

MIT Open Access Articles

Junta Correlation is Testable

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: De, Anindya, Mossel, Elchanan and Neeman, Joe. 2019. "Junta Correlation is Testable." Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS, 2019-November.

As Published: 10.1109/FOCS.2019.00090

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/137504>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Junta-Correlation is Testable

Anindya De*
University of Pennsylvania
anindyad@cis.upenn.edu

Elchanan Mossel†
MIT
elmos@mit.edu

Joe Neeman‡
UT Austin
jneeman@math.utexas.edu

April 9, 2019

Abstract

The problem of tolerant junta testing is a natural and challenging problem which asks if the property of a function having some specified correlation with a k -Junta is testable. In this paper we give an affirmative answer to this question: We show that given distance parameters $\frac{1}{2} > c_u > c_\ell \geq 0$, there is a tester which given oracle access to $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, with query complexity $2^k \cdot \text{poly}(k, 1/|c_u - c_\ell|)$ and distinguishes between the following cases:

1. The distance of f from any k -junta is at least c_u ;
2. There is a k -junta g which has distance at most c_ℓ from f .

This is the first non-trivial tester (i.e., query complexity is independent of n) which works for all $1/2 > c_u > c_\ell \geq 0$. The best previously known results by Blais *et al.*, required $c_u \geq 16c_\ell$. In fact, with the same query complexity, we accomplish the stronger goal of identifying the most correlated k -junta, up to permutations of the coordinates.

We can further improve the query complexity to $\text{poly}(k, 1/|c_u - c_\ell|)$ for the (weaker) task of distinguishing between the following cases:

1. The distance of f from any k' -junta is at least c_u .
2. There is a k -junta g which is at a distance at most c_ℓ from f .

Here $k' = O(k^2/|c_u - c_\ell|^2)$.

Our main tools are Fourier analysis based algorithms that simulate oracle access to influential coordinates of functions.

*Supported by NSF grant CCF-1814706

†Supported by ONR grant N00014-16-1-2227 and NSF grant CCF 1320105.

‡Supported by the Alfred P. Sloan Foundation

1 Introduction

Juntas are a fundamental class of functions in Boolean function analysis. A function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ is said to be a k -junta if there are some k -coordinates $i_1, \dots, i_k \in [n]$ such that $f(x)$ only depends on x_{i_1}, \dots, x_{i_k} . In particular, special attention has been devoted to the problem of testing juntas.

We recall that a property testing algorithm for a class of functions \mathcal{C} is an algorithm which given oracle access to an $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and a distance parameter $\epsilon > 0$, satisfies

1. If $f \in \mathcal{C}$, then the algorithm accepts with probability at least $2/3$;
2. If $\text{dist}(f, g) \geq \epsilon$ for every $g \in \mathcal{C}$, then the algorithm rejects with probability at least $2/3$. Here $\text{dist}(f, g) = \Pr_{x \in \{-1, 1\}^n} [f(x) \neq g(x)]$.

The principal measure of the efficiency of the algorithm is its *query complexity*. Also, the precise value of the confidence parameter is irrelevant and $2/3$ can be replaced by any constant $1/2 < c < 1$.

Fischer *et al.* [14] were the first to study the problem of testing k -juntas and showed that k -juntas can be tested with query complexity $\tilde{O}(k^2/\epsilon)$. The crucial feature of their algorithm is that the query complexity is independent of the ambient dimension n . Since then, there has been a long line of work on testing juntas [3, 2, 27, 10, 9] and it continues to be of interest down to the present day. The flagship result here is that k -juntas can be tested with $\tilde{O}(k/\epsilon)$ queries and this is tight [3, 10]. While the initial motivation to study this problem came from long-code testing [1, 25] (related to PCPs and inapproximability), another strong motivation comes from the *feature selection* problem in machine learning (see, e.g. [6, 7]).

Tolerant testing The definition of property tester above requires the algorithm to accept if and only if $f \in \mathcal{C}$. However, for many applications, it is important consider a *noise-tolerant* definition of property testing. In particular, Parnas, Ron and Rubinfeld [24] introduced the following definition of noise tolerant testers.

Definition 1.1. For constants $1/2 > c_u > c_\ell \geq 0$ and a function class \mathcal{C} , a (c_u, c_ℓ) -noise tolerant tester for \mathcal{C} is an algorithm which given oracle access to a function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$

1. accepts with probability at least $2/3$ if $\min_{g \in \mathcal{C}} \text{dist}(f, g) \leq c_\ell$.
2. rejects with probability at least $2/3$ if $\min_{g \in \mathcal{C}} \text{dist}(f, g) \geq c_u$.

We observe that we restrict $c_u, c_\ell < 1/2$. This is because most natural classes \mathcal{C} are closed under complementation – i.e., if $g \in \mathcal{C}$, then $-g \in \mathcal{C}$. For such a class \mathcal{C} and for any f , $\min_{g \in \mathcal{C}} \text{dist}(f, g) \leq 1/2$. Further, note that the standard notion of property testing corresponds to a $(\epsilon, 0)$ -noise tolerant tester.

The problem of testing juntas becomes quite challenging in the presence of noise. Parnas *et al.* [24] observed that any tester whose (individual) queries are uniformly distributed are inherently noise tolerant in a very weak sense. In particular, [13] used this observation to show that the junta tester of [14] is in fact a $(\epsilon, \text{poly}(\epsilon/k))$ -noise tolerant tester for k -juntas – note that c_ℓ is quite small, namely $\text{poly}(\epsilon/k)$. Later, Chakraborty *et al.* [8] showed that the tester of Blais [3] yields a $(C\epsilon, \epsilon)$ tester (for some large but fixed $C > 1$) with query complexity $\exp(k/\epsilon)$. Recently, there has been a surge of interest in tolerant junta testing. On one hand, Levi and Waingarten showed that there are constants $1/2 > \epsilon_1 > \epsilon_2 > 0$ such that any non-adaptive (ϵ_1, ϵ_2) tester requires $\tilde{\Omega}(k^2)$ non-adaptive queries. Contrast this with the result of Blais [3] who showed that there is non-adaptive tester for k -juntas with $O(k^{3/2})$ queries when there is no noise. In particular, this shows a gap between testing in the noisy and noiseless case.

In the opposite (i.e., algorithmic) direction, Blais *et al.* [4] proved the following theorem.

Theorem 1.2. *There is an algorithm which for any $\rho \in (0, 1)$, $\epsilon \in (0, 1)$ and parameter $k \in \mathbb{N}$, is a $(\epsilon, \frac{\rho\epsilon}{16})$ -noise tolerant tester for k -juntas. The query complexity of the tester is $O(\frac{k \log k}{\epsilon\rho(1-\rho)^k})$.*

Note that for any $C > 16$, this yields an $(\epsilon, \epsilon/C)$ -tolerant tester for k -juntas with $\exp(k)$ query complexity. Thus, it improves on the result of [8] who showed the same result for an unspecified large constant C .

To understand the main shortcoming of [4], note that this algorithm does not yield a (c_u, c_ℓ) noise tolerant tester once $c_\ell > \frac{1}{32}$ – e.g, no setting of parameters in the tester of [4] can yield (say) a $(0.1, 0.05)$ noise-tolerant tester for k -juntas. Naturally, one would like to obtain (c_u, c_ℓ) testers for any $1/2 > c_u > c_\ell$. The main result of this paper accomplishes this goal. Below we formalize and state the main results of the paper.

We will use $\mathcal{J}_{n,k}$ to denote the class of k -juntas on n variables. Also, for a subset $S \subseteq [n]$, we let $\mathcal{J}_{S,k}$ denote the class of k -juntas on the variables in S . Further, unless indicated otherwise, all expectations are taken over uniformly random elements of $\{-1, 1\}^n$ where the ambient dimension n will be clear from the context. Our first result constructs (c_u, c_ℓ) testers for any $1/2 > c_u > c_\ell$.

