Distribution-Free Junta Testing*

Zhengyang Liu
Shanghai Jiao Tong University
Shanghai, China
lzy5118@sjtu.edu.cn

Xi Chen
Columbia University
New York, USA
xichen@cs.columbia.edu

Rocco A. Servedio
Columbia University
New York, USA
rocco@cs.columbia.edu

Ying Sheng
Columbia University
New York, USA
ys2982@columbia.edu

Jinyu Xie
Columbia University
New York, USA
jinyu@cs.columbia.edu

ABSTRACT

We study the problem of testing whether an unknown n -variable Boolean function is a k -junta in the *distribution-free* property testing model, where the distance between functions is measured with respect to an arbitrary and unknown probability distribution over $\{0, 1\}^n$. Our first main result is that distribution-free k -junta testing can be performed, with one-sided error, by an adaptive algorithm that uses $\tilde{O}(k^2)/\epsilon$ queries (independent of n). Complementing this, our second main result is a lower bound showing that any *non-adaptive* distribution-free k -junta testing algorithm must make $\Omega(2^{k/3})$ queries even to test to accuracy $\epsilon = 1/3$. These bounds establish that while the optimal query complexity of non-adaptive k -junta testing is $2^{\Theta(k)}$, for adaptive testing it is $\text{poly}(k)$, and thus show that adaptivity provides an exponential improvement in the distribution-free query complexity of testing juntas.

CCS CONCEPTS

• **Theory of computation** → **Streaming, sublinear and near linear time algorithms**;

KEYWORDS

Property testing, juntas, distribution-free testing

ACM Reference Format:

Zhengyang Liu, Xi Chen, Rocco A. Servedio, Ying Sheng, and Jinyu Xie. 2018. Distribution-Free Junta Testing. In *Proceedings of 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC'18)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3188745.3188842>

1 INTRODUCTION

Property testing of Boolean functions was first considered in the seminal works of Blum, Luby and Rubinfeld [10] and Rubinfeld

and Sudan [39] and has developed into a robust research area at the intersection of sub-linear algorithms and complexity theory. Roughly speaking, a property tester for a class C of functions from $\{0, 1\}^n$ to $\{0, 1\}$ is a randomized algorithm that is given some form of access to the (unknown) input Boolean function f , and must with high probability distinguish the case that $f \in C$ versus the case that f is ϵ -far from every function $g \in C$. In the usual (uniform-distribution) property testing scenario, the testing algorithm may access f by making black-box queries on inputs $x \in \{0, 1\}^n$, and the distance between two functions f and g is measured with respect to the uniform distribution on $\{0, 1\}^n$; the goal is to develop algorithms that make as few queries as possible. Many different classes of Boolean functions have been studied from this perspective, see [1, 3–5, 8, 9–10, 12–14, 16, 18, 19, 22, 23, 25, 28, 30, 32–36] and other works referenced in the surveys [27, 37, 38]. Among these, the class of *k-juntas* — Boolean functions that depend only on (an unknown set of) at most k of their n input variables — is one of the best-known and most intensively investigated such classes [6, 7, 11, 21, 24, 40], with ongoing research on junta testing continuing right up to the present [17].

The query complexity of junta testing in the uniform distribution framework is now well understood. Improving on $\text{poly}(k)/\epsilon$ -query algorithms given in [24] (which introduced the junta testing problem), in [6] Blais gave a non-adaptive algorithm that makes $\tilde{O}(k^{3/2})/\epsilon$ queries, and in [7] Blais gave an $O(k \log k + k/\epsilon)$ -query adaptive algorithm. On the lower bounds side, Fischer et al. [24] initially gave an $\Omega(\sqrt{k})$ lower bound for non-adaptively testing k -juntas, which also implies an $\Omega(\log k)$ lower bound for adaptive testing. Chockler and Gutfreund improved the adaptive lower bound to $\Omega(k)$ in [21], and very recently Chen et al. [17] gave an $\tilde{\Omega}(k^{3/2})/\epsilon$ non-adaptive lower bound. Thus in both the adaptive and non-adaptive uniform distribution settings, the query complexity of k -junta testing has now been pinned down to within logarithmic factors.

Distribution-Free Property Testing. This work studies the junta testing problem in the *distribution-free* property testing model that was first introduced by Goldreich et al in [29]. In this model the distance between Boolean functions is measured with respect to a distribution $\mathcal{D}$ over $\{0, 1\}^n$ which is arbitrary and unknown to the testing algorithm. Since the distribution is unknown, in this model the testing algorithm is allowed (in addition to making black-box queries) to draw random labeled samples $(x, f(x))$ where each x is independently distributed according to $\mathcal{D}$. The query complexity

* Author ordering reflects administrative constraints. This work is supported by NSF CCF-1703925, NSF CCF-1420349, and NSF CCF-1563155.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC'18, June 25–29, 2018, Los Angeles, CA, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5559-9/18/06...\$15.00

<https://doi.org/10.1145/3188745.3188842>

of an algorithm in this framework is the worst-case total number of black-box oracle calls plus random labeled samples that are used, across all possible distributions. (It follows that distribution-free testing of a class C requires at least as many queries as testing C in the standard uniform-distribution model.)

Distribution-free property testing is in the spirit of similar distribution-free models in computational learning theory such as the celebrated PAC learning model of Valiant [41]. Such models are attractive because of their minimal assumptions; they are well motivated both because in many natural settings the uniform distribution over $\{0, 1\}^n$ may not be the best way to measure distances, and because they capture the notion of an algorithm dealing with an unknown and arbitrary environment (modeled here by the unknown and arbitrary distribution $\mathcal{D}$ over $\{0, 1\}^n$ and the unknown and arbitrary Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$). Researchers have studied distribution-free testing of a number of Boolean function classes, including monotone functions, low-degree polynomials, dictators (1-juntas) and k -juntas [31], disjunctions and conjunctions (monotone and non-monotone), decision lists, and linear threshold functions [20, 23, 26]. Since depending on few variables is an appealingly flexible “real-world” property in comparison with more highly structured syntactically defined properties, we feel that junta testing is a particularly natural task to study in the distribution-free model.

Prior Results on Distribution-Free Junta Testing. Given how thoroughly junta testing has been studied in the uniform distribution model, surprisingly little was known in the distribution-free setting. The adaptive $\Omega(k)$ and non-adaptive $\tilde{\Omega}(k^{3/2})/\epsilon$ uniform-distribution lower bounds from [17, 21] mentioned earlier trivially extend to the distribution-free model, but no other lower bounds on distribution-free junta testing were known prior to this work. On the positive side, Halevy and Kushilevitz showed in [31] that any class C that has (i) a one-sided error uniform-distribution testing algorithm and (ii) a self-corrector, has a one-sided error distribution-free testing algorithm. As $\text{poly}(k)/\epsilon$ -query one-sided junta testers were given in [24], and k -juntas have $O(2^k)$ -query self-correctors [2], this yields a one-sided non-adaptive distribution-free junta tester with query complexity $O(2^k/\epsilon)$. No other results were known.

Thus, prior to this work there were major gaps in our understanding of distribution-free k -junta testing: is the query complexity of this problem polynomial in k , exponential in k , or somewhere in between? Does adaptivity confer an exponential advantage, a sub-exponential advantage, or no advantage at all? Our results, described below, answer both these questions.

1.1 Our Results

Our main positive result is a $\text{poly}(k)/\epsilon$ -query one-sided adaptive algorithm for distribution-free k -junta testing:

THEOREM 1.1 (UPPER BOUND). *For any parameter $\epsilon > 0$, there is a one-sided distribution-free adaptive algorithm for ϵ -testing k -juntas with $\tilde{O}(k^2)/\epsilon$ queries.*

Theorem 1.1 shows that k -juntas stand in interesting contrast with many other well-studied classes of Boolean functions in property testing such as conjunctions, decision lists, linear threshold functions, and monotone functions. For each of these classes of

Boolean functions, distribution-free testing requires dramatically more queries than uniform-distribution testing: for the first three classes the separation is $\text{poly}(1/\epsilon)$ queries in the uniform setting [34, 36] versus $n^{\Omega(1)}$ queries in the distribution-free setting [20, 26]; for n -variable monotone functions $\text{poly}(n)$ queries suffice in the uniform setting [28, 32] whereas [31] shows that $2^{\Omega(n)}$ many queries are required in the distribution-free setting. In contrast, Theorem 1.1 implies that for k -juntas the query complexities of uniform-distribution and distribution-free testing are polynomially related (indeed, within at most a quadratic factor).

Complementing the strong upper bound which Theorem 1.1 gives for adaptive testers, our main negative result is an $\Omega(2^{k/3})$ -query lower bound for non-adaptive testers:

THEOREM 1.2 (LOWER BOUND). *For $k \leq n/200$, any non-adaptive algorithm that distribution-free ϵ -tests k -juntas over $\{0, 1\}^n$, for $\epsilon = 1/3$, must have query complexity $\Omega(2^{k/3})$.*

Theorems 1.1 and 1.2 together show that adaptivity enables an exponential improvement in the distribution-free query complexity of testing juntas. This is in sharp contrast with uniform-distribution junta testing, where the adaptive and non-adaptive query complexities are polynomially related (with an exponent of only $3/2$). To the best of our knowledge, this is the first example of an exponential separation between adaptive and nonadaptive distribution-free testers.

