline{M} and then drawing a random subset of $\mathbf{A} \cap T_\ell$ independently by including each element of $\mathbf{A} \cap T_\ell$ with the same probability ϵ/\sqrt{n} (recall in particular equation (4) and step 3 of the randomized procedure specifying \mathcal{D}_{yes} in Section 2). Since $\mathbf{f}_{M, \mathbf{A}, \mathbf{H}}(x)$ for $x \in X_\ell$ is determined by a random Boolean function $\mathbf{h}_{\rho(\ell)}$ from $\{0, 1\}^{\mathbf{S}_{\rho(\ell)}}$ to $\{0, 1\}$, and since all the queries in X_ℓ only differ by coordinates in T_ℓ , the distribution of the q bits that g feeds to h when $\mathbf{v} \sim \mathcal{V}_{\text{yes}}$ is the same as the distribution of $(\mathbf{f}(x) : x \in X)$ when $\mathbf{f} \sim \mathcal{E}_{\text{yes}}(M)$. This finishes the proof of (9), and concludes our reduction argument.

5 A lower bound on the non-adaptive query complexity of SSSQ

We will prove Lemma 9 by first giving a reduction from an even simpler algorithmic task, which we describe next in Section 5.1. We will then prove a lower bound for the simpler task in Section 5.2.

5.1 Set-Size-Element-Queries (SSEQ)

Recall the parameters m, p, q and ϵ and the two distributions \mathcal{A}_{yes} and \mathcal{A}_{no} used in the definition of problem SSSQ. We now introduce a simpler algorithmic task called the Set-Size-Element-Queries (SSEQ) problem using the same parameters and distributions.

Let A be a subset of $[m]$ hidden in an oracle. An algorithm accesses the oracle to tell whether it is drawn from \mathcal{A}_{yes} or \mathcal{A}_{no} . The difference between SSSQ and SSEQ is the way an algorithm accesses A . In SSEQ, an algorithm $\text{ALG}' = (h, \ell)$ submits a vector $\ell = (\ell_1, \dots, \ell_m)$ of nonnegative integers.

- On receiving ℓ , the oracle returns a random response vector $\mathbf{b} \in \{0, 1\}^m$, where each entry b_i is distributed independently as follows: if $i \notin A$ then $b_i = 0$, and if $i \in A$ then

$$b_i = \begin{cases} 1 & \text{with probability } \lambda(\ell_i) \\ 0 & \text{with probability } 1 - \lambda(\ell_i) \end{cases}, \quad \text{where } \lambda(\ell_i) = 1 - \left(1 - \frac{\epsilon}{\sqrt{n}}\right)^{\ell_i}.$$

Equivalently, for each $i \in A$, the oracle independently flips ℓ_i coins, each of which is 1 with probability ϵ/\sqrt{n} , and at the end returns $b_i = 1$ to the algorithm if and only if at least one of the coins is 1. Thus, we refer to each ℓ_i as ℓ_i *element-queries* for the i th element.

- After receiving the vector \mathbf{b} from the oracle, ALG' returns the value $h(\mathbf{b}) \in \{\text{“yes”}, \text{“no”}\}$. Here h is a deterministic map from $\{0, 1\}^m$ to $\{\text{“yes”}, \text{“no”}\}$.

Similar to before, the performance of ALG' is measured by its query complexity and its advantage:

- The query complexity of $\text{ALG}' = (h, \ell)$ is defined as $\|\ell\|_1 = \sum_{i=1}^m \ell_i$. For its advantage, we let \mathcal{B}_{yes} denote the distribution of response vectors \mathbf{b} to query ℓ when $\mathbf{A} \sim \mathcal{A}_{\text{yes}}$, and \mathcal{B}_{no} denote the distribution when $\mathbf{A} \sim \mathcal{D}_{\text{no}}$. The advantage of $\text{ALG}' = (h, \ell)$ is then defined as

$$\Pr_{\mathbf{b} \sim \mathcal{B}_{\text{yes}}} [h(\mathbf{b}) = \text{“yes”}] - \Pr_{\mathbf{b} \sim \mathcal{B}_{\text{no}}} [h(\mathbf{b}) = \text{“yes”}].$$