Theorem 1.3. *There is an algorithm **Maximum-correlation-junta** which takes as input parameters $k \in \mathbb{N}$, distance parameter $\epsilon > 0$, oracle access to function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and has the following guarantee: With probability $2/3$, it outputs a number $\widehat{\text{Corr}}_{f,k}$ such that*

$$|\widehat{\text{Corr}}_{f,k} - \max_{\ell \in \mathcal{J}_{n,k}} \mathbf{E}_{\mathbf{x}}[\ell(\mathbf{x}) \cdot f(\mathbf{x})]| \leq \epsilon.$$

It also outputs a function $h : \{-1, 1\}^k \rightarrow \{-1, 1\}$ such that there is a set of coordinates $\mathcal{T} = \{i_1, \dots, i_k\} \subseteq [n]$ and

$$|\max_{\ell \in \mathcal{J}_{n,k}} \mathbf{E}_{\mathbf{x}}[\ell(\mathbf{x}) \cdot f(\mathbf{x})] - \mathbf{E}_{\mathbf{x}}[h(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k}) \cdot f(\mathbf{x})]| \leq \epsilon.$$

The query complexity of the algorithm is $2^k \cdot \text{poly}(k, 1/\epsilon)$.

Note that the above algorithm is doing something stronger than “merely” computing correlation of f with k -juntas – in fact, the algorithm also outputs the a k -junta that is most correlated up to ϵ . Note that the algorithm cannot identify the actual subset of the coordinates of f that maps to those of the junta, as an standard information theory argument shows that this requires the number of queries to depend on n , even without noise. Further, for the task of approximately finding the most correlated k -junta, our query complexity is essentially optimal, since even giving the description of the most correlated k -junta takes 2^k bits. An immediate corollary of Theorem 1.3 is the existence of a noise tolerant tester for k -juntas.

Corollary 1.4. *For any constant $\frac{1}{2} > c_u > c_\ell \geq 0$ and $k \in \mathbb{N}$, there is (c_u, c_ℓ) -noise tolerant tester for k -juntas with query complexity $2^k \cdot \text{poly}(k, 1/|c_u - c_\ell|)$.*

Proof. Let $\epsilon = \frac{c_u - c_\ell}{2}$. Run the algorithm **Maximum-correlation-junta** with distance parameter ϵ . Let the output be $\widehat{\text{Corr}}_{f,k}$. Set $\text{Thr} = 1 - 2c_\ell - 2\epsilon = 1 - 2c_u + 2\epsilon$. The rest of the algorithm is

1. If $\widehat{\text{Corr}}_{f,k} \geq \text{Thr}$, then the algorithm accepts.
2. The algorithm rejects otherwise.

Note that if there is a k -junta g such that $\text{dist}(f, g) \leq c_\ell$, then $\max_{g \in \mathcal{J}_{n,k}} \mathbf{E}_{\mathbf{x}}[g(\mathbf{x}) \cdot f(\mathbf{x})] \geq 1 - 2c_\ell$. Thus, $\widehat{\text{Corr}}_{f,k} \geq 1 - 2c_\ell - \epsilon$ (w.p. $2/3$), and so the algorithm will accept.

On the other hand, if $\text{dist}(f, g) \geq c_u$ for every $g \in \mathcal{J}_{n,k}$ then $\max_{g \in \mathcal{J}_{n,k}} \mathbf{E}_{\mathbf{x}}[g(\mathbf{x})f(\mathbf{x})] \leq 1 - 2c_u = \text{Thr} - 2\epsilon$, meaning that the algorithm will reject with probability at least $2/3$. \square

We also remark here that the algorithm Maximum-correlation-junta can be modified in a straightforward manner to yield a noise tolerant tester against any *subclass of juntas*, including any *specific junta* – e.g., for any $1/2 > c_u > c_\ell \geq 0$, we can obtain a (c_u, c_ℓ) -tester for k -linear functions [5, 26] with query complexity $\text{poly}(k, 1/|c_u - c_\ell|)$. We leave the proof to the interested reader.

Finally, we can also improve the query complexity to have a polynomial dependence on k at the cost of achieving a weaker guarantee.

Theorem 1.5. *There is an algorithm Maximum-correlation-gap-junta which takes as input parameters $k \in \mathbb{N}$, distance parameter $\epsilon > 0$, oracle access to a function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and has the following guarantee: With probability $2/3$, it outputs a number $\widehat{\text{Corr}}_{f, \text{gap}, k}$ satisfying*

$$\max_{g \in \mathcal{J}_{n, k}} \mathbf{E}_{\mathbf{x}}[g(\mathbf{x})f(\mathbf{x})] - \epsilon \leq \widehat{\text{Corr}}_{f, \text{gap}, k} \leq \max_{g \in \mathcal{J}_{n, k^2/\epsilon^2}} \mathbf{E}[g(\mathbf{x})f(\mathbf{x})] + \epsilon.$$

The query complexity of the algorithm is $\text{poly}(k, 1/\epsilon)$.

Analogous to Corollary 1.4, we get the following corollary by applying Theorem 1.5.

Corollary 1.6. *For any constant $\frac{1}{2} > c_u > c_\ell \geq 0$ and $k \in \mathbb{N}$, there is an algorithm with query complexity $\text{poly}(k, 1/|c_u - c_\ell|)$ which with probability $2/3$ can distinguish between the following cases:*

1. $\min_{g \in \mathcal{J}_{n, k}} \text{dist}(g, f) \leq c_\ell$.
2. $\min_{g \in \mathcal{J}_{n, k'}} \text{dist}(g, f) \geq c_u$, where $k' = k^2/(c_u - c_\ell)^2$.

We remark that [4] contains a result along the same lines as Corollary 1.6, but with $k' = 4k$ and $c_u = 16c_\ell$. That is, compared to the result of [4], Corollary 1.6 has a worse k' , but allows for arbitrarily good noise tolerance.

1.1 Overview of techniques

One of our main contributions is a *Fourier based algorithm* to simulate *oracles* to interesting coordinates of f . In particular, the first step in both Maximum-correlation-junta and Maximum-correlation-gap-junta is to obtain oracle access to the functions $x \mapsto x_\ell$ for all ℓ with large low-degree influence in f . This idea previously appeared (but in a real-valued setting) in [12], and it may be of independent interest. In particular, it is a substantial departure from previous approaches to tolerant junta testing.

1.1.1 Sketch of algorithm Maximum-correlation-junta. We begin by giving the high level overview of the algorithm Maximum-correlation-junta (from Theorem 1.3). Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and let $g : \{-1, 1\}^n \rightarrow \{-1, 1\}$ be a maximally correlated k -junta – for this description, assume that $\mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot g(\mathbf{x})] = \Omega(1)$. The main steps in our algorithm is as follows.

1. First, we show that up to a small loss in correlation, we may assume that every variable in g has at least $k^{-\Theta(1)}$ influence in f – see Claim 3.8 for the precise quantitative parameters. We call these the “interesting variables,” and our first main goal is to obtain oracle access to them.
2. Suppose that x_ℓ is an interesting variable. We show that by randomly (in a carefully chosen sense) restricting certain variables of f , with probability $k^{-O(1)}$ we obtain a function (call it $f_{\uparrow \rho}$) such that $|\widehat{f_{\uparrow \rho}}(\ell)| \geq k^{-O(1)}$. (See Claim 3.7 for the precise details.) In other words, influential coordinates of f end up with large Fourier coefficients under random restrictions.