1.2 Ideas and Techniques

The Algorithm. As a first step toward our main $\tilde{O}(k^2)/\epsilon$ -query algorithm, in Section 3 we present a simple one-sided adaptive algorithm, which we call **SimpleDJunta**, that distribution-free tests k -juntas using $O((k/\epsilon) + k \log n)$ queries. **SimpleDJunta** uses binary search and is an adaptation to the distribution-free setting of the $O((k/\epsilon) + k \log n)$ -query uniform-distribution algorithm which is implicit in [7]. The algorithm maintains a set I of *relevant* variables: a string $x \in \{0, 1\}^n$ has been found for each $i \in I$ such that $f(x) \neq f(x^{(i)})$ (we use $x^{(i)}$ to denote the string obtained by flipping the i -th bit of x), and the algorithm rejects only when $|I|$ becomes larger than k . In each round, the algorithm samples a string $x \leftarrow \mathcal{D}$ and a subset R of $\bar{I} := [n] \setminus I$ uniformly at random. A simple lemma, Lemma 3.2, states that if f is far from every k -junta with respect to $\mathcal{D}$, then we must have

$$f(x) \neq f(x^{(R)})$$

with at least some moderately large probability as long as $|I| \leq k$, where we use $x^{(R)}$ to denote the string obtained from x by flipping every coordinate in R . With $(x, x^{(R)})$ in hand, it is straightforward to find a new relevant variable using binary search over coordinates in R (see Figure 1, where we use $\text{diff}(x, y)$ to denote the set of $i \in [n]$ with $x_i \neq y_i$), with at most $\log n$ additional queries.

In order to achieve a query complexity that is independent of n , one must employ a more efficient approach than binary search over $\Omega(n)$ coordinates (since most likely the set R has size $\Omega(n)$ for the range of k we are interested in). In the uniform-distribution setting this is accomplished in [7] by first randomly partitioning the variable space $[n]$ into $s = \text{poly}(k/\epsilon)$ disjoint blocks $B_1, \dots, B_s$ of variables and carrying out binary search over blocks (see Figure

Procedure BinarySearch(f, x, y)**Input:** Query access to $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and two strings $x, y \in \{0, 1\}^n$ with $f(x) \neq f(y)$.**Output:** Two strings $x', y' \in \{0, 1\}^n$ with $f(x') \neq f(y')$ and $x' = y'^{(i)}$ for some $i \in \text{diff}(x, y)$.

- (1) Let $B \subseteq [n]$ be the set such that $x = y^{(B)}$.
- (2) If $|B| = 1$, return x and y .
- (3) Partition (arbitrarily) B into B_1 and B_2 of size $\lfloor |B|/2 \rfloor$ and $\lceil |B|/2 \rceil$, respectively.
- (4) Query $f(x^{(B_1)})$.
- (5) If $f(x) \neq f(x^{(B_1)})$, return **BinarySearch**($f, x, x^{(B_1)}$).
- (6) Otherwise, return **BinarySearch**($f, x^{(B_1)}, y$).

Figure 1: Description of the standard binary search procedure.**Procedure BlockBinarySearch**($f, x, y, B_1, \dots, B_r$)**Input:** Query access to $f: \{0, 1\}^n \rightarrow \{0, 1\}$, two strings $x, y \in \{0, 1\}^n$ with $f(x) \neq f(y)$, and a sequence of pairwise disjoint blocks $B_1, \dots, B_r$ for some $r \geq 1$ with $\text{diff}(x, y) \subseteq B_1 \cup \dots \cup B_r$.**Output:** Two strings $x', y' \in \{0, 1\}^n$ with $f(x') \neq f(y')$ and $\text{diff}(x, y) \subseteq B_i$ for some $i \in [r]$.

- (1) If $r = 1$, return x and y .
- (2) Let $t = \lfloor r/2 \rfloor$ and B be the intersection of $\text{diff}(x, y)$ and $B_1 \cup \dots \cup B_t$.
- (3) Query $f(x^{(B)})$.
- (4) If $f(x) \neq f(x^{(B)})$, return **BlockBinarySearch**($f, x, x^{(B)}; B_1, \dots, B_t$).
- (5) Otherwise, return **BlockBinarySearch**($f, x^{(B)}, y; B_{t+1}, \dots, B_r$).

Figure 2: Description of the blockwise version of the binary search procedure.

2) rather than over individual coordinates; this reduces the cost of each binary search to $\log(k/\epsilon)$ rather than $\log n$. The algorithm maintains a set of *relevant blocks*: two strings $x, y \in \{0, 1\}^n$ have been found for each such block B which satisfy $f(x) \neq f(y)$ and $y = x^{(S)}$ with $S \subseteq B$, and the algorithm rejects when more than k relevant blocks have been found. In each round the algorithm samples two strings x, y uniformly at random conditioned on their agreeing with each other on the relevant blocks that have already been found in previous rounds; if $f(x) \neq f(y)$, then the binary search over blocks is performed to find a new relevant block. To establish the correctness of this approach [7] employs a detailed and technical analytic argument based on the influence of coordinates and the Efron-Stein orthogonal decomposition of functions over product spaces. This machinery is well suited for dealing with product distributions, and indeed the analysis of [7] goes through for any product distribution over $\{0, 1\}^n$ (and even for more general finite domains and ranges). However, it is far from clear how to extend this machinery to work for the completely unstructured distributions $\mathcal{D}$ that must be handled in the distribution-free model.

Our main distribution-free junta testing algorithm, which we denote **MainDJunta**, draws ideas from both **SimpleDJunta** (mainly Lemma 3.2) and the uniform distribution tester of [7]. To avoid the $\log n$ cost, the algorithm carries out binary search over blocks rather than over individual coordinates, and maintains a set of disjoint relevant blocks $B_1, \dots, B_\ell$, i.e., for each B_j a pair of strings x^j and y^j have been found such that they agree with each other over $\overline{B_j}$ and satisfy $f(x^j) \neq f(y^j)$. Let w^j be the projection of x^j (and y^j) over $\overline{B_j}$ and let g_j be the Boolean function over $\{0, 1\}^{B_j}$ obtained from f by setting variables in $\overline{B_j}$ to w^j . For clarity we assume further that

every function g_j is very close to a *literal* (i.e. for some $\tau \in \{x_{i_j}, \overline{x_{i_j}}\}$ we have $g_j(x) = \tau$ for all $x \in \{0, 1\}^{B_j}$ for some $i_j \in B_j$) under the *uniform* distribution. (To justify this assumption we note that if g_j is far from every literal under the uniform distribution, then it is easy to split B_j further into two relevant blocks using the uniform distribution algorithm of [7].) Let $I = \{i_j : j \in [\ell]\}$. Even though the algorithm does not know I , there is indeed a way to draw uniformly random subsets $\mathbf{R}$ of $\overline{I}$. First we draw a partition of B_j into $\mathbf{P}_j$ and $\mathbf{Q}_j$ uniformly at random, for each j . Since g_j is close to a literal, it is not difficult to figure out whether $\mathbf{P}_j$ or $\mathbf{Q}_j$ contains the hidden i_j , say it is $\mathbf{P}_j$ for every j . Then the union of all $\mathbf{Q}_j$'s together with a uniformly random subset of $\overline{B_1 \cup \dots \cup B_\ell}$, denoted by $\mathbf{R}$, turns out to be a uniformly random subset of $\overline{I}$. With $\mathbf{R}$ in hand, Lemma 3.2 implies that $f(x) \neq f(x^{(\mathbf{R})})$ with high probability when $x \leftarrow \mathcal{D}$, and when this happens, one can carry out binary search over blocks to increase the number of relevant blocks by one. In Section 4.1 we explain the intuition behind the main algorithm in more detail.

The Lower Bound. A q -query non-adaptive distribution-free tester is a randomized algorithm A that works as follows. When A is run on a pair $(\phi, \mathcal{D})$ ¹, it is first given the result $(y^1, \phi(y^1)), \dots, (y^q, \phi(y^q))$ of q queries from the sampling oracle. Based on them, it queries the black-box oracle q times on strings $z^1, \dots, z^q$. The z^j 's may depend on the random pairs $(y^i, \phi(y^i))$ received from the sampling oracle, but the j -th black-box query string z^j may not depend on the responses $\phi(z^1), \dots, \phi(z^{j-1})$ to any of the $j-1$ earlier black-box queries.

¹For clarity, throughout our discussion of lower bounds we write ϕ to indicate a function which may be either a “yes-function” or a “no-function”, f to denote a “yes-function” and g to denote a “no-function”.