► **Remark 11.** It is worth pointing out (we will use it later) that the highest possible advantage over all deterministic maps h is a monotonically non-decreasing function of the coordinates of ℓ . To see this, let A be the underlying set and let ℓ and ℓ' be two vectors with $\ell_i \leq \ell'_i$ for every $i \in [m]$. Let \mathbf{b} and \mathbf{b}' be the random vectors returned by the oracle upon ℓ and ℓ' . Then we can define \mathbf{b}^* using \mathbf{b}' as follows: $b_i^* = 0$ if $b'_i = 0$; otherwise when $b'_i = 1$, we set

$$b_i^* = \begin{cases} 1 & \text{with probability } \lambda(\ell_i)/\lambda(\ell'_i) \\ 0 & \text{with probability } 1 - \lambda(\ell_i)/\lambda(\ell'_i) \end{cases}.$$

One can verify that the distribution of \mathbf{b} is exactly the same as the distribution of \mathbf{b}^* . Hence there is a randomized map h' such that the advantage of (h', ℓ') is at least as large as the highest possible advantage achievable using ℓ . The remark now follows by our earlier observation in Remark 8 that the highest possible advantage using ℓ' is always achieved by a deterministic h' .

The following lemma reduces the proof of Lemma 9 to proving a lower bound for SSEQ.

► **Lemma 12.** *Given any deterministic and non-adaptive algorithm $\text{ALG} = (g, T)$ for SSSQ, there is a deterministic and non-adaptive algorithm $\text{ALG}' = (h, \ell)$ for SSEQ with the same query complexity as ALG and advantage at least as large as that of ALG .*

Proof. We show how to construct $\text{ALG}' = (h, \ell)$ from $\text{ALG} = (g, T)$, where h is a randomized map, such that ALG' has exactly the same query complexity and advantage as those of ALG . The lemma then follows from the observation we made earlier in Remark 8.

We define ℓ first. Given $T = (T_1, \dots, T_d)$ for some $d \geq 1$, $\ell = (\ell_1, \dots, \ell_m)$ is defined as

$$\ell_j = |\{i \in [d] : j \in T_i\}|.$$

So $\|\ell\|_1 = \sum_{i=1}^d |T_i|$. To define h we describe a randomized procedure P that, given any $b \in \{0, 1\}^m$, outputs a sequence of random vectors $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_d)$ such that the following claim holds.

► **Claim 13.** *If $\mathbf{b} \sim \mathcal{B}_{\text{yes}}$ (or \mathcal{B}_{no}), then $P(\mathbf{b})$ is distributed the same as \mathcal{V}_{yes} (or \mathcal{V}_{no} , respectively).*

Assuming Claim 13, we can set $h = g \circ P$ and the advantage of ALG' would be the same as that of ALG . In the rest of the proof, we describe the randomized procedure P and prove Claim 13.

Given $b \in \{0, 1\}^m$, P outputs a sequence of random vectors $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_d)$ as follows:

- If $b_j = 0$, then for each $i \in [d]$ with $j \in T_i$, P sets $\mathbf{v}_{i,j} = 0$.
- If $b_j = 1$ (this implies that $\ell_j > 0$ and $j \in T_i$ for some $i \in [d]$), P sets $(\mathbf{v}_{i,j} : i \in [d], j \in T_i)$ to be a length- r , where $r = |\{i \in [d] : j \in T_i\}|$, binary string in which each bit is independently 1 with probability ϵ/\sqrt{n} and 0 with probability $1 - \epsilon/\sqrt{n}$, conditioned on its not being 0^r .

Proof of Claim 13. It suffices to prove that, fixing any $A \subseteq [m]$ as the underlying set hidden in the oracle, the distribution of \mathbf{v} is the same as the distribution of $P(\mathbf{b})$. The claim then follows since in the definitions of both \mathcal{B}_{yes} and \mathcal{V}_{yes} (or \mathcal{B}_{no} and \mathcal{V}_{no}), A is drawn from \mathcal{A}_{yes} (or \mathcal{A}_{no} , respectively).