3. Assuming that $|\widehat{f}_{|\rho}(\ell)| \geq k^{-O(1)}$, we construct a (randomized) operator on the function $f_{|\rho}$ which, with probability $k^{-O(1)}$, gives us an oracle to the variable x_ℓ . This operator is a variant of the operator used by Håstad [19]) in the context of dictatorship testing and in turn uses a modified version of the standard Bonami-Beckner noise. The details of this are in Section 3.1.
4. Having obtained an oracle for one particular variable x_ℓ , we can just repeat steps 2 and 3 $k^{\Theta(1)}$ times to obtain a set \mathcal{S} of (oracles to) variables that contains all of the interesting variables. This reduces the original problem (estimating $\max_{g \in \mathcal{J}_{n,k}} \mathbf{E}_{\mathbf{x}}[g(\mathbf{x})f(\mathbf{x})]$) to the problem of estimating $\max_{g \in \mathcal{J}_{\mathcal{S},k}} \mathbf{E}_{\mathbf{x}}[g(\mathbf{x})f(\mathbf{x})]$. We do this via a simple sampling based algorithm in Find-best-fit. The query complexity of this routine is $2^k \cdot \text{poly}(k)$ and is the bottleneck for Maximum-correlation-junta.

1.1.2 Sketch of Algorithm Maximum-correlation-gap-junta The difference in the proof of Theorem 1.5 vis-a-vis Theorem 1.3 lies in Step 4 of the above overview. Namely, having obtained the set \mathcal{S} , our goal is find smaller subset $\mathcal{S}' \subset \mathcal{S}$ of size $O(k^2/\epsilon^2)$ of the variables that achieves the same correlation and moreover find this correlation.

1. The novel idea of the proof is to find a polynomial algorithm that is able to compute the function

$$f_{\text{avg},\mathcal{S}}(x) := \mathbf{E}_{\mathbf{y} \in \{-1,1\}^{[n] \setminus \mathcal{S}}} [f(x, \mathbf{y})]$$

while only having oracle access to the variables in \mathcal{S} . The details of this algorithm are in Section 5.1, and we will give an outline shortly. Note that $f_{\text{avg},\mathcal{S}}$ depends only on the variables in \mathcal{S} (of which there are $\text{poly}(k)$), and among all such functions it has the highest correlation with f . Further, $\max_{g \in \mathcal{J}_{\mathcal{S},k}} \mathbf{E}_{\mathbf{x}}[g(\mathbf{x})f(\mathbf{x})] = \max_{g \in \mathcal{J}_{\mathcal{S},k}} \mathbf{E}_{\mathbf{x}}[g(\mathbf{x})f_{\text{avg},\mathcal{S}}(\mathbf{x})]$.

2. We replace $f_{\text{avg},\mathcal{S}}$ by $T_\rho f_{\text{avg},\mathcal{S}}$ (for $\rho = 1 - O(\epsilon/k)$), incurring an $O(\epsilon)$ error in the correlation. The advantage of $T_\rho f_{\text{avg},\mathcal{S}}$ over $f_{\text{avg},\mathcal{S}}$ is that it has at most $O(k^2/\epsilon^2)$ high-influence (meaning, influence $\Omega(\epsilon/k)$) variables, and that restricting our attention to juntas on these variables only incurs another $O(\epsilon)$ error in the correlation. It is also easy to produce an oracle to $T_\rho f_{\text{avg},\mathcal{S}}$ from $f_{\text{avg},\mathcal{S}}$ with polynomially many samples.
3. Our next step is to estimate the influence (in the function $T_\rho f_{\text{avg},\mathcal{S}}$) of all the variables in \mathcal{S} . We can do so by sampling correlated pairs (x, y) repeatedly until we obtain pairs that differ in one coordinate and then checking the effect on $T_\rho f_{\text{avg},\mathcal{S}}$. Having estimated the influences, we let $\mathcal{S}' \subseteq \mathcal{S}$ be the set of high-influence variables. The problem is reduced to that of estimating $\max_{g \in \mathcal{J}_{\mathcal{S}',k}} \mathbf{E}_{\mathbf{x}}[g(\mathbf{x})f(\mathbf{x})]$.
4. The final output of the algorithm is an estimate for the correlation of f with the best function depending only on variables of \mathcal{S}' . This is just

$$\mathbf{E}_{\mathbf{x}}[|f_{\text{avg},\mathcal{S}'}(\mathbf{x})|],$$

which can be estimated using the same averaging procedure that we mentioned in the first step.

It remains to explain how to carry out the averaging procedure needed for the first and last steps of the outline above: how do we estimate $\mathbf{E}_{\mathbf{y} \in \{\pm 1\}^{[n] \setminus \mathcal{S}}} [f(x, \mathbf{y})]$, given only oracle access to the variables in \mathcal{S} ? The basic idea is to perform a random walk on the the subset of y 's that agree with x on all the elements of \mathcal{S} . We let $y^{(0)} = x$. Given $y^{(i)}$ we sample $y^{(i+1)}$ to be a noisy version of $y^{(i)}$, where each coordinate is flipping with probability about $1/|\mathcal{S}|$. We *accept* $y^{(i+1)}$ if the value of all oracle functions in \mathcal{S} is identical for $y^{(i+1)}$ and $y^{(i)}$. If we reject the current proposal of $y^{(i+1)}$ we try again an independent noisy $y^{(i+1)}$.

Thus we effectively perform a random walk on the noisy hyper-cube on the coordinates in $[n] \setminus \mathcal{S}$. The spectral gap of the random walk is inverse polynomial in k , and hence by taking $\text{poly}(k)$ steps of this random walk, we can essentially independently resample those coordinates of x that do not belong to \mathcal{S} . By repeating this, we can estimate $f_{\text{avg}, \mathcal{S}}(x)$.

2 Preliminaries

We begin with the basics of Fourier analysis, in particular the notion of Fourier expansion of functions.

Definition 2.1. For any subset $S \subseteq [n]$, we define $\chi_S : \{-1, 1\}^n \rightarrow \{-1, 1\}$ as $\chi_S(x) = \prod_{i \in S} x_i$. Any function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ can be expressed as a linear combination of $\{\chi_S(x)\}_{S \subseteq [n]}$ (as follows):

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \chi_S(x).$$

This is referred to as the Fourier expansion of f and the coefficients $\{\widehat{f}(S)\}_{S \subseteq [n]}$ are referred to as the Fourier coefficients of f .

We next define the concept of influence of variables in $f : \{-1, 1\}^n \rightarrow \mathbb{R}$

Definition 2.2. For any function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and any $i \in [n]$, $\text{Inf}_i(f) = \Pr_{x \in \{-1, 1\}^n} [f(x) \neq f(x^{\oplus i})]$ (where $x^{\oplus i}$ differs from x exactly in the i^{th} position). In terms of Fourier coefficients, $\text{Inf}_i(f) = \sum_{S \ni i} \widehat{f}^2(S)$; in the case of a real-valued function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$, we take this latter formula as the definition of $\text{Inf}_i(f)$.

For a number $k \leq n$, we let $\text{Inf}_i^{\leq k}(f)$ denote the quantity $\text{Inf}_i^{\leq k}(f) = \sum_{S \ni i: |S| \leq k} \widehat{f}^2(S)$. We also define the total influence of f , denoted by $\text{Inf}(f)$, as $\sum_S |S| \widehat{f}^2(S)$.

We now define the Bonami-Beckner noise operator on the space of functions on $\{-1, 1\}^n$. To do this, we first define a general notion of noise distribution on $\{-1, 1\}^n$. For $\eta \in [-1, 1]^n$, we let \mathbf{Z}_η denote the product distribution on $\{-1, 1\}^n$ where the expectation of the i^{th} bit is η_i .