Figure 3: A schematic depiction of how $\{0, 1\}^n$ is labeled by a function g from NO . The domain $\{0, 1\}^n$ is partitioned into 2^k sections corresponding to different settings of the variables in J ; each section is a vertical strip in the figure. Shaded regions correspond to strings where g evaluates to 1 and unshaded regions to strings where g evaluates to 0. Each string in S is a black dot and the value of g on each such string is chosen uniformly at random. Since in this figure the truncated circles are disjoint, the tie-breaking rule does not come into effect, and for each $z \in S$ all strings in its section within distance at most $0.4n$ (the points in the truncated circle around z) have the same value as z . The value of g on other points is determined by the background junta h which assigns a uniform random bit to each section.

As is standard in property testing lower bounds, our argument employs a distribution $\mathcal{YES}$ over yes-instances and a distribution NO over no-instances. Here $\mathcal{YES}$ is a distribution over (function, distribution) pairs $(f, \mathcal{D})$ in which f is guaranteed to be a k -junta; NO is a distribution over pairs $(g, \mathcal{D})$ such that with probability $1 - o(1)$, g is $1/3$ -far from every k -junta with respect to $\mathcal{D}$. To prove the desired lower bound against non-adaptive distribution-free testers, it suffices to show that for $q = 2^{k/3}$, any deterministic non-adaptive algorithm A as described above is roughly equally likely to accept whether it is run on an input drawn from $\mathcal{YES}$ or from NO .

Our construction of the $\mathcal{YES}$ and NO distributions is essentially as follows. In making a draw either from $\mathcal{YES}$ or from NO , first $m = \Theta(2^k \log n)$ strings are selected uniformly at random from $\{0, 1\}^n$ to form a set S , and the distribution $\mathcal{D}$ in both $\mathcal{YES}$ and NO is set to be the uniform distribution over S . Also in both $\mathcal{YES}$ and NO , a “background” k -junta h is selected uniformly at random by first picking a set J of k variables at random and then a random truth table for h over the variables in J . We view the variables in J as partitioning $\{0, 1\}^n$ into 2^k disjoint “sections” depending on how they are set.

In the case of a draw from $\mathcal{YES}$, the function f that goes with the above-described $\mathcal{D}$ is simply the background junta $f = h$. In the case of a draw from NO , the function g that goes with $\mathcal{D}$ is formed by modifying the background junta h in the following way

(roughly speaking; see Section 5.1 in the full version [15] for precise details): for each $z \in S$, we toss a fair coin $b(z)$ and set the value of all the strings in z ’s section that lie within Hamming distance $0.4n$ from z (including z itself) to $b(z)$ (see Figure 3). Note that the value of g at each string in S is a fair coin toss, which is completely independent of the background junta h . Using the choice of m it can be argued that with high probability g is $1/3$ -far from every k -junta with respect to $\mathcal{D}$.

The rough idea of why a pair $(f, \mathcal{D}) \leftarrow \mathcal{YES}$ is difficult for a $(q = 2^{k/3})$ -query non-adaptive algorithm A to distinguish from a pair $(g, \mathcal{D}) \leftarrow NO$ is as follows. Intuitively, in order for A to distinguish the no-case from the yes-case, it must obtain two strings x^1, x^2 that belong to the same section but are labeled differently. Since there are 2^k sections but q is only $2^{k/3}$, by the birthday paradox it is very unlikely that A obtains two such strings among the q samples $y^1, \dots, y^q$ that it is given from the distribution $\mathcal{D}$. In fact, in both the yes- and no- cases, writing $(\phi, \mathcal{D})$ to denote the (function, distribution) pair, the distribution of the q pairs

$$(y^1, \phi(y^1)), \dots, (y^q, \phi(y^q))$$

will be statistically very close to

$$(x^1, b_1), \dots, (x^q, b_q),$$

where each pair (x^j, b_j) is drawn independently and uniformly from $\{0, 1\}^n \times \{0, 1\}$. Intuitively, this translates into the examples $(y^i, \phi(y^i))$ from the sampling oracle “having no useful information” about the set J of variables that the background junta depends on.

What about the q strings $z^1, \dots, z^q$ that A feeds to the black-box oracle? It is also unlikely that any two elements of $y^1, \dots, y^q, z^1, \dots, z^q$ belong to the same section but are labeled differently. Fix an $i \in [q]$; we give some intuition here as to why it is very unlikely that there is any j such that z^i lies in the same section as y^j but has $f(z^i) \neq f(y^j)$ (via a union bound, the same intuition handles all $i \in [q]$). Intuitively, since the random examples from the sampling oracle provide no useful information about the set J defining the background junta, the only thing that A can do in selecting z^i is to choose how far it lies, in terms of Hamming distance, from the points in $y^1, \dots, y^q$ (which, recall, are uniform random). Fix $j \in [q]$: if z^i is within Hamming distance $0.4n$ from y^j , then even if z^i lies in the same section as y^j it will be labeled the same way as y^j whether we are in the yes- case or the no- case. On the other hand, if z^i is farther than $0.4n$ in Hamming distance from y^j , then it is overwhelmingly likely that z^i will lie in a different section from y^j (since it is very unlikely that all $0.4n$ of the flipped coordinates avoid the k -element set J). We prove Theorem 1.2 in the full version [15] via a formal argument that proceeds somewhat differently from but is informed by the above intuitions.

Organization. In Section 2 we define the distribution-free testing model and introduce some useful notation. In Section 3 we present **SimpleDJunta** and prove Lemma 3.2 in its analysis. In Section 4 we present our main algorithm **MainDJunta** and prove Theorem 1.1.

2 PRELIMINARIES

Notation. We use $[n]$ to denote $\{1, \dots, n\}$. We use f and g to denote Boolean functions, which are maps from $\{0, 1\}^n$ to $\{0, 1\}$ for some positive integer n . We use the calligraphic font (e.g., $\mathcal{D}$ and NO)

to denote probability distributions, boldface letters such as $\mathbf{x}$ to denote random variables, and write “ $\mathbf{x} \leftarrow \mathcal{D}$ ” to indicate that $\mathbf{x}$ is a random variable drawn from a distribution $\mathcal{D}$. We also write $\mathbf{x} \leftarrow \{0, 1\}^n$ to denote that $\mathbf{x}$ is a string drawn uniformly at random. Given $S \subseteq [n]$, we write $\mathbf{R} \leftarrow S$ to indicate that $\mathbf{R}$ is a subset of S drawn uniformly at random, i.e., each index $i \in S$ is included in $\mathbf{R}$ independently with probability $1/2$.

Given a subset $B \subseteq [n]$, we use $\bar{B}$ to denote its complement with respect to $[n]$, and $\{0, 1\}^B$ to denote the set of all binary strings of length $|B|$ with coordinates indexed by $i \in B$. Given an $x \in \{0, 1\}^n$ and a $B \subseteq [n]$, we write $x_B \in \{0, 1\}^B$ to denote the projection of x over coordinates in B and $x^{(B)} \in \{0, 1\}^n$ to denote the string obtained from x by flipping coordinates in B . Given $x \in \{0, 1\}^B$ and $y \in \{0, 1\}^{\bar{B}}$, we write $x \circ y \in \{0, 1\}^n$ to denote their concatenation, a string that agrees with x over coordinates in B and agrees with y over $\bar{B}$. (As an example of the notation, given $x, y \in \{0, 1\}^n$ and $B \subseteq [n]$, $x_B \circ y_{\bar{B}}$ denotes the string that agrees with x over B and with y over $\bar{B}$.) Given $x, y \in \{0, 1\}^n$, we write $d(x, y)$ to denote the Hamming distance between x and y , and $\text{diff}(x, y) \subseteq [n]$ to denote the set of coordinates $i \in [n]$ with $x_i \neq y_i$.

Given $f, g : \{0, 1\}^n \rightarrow \{0, 1\}$ and a distribution $\mathcal{D}$ over $\{0, 1\}^n$, we denote the *distance* between f and g with respect to $\mathcal{D}$ by

$$\text{dist}_{\mathcal{D}}(f, g) := \Pr_{z \leftarrow \mathcal{D}} [f(z) \neq g(z)].$$

Given a class $\mathcal{C}$ of Boolean functions,

$$\text{dist}_{\mathcal{D}}(f, \mathcal{C}) := \min_{g \in \mathcal{C}} (\text{dist}_{\mathcal{D}}(f, g))$$

denotes the *distance* between f and $\mathcal{C}$ with respect to $\mathcal{D}$, where the minimum is taken over g with the same number of variables as f . We say f is ϵ -far from $\mathcal{C}$ with respect to $\mathcal{D}$ if $\text{dist}_{\mathcal{D}}(f, \mathcal{C}) \geq \epsilon$.

We will often work with restrictions of Boolean functions. Given a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, $S \subseteq [n]$ and $z \in \{0, 1\}^{\bar{S}}$, the *restriction of f over S by z* , denoted by $f|_z$, is the Boolean function $g : \{0, 1\}^S \rightarrow \{0, 1\}$ defined by $g(x) = f(x \circ z)$ for all $x \in \{0, 1\}^S$.