Consider a sequence v of d vectors v_1, \dots, v_d with $v_i \in \{0, 1\}^{T_i}$ for each $i \in [d]$, and let

$$n_{j,1} = |\{i \in [d] : j \in T_i \text{ and } v_{i,j} = 1\}| \quad \text{and} \quad n_{j,0} = |\{i \in [d] : j \in T_i \text{ and } v_{i,j} = 0\}|,$$

for each $j \in [m]$. Then the \mathbf{v} returned by the oracle (in SSSQ) is equal to v with probability:

$$\mathbf{1}\{\forall j \notin A, n_{j,1} = 0\} \cdot \prod_{j \in A} \left(\frac{\epsilon}{\sqrt{n}}\right)^{n_{j,1}} \left(1 - \frac{\epsilon}{\sqrt{n}}\right)^{n_{j,0}}, \quad (11)$$

since all coordinates $\mathbf{v}_{i,j}$ are independent. On the other hand, the probability of $P(\mathbf{b}) = v$ is

$$\mathbf{1}\{\forall j \notin A, n_{j,1} = 0\} \cdot \prod_{j \in A} \left(\mathbf{1}\{n_{j,0} = \ell_j\} \cdot \left(1 - \frac{\epsilon}{\sqrt{n}}\right)^{\ell_j} + \mathbf{1}\{n_{j,1} \geq 1\} \cdot \left(\frac{\epsilon}{\sqrt{n}}\right)^{n_{j,1}} \left(1 - \frac{\epsilon}{\sqrt{n}}\right)^{n_{j,0}} \right),$$

which is exactly the same as the probability of $\mathbf{v} = v$ in (11). ◀

This finishes the proof of Lemma 12. ◀

5.2 A lower bound for SSEQ

We prove the following lower bound for SSEQ, from which Lemma 9 follows:

► **Lemma 14.** *Any deterministic, non-adaptive ALG' for SSEQ with advantage at least $2/3$ satisfies*

$$\|\ell\|_1 > s \stackrel{\text{def}}{=} \frac{n^{3/2}}{\epsilon \cdot \log^3 n \cdot \log^2(n/\epsilon)}.$$

Proof. Assume for contradiction that there is an algorithm $\text{ALG}' = (h, \ell)$ with $\|\ell\|_1 \leq s$ and advantage at least $2/3$. Let ℓ^* be the vector obtained from ℓ by rounding each positive ℓ_i to the smallest power of 2 that is at least as large as ℓ_i (and taking $\ell_i^* = 0$ if $\ell_i = 0$). From Remark 11, there must be a map h^* such that (h^*, ℓ^*) also has advantage at least $2/3$ but now we have 1) $\|\ell^*\|_1 \leq 2s$ and 2) every positive entry of ℓ^* is a power of 2. Below we abuse notation and still use $\text{ALG}' = (h, \ell)$ to denote (h^*, ℓ^*) : $\text{ALG}' = (h, \ell)$ satisfies $\|\ell\|_1 \leq 2s$, every

positive entry of ℓ is a power of 2, and has advantage at least $2/3$. We obtain a contradiction below by showing that any such ℓ can only have an advantage of $o(1)$.

Let $L = \lceil \log(2s) \rceil = O(\log(n/\epsilon))$. Given that $\|\ell\|_1 \leq 2s$ we can partition $\{i \in [m] : \ell_i > 0\}$ into $L + 1$ sets C_0, \dots, C_L , where bin C_j contains those coordinates $i \in [m]$ with $\ell_i = 2^j$. We may make two further assumptions on $\text{ALG}' = (h, \ell)$ that will simplify the lower bound proof:

- We may reorder the entries in decreasing order and assume without loss of generality that

$$\ell = \left(\underbrace{2^L, \dots, 2^L}_{c_L}, \underbrace{2^{L-1}, \dots, 2^{L-1}}_{c_{L-1}}, \dots, \underbrace{1, \dots, 1}_{c_0}, 0, \dots, 0 \right), \quad (12)$$

where $c_j = |C_j|$ satisfies $\sum_j c_j \cdot 2^j \leq 2s$. This is without loss of generality since \mathcal{A}_{yes} and \mathcal{A}_{no} are symmetric in the coordinates (and so are \mathcal{B}_{yes} and \mathcal{B}_{no}).