Definition 2.3. For any $\rho \in [-1, 1]$, let $\bar{\rho} \in [-1, 1]^n$ denote the vector all of whose coordinates are ρ . The Bonami-Beckner noise operator (denoted by T_ρ) operates on $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ as

$$T_\rho f(x) = \mathbf{E}_{\mathbf{y} \sim \mathbf{Z}_{\bar{\rho}}} [f(x \cdot \mathbf{y})].$$

We let $x \cdot y$ denote the coordinate wise product of x and y .

A standard fact about the operator T_ρ is its action on the Fourier expansion of f (see [23] for details).

$$T_\rho f(x) = \sum_{S \subseteq [n]} \rho^{|S|} \widehat{f}(S) \chi_S(x).$$

Bonami-Beckner noise operator as a Markov chain It will be useful for us to view the Bonami-Beckner noise operator as a Markov chain. We recall the definition of a Markov chain (on a finite set).

Definition 2.4. Let G be a finite set and let $P \in \mathbb{R}^{G \times G}$ be a stochastic matrix. The random variables (taking values in G) $(\mathbf{x}_i)_{i=1}^T$ are said to follow the Markov chain \mathcal{M}_P (with transition matrix given by P) if for any $T \geq j > 1$ and any $g_1, \dots, g_j \in G$,

$$\Pr[\mathbf{x}_j = g_j | \mathbf{x}_{j-1} = g_{j-1}, \dots, \mathbf{x}_1 = g_1] = \Pr[\mathbf{x}_j = g_j | \mathbf{x}_{j-1} = g_{j-1}] = P(g_{j-1}, g_j).$$

We refer the reader to the book by Levin and Peres [21] for definitions of standard notions such as ergodicity, aperiodicity and stationary distributions.

Now, consider any $\rho \in [-1, 1]$ and define the stochastic matrix P_ρ (whose rows and columns are indexed by $\{-1, 1\}^n$) such that

$$P_\rho(x, y) = \Pr_{\mathbf{z} \sim \mathbf{Z}_\rho} [x \cdot \mathbf{z} = y].$$

It is easy to see that the matrix P_ρ is a symmetric matrix and further, the second largest eigenvalue of the matrix P_ρ is at most ρ . The matrix P_ρ also defines a corresponding Markov chain \mathcal{M}_ρ (i.e., the transition matrix of \mathcal{M}_ρ is P_ρ). Markov chains have a certain ‘‘averaging property’’ which is particularly useful for us and is stated below. We will instantiate it to the Markov chain \mathcal{M}_ρ in Section 5 later. We now state the following result due to Lezaud [22] (Theorem 1.1 in the paper) which applies to ergodic and reversible Markov chains. Similar results which apply to the special case of random walks on undirected graphs have found many applications in computer science [17, 20, 28, 16].

Lemma 2.5. *For $\rho \in (-1, 1)$, let $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots$ follow the Bonami-Beckner Markov chain \mathcal{M}_ρ with an arbitrary initial value $\mathbf{x}^{(1)} = x \in \{\pm 1\}^n$. There is a constant C such that for any $f : \{\pm 1, 1\}^n \rightarrow [-1, 1]$ and any $\gamma, \delta > 0$, if $T \geq \frac{C \log(1/\delta)}{\gamma^2 \cdot (1-\rho)}$ then*

$$\Pr \left[\left| \frac{f(\mathbf{x}_1) + \dots + f(\mathbf{x}_T)}{T} - \mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [f(\mathbf{x})] \right| > \gamma \right] \leq \delta.$$

2.1 Random restrictions

A crucial role in our algorithm will be played by the notion of random restrictions from circuit complexity [15, 18].

Definition 2.6. *For any $\mu \in [0, 1]$, we let $\mathcal{R}_\mu \in \{-1, 1, *\}^n$ denote the product distribution where each coordinate is $*$ with probability μ , ± 1 with probability $(1 - \mu)/2$ each. Further, for $f : \{\pm 1\}^n \rightarrow \{\pm 1\}$ and $\xi \in \{-1, 0, 1\}^n$, we let $f_{|\xi} : \{\pm 1, 1\}^S \rightarrow \{\pm 1, 1\}$ where*

1. $S = \{i \in [n] : \xi_i = 0\}$. The set of variables in S are said to survive in $f_{|\xi}$.
2. For $x \in \{\pm 1\}^S$, we $f_{|\xi}(x) = f(z)$ where $z_S = x$ and $z_j = \xi_j$ for $j \notin S$.

3 Construction of coordinate oracles

The main result of this section is an algorithm for constructing a set of ‘‘oracles’’ to all of the interesting coordinates of a function, assuming that there are not too many of them. The basic definition is the following: For $1 \leq i \leq n$, let $\text{Dict}_i : \{\pm 1, 1\}^n \rightarrow \{\pm 1, 1\}$ be defined as $\text{Dict}_i : x \mapsto x_i$.

Definition 3.1. *Let \mathcal{D} be a set of functions from $\{\pm 1, 1\}^n$ to $\{\pm 1, 1\}$. We say that \mathcal{D} is an oracle for the coordinates $S \subset [n]$ if*

- for every $g \in \mathcal{D}$, there is some $i \in S$ such that $g = \text{Dict}_i$ or $g = -\text{Dict}_i$; and
- for every $i \in S$, there is some $g \in \mathcal{D}$ such that $g = \text{Dict}_i$ or $g = -\text{Dict}_i$.

In other words, \mathcal{D} is an oracle for S if $\mathcal{D} = \{\text{Dict}_i : i \in S\}$ ‘‘up to sign’’.

Due to our constraints, we will not be able to produce coordinate oracles exactly according to the definition above, so we will relax it slightly. Recall that we have fixed an underlying function $f : \{\pm 1\}^n \rightarrow \{-1, 1\}$ and a parameter $k \in \mathbb{N}$.

Definition 3.2. Let \mathcal{D} be a set of functions from $\{-1, 1\}^n$ to $\{-1, 1\}$. For $\epsilon \leq \frac{1}{8}$, we say that \mathcal{D} is an ν -oracle for $\mathcal{S} \subset [n]$ if

- for every $g \in \mathcal{D}$, there is some $i \in \mathcal{S}$ such that g is ν -close to Dict_i or $-\text{Dict}_i$ (necessarily only one, since $\nu \leq \frac{1}{8}$);
- for every $i \in \mathcal{S}$, there is exactly one $g \in \mathcal{D}$ that is ν -close to Dict_i or $-\text{Dict}_i$; and
- for every $x \in \{\pm 1\}^n$, every $g \in \mathcal{D}$, and every $\delta > 0$, there is a randomized algorithm to compute $g(x)$ correctly with probability at least $1 - \delta$, using $\text{poly}(k, \log \frac{1}{\delta})$ queries to f .

While the definition of \mathcal{D} involves both ν and k , since the latter will remain fixed throughout, the above definition is only quantified in terms of ν . The parameters ν and δ that we choose will essentially allow us to pretend that an ν -oracle is an oracle. In particular, we will fix $\delta = 2^{-\omega(k)}$ when evaluating coordinate oracles at a point. This will preserve the $\text{poly}(k)$ query complexity of each oracle query, while ensuring that (with high probability) every query that we make to a coordinate oracle will be computed correctly (since in all of our algorithms, we will make no more than $2^k \cdot \text{poly}(k)$ queries). Our choice of ν will depend on the setting: in the setting of Theorem 1.3, we will make at most $2^k \cdot \text{poly}(k)$ queries to each coordinate oracle, so we will take $\nu = (2^k \cdot \text{poly}(k))^{-1}$. This means that each oracle query requires at most $\text{poly}(k)$ queries to f , while ensuring that each coordinate oracle is so close to a dictator (or anti-dictator) that we will (with high probability) not observe the difference. In the setting of Theorem 1.5, we will set take $\nu = \text{poly}(1/k)$ and make at most $\text{poly}(k)$ queries to each coordinate oracle; this requires $\text{poly}(k)$ queries to f and ensures that (with high probability) we will not observe the difference between any coordinate oracle and its corresponding dictator. With this in mind, and to prevent a proliferation of parameters, we will often pretend that we have access to an oracle in the sense of Definition 3.1 when we really have access to an ν -oracle in the sense of Definition 3.2.