Distribution-Free Property Testing. Now we define distribution-free property testing:

DEFINITION 2.1. We say an algorithm A has oracle access to a pair $(f, \mathcal{D})$, where $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is an unknown Boolean function and $\mathcal{D}$ is an unknown probability distribution over $\{0, 1\}^n$, if it can (1) access f via a black-box oracle that returns $f(x)$ when a string $x \in \{0, 1\}^n$ is queried, and (2) access $\mathcal{D}$ via a sampling oracle that, upon each request, returns a pair $(x, f(x))$ where $x \leftarrow \mathcal{D}$ independently.

Let $\mathcal{C}$ be a class of Boolean functions. A distribution-free testing algorithm A for $\mathcal{C}$ is a randomized algorithm that, given as input a distance parameter $\epsilon > 0$ and oracle access to a pair $(f, \mathcal{D})$, accepts with probability at least $2/3$ if $f \in \mathcal{C}$ and rejects with probability at least $2/3$ if f is ϵ -far from $\mathcal{C}$ with respect to $\mathcal{D}$. We say A is one-sided if it always accepts when $f \in \mathcal{C}$. The query complexity of a distribution-free testing algorithm is the number of queries made on f plus the number of samples drawn from $\mathcal{D}$.

We often use the term “block” to refer to a nonempty subset of $[n]$, which should be interpreted as a nonempty subset of the n variables of a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$. The following

definition of distinguishing pairs and relevant blocks will be heavily used in our algorithms.

DEFINITION 2.2 (DISTINGUISHING PAIRS AND RELEVANT BLOCKS). Given two strings $x, y \in \{0, 1\}^n$ and a block $B \subseteq [n]$, we say that (x, y) is a distinguishing pair for B if $x_{\bar{B}} = y_{\bar{B}}$ and $f(x) \neq f(y)$. We say B is a relevant block of f if such a distinguishing pair exists for B (or equivalently, the influence of B in f is positive).

When $\{i\}$ is a relevant block we simply say that the i -th variable is relevant to f .

As will become clear later, all our algorithms reject a Boolean function f only when they have found $k + 1$ pairwise disjoint blocks $B_1, \dots, B_{k+1}$ and a distinguishing pair for each B_i . When this occurs, it means that $B_1, \dots, B_{k+1}$ are pairwise disjoint relevant blocks of f , which implies that f cannot be a k -junta. As a result, our algorithms are one-sided. To prove their correctness, it suffices to show that they reject with probability at least $2/3$ when f is ϵ -far from k -juntas with respect to $\mathcal{D}$.

For the standard property testing model under the uniform distribution, Blais [7] obtained a nearly optimal algorithm:

THEOREM 2.3 ([7]). There exists a one-sided, $O((k/\epsilon) + k \log k)$ -query algorithm **UniformJunta** (f, k, ϵ) that rejects f with probability at least $2/3$ when it is ϵ -far from k -juntas under the uniform distribution. Moreover, it rejects only when it has found $k + 1$ pairwise disjoint blocks and a distinguishing pair of f for each of them.

Binary Search. The standard binary search procedure (see Figure 1) takes as input two strings $x, y \in \{0, 1\}^n$ with $f(x) \neq f(y)$, makes $O(\log n)$ queries on f , and returns a pair of strings $x', y' \in \{0, 1\}^n$ with $f(x') \neq f(y')$ and $x' = y'^{(i)}$ for some $i \in \text{diff}(x, y)$, i.e., a distinguishing pair for the i -th variable for some $i \in \text{diff}(x, y)$.

However, we cannot afford to use the standard binary search procedure directly in our main algorithm due to its query complexity of $O(\log n)$; recall our goal is to have the query complexity depend on k only. Instead we will employ a blockwise version of the binary search procedure, as described in Figure 2. It takes as input two strings $x, y \in \{0, 1\}^n$ with $f(x) \neq f(y)$ as well as a sequence of pairwise disjoint blocks $B_1, \dots, B_r$ such that $\text{diff}(x, y) \subseteq B_1 \cup \dots \cup B_r$ (i.e., (x, y) is a distinguishing pair for $B_1 \cup \dots \cup B_r$), makes $O(\log r)$ queries on f , and returns two strings $x', y' \in \{0, 1\}^n$ satisfying $f(x') \neq f(y')$ and $\text{diff}(x', y') \subseteq B_i$ for some $i \in [r]$ (i.e., (x', y') is a distinguishing pair for one of the blocks B_i in the input).

3 WARMUP: A TESTER WITH $O((k/\epsilon) + k \log n)$ QUERIES

As a warmup, we first present a simple, one-sided distribution-free algorithm for testing k -juntas (**SimpleDJunta**, where the capital letter D is a shorthand for distribution-free). It uses $O((k/\epsilon) + k \log n)$ queries, where n as usual denotes the number of variables of the function being tested. The idea behind **SimpleDJunta** and its analysis (Lemma 3.2) will be useful in the next section where we present our main algorithm to remove the dependency on n .

The algorithm **SimpleDJunta** maintains a set $I \subseteq [n]$ which is such that a distinguishing pair has been found for each $i \in I$ (i.e., I is a set of relevant variables of f discovered so far). The algorithm sets $I = \emptyset$ at the beginning and rejects only when $|I|$

Algorithm SimpleDJunta($f, \mathcal{D}, k, \epsilon$)

Input: Oracle access to a Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and a probability distribution $\mathcal{D}$ over $\{0, 1\}^n$, a positive integer k , and a distance parameter $\epsilon > 0$.

Output: Either “accept” or “reject.”

- (1) Set $I = \emptyset$.
- (2) Repeat $8(k+1)/\epsilon$ times:
 - (3) Sample $\mathbf{x} \leftarrow \mathcal{D}$ and a subset $\mathbf{R}$ of $\bar{I}$ uniformly at random. Set $\mathbf{y} = \mathbf{x}^{(\mathbf{R})}$.
 - (4) If $f(\mathbf{x}) \neq f(\mathbf{y})$, then run the standard binary search on $\mathbf{x}, \mathbf{y}$ to find a distinguishing pair for a new relevant variable $i \in \mathbf{R} \subseteq \bar{I}$. Set $I = I \cup \{i\}$.
 - (5) If $|I| > k$, then halt and output “reject.”
 - (6) Halt and output “accept.”

Figure 4: Description of the distribution-free testing algorithm SimpleDJunta for k -juntas.

reaches $k+1$, which implies immediately that the algorithm is one-sided. **SimpleDJunta** proceeds round by round. In each round it draws a pair of random strings $\mathbf{x}$ and $\mathbf{y}$ with $\mathbf{x}_I = \mathbf{y}_I$. If $f(\mathbf{x}) \neq f(\mathbf{y})$, the standard binary search procedure is used on $\mathbf{x}$ and $\mathbf{y}$ to find a distinguishing pair for a new variable $i \in \bar{I}$, which is then added to I . The detailed description of the algorithm is given in Figure 4.

The following theorem establishes correctness of the algorithm.

THEOREM 3.1. (i) **SimpleDJunta** makes $O((k/\epsilon) + k \log n)$ queries and always accepts when f is a k -junta. (ii) It rejects with probability at least $2/3$ if f is ϵ -far from k -juntas with respect to $\mathcal{D}$.

PROOF. For part (i) of the theorem, note that the algorithm only runs binary search (and spends $O(\log n)$ queries) when $f(\mathbf{x}) \neq f(\mathbf{y})$ and this happens at most $k+1$ times (even though the algorithm has $O(k/\epsilon)$ rounds). The rest of part (i) is immediate from the description of the algorithm.

For part (ii), it suffices to show that when $|I| \leq k$ at the beginning of a round, a new relevant variable is discovered in this round with at least a modestly large probability. For this purpose we use the following simple but crucial lemma and note the fact that $\mathbf{x}$ and $\mathbf{y}$ on line 3 can be equivalently drawn by first sampling $\mathbf{x} \leftarrow \mathcal{D}$ and $\mathbf{w} \leftarrow \{0, 1\}^n$ and then setting $\mathbf{y} = \mathbf{x}_I \circ \mathbf{w}_{\bar{I}}$ (the way we draw $\mathbf{x}$ and $\mathbf{y}$ in Figure 4 via $\mathbf{R} \leftarrow \bar{I}$ makes it easier to connect with the main algorithm in the next section).