- For the same reason we may assume that the map $h(b)$ depends only on the number of 1's of b in each set C_j , which we refer to as the *summary* $S(b)$ of b :

$$S(b) \stackrel{\text{def}}{=} \left(\|b_{|C_L}\|_1, \|b_{|C_{L-1}}\|_1, \dots, \|b_{|C_0}\|_1 \right) \in \mathbb{Z}_{\geq 0}^{L+1}.$$

To see that this is without loss of generality, consider a randomized procedure P that, given $b \in \{0, 1\}^m$, applies an independent random permutation over the entries of C_j for each bin $j \in [0 : L]$. One can verify that the random map $h' = h \circ P$ only depends on the summary $S(b)$ of b but achieves the same advantage as h .

Given a query ℓ as in (12), we define \mathcal{S}_{yes} to be the distribution of $S(\mathbf{b})$ for $\mathbf{b} \sim \mathcal{B}_{\text{yes}}$ (recall that \mathcal{B}_{yes} is the distribution of the vector \mathbf{b} returned by the oracle upon the query ℓ when $\mathbf{A} \sim \mathcal{A}_{\text{yes}}$). Similarly we define \mathcal{S}_{no} as the distribution of $S(\mathbf{b})$ for $\mathbf{b} \sim \mathcal{B}_{\text{no}}$. As h only depends on the summary the advantage is at most $d_{\text{TV}}(\mathcal{S}_{\text{yes}}, \mathcal{S}_{\text{no}})$, which we upper bound below by $o(1)$.

From the definition of \mathcal{B}_{yes} (or \mathcal{B}_{no} , respectively) and the fact that \mathcal{A}_{yes} (or \mathcal{A}_{no} , respectively) is symmetric over the m coordinates, we have that the $L + 1$ entries of \mathcal{S}_{yes} (of \mathcal{S}_{no} , respectively) are mutually independent, and that their entries for each C_j , $j \in [0 : L]$, are distributed as $\text{Bin}(c_j, p\lambda_j)$ (as $\text{Bin}(c_j, q\lambda_j)$, respectively), where we have $\lambda_j = 1 - (1 - (\epsilon/\sqrt{n}))^{2^j}$.

In order to prove that $d_{\text{TV}}(\mathcal{S}_{\text{yes}}, \mathcal{S}_{\text{no}}) = o(1)$ and achieve the desired contradiction, we will give upper bounds on the total variation distance between their C_j -entries for each $j \in \{0, \dots, L\}$.

- **Claim 15.** *Let $\mathbf{X} \sim \text{Bin}(c_j, p\lambda_j)$ and $\mathbf{Y} \sim \text{Bin}(c_j, q\lambda_j)$. Then $d_{\text{TV}}(\mathbf{X}, \mathbf{Y}) \leq o(1/L)$.*

We delay the proof of Claim 15, but assuming it we may simply apply the following well-known proposition to conclude that $d_{\text{TV}}(\mathcal{S}_{\text{yes}}, \mathcal{S}_{\text{no}}) = o(1)$.

- **Proposition 16** (Subadditivity of total variation distance). *Let $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_k)$ and $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_k)$ be two tuples of independent random variables. Then $d_{\text{TV}}(\mathbf{X}, \mathbf{Y}) \leq \sum_{i=1}^k d_{\text{TV}}(\mathbf{X}_i, \mathbf{Y}_i)$.*

This gives us a contradiction and finishes the proof of Lemma 14. ◀

Below we prove Claim 15.

Proof of Claim 15. The claim is trivial when $c_j = 0$ so we assume below that $c_j > 0$.

Let $r = p\lambda_j$ and $x = \log n \cdot \lambda_j / \sqrt{n}$. Then $\mathbf{X} \sim \text{Bin}(c_j, r)$ and $\mathbf{Y} \sim \text{Bin}(c_j, r + x)$. As indicated in Equation (2.15) of [1], Equation (15) of [37] gives

$$d_{\text{TV}}(\mathbf{X}, \mathbf{Y}) \leq O\left(\frac{\tau(x)}{(1 - \tau(x))^2}\right), \quad \text{where } \tau(x) \stackrel{\text{def}}{=} x \sqrt{\frac{c_j + 2}{2r(1 - r)}}, \quad (13)$$

whenever $\tau(x) < 1$. Substituting for x and r , we have (using $c_j \geq 1$, $r \leq 1/2$ and $p = 1/2$)