3.1 A single oracle

We begin by describing how to construct an oracle to a single coordinate. The basic step notion is the following operator, which is related to one used by Håstad in the context of dictatorship testing [19].

Definition 3.3. Let $\eta \in [-1, 1]^n$ and \mathbf{Z}_η denote the product distribution on $\{-1, 1\}^n$ where the expectation of the i^{th} bit is η_i . For any $f : \{-1, 1\}^n \rightarrow \mathbb{R}$, we define the operator

$$\text{Has}_\eta f(x) = \mathbf{E}_{\mathbf{y}_1, \mathbf{y}_2 \in \{-1, 1\}^n, \mathbf{y}_3 \in \mathbf{Z}_\eta} [f(\mathbf{y}_1) f(\mathbf{y}_2) f(x \oplus \mathbf{y}_1 \oplus \mathbf{y}_2 \oplus \mathbf{y}_3)].$$

In terms of the Fourier expansion, it is easy to check that

$$\text{Has}_\eta f(x) = \sum_S \hat{f}^3(S) \chi_S(x) \eta^S,$$

where $\eta^S = \prod_{j \in S} \eta_j$. A consequence of this expansion is that for the right choice of η , $\text{Has}_\eta f$ is a good approximation to a certain dictator function (depending on η).

Lemma 3.4. Suppose that $|\hat{f}(1)| \geq \kappa$, where $\kappa \in (0, 1)$, and let $\alpha = \frac{\kappa^3}{16}$. Choose $\eta \in \{0, \alpha\}^n$ randomly so that $\Pr[\eta_i = \alpha] = \kappa^6/16$, independently for every i . Then with probability at least $\Omega(\kappa^6)$, for every $x \in \{\pm 1\}^n$,

$$|\text{Has}_\eta f(x) - \hat{f}^3(0) - \alpha \hat{f}^3(1)x_1| \leq \frac{\alpha}{4} |\hat{f}(1)|^3.$$

Proof. Let $p = \kappa^6/16 = \Pr[\eta_i = \alpha]$. Let $\Gamma = \{i : |\hat{f}(i)| \geq \kappa^3/8\}$; since $\sum_S \hat{f}(S)^2 \leq 1$, we have $|\Gamma| \leq 64\kappa^{-6}$. Let E be the event that $\eta_1 = \alpha$ and $\eta_j = 0$ for all $j \in \Gamma \setminus \{1\}$. Then $\Pr[E] = p(1-p)^{|\Gamma|-1} \geq p(1-p)^{64\kappa^{-6}} = \Omega(\kappa^6)$, and we will show that the claimed inequality happens on E . Indeed, the Fourier expansion above implies that

$$\begin{aligned} \text{Has}_\eta f(x) - \hat{f}^3(0) - \alpha \hat{f}^3(1)x_1 &= \sum_{1 < j \leq n} \eta_j \hat{f}^3(j)x_j + \sum_{|S| > 1} \eta^S \hat{f}^3(S)\chi_S(x) \\ &= \sum_{j \notin \Gamma} \eta_j \hat{f}^3(j)x_j + \sum_{|S| > 1} \eta^S \hat{f}^3(S)\chi_S(x), \end{aligned}$$

where the second equality holds on the event E . Now,

$$\sum_{j \notin \Gamma} |\eta_j \hat{f}^3(j)| \leq \alpha \frac{\kappa^3}{8} \sum_{j \notin \Gamma} \hat{f}^2(j) \leq \frac{\alpha \kappa^3}{8}$$

and

$$\sum_{|S| > 1} |\eta^S \hat{f}^3(S)| \leq \sum_{|S| > 1} \alpha^2 \hat{f}^2(S) \leq \alpha^2,$$

and so the triangle inequality gives

$$|\text{Has}_\eta f(x) - \hat{f}^3(0) - \alpha \hat{f}^3(1)x_1| \leq \frac{\alpha \kappa^3}{8} + \alpha^2 \leq \frac{\alpha \kappa^3}{4} \leq \frac{\alpha}{4} |\hat{f}(1)|^3. \quad \square$$

Lemma 3.4 gives us a natural algorithm for computing something that might be a coordinate oracle: the basic idea is to sample η and then to define

$$g_\eta(x) = \text{sgn}(\text{Has}_\eta f(x) - \hat{f}^3(0)) = \text{sgn}(\text{Has}_\eta f(x) - \mathbf{E}[f]^3).$$

Note that computing the function g_η requires randomness (to estimate $\text{Has}_\eta f(x)$ and $\mathbf{E}[f]$). However, a straightforward Chernoff bound implies that with $O(\kappa^{-12} \log \frac{1}{\delta})$ queries to f , we can estimate both $\mathbf{E}[f]$ and $\text{Has}_\eta f(x)$ to additive accuracy $O(\kappa^6)$, with probability $1 - \delta$; according to Lemma 3.4, this is sufficient to correctly evaluate $g_\eta(x)$ with probability $1 - \delta$. This is the same sort of guarantee required in Definition 3.2; as discussed there, we can choose $\delta = 2^{-k^2}$ so that with high probability, every evaluation of g_η will be correct.

Now, Lemma 3.4 only provides us with a small probability of finding a good g_η . To filter out the bad ones, we add in a dictatorship test: for $i \leq n$, let $\text{Dict}_i : \{-1, 1\}^n \rightarrow \{-1, 1\}$ denote the function $\text{Dict}_i(x) = x_i$. We call Dict_i a dictator function, and $-\text{Dict}_i$ an anti-dictator function. There is an algorithm for testing whether a function is a dictator function (see Chapter 7 of [23]):

Theorem 3.5. *There is an algorithm Dictator-test which given an error parameter $\nu > 0$ and confidence parameter $\tilde{\delta} > 0$, makes $O(\nu^{-1} \log \tilde{\delta}^{-1})$ queries to $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and has the following properties:*

1. If $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ is a dictator or an anti-dictator, then it accepts with probability 1.

2. Any $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ which is ν -far from every dictator and anti-dictator is accepted with probability at most $\tilde{\delta}$.

Algorithm Construct-coordinate-oracle gives the algorithm for constructing coordinate oracles.

Input: f (target function), k (arity of Junta),
 δ (confidence parameter), $\nu \leq 1/8$ (first accuracy parameter), τ (second accuracy parameter)
Output: an oracle \mathcal{D}

```

// Construct the initial oracles
1 Let  $T = Ck^5\tau^{-5}\log(1/\delta)$  and let  $M = Ck^7\tau^{-7}\log(1/\delta)$ ;
2 Let  $\tilde{\delta} = \delta/(MT)$ ;
3 Initialize  $\mathcal{D} = \emptyset$ ;
4 repeat  $T$  times
5   Sample  $\rho$  according to  $\mathcal{R}_{1/k}$  (as in Definition 2.6);
6   repeat  $M$  times
7     Sample  $\eta$  as in Lemma 3.4;
8     Let  $g(x) = \text{sgn}(\text{Has}_\eta f_{\uparrow\rho}(x) - \mathbf{E}[f_{\uparrow\rho}]^3)$ ;
9     Apply Dictator-test to  $g$  with confidence  $\tilde{\delta}$  and accuracy  $\nu$ ;
10    if Dictator-test accepts then
11      Add  $g$  to  $\mathcal{D}$ ;
// Clean out duplicates
12 Let  $N = C \log(MT/\delta)$ , and sample  $\mathbf{x}^{(1)} \dots, \mathbf{x}^{(N)} \in \{\pm 1\}^n$  independently and uniformly;
13 while there exist  $g \neq h \in \mathcal{D}$  such that  $|N^{-1} \sum_i g(\mathbf{x}^{(i)})h(\mathbf{x}^{(i)})| \geq \frac{1}{2}$  do
14   Remove  $g$  from  $\mathcal{D}$ ;
15 return  $\mathcal{D}$ 