LEMMA 3.2. If f is ϵ -far from k -juntas with respect to $\mathcal{D}$, then for any $I \subset [n]$ of size at most k , we have

$$\Pr_{\mathbf{x} \leftarrow \mathcal{D}, \mathbf{w} \leftarrow \{0, 1\}^n} [f(\mathbf{x}) \neq f(\mathbf{x}_I \circ \mathbf{w}_{\bar{I}})] \geq \epsilon/2. \quad (1)$$

Before proving Lemma 3.2, we use it to finish the proof of part (ii). Assuming Lemma 3.2 and that f is ϵ -far from k -juntas with respect to $\mathcal{D}$, for each round in which $|I| \leq k$ the algorithm finds a new relevant variable with probability at least $\epsilon/2$. By a coupling argument, the probability that the algorithm rejects f (i.e., $|I|$ reaches $k+1$ during the $8(k+1)/\epsilon$ rounds) is at least the probability that

$$\sum_{i=1}^{8(k+1)/\epsilon} \mathbf{Z}_i \geq k+1,$$

where $\mathbf{Z}_i$'s are i.i.d. $\{0, 1\}$ -variables that are 1 with probability $\epsilon/2$. It follows from the Chernoff bound that the latter probability is at least $2/3$. This finishes the proof of the theorem. $\square$

PROOF OF LEMMA 3.2. Let I be a subset of $[n]$ of size at most k . To prove (1) for I , we use I to define the following Boolean function $h: \{0, 1\}^n \rightarrow \{0, 1\}$ over n variables: for each $\mathbf{x} \in \{0, 1\}^n$ we set

$$h(\mathbf{x}) := \arg \max_{b \in \{0, 1\}} \left\{ \Pr_{\mathbf{w} \leftarrow \{0, 1\}^n} [f(\mathbf{x}_I \circ \mathbf{w}_{\bar{I}}) = b] \right\},$$

where we break ties arbitrarily. Then for any $\mathbf{x} \in \{0, 1\}^n$, we have

$$\Pr_{\mathbf{w} \leftarrow \{0, 1\}^n} [f(\mathbf{x}_I \circ \mathbf{w}_{\bar{I}}) = h(\mathbf{x})] \geq 1/2. \quad (2)$$

Furthermore, we have

$$\begin{aligned} & \Pr_{\mathbf{x} \leftarrow \mathcal{D}, \mathbf{w} \leftarrow \{0, 1\}^n} [f(\mathbf{x}) \neq f(\mathbf{x}_I \circ \mathbf{w}_{\bar{I}})] \\ &= \sum_{z \in \{0, 1\}^n} \Pr_{\mathbf{x} \leftarrow \mathcal{D}} [\mathbf{x} = z] \cdot \Pr_{\mathbf{w} \leftarrow \{0, 1\}^n} [f(z) \neq f(z_I \circ \mathbf{w}_{\bar{I}})] \\ &\geq \sum_{z \in \{0, 1\}^n} \Pr_{\mathbf{x} \leftarrow \mathcal{D}} [\mathbf{x} = z] \cdot \left((1/2) \cdot \mathbf{1}[f(z) \neq h(z)] \right) \\ &= (1/2) \cdot \Pr_{\mathbf{x} \leftarrow \mathcal{D}} [f(\mathbf{x}) \neq h(\mathbf{x})] \geq \epsilon/2, \end{aligned}$$

where the first inequality follows from (2) and the second inequality follows from the assumption that f is ϵ -far from every k -junta with respect to $\mathcal{D}$ and the fact that h is a k -junta (since it only depends on variables in I and $|I| \leq k$). This finishes the proof of the lemma. $\square$

4 PROOF OF THEOREM 1.1: A TESTER WITH $\tilde{O}(k^2)/\epsilon$ QUERIES

In this section, we present our main $\tilde{O}(k^2)/\epsilon$ -query algorithm for the distribution-free testing of k -juntas. We start with some intuition behind the algorithm.

4.1 Intuition

Recall that the factor of $\log n$ in the query complexity of **SimpleDJunta** from the previous section is due to the use of the standard binary search procedure. To avoid it, one could choose to terminate each call to binary search early but this ends up giving us relevant

blocks of variables instead of relevant variables. To highlight the challenge, imagine that the algorithm has found so far $\ell \leq k$ many pairwise disjoint relevant blocks B_j , $j \in [\ell]$, i.e., it has found a distinguishing pair for each block B_j . By definition, each B_j must contain at least one relevant variable $i_j \in B_j$. However, we do not know exactly which variable in B_j is i_j , and thus it is not clear how to draw a set $\mathbf{R}$ from $\bar{I}$ uniformly at random, where $I = \{i_j : j \in [\ell]\}$, as on line 3 of **SimpleDJunta**, in order to apply Lemma 3.2 to discover a new relevant block. It seems that we are facing a dilemma when trying to improve **SimpleDJunta** to remove the $\log n$ factor: unless we pin down a set of relevant variables, it is not clear how to draw a random set from their complement, but pinning down a relevant variable using the standard binary search would already cost $\log n$ queries.

To explain the main idea behind our $\tilde{O}(k^2)/\epsilon$ -query algorithm, let's assume again that $\ell \leq k$ many disjoint relevant blocks B_j have been found so far, with a distinguishing pair $(x^{[j]}, y^{[j]})$ for each B_j (satisfying that $\text{diff}(x^{[j]}, y^{[j]}) \subseteq B_j$ and $f(x^{[j]}) \neq f(y^{[j]})$ by definition). Let

$$w^{[j]} = (x^{[j]})_{\overline{B_j}} = (y^{[j]})_{\overline{B_j}} \in \{0, 1\}^{\overline{B_j}}.$$

Next let us assume further that the function $g_j := f|_{w^{[j]}}$, for each $j \in [\ell]$, is a *literal*, i.e. either $g_j(z) = z_{i_j}$ for all $z \in \{0, 1\}^{B_j}$ or $g_j(z) = \bar{z}_{i_j}$ for all $z \in \{0, 1\}^{B_j}$, for some variable $i_j \in B_j$, but the variable i_j is of course unknown to the algorithm. (While this may seem very implausible, we make this assumption for now and explain below why it is not too far from real situations.)

To make progress, we draw a random two-way partition of each B_j into P_j and Q_j , i.e., each $i \in B_j$ is added to either P_j or Q_j with probability $1/2$ (so they are disjoint and $B_j = P_j \cup Q_j$). We make three simple but crucial observations to increase the number of disjoint relevant blocks by one.

- (1) Since g_j is assumed to be a literal on the i_j -th variable (and by the definition of g_j we have query access to g_j), it is easy to tell whether $i_j \in P_j$ or $i_j \in Q_j$, by picking an arbitrary string $x \in \{0, 1\}^{B_j}$ and then comparing $g_j(x)$ with $g_j(x^{(P_j)})$. Below we assume that the algorithm correctly determines whether i_j is in P_j or Q_j for all $j \in [\ell]$. We let S_j denote the element of $\{P_j, Q_j\}$ that contains i_j and let T_j denote the other one. We also assume below that the algorithm has obtained a distinguishing pair of g_j for each block S_j .
- (2) Next we draw a subset $\mathbf{T}$ of $\overline{B_1 \cup \dots \cup B_\ell}$ uniformly at random. Crucially, the way that P_j and Q_j were drawn, and the above assumption that S_j contains i_j , implies that $\mathbf{R} := \mathbf{T} \cup \mathbf{T}_1 \cup \dots \cup \mathbf{T}_\ell$ is indeed a subset of $\bar{I}$ drawn uniformly at random (recall that $I = \{i_j : j \in [\ell]\}$) since other than those in I , each variable is included in $\mathbf{R}$ independently with probability $1/2$. If we draw a random string $\mathbf{x} \leftarrow \mathcal{D}$, then Lemma 3.2 implies that $f(\mathbf{x}) \neq f(\mathbf{y})$, where $\mathbf{y} = \mathbf{x}^{(\mathbf{R})}$, with probability at least $\epsilon/2$.
- (3) Finally, assuming that $f(\mathbf{x}) \neq f(\mathbf{y})$ (with $\text{diff}(\mathbf{x}, \mathbf{y}) = \mathbf{R}$), running the blockwise binary search on $\mathbf{x}, \mathbf{y}$ and blocks $\mathbf{T}, \mathbf{T}_1, \dots, \mathbf{T}_\ell$ will lead to a distinguishing pair for one of these blocks and will only require $O(\log \ell) \leq O(\log k)$ queries. If it is a distinguishing pair for $\mathbf{T}$, then we can add

$\mathbf{T}$ to the list of relevant blocks $B_1, \dots, B_\ell$ and they remain pairwise disjoint. If it is $\mathbf{T}_j$ for some $j \in [\ell]$, then we can replace B_j in the list by S_j and $\mathbf{T}_j$, for each of which we have found a distinguishing pair (recall that a distinguishing pair has already been found for each S_j in the first step). In either case we have that the number of pairwise disjoint relevant blocks grows by one.

Coming back to the assumption we made earlier, although g_j is very unlikely to be a literal, it must fall into one of the following three cases: (1) close to a literal; (2) close to a (all-0 or all-1) constant function; or (3) far from 1-juntas. Here in all cases “close” and “far” means with respect to the *uniform distribution* over $\{0, 1\}^{B_j}$. As we discuss in more detail in the rest of the section, with some more careful probability analysis the above arguments generalize to the case in which every g_j is only close to (rather than exactly) a literal. On the other hand, if one of the blocks B_j is in case (2) or (3), then (using the fact that we have a distinguishing pair for B_j) it is easy to split B_j into two blocks and find a distinguishing pair for each of them. (For example, for case (3) this can be done by running Blais’s uniform distribution junta testing algorithm.) As a result, we can make progress by increasing the number of pairwise disjoint relevant blocks by one. Our algorithm basically keep repeating these steps until the number of such blocks reaches $k + 1$.