$$\tau(x) = O\left(\frac{\log n \cdot \lambda_j}{\sqrt{n}} \cdot \sqrt{\frac{c_j}{r}}\right) = O\left(\log n \cdot \sqrt{\frac{\lambda_j \cdot c_j}{n}}\right) = O\left(\frac{1}{L} \cdot \sqrt{\frac{n^{1/2} \cdot \lambda_j}{2^j \cdot \epsilon \cdot \log n}}\right),$$

where the last inequality follows from

$$c_j \cdot 2^j \leq 2s \leq O\left(\frac{n^{3/2}}{\epsilon \cdot \log^3 n \cdot L^2}\right).$$

Finally, note that (using $1 - x > e^{-2x}$ for small positive x and $1 - x \leq e^{-x}$ for all x):

$$1 - \lambda_j = \left(1 - \frac{\epsilon}{\sqrt{n}}\right)^{2^j} \geq \left(e^{-2\epsilon/\sqrt{n}}\right)^{2^j} = e^{-2^{j+1}\epsilon/\sqrt{n}} \geq 1 - O(2^j\epsilon/\sqrt{n})$$

and $\frac{\sqrt{n} \cdot \lambda_j}{2^j \cdot \epsilon} = O(1)$. This implies $\tau(x) = o(1/L) = o(1)$. The claim then follows from (13). \blacktriangleleft

References

- 1 José A Adell and Pedro Jodrá. Exact kolmogorov and total variation distances between some familiar discrete distributions. *Journal of Inequalities and Applications*, 2006(1):1–8, 2006.
- 2 N. Alon, T. Kaufman, M. Krivelevich, S. Litsyn, and D. Ron. Testing Reed-Muller Codes. *IEEE Transactions on Information Theory*, 51(11):4032–4039, 2005.
- 3 Rokhsana Baleshazar, Meiram Murzabulatov, Ramesh Krishnan S. Pallavoor, and Sofya Raskhodnikova. Testing unateness of real-valued functions. *CoRR*, abs/1608.07652, 2016.
- 4 A. Belovs and E. Blais. A polynomial lower bound for testing monotonicity. In *Proceedings of the 48th ACM Symposium on Theory of Computing*, pages 1021–1032, 2016.
- 5 Arthur J. Bernstein. Maximally connected arrays on the n -cube. *SIAM J. Appl. Math.*, 15(6):1485–1489, 1967.
- 6 Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of reed-muller codes. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 488–497, 2010.
- 7 Eric Blais. Improved bounds for testing juntas. In *Proc. RANDOM*, pages 317–330, 2008.
- 8 Eric Blais. Testing juntas nearly optimally. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 151–158, 2009. doi:10.1145/1536414.1536437.
- 9 Eric Blais, Joshua Brody, and Kevin Matulef. Property testing lower bounds via communication complexity. In *CCC*, pages 210–220, 2011.
- 10 Eric Blais and Daniel M. Kane. Tight bounds for testing k -linearity. In *RANDOM*, pages 435–446, 2012.
- 11 M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47:549–595, 1993. Earlier version in STOC’90.