```

Algorithm 1: Construct-coordinate-oracle

Lemma 3.6. *The algorithm Construct-coordinate-oracle has the following guarantee:*

1. As input, it gets oracle access to $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$, an arity parameter k , two accuracy parameters τ and $\nu \leq 1/8$ and a confidence parameter $\delta > 0$.
2. With probability $1 - \delta$, there is some set $\mathcal{S} \supseteq \{i : \text{Inf}_i^{\leq k}(f) \geq \tau^2/k^2\}$ such that the output of Algorithm 1 is an ν -oracle to \mathcal{S} .
3. The number of oracles in \mathcal{D} is at most $\text{poly}(k, \tau^{-1}, \log(1/\delta))$.
4. The query complexity of the procedure is $\nu^{-1} \cdot \text{poly}(k, \tau^{-1}, \log(1/\delta))$.

Proof. Note that the number of oracles in \mathcal{D} is at most MT . This immediately gives Item 3. Similarly, the query complexity of the algorithm is $M \cdot T$ times the cost of applying the Dictator-test in Step 9. By plugging the bound from Theorem 3.5, we get Item 4. This leaves us with Item 2. Like the algorithm itself, there are two steps in this proof. The first step is to show that after executing the loop on line 4, the set \mathcal{D} contains only ν -approximate dictators, and it contains at least one dictator for every coordinate i with large $\text{Inf}_i^{\geq k}(f)$. The second step is to show that by the end of the algorithm, each coordinate that was represented by a dictator after executing the loop on line 4 is now represented by a unique dictator. Actually, this second step is trivial: if g and h are ν -approximate dictators for $\nu \leq \frac{1}{8}$ then $|\mathbf{E}[gh]| \leq \frac{1}{4}$ if they represent different coordinates, while $|\mathbf{E}[gh]| \geq \frac{3}{4}$ if they represent the same coordinate. By a union bound over the $O(M^2T^2)$ pairs of elements in \mathcal{D} , $\log(CMT/\delta)$ samples are enough to estimate all of these correlations to accuracy $\frac{1}{4}$ with confidence $1 - \delta/4$, meaning that

the loop starting on line 13 correctly (with probability at least $1 - \delta/4$) chooses exactly one dictator function to represent each coordinate.

We turn to the correctness of the first loop: it should find a dictator for every influential coordinate. The analysis of this part is itself divided into two parts, corresponding to the two nested loops: we will argue that for every influential coordinate i , at least one iteration of the outer loop will sample ρ for which $\widehat{f}_{|\rho}(i)$ is large. For this execution of the outer loop, we will argue that at least one iteration of the inner loop will find an oracle for coordinate i .

We will write the first of these parts as a separate claim, and prove it later:

Claim 3.7. *If $\text{Inf}_i^{\leq k}(f) \geq \tau^2/k^2$ and $\rho \sim \mathcal{R}_{1/k}$ then with probability $\Omega(k^{-4}\tau^4)$, $|\widehat{f}_{|\rho}(i)| \geq \tau/(4k)$.*

With our choice of T , Claim 3.7 and a union bound imply that with probability at least $1 - \delta/3$, for every i with $\text{Inf}_i^{\leq k}(f) \geq \tau^2/k^2$ there is at least one iteration of the outer loop for which $|\widehat{f}_{|\rho}(i)| \geq \frac{\tau}{4k}$. We will fix this iteration, and examine the inner loop: by Lemma 3.4 with $\kappa = \frac{\tau}{4k}$, each iteration has $\Omega(\tau^6 k^{-6})$ probability of producing g that is equal to $\pm \text{Dict}_i$. By a union bound, with probability at least $1 - \delta/3$, the inner loop will succeed in adding $\pm \text{Dict}_i$ to \mathcal{D} .

Finally, a union bound implies that with probability at least $1 - \delta/3$, every g that passes the dictatorship test is in fact ϵ -close to some dictator. \square

Proof of Claim 3.7. For $i \in [n]$, define the polynomial $D_i f : \mathbf{R}^n \rightarrow \mathbf{R}$ by

$$(D_i f)(x) = \frac{1}{2}(f(x) - f(x_{-,i})),$$

where $x_{-,i}$ is equal to x , except that the i th coordinate is negated. In terms of Fourier coefficients, it is easy to check that

$$(D_i f)(x) = x_i^2 \sum_{S \subseteq [n] \setminus \{i\}} \widehat{f}(S \cup \{i\}) \chi_S(x).$$

On the other hand, it is also immediate that for every $x \in \{-1, 0, 1\}^n$, and it is easy to check that $\widehat{f}_{|\rho}(i) = D_i f(\rho)$. In particular,

$$\begin{aligned} \text{Var}_\rho[\widehat{f}_{|\rho}(i)] &= \text{Var}_\rho[D_i f(\rho)] \\ &= \sum_{S \subseteq [n] \setminus \{i\}} \widehat{f}^2(S \cup \{i\}) (1 - 1/k)^{|S|} \\ &\geq (1 - 1/k)^k \sum_{\substack{S \subseteq [n] \setminus \{i\} \\ |S| \leq k}} \widehat{f}^2(S \cup \{i\}) \\ &\geq \frac{1}{4} \text{Inf}_i^{\leq k}(f) \\ &\geq \frac{\eta^2}{4k^2}. \end{aligned}$$

We will apply an anti-concentration inequality to the random variable $D_i f(\rho) = \widehat{f}_{|\rho}(i)$ (see Lemma 13 in [11]): for any real-valued random variable \mathbf{X} with variance at least σ^2 and central fourth moment at most $t^4 \sigma^4$,

$$\Pr \left[|\mathbf{X}| > \frac{\sigma}{2} \right] \geq \frac{9}{128(t+2)^4}.$$

Immediately from the definition of $D_i f$, we see that the central fourth moment of $D_i f(\rho)$ is at most 16. Setting $\sigma = \eta/(2k)$ and $t = 4k/\eta$, we have

$$\Pr \left[|\widehat{f}_{|\rho}(i)| > \frac{\eta}{4k} \right] \geq \Omega(k^{-4}\eta^4),$$

as claimed. \square

Finally, the following claim, which will be useful in both Section 4 and Section 5, says that if f correlates with a k -junta, then it also correlates nearly as well with a junta on some set S such that for every variable $j \in S$, $\text{Inf}_j^{\leq k}(f)$ is large. This is useful because Lemma 3.6 can then be used to construct oracles to all these *relevant variables*.