4.2 Description of the Main Algorithm and the Proof of Correctness

Our main algorithm **MainDJunta**($f, \mathcal{D}, k, \epsilon$) is described in Figure 5. It maintains two collections of blocks $V = \{B_1, \dots, B_v\}$ (V for “verified”) and $U = \{C_1, \dots, C_u\}$ (U for “unverified”) for some nonnegative integers v and u . They are set to be $\emptyset$ at initialization and always satisfy:

- (A). $B_1, \dots, B_v, C_1, \dots, C_u \subseteq [n]$ are pairwise disjoint (nonempty) blocks of variables; and
- (B). A distinguishing pair has been found for each of these blocks. For notational convenience we use $(x^{[j]}, y^{[j]})$ to denote the distinguishing pair for each B_j and $(x^{[C]}, y^{[C]})$ to denote the distinguishing pair for each block $C \in U$. We also use the notation

$$w^{[j]} := (x^{[j]})_{\overline{B_j}} = (y^{[j]})_{\overline{B_j}} \in \{0, 1\}^{\overline{B_j}} \quad \text{and}$$

$$w^{[C]} := (x^{[C]})_{\overline{C}} = (y^{[C]})_{\overline{C}} \in \{0, 1\}^{\overline{C}},$$

and we let $g_j := f|_{w^{[j]}}$ and $g_C := f|_{w^{[C]}}$, functions over $\{0, 1\}^{B_j}$ and $\{0, 1\}^C$, respectively.

The algorithm rejects only when the total number of blocks $v + u \geq k + 1$ so it is one-sided.

Throughout the algorithm and its analysis, we set a key parameter $\gamma := 1/(8k)$. Blocks in V are intended to be those that have been “verified” to satisfy the condition that g_j is γ -close to a literal (for some unknown variable $i_j \in B_j$) under the uniform distribution, while blocks in U have not been verified yet so they may or may not satisfy the condition. More formally, at any point in the execution of the algorithm we say that the algorithm is in *good condition* if its

Algorithm MainDJunta($f, \mathcal{D}, k, \epsilon$) with the same input / output as **SimpleDJunta** in Figure 4.

- (1) Initialization: Set $V = U = \emptyset$, $r_1 = 64k/\epsilon$ and $r_2 = 3(k+1)$.
- (2) While $r_1 > 0$ and $r_2 > 0$ do (letting $V = \{B_1, \dots, B_v\}$ and $U = \{C_1, \dots, C_u\}$)
- (3) If $u = 0$, then
- (4) Set r_1 to be $r_1 - 1$.
- (5) For $j = 1$ to v do $((x^{[j]}, y^{[j]}):$ distinguishing pair for B_j , $w^{[j]} = (x^{[j]})_{\overline{B_j}}$ and $g_j = f|_{w^{[j]}})$
- (6) Draw a random partition P_j, Q_j of B_j and run **WhereIsTheLiteral**(g_j, P_j, Q_j).
- (7) If it returns a distinguishing pair of g_j for P_j , set $S_j = P_j$ and $T_j = Q_j$;
- (8) Else if it returns a distinguishing pair of g_j for Q_j , set $S_j = Q_j$ and $T_j = P_j$;
- (9) Else (it returns “fail”), skip this round and go back to line 2.
- (10) Draw $\mathbf{x} \leftarrow \mathcal{D}$ and a subset $\mathbf{T}$ of $B_1 \cup \dots \cup B_v$ uniformly at random.
- (11) If $f(\mathbf{x}) \neq f(\mathbf{y})$, where $\mathbf{y} = \mathbf{x}^{(\mathbf{R})}$ with $\mathbf{R} = \mathbf{T} \cup \mathbf{T}_1 \cup \dots \cup \mathbf{T}_v$, then
- (12) Run the blockwise binary search on $\mathbf{x}$ and $\mathbf{y}$ with blocks $\mathbf{T}, \mathbf{T}_1, \dots, \mathbf{T}_v$.
- (13) If a distinguishing pair of f for $\mathbf{T}$ is found, add $\mathbf{T}$ to U .
- (14) Else (a distinguishing pair of f for $\mathbf{T}_{j^*}$, for some $j^* \in [v]$, is found)
- (15) Concatenate $w^{[j^*]}$ to the distinguishing pair of g_{j^*} for S_{j^*} found on line 7-8.
- (16) This gives us a distinguishing pair of f for S_{j^*} .
- (17) Remove B_{j^*} from V and add both S_{j^*} and T_{j^*} to U .
- (18) Else (i.e., $u > 0$)
- (19) Set r_2 to be $r_2 - 1$.
- (20) Pick a $C \in U$ arbitrarily; let (x, y) be its distinguishing pair, $w = x_{\overline{C}}$ and $g = f|_w$.
- (21) If **Literal**(g) returns “true,” remove C from U and add it to V .
- (22) Else (it returns disjoint subsets C', C^* of C and each a distinguishing pair of g_C)
- (23) Concatenate w to obtain a distinguishing pair of f for each of C' and C^*
- (24) Remove C from U and add both C' and C^* to U .
- (25) If $|V| + |U| \geq k + 1$, then halt and output “reject.”
- (26) Halt and output “accept.”

Figure 5: Description of the distribution-free testing algorithm MainDJunta for k -juntas.

current collections V and U satisfy conditions (A), (B) and (C): Every $g_j, j \in [v]$, is γ -close to a literal under the uniform distribution over $\{0, 1\}^{B_j}$.

The algorithm **MainDJunta**(f, k, ϵ) starts with $V = U = \emptyset$ and proceeds round by round. For each round, we consider two different types that the round may have: type 1 is that $u = 0$, and type 2 is that $u > 0$. In a type-1 round (with $u = 0$) we follow the idea sketched in Section 4.1 to increase the total number of disjoint relevant blocks by one. We prove the following lemma for this case in Section 4.3.

LEMMA 4.1. *Assume that MainDJunta is in good condition at the beginning of a round, with $u = 0$ and $v \leq k$. Then it must remain in good condition at the end of this round. Moreover, letting V' and U' be the two collections of blocks at the end of this round, we have either $|V'| = v$ and $|U'| = 1$, or $|V'| = v - 1$ and $|U'| = 2$ with probability at least $\epsilon/4$.*

In a type-2 round (with $u \geq 1$), we pick an arbitrary block C from U and check whether g_C is close to a literal under the uniform distribution. The following lemma, which we prove in Section 4.4, shows that with high probability, either C is moved to collection

V and the algorithm remains in good condition, or the algorithm finds two disjoint subsets of C and a distinguishing pair for each of them so that V stays the same but $|U|$ goes up by one (we add these two blocks to U since they have not yet been verified).

LEMMA 4.2. *Assume that MainDJunta is in good condition at the beginning of a round, with $u > 0$ and $v + u \leq k$. Then with probability at least $1 - 1/(64k)$, one of the following two events occurs at the end of this round (letting V' and U' be the two collections of blocks at the end of this round):*

- (1) *The algorithm remains in good condition with $|V'| = |V| + 1$ and $|U'| = |U| - 1$; or*
- (2) *The algorithm remains in good condition with $V' = V$ and $|U'| = |U| + 1$.*

Assuming Lemma 4.1 and Lemma 4.2, we are ready to prove the correctness of **MainDJunta**.

THEOREM 4.3. (i) **MainDJunta** makes $\tilde{O}(k^2)/\epsilon$ many queries and always accepts f when it is a k -junta. (ii) It rejects with probability at least $2/3$ when f is ϵ -far from every k -junta with respect to $\mathcal{D}$.

PROOF OF THEOREM 4.3 ASSUMING LEMMA 4.1 AND LEMMA 4.2.

MainDJunta is one-sided since it rejects f only when it has found $k + 1$ pairwise disjoint relevant blocks of f . Its query complexity is

$$\begin{aligned} & (\# \text{ type-1 rounds}) \cdot (\# \text{ queries per type-1 round}) \\ & + (\# \text{ type-2 rounds}) \cdot (\# \text{ queries per type-2 round}) \\ & = O(k/\epsilon) \cdot (O(k) + O(\log k)) + O(k) \cdot O(\log k) \cdot O(k) \\ & = O(k^2/\epsilon) + O(k^2 \log k) = O(k^2 \log k)/\epsilon. \end{aligned}$$

In the rest of the proof we show that it rejects f with probability at least $2/3$ when f is ϵ -far from every k -junta with respect to $\mathcal{D}$.