- 12 Harry Buhrman, David García-Soriano, Arie Matsliah, and Ronald de Wolf. The non-adaptive query complexity of testing k -parities. *Chicago Journal of Theoretical Computer Science*, 2013, 2013.
- 13 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.
- 14 Deeparnab Chakrabarty and C. Seshadhri. A $\tilde{O}(n)$ non-adaptive tester for unateness. *CoRR*, abs/1608.06980, 2016.
- 15 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 47th ACM Symposium on Theory of Computing*, pages 519–528, 2015.
- 16 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.
- 17 H. Chockler and D. Gutfreund. A lower bound for testing juntas. *Information Processing Letters*, 90(6):301–305, 2004.
- 18 I. Diakonikolas, H. Lee, K. Matulef, K. Onak, R. Rubinfeld, R. Servedio, and A. Wan. Testing for concise representations. In *Proc. 48th Ann. Symposium on Computer Science (FOCS)*, pages 549–558, 2007.
- 19 E. Fischer, G. Kindler, D. Ron, S. Safra, and A. Samorodnitsky. Testing juntas. *J. Computer & System Sciences*, 68(4):753–787, 2004.
- 20 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.
- 21 Peter Frankl. On the trace of finite sets. *J. Comb. Theory, Ser. A*, 34(1):41–45, 1983.
- 22 O. Goldreich, editor. *Property Testing: Current Research and Surveys*. Springer, 2010. LNCS 6390.
- 23 O. Goldreich, S. Goldwasser, E. Lehman, D. Ron, and A. Samordinsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000.
- 24 P. Gopalan, R. O’Donnell, R. Servedio, A. Shpilka, and K. Wimmer. Testing Fourier dimensionality and sparsity. *SIAM J. on Computing*, 40(4):1075–1100, 2011.
- 25 Larry H. Harper. Optimal assignments of numbers to vertices. *SIAM J. Appl. Math.*, 12(1):131–135, 1964.
- 26 Sergiu Hart. A note on the edges of the n -cube. *Disc. Math.*, 14:157–163, 1976.
- 27 Subhash Khot, Dor Minzer, and Muli Safra. On monotonicity testing and boolean isoperimetric type theorems. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science*, pages 52–58, 2015.
- 28 Subhash Khot and Igor Shinkar. An $o(n)$ queries adaptive tester for unateness. In *Approximation, Randomization, and Combinatorial Optimization Algorithms and Techniques*, 2016.
- 29 J. H. Lindsey. Assignment of numbers to vertices. *Amer. Math. Monthly*, 71:508–516, 1964.
- 30 K. Matulef, R. O’Donnell, R. Rubinfeld, and R. Servedio. Testing halfspaces. *SIAM J. on Comput.*, 39(5):2004–2047, 2010.
- 31 Kevin Matulef, Ryan O’Donnell, Ronitt Rubinfeld, and Rocco A. Servedio. Testing ± 1 -weight halfspace. In *APPROX-RANDOM*, pages 646–657, 2009. doi:10.1007/978-3-642-03685-9_48.
- 32 M. Parnas, D. Ron, and A. Samorodnitsky. Testing Basic Boolean Formulae. *SIAM J. Disc. Math.*, 16:20–46, 2002.
- 33 D. Ron. Property Testing: A Learning Theory Perspective. *Foundations and Trends in Machine Learning*, 1(3):307–402, 2008.

- 34 D. Ron. Algorithmic and analysis techniques in property testing. *Foundations and Trends in Theoretical Computer Science*, 5:73–205, 2010.
- 35 D. Ron and R. Servedio. Exponentially improved algorithms and lower bounds for testing signed majorities. In *SODA*, pages 1319–1336, 2013.
- 36 Dana Ron and Gilad Tsur. Testing computability by width-two obdds. Technical Report 11(041), ECCCC, 2011. available at <http://ecccc.hpi-web.de/report/2011/041/>.
- 37 B. Roos. Binomial approximation to the Poisson binomial distribution: The Krawtchouk expansion. *Theory Probab. Appl.*, 45:328–344, 2000.
- 38 Rocco Servedio, Li-Yang Tan, and John Wright. Adaptivity helps for testing juntas. In *Proceedings of the 30th IEEE Conference on Computational Complexity*, pages 264–279, 2015. volume 33 of LIPIcs.

A Proof of Theorem 1 assuming Theorem 2

We prove the following claim in Appendix A.1.

► **Claim 17.** *Let $\epsilon(n)$ be a function that satisfies $2^{-n} \leq \epsilon(n) \leq 1/5$ for sufficiently large n . Then any non-adaptive algorithm that accepts the all-0 function with probability at least $5/6$ and rejects every function that is ϵ -far from $(n-1)$ -juntas with probability at least $5/6$ must make $\Omega(1/\epsilon)$ queries.*

Next let $k(n)$ and $\epsilon(n)$ be the pair of functions from the statement of Theorem 1. We consider a sufficiently large n (letting $k = k(n)$ and $\epsilon = \epsilon(n)$ below) and separate the proof into two cases:

$$2^{-(2\alpha-1)k/(2\alpha)} \leq \epsilon \leq 1/6 \quad \text{and} \quad 2^{-n} \leq \epsilon < 2^{-(2\alpha-1)k/(2\alpha)}.$$