Claim 3.8. *Let $f : \{\pm 1\}^n \rightarrow \{\pm 1\}$ and let $g : \{\pm 1\}^n \rightarrow \{\pm 1\}$ be a k -junta on the variables $\{1, \dots, k\}$ such that $\mathbf{E}[f \cdot g] \geq c$. Then, for any $\tau > 0$, there is a set $S \subseteq [k]$ of variables such that*

1. *For $i \in S$, $\text{Inf}_i^{\leq k}(f) \geq \frac{\tau^2}{k^2}$.*
2. *There is a junta on S with correlation $c - \tau$ with f .*

Proof. Let us first start with $S = [k]$. If all variables $i \in S$ satisfy $\text{Inf}_i^{\leq k}(f) \geq \frac{\tau^2}{k^2}$, we are done. Else, let j be any variable such that $\text{Inf}_j^{\leq k}(f) \leq \frac{\tau^2}{k^2}$. For this variable j , define $g_j : \{\pm 1\}^n \rightarrow [-1, 1]$ as the junta on the set $[k] \setminus \{j\}$

$$g_j(x) = \frac{1}{2}(g(x) + g(x^{\oplus j})),$$

where $x^{\oplus j}$ is the same as x with the coordinate at j flipped. Observe that for any S , $\widehat{g}_j(S) = \widehat{g}(S)$ if $j \notin S$ and 0 otherwise. Note that

$$|\mathbf{E}_x[f \cdot g_j(x)] - \mathbf{E}_x[f \cdot g(x)]| = \left| \sum_{S \ni j} \widehat{g}(S) \widehat{f}(S) \right| = \left| \sum_{S \ni j: |S| \leq k} \widehat{g}(S) \widehat{f}(S) \right| \leq \sqrt{\sum_{S \ni j: |S| \leq k} \widehat{f}^2(S)} \leq \sqrt{\text{Inf}_j^{\leq k}(f)}.$$

Thus, $\mathbf{E}[f \cdot g_j(x)] \geq c - \frac{\tau}{k}$. By doing a simple randomized rounding step, we can in fact, assume that the range of g_j is ± 1 . We now set $S \leftarrow [k] \setminus j$ and $g = g_j$ and inductively repeat the argument. This finishes the proof. \square

4 Description of Maximum-correlation-junta

In this section, we want to prove Theorem 1.3. The final ingredient required to describe the algorithm Maximum-correlation-junta (from Theorem 1.3) is the routine Find-best-fit. Here, we state the algorithmic guarantee for this routine. The description and its proof of correctness is deferred to Appendix A.

Lemma 4.1. *There is an algorithm Find-best-fit with the following guarantee:*

1. *The algorithm gets as input oracle access to a function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ as well as a set of oracles $\mathcal{S} \subseteq \{\text{Dict}_1, \dots, \text{Dict}_n\}$. We clarify that algorithm is only given the oracle Dict_i (for $\text{Dict}_i \in \mathcal{S}$) but not i .*
2. *The algorithm gets as input error parameter $\epsilon > 0$, arity parameter k and confidence parameter $\delta > 0$.*
3. *The algorithm makes $N(k, |\mathcal{S}|, \epsilon, \delta) = O(2^k/\epsilon^2 \cdot |\mathcal{S}| \cdot (\log(1/\delta) + k^2 + |\mathcal{S}|))$ queries with probability $1 - \delta$.*
4. *Each query point (to either f or oracle in \mathcal{S}) is distributed as a uniformly random element of $\{-1, 1\}^n$.*
5. *With probability $1 - \delta$, the algorithm outputs a number $\widehat{\text{Corr}}_{f, \mathcal{S}, k}$ and $h_{\mathcal{S}, k} : \{-1, 1\}^k \rightarrow \{-1, 1\}$ such that there exists oracles $\text{Dict}_{i_1}, \dots, \text{Dict}_{i_k} \in \mathcal{S}$*

$$\mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot h_{\mathcal{S}, k}(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k})] \geq \max_{\ell \in \mathcal{J}_{\mathcal{S}, k}} \mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot \ell(\mathbf{x})] - \epsilon \text{ and } |\widehat{\text{Corr}}_{f, \mathcal{S}, k} - \max_{\ell \in \mathcal{J}_{\mathcal{S}, k}} \mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot \ell(\mathbf{x})]| \leq \epsilon.$$

We next describe the algorithm Maximum-correlation-junta (Algorithm 2).

Input: f (target function), k (arity of Junta), ϵ (distance parameter)
Output: $\widehat{\text{Corr}}_{f,k} \in [0, 1]$ and $h : \{-1, 1\}^k \rightarrow \{-1, 1\}$
// Set parameters for the algorithm
1 Let $\tau = \frac{\epsilon}{4}$, $\delta = 1/3$;
2 Let $R = \text{poly}(k, \tau^{-1}, \log(1/\delta))$ – the upper bound on output size from Item 3 of Lemma 3.6 ;
3 Set $N = N(k, R, \epsilon/4, \delta/4)$ where $N(\cdot)$ is the function in Item 3 of Lemma 4.1 ;
4 Set $\nu = \delta/4N$;
// Construct the initial oracles
5 Run **Construct-coordinate-oracle** with input function f , junta arity parameter k , confidence parameter $\delta/4$, first accuracy parameter ν and second accuracy parameter τ ;
6 Let \mathcal{D} be the set of returned oracles ;
// Use the oracles to find maximally correlated junta
7 Error parameter for queries to any oracle $g \in \mathcal{D}$ is set to ν ;
8 Run **Find=best-fit** with target function f , oracles given by \mathcal{D} , confidence parameter $\delta/4$, distance parameter $\epsilon/4$ and arity parameter k . ;
9 Let the output of this routine be $\widehat{\text{Corr}}_{f,\mathcal{D},k}$ and $h_{\mathcal{D},k} : \{-1, 1\}^k \rightarrow \{-1, 1\}$. ;
10 Set $h \leftarrow h_{\mathcal{D},k}$ and $\widehat{\text{Corr}}_{f,k} \leftarrow \widehat{\text{Corr}}_{f,\mathcal{D},k}$;
11 **return** $(\widehat{\text{Corr}}_{f,k}, h)$

Algorithm 2: Maximum-correlation-junta

Proof of Theorem 1.3: We begin with the following claim.

Claim 4.2. *The query complexity of the algorithm is $2^k \cdot \text{poly}(k, 1/\epsilon)$.*

Proof. By just plugging the bounds, we see that R defined in Step 2 of the algorithm is $\text{poly}(k/\epsilon)$; N defined in Step 3 is $2^k \cdot \text{poly}(k/\epsilon)$ and ν defined in Step 4 is $2^{-k} \cdot \text{poly}(\epsilon/k)$.

Now, the query complexity of Step 5, i.e., **Construct-coordinate-oracle** is then $\nu^{-1} \cdot \text{poly}(k, \tau^{-1}) = 2^k \cdot \text{poly}(k/\epsilon)$ (Item 4 of Lemma 3.6). Next, note that $|\mathcal{D}| \leq \text{poly}(k/\epsilon)$ (by plugging the bound from Item 3 of Lemma 3.6). However, this means that the algorithm **Find-best-fit** (from Step 3 of Lemma 4.1) makes at most $2^k \cdot \text{poly}(k, |\mathcal{D}|) = 2^k \cdot \text{poly}(k, 1/\epsilon)$ queries where each query is either to f or to an oracle in \mathcal{D} . However, each call to \mathcal{D} is made with error ν , the query complexity of making each such call is $\text{poly}(k, \log(1/\nu)) = \text{poly}(k)$. Thus, the total query complexity is $2^k \cdot \text{poly}(k/\epsilon)$. \square

Having proven a bound on the query complexity, we now turn to the proof of correctness of this algorithm. Note that every oracle $g \in \mathcal{D}$ is ν close to Dict_i for some $i \in [n]$. Further, at point x , (by definition of a ν -oracle), we have an algorithm, which returns the value $g(x)$ with probability at least $1 - \nu$. We say that the evaluation of g at point x is *good*, if we get the value of $g(x) = \text{Dict}_i(x)$. Note that a randomly chosen point $x \in \{-1, 1\}^n$,

$$\begin{aligned} \Pr_{x \in \{-1, 1\}^n} [\text{Evaluation of } g \text{ is not good}] &\leq \Pr_{x \in \{-1, 1\}^n} [g(x) \neq \text{Dict}_i(x)] + \Pr[\text{Evaluation of } g \text{ is incorrect}] \\ &\leq \nu + \nu = 2\nu. \end{aligned}$$

Since the query by the algorithm **Find-best-fit** to each oracle in \mathcal{D} is a random point in $\{-1, 1\}^n$ (Item 4 in Lemma 4.1) and the total number of queries to **Find-best-fit** is N , hence the probability the evaluation of $g \in \mathcal{D}$ is not good at any point is at most $2N\nu = \delta/2$. Thus, from now on, we will assume that $g \in \mathcal{D}$ is ν -close to Dict_i , then the algorithm **Find-best-fit** has exact access to Dict_i . However, with this the following claim is immediate from Item 5 of Lemma 4.1.