For this purpose we introduce a simple potential function

$$F(V, U) := 3|V| + 2|U|$$

to measure the progress: Each round of the algorithm is either of type-1 (when $|U| = 0$) or of type-2 (when $|U| > 0$). By Lemma 4.1, if the algorithm is in good condition at the beginning of a type-1 round, then the algorithm ends the round in good condition and the potential function F goes up by at least one with probability at least $\epsilon/4$ (in which case we say that the algorithm succeeds in this type-1 round). By Lemma 4.2, if the algorithm is in good condition at the beginning of a type-2 round, then the algorithm ends the round in good condition and F goes up by at least one with probability at least $1 - 1/(32k)$ (in which case we say it succeeds in this type-2 round).

Note that F is 0 at the beginning ($V = U = \emptyset$) and that we must have $|U| + |V| \geq k + 1$ (and thus, the algorithm rejects) when the potential function F reaches $3(k + 1)$ or above. As a result, a necessary condition for the algorithm to accept is that one of the following two events occurs:

E_1 : At least one of the (no more than $3(k + 1)$ many) type-2 rounds fails.

E_2 : E_1 does not occur (so every round ends in good condition, and the reason that the algorithm accepts cannot be that it uses up all the $3(k + 1)$ many type-2 rounds), and the algorithm uses up all the $64k/\epsilon$ many type-1 rounds but at most $3k + 2$ of them succeed.

By a union bound the probability of E_1 is at most

$$3(k + 1) \cdot 1/(64k) \leq 6k \cdot 1/(64k) < 1/8.$$

As the algorithm ends every round in good condition, it follows from Lemma 4.1 and a coupling argument that the probability of E_2 is at most the probability that

$$\sum_{i=1}^{64k/\epsilon} Z_i \leq 3k + 2,$$

where Z_i 's are i.i.d. $\{0, 1\}$ -valued random variables that take 1 with probability $\epsilon/4$. It follows from the Chernoff bound the probability is at most (using $3k + 2 \leq 5k$)

$$\exp\left(-\left(\frac{11}{16}\right)^2 \cdot \frac{16k}{2}\right) = \exp\left(-\frac{121k}{32}\right) < \exp(-3) < 1/8.$$

Finally it follows from a union bound that the algorithm rejects with probability at least $3/4$. $\square$

4.3 Proof of Lemma 4.1

We start with a lemma for the subroutine **WhereIsTheLiteral** in **MainDJunta**, which is described in Figure 6.

LEMMA 4.4. Assume that a Boolean function $g : \{0, 1\}^B \rightarrow \{0, 1\}$ is γ -close (with respect to the uniform distribution) to a literal x_i or $\bar{x}_i$ for some $i \in B$. If $i \in P$, then **WhereIsTheLiteral**(g, P, Q) returns a distinguishing pair of g for P with probability at least $1 - 4\gamma$; If $i \in Q$, then it returns a distinguishing pair of g for Q with probability at least $1 - 4\gamma$.

PROOF. Let K be the set of strings $x \in \{0, 1\}^B$ such that $g(x)$ disagrees with the literal to which it is γ -close (so $|K| \leq \gamma \cdot 2^{|B|}$). We work on the case when $i \in Q$; the case when $i \in P$ is similar.

By the description of **WhereIsTheLiteral**, it returns a distinguishing pair for Q if

$$g(w \circ z) = g(w \circ z^{(P)}) \quad \text{and} \quad g(w' \circ z') \neq g(w' \circ z'^{(Q)}).$$

Note that this holds if all four strings fall outside of K and thus, the probability that it does not hold is at most the probability that at least one of these four strings falls inside K . The latter by a union bound is at most 4γ since each of these four strings is drawn uniformly at random from $\{0, 1\}^B$ by itself. This finishes the proof of the lemma. $\square$

We are now ready to prove Lemma 4.1.

PROOF OF LEMMA 4.1. First, it is easy to verify that if the algorithm starts a round in good condition, then it ends it in good condition. This is because whenever a block is added to U , it is disjoint from other blocks and we have found a distinguishing pair for it.

Next it follows directly from Lemma 4.4 and a union bound that, for any sequence of partitions P_j and Q_j of B_j picked on line 6, the probability that the for-loop correctly sets S_j to be the one that contains the special variable i_j for all $j \in [v]$ is at least (recalling that $\gamma = 1/(8k)$)

$$1 - 4\gamma \cdot v \geq 1 - 4\gamma \cdot k = 1/2.$$

Now we can view the process equivalently as follows. First we draw $\mathbf{x} \leftarrow \mathcal{D}$, $\mathbf{T} \leftarrow B_1 \cup \dots \cup B_v$, and random partitions P_j, Q_j of each B_j . If we let $\mathbf{T}_j^*$ be the set in P_j, Q_j that does not contain the special variable, then we have that

$$\mathbf{R}^* = \mathbf{T} \cup \mathbf{T}_1^* \cup \dots \cup \mathbf{T}_v^*$$

is a set drawn uniformly at random from $\bar{I}$, where $I = \{i_j : j \in [v]\}$ consists of the special variables. Therefore, it follows from Lemma 3.2 that $f(\mathbf{x}) \neq f(\mathbf{x}^{(\mathbf{R}^*)})$ with probability at least $\epsilon/2$. Since with probability at least $1/2$, the set $\mathbf{R}$ on line 11 agrees with $\mathbf{R}^*$, we have that the algorithm reaches line 12 with $f(\mathbf{x}) \neq f(\mathbf{y})$ with probability at least $\epsilon/4$. Given this, the lemma is immediate by inspection of lines 12-17 of the algorithm. $\square$

4.4 Proof of Lemma 4.2

First it follows from the description of the subroutine **Literal**(g) that it either returns “true” or a pair of nonempty disjoint subsets C', C^* of C and a distinguishing pair of g for each of them (see Theorem 2.3). Next, let $C \in V$ be the block picked in line 20. If g is γ -close to a literal, then it is easy to verify that one of the two

Subroutine WhereIsTheLiteral(g, P, Q)**Input:** Oracle access to a Boolean function g over $\{0, 1\}^B$ with P, Q being a partition of B .**Output:** Either a distinguishing pair for P , a distinguishing pair for Q , or “fail.”

- (1) Draw $\mathbf{w} \leftarrow \{0, 1\}^Q$ and $\mathbf{z} \leftarrow \{0, 1\}^P$ independently and uniformly at random.
- (2) If $g(\mathbf{w} \circ \mathbf{z}) \neq g(\mathbf{w} \circ \mathbf{z}^{(P)})$, return $(\mathbf{w} \circ \mathbf{z}, \mathbf{w} \circ \mathbf{z}^{(P)})$ as a distinguishing pair for P .
- (3) Draw $\mathbf{w}' \leftarrow \{0, 1\}^P$ and $\mathbf{z}' \leftarrow \{0, 1\}^Q$ independently and uniformly at random.
- (4) If $g(\mathbf{w}' \circ \mathbf{z}') \neq g(\mathbf{w}' \circ \mathbf{z}'^{(Q)})$, return $(\mathbf{w}' \circ \mathbf{z}', \mathbf{w}' \circ \mathbf{z}'^{(Q)})$ as a distinguishing pair for Q .
- (5) Return “fail.”

Figure 6: Description of the subroutine WhereIsTheLiteral.**Subroutine Literal**(g)**Input:** Oracle access to a Boolean function g over $\{0, 1\}^C$ where C has a distinguishing pair.**Output:** “True” or disjoint nonempty subsets C', C^* of C and a distinguishing pair for each.

- (1) Repeat $\log k + 6$ times:
 - (2) If **UniformJunta**($g, 1, \gamma$) rejects, then
 - (3) Return the two disjoint blocks it has found and a distinguishing pair for each.
- (4) Let (x, y) be the distinguishing pair for C .
- (5) Repeat $\log k + 3$ times:
 - (6) Draw a random partition C', C^* of C and query $g(x^{(C')})$, $g(x^{(C^*)})$, $g(y^{(C')})$, $g(y^{(C^*)})$.
 - (7) If $g(x^{(C')}) = g(x^{(C^*)}) \neq g(x)$, then
 - (8) Return C', C^* and $(x, x^{(C')})$ and $(x, x^{(C^*)})$ as their distinguishing pairs.
 - (9) If $g(y^{(C')}) = g(y^{(C^*)}) \neq g(y)$, then
 - (10) Return C', C^* and $(y, y^{(C')})$ and $(y, y^{(C^*)})$ as their distinguishing pairs.
- (11) Return “true.”

Figure 7: Description of the subroutine Literal.

events described in Lemma 4.2 must hold (using the property of **Literal**(g) above). So we focus on the other two cases in the rest of the proof: g is γ -far from 1-juntas or g is γ -close to a (all-1 or all-0) constant function. In both cases we show below that the second event happens with high probability.

When g is γ -far from 1-juntas under the uniform distribution, we have that one of the $\log k + 6$ calls to **UniformJunta** in **Literal** rejects with probability at least

$$1 - (1/3)^{\log k + 6} > 1 - 1/(64k).$$

The second event in Lemma 4.2 occurs when this happens.