For the first case, if $k = O(1)$ then the bound we aim for is simply $\tilde{\Omega}(1/\epsilon)$, which follows trivially from Claim 17 (since $k \leq \alpha n < n-1$ and the all-0 function is a k -junta). Otherwise we combine the following reduction with Theorem 2: any ϵ -tester for k -juntas over n -variable functions can be used to obtain an ϵ -tester for k -juntas over (k/α) -variable functions. This can be done by adding $n - k/\alpha$ dummy variables to any (k/α) -variable function to make the number of variables n (as $k \leq \alpha n$). The lower bound then follows from Theorem 2 since α is a constant. For the second case, the lower bound claimed in Theorem 1 is $\tilde{\Omega}(1/\epsilon)$, which follows again from Claim 17. This concludes the proof of Theorem 1 given Theorem 2 and Claim 17. ◀

A.1 Proof of Claim 17

Let C be a sufficiently large constant. We prove Claim 17 by considering two cases:

$$\epsilon \geq \frac{C \log n}{2^n} \quad \text{and} \quad \epsilon < \frac{C \log n}{2^n}.$$

For the first case of $2^n \epsilon \geq C \log n$, we use \mathcal{D}_1 to denote the following distribution over n -variable Boolean functions: to draw $\mathbf{g} \sim \mathcal{D}_1$, independently for each $x \in \{0, 1\}^n$ the value of $\mathbf{g}(x)$ is set to 0 with probability $1 - 3\epsilon$ (recall that $\epsilon \leq 1/5$) and 1 with probability 3ϵ .

We prove the following lemma for the distribution \mathcal{D}_1 :

► **Lemma 18.** *With probability at least $1 - o(1)$, $\mathbf{g} \sim \mathcal{D}_1$ is ϵ -far from every $(n-1)$ -junta.*

Proof. Note that every $(n - 1)$ -junta is such that for some $i \in [n]$, the function does not depend on the i -th variable; we refer to such a function as a type- i junta. An easy lower bound for the distance from a function g to all type- i juntas is the number of g -bichromatic edges $(x, x^{(i)})$ divided by 2^n . When $\mathbf{g} \sim \mathcal{D}_1$ each edge $(x, x^{(i)})$ is independently \mathbf{g} -bichromatic with probability $6\epsilon(1 - 3\epsilon) \geq 12\epsilon/5$ (as $\epsilon \leq 1/5$). Thus when $2^n\epsilon \geq C \log n$, the expected number of such edges is at least

$$2^{n-1} \cdot (12\epsilon/5) \geq (6/5) \cdot 2^n\epsilon \geq (6/5) \cdot C \log n.$$

Using a Chernoff bound, the probability of having fewer than $2^n\epsilon$ bichromatic edges along direction i is at most $1/n^2$ when C is sufficiently large. The lemma follows from a union bound over i . ◀

As a result, when $2^n\epsilon \geq C \log n$, if \mathcal{A} is a non-adaptive algorithm with the property described in Claim 17, then \mathcal{A} must satisfy

$$\Pr[\mathcal{A} \text{ accepts the all-0 function}] - \Pr_{\mathbf{g} \sim \mathcal{D}_1}[\mathcal{A} \text{ accepts } \mathbf{g}] \geq 2/3 - o(1).$$

But any such non-adaptive algorithm must make $\Omega(1/\epsilon)$ queries as otherwise with high probability all of its queries to $\mathbf{g} \sim \mathcal{D}_1$ would be answered 0, and hence its behavior would be the same as if it were running on the all-0 function.

Finally we work on the case when $1 \leq 2^n\epsilon = O(\log n)$. The proof is the same except that we let \mathbf{g} be drawn from \mathcal{D}_2 , which we define to be the distribution where all entries of $\mathbf{g} \sim \mathcal{D}_2$ are 0 except for exactly $2^n\epsilon$ of them picked uniformly at random. The claim follows from the following lemma:

► **Lemma 19.** *With probability at least $1 - o(1)$, $\mathbf{g} \sim \mathcal{D}_2$ is ϵ -far from every $(n - 1)$ -junta.*

Proof. This follows from the observation that, with probability $1 - o(1)$, no two points picked form an edge. When this happens, we have $2^n\epsilon$ bichromatic edges along the i th direction for all i . ◀