When g is γ -close to a constant function (say the all-1 function), we have that either string x or y in the distinguishing pair for C disagrees with the function (say $g(x) = 0$, since $g(x) \neq g(y)$). Let K be the set of strings in $\{0, 1\}^C$ that disagree with the all-1 function. Then line 7 of **Literal**(g) does not hold only when one of $x^{(C')}$ or $x^{(C^*)}$ lies in K . As both strings are distributed uniformly over $\{0, 1\}^C$ by themselves, this happens with probability at most 2γ by a union bound. Thus, the probability that line 7 holds at least once is at least

$$\begin{aligned} 1 - (2\gamma)^{\log k + 3} &= 1 - (1/(4k))^{\log k + 3} \\ &> 1 - (1/4)^{\log k + 3} = 1 - 1/(64k^2). \end{aligned}$$

Therefore the second event in Lemma 4.2 occurs with probability at least $1 - 1/(64k^2)$.

This finishes the proof of Lemma 4.2.

REFERENCES

- [1] N. Alon, T. Kaufman, M. Krivelevich, S. Litsyn, and D. Ron. 2005. Testing Reed-Muller Codes. *IEEE Transactions on Information Theory* 51, 11 (2005), 4032–4039.
- [2] Noga Alon and Amit Weinstein. 2012. Local correction of juntas. *Inf. Process. Lett.* 112, 6 (2012), 223–226.
- [3] Roksana Baleshazar, Meiram Murzabulatov, Ramesh Krishnan S. Pallavoor, and Sofya Raskhodnikova. 2016. Testing Unateness of Real-Valued Functions. *CoRR* abs/1608.07652 (2016).
- [4] A. Belovs and E. Blais. 2016. A Polynomial Lower Bound for Testing Monotonicity. In *Proceedings of the 48th ACM Symposium on Theory of Computing*. 1021–1032.
- [5] Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. 2010. Optimal Testing of Reed-Muller Codes. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*. 488–497.
- [6] Eric Blais. 2008. Improved bounds for testing juntas. In *Proc. RANDOM*. 317–330.
- [7] Eric Blais. 2009. Testing juntas nearly optimally. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC)*. 151–158.
- [8] Eric Blais, Joshua Brody, and Kevin Matulef. 2011. Property Testing Lower Bounds via Communication Complexity. In *CCC*. 210–220.
- [9] Eric Blais and Daniel M. Kane. 2012. Tight Bounds for Testing k -Linearity. In *RANDOM*. 435–446.
- [10] M. Blum, M. Luby, and R. Rubinfeld. 1993. Self-testing/correcting with applications to numerical problems. *J. Comput. System Sci.* 47 (1993), 549–595. Earlier version in STOC'90.
- [11] Harry Buhrman, David García-Soriano, Arie Matsliah, and Ronald de Wolf. 2013. The non-adaptive query complexity of testing k -parities. *Chicago Journal of Theoretical Computer Science* 2013 (2013).
- [12] Deeparnab Chakrabarty and C. Seshadhri. 2013. A $o(n)$ monotonicity tester for boolean functions over the hypercube. In *Proceedings of the 45th ACM Symposium on Theory of Computing*. 411–418.
- [13] Deeparnab Chakrabarty and C. Seshadhri. 2016. A $\tilde{O}(n)$ non-adaptive tester for unateness. *CoRR* abs/1608.06980 (2016).
- [14] Xi Chen, Anindya De, Rocco A. Servedio, and Li-Yang Tan. 2015. Boolean function monotonicity testing requires (almost) $n^{1/2}$ non-adaptive queries. In *Proceedings of the 47th ACM Symposium on Theory of Computing*. 519–528.

- [15] Xi Chen, Zhengyang Liu, Rocco A. Servedio, Ying Sheng, and Jinyu Xie. 2018. Distribution-free Junta Testing. *arXiv:1802.04859* (2018).
- [16] X. Chen, R.A. Servedio, and L.-Y. Tan. 2014. New Algorithms and Lower Bounds for Monotonicity Testing. In *Proc. 55th IEEE Symposium on Foundations of Computer Science (FOCS)*. 286–295.
- [17] Xi Chen, Rocco A. Servedio, Li-Yang Tan, Erik Waingarten, and Jinyu Xie. 2017. Settling the Query Complexity of Non-Adaptive Junta Testing. In *32nd Computational Complexity Conference (CCC)*. 26:1–26:19.
- [18] Xi Chen, Erik Waingarten, and Jinyu Xie. 2017. 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 2017*. 523–536.
- [19] Xi Chen, Erik Waingarten, and Jinyu Xie. 2017. Boolean Unateness Testing with $O(n^{3/4})$ Adaptive Queries. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 868–879.
- [20] Xi Chen and Jinyu Xie. 2016. Tight Bounds for the Distribution-Free Testing of Monotone Conjunctions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*. 54–71.
- [21] H. Chockler and D. Gutfreund. 2004. A lower bound for testing juntas. *Inform. Process. Lett.* 90, 6 (2004), 301–305.
- [22] I. Diakonikolas, H. Lee, K. Matulef, K. Onak, R. Rubinfeld, R. Servedio, and A. Wan. 2007. Testing for Concise Representations. In *Proc. 48th Ann. Symposium on Computer Science (FOCS)*. 549–558.
- [23] E. Dolev and D. Ron. 2011. Distribution-Free Testing for Monomials with a Sublinear Number of Queries. *Theory of Computing* 7, 1 (2011), 155–176.
- [24] E. Fischer, G. Kindler, D. Ron, S. Safra, and A. Samorodnitsky. 2004. Testing juntas. *J. Computer & System Sciences* 68, 4 (2004), 753–787.
- [25] E. Fischer, E. Lehman, I. Newman, S. Raskhodnikova, R. Rubinfeld, and A. Samorodnitsky. 2002. Monotonicity Testing Over General Poset Domains. In *Proc. 34th Annual ACM Symposium on the Theory of Computing*. 474–483.
- [26] Dana Glasner and Rocco A. Servedio. 2009. Distribution-Free Testing Lower Bound for Basic Boolean Functions. *Theory of Computing* 5, 1 (2009), 191–216.
- [27] O. Goldreich (Ed.). 2010. *Property Testing: Current Research and Surveys*. Springer. LNCS 6390.
- [28] O. Goldreich, S. Goldwasser, E. Lehman, D. Ron, and A. Samordinsky. 2000. Testing Monotonicity. *Combinatorica* 20, 3 (2000), 301–337.
- [29] O. Goldreich, S. Goldwasser, and D. Ron. 1998. Property testing and its connection to learning and approximation. *J. ACM* 45 (1998), 653–750.
- [30] P. Gopalan, R. O'Donnell, R. Servedio, A. Shpilka, and K. Wimmer. 2011. Testing Fourier dimensionality and sparsity. *SIAM J. on Computing* 40, 4 (2011), 1075–1100.
- [31] S. Halevy and E. Kushilevitz. 2007. Distribution-Free Property Testing. *SIAM J. Comput.* 37, 4 (2007), 1107–1138.
- [32] Subhash Khot, Dor Minzer, and Muli Safra. 2015. On Monotonicity Testing and Boolean Isoperimetric type Theorems. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science*. 52–58.
- [33] Subhash Khot and Igor Shinkar. 2016. An $O(n)$ queries adaptive tester for unateness. In *Approximation, Randomization, and Combinatorial Optimization Algorithms and Techniques*.
- [34] K. Matulef, R. O'Donnell, R. Rubinfeld, and R. Servedio. 2010. Testing Halfspaces. *SIAM J. on Comput.* 39, 5 (2010), 2004–2047.
- [35] Kevin Matulef, Ryan O'Donnell, Ronitt Rubinfeld, and Rocco A. Servedio. 2009. Testing ± 1 -weight halfspace. In *APPROX-RANDOM*. 646–657.
- [36] M. Parnas, D. Ron, and A. Samorodnitsky. 2002. Testing Basic Boolean Formulae. *SIAM J. Disc. Math.* 16 (2002), 20–46.
- [37] D. Ron. 2008. Property Testing: A Learning Theory Perspective. *Foundations and Trends in Machine Learning* 1, 3 (2008), 307–402.
- [38] D. Ron. 2010. Algorithmic and Analysis Techniques in Property Testing. *Foundations and Trends in Theoretical Computer Science* 5 (2010), 73–205. Issue 2.
- [39] Ronitt Rubinfeld and Madhu Sudan. 1996. Robust characterizations of polynomials with applications to program testing. *SIAM J. Comput.* 25, 2 (1996), 252–271.
- [40] Rocco Servedio, Li-Yang Tan, and John Wright. 2015. Adaptivity helps for testing juntas. In *Proceedings of the 30th IEEE Conference on Computational Complexity*. 264–279. volume 33 of LIPIcs.
- [41] L. Valiant. 1984. A theory of the learnable. *Commun. ACM* 27, 11 (1984), 1134–1142.