Claim 4.3. Suppose the algorithm outputs $(\widehat{\text{Corr}}_{f,k}, h)$. Let $V = \{i \in [n] : \text{Dict}_i \text{ is } \nu\text{-close to some } g \in \mathcal{D}\}$. Then, with probability $1 - \delta/4$, there is some $\{i_1, \dots, i_k\} \subseteq V$ such that

$$\mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot h(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k})] \geq \max_{\ell \in \mathcal{J}_{V,k}} \mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot \ell(\mathbf{x})] - \frac{\epsilon}{4} \text{ and } |\widehat{\text{Corr}}_{f,k} - \max_{\ell \in \mathcal{J}_{V,k}} \mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot \ell(\mathbf{x})]| \leq \frac{\epsilon}{4}.$$

Finally, we have the following claim.

Claim 4.4. Let $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ and assume that there exists a k -junta $g : \{-1, 1\}^n \rightarrow \{-1, 1\}$ such that $\mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot g(\mathbf{x})] \geq \text{Corr}_{f,g}$. Let \mathcal{D} be the set of returned oracles in Step 6. Let $V = \{i \in [n] : \text{Dict}_i \text{ is } \nu\text{-close to some } g \in \mathcal{D}\}$. Then, with probability $1 - \delta/4$, there exists $T \subseteq V$, $|T| = k$ and a junta on T with correlation at least $\text{Corr}_{f,g} - \epsilon/4$ with f .

Proof. Let W be the set of variables appearing g . By Claim 3.8 (setting $\tau = \epsilon/4$), there is subset $W' \subseteq W$ and a W' -junta $r : \{-1, 1\}^n \rightarrow \{-1, 1\}$ with the following properties: (a) $\mathbf{E}_{\mathbf{x}}[r(\mathbf{x}) \cdot f(\mathbf{x})] \geq \text{Corr}_{f,g} - \frac{\epsilon}{4}$. (b) For every $i \in W'$, $\text{Inf}_i^{\leq k}(f) \geq \frac{\epsilon^2}{16k^2}$. Now, with the same value of τ , applying Lemma 3.6 implies that $W' \subseteq V$ with probability $1 - \delta/4$. This finishes the claim. \square

The proof of correctness is now immediate from Claim 4.4 and Claim 4.3. \square

5 The polynomial-query gap tester

Recall the context of the relaxed tester compared to the original one: we have already identified (using Lemma 3.6) a subset \mathcal{S} of size $\text{poly}(k, 1/\epsilon)$ such that all potentially interesting coordinates of f are contained in \mathcal{S} (in the sense of Claim 3.8). Recall, moreover, that we do not have an explicit representation of the coordinates \mathcal{S} ; we only have access to coordinate oracles in the sense of Definition 3.1. Nevertheless, let us first pretend that we do have explicit access to \mathcal{S} , in order to explain roughly how our algorithm works.

The first step of our algorithm is to replace f by $f_{\text{avg},\mathcal{S}}$, defined by $f_{\text{avg},\mathcal{S}}(x) = \mathbf{E}_{\mathbf{y} \sim \text{Unif}}[f(\mathbf{y}) \mid \mathbf{y}_{\mathcal{S}} = x_{\mathcal{S}}]$. This step turns out to be unnecessary if we have explicit access to \mathcal{S} , but on the other hand it is also clearly harmless, because $f_{\text{avg},\mathcal{S}}$ is the orthogonal projection of f onto the space of \mathcal{S} -juntas, and hence $\mathbf{E}[fg] = \mathbf{E}[f_{\text{avg},\mathcal{S}}g]$ for every $g \in \mathcal{J}_{\mathcal{S},k}$.

The second step of our algorithm is to replace $f_{\text{avg},\mathcal{S}}$ with f_{smooth} , defined by $f_{\text{smooth}} = T_{1-\epsilon/(2k)} f_{\text{avg},\mathcal{S}}$. Applying this noise has two nice features: it approximately preserves correlation with k -Juntas, and it allows us to bound total influence. First, we will show that it preserves correlation:

Lemma 5.1. For every k -Junta g , if $s = \epsilon/(2k)$ then

$$\|g - T_{1-s}g\|_2 \leq \frac{\epsilon}{2} \|g\|_2.$$

In particular, for all $f : \{\pm 1\}^n \rightarrow [-1, 1]$ and $g \in \mathcal{J}_{\mathcal{S},k}$

$$|\mathbf{E}[(T_{1-s}f)g] - \mathbf{E}[fg]| \leq \epsilon/2.$$

Consequently, $\max_{g \in \mathcal{J}_{\mathcal{S},k}} \mathbf{E}[f_{\text{smooth}}g]$ is essentially the same as $\max_{g \in \mathcal{J}_{\mathcal{S},k}} \mathbf{E}[f_{\text{avg},\mathcal{S}}g]$.

Proof. If g is a k -Junta then $\hat{g}_{\mathcal{S}} = 0$ whenever $|\mathcal{S}| > k$, and so

$$\|g - T_{1-s}g\|_2^2 = \sum_{|\mathcal{S}| \leq k} (1 - (1-s)^{|\mathcal{S}|})^2 \hat{g}^2(\mathcal{S}) \leq \frac{\epsilon^2}{4} \sum_{\mathcal{S}} \hat{g}^2(\mathcal{S}) \leq \frac{\epsilon^2}{4} \|g\|_2^2.$$

The second claim follows from the fact that $\mathbf{E}[(T_{1-s}f)g] = \mathbf{E}[f(T_{1-s}g)]$, and so $\mathbf{E}[(T_{1-s}f - f)g] = \mathbf{E}[f(T_{1-s}g - g)]$. Then apply Cauchy-Schwarz and the first claim (together with the fact that $\|f\|_2$ and $\|g\|_2$ are at most 1). \square

As we claimed above, we can also bound the total influence of a smoothed function:

Lemma 5.2. *If $f : \{\pm 1\}^n \rightarrow [-1, 1]$ then $\text{Inf}(T_{1-s}f) \leq \frac{1}{2s}$. In particular,*

$$|\{1 \leq j \leq n : \text{Inf}_j(T_{1-s}f) \geq 2s\}| \leq \frac{1}{4s^2}.$$

Proof. Since f takes values in $[-1, 1]$, $\mathbf{E}[f^2] \leq 1$ and so $\sum_S \hat{f}^2(S) \leq 1$. On the other hand,

$$\text{Inf}(T_{1-s}f) = \sum_S |S|(1-s)^{|S|} \hat{f}^2(S) \leq \max\{j(1-s)^j : 1 \leq j \leq n\} \leq \frac{1}{es} \leq \frac{1}{2s}$$

This proves the first claim. The second follows from Markov's inequality. \square

According to Lemma 5.2, if $S' = \{i \in S : \text{Inf}_i(f_{\text{smooth}}) \geq \epsilon/k\}$ then $|S'| \leq k^2/\epsilon^2 = k'$. On the other hand, we can easily see that coordinates not belonging to S' are irrelevant when it comes to approximating f_{smooth} by a k -Junta:

Lemma 5.3. *For $f : \{\pm 1\}^n \rightarrow [-1, 1]$, let $S' \subset S \subset [n]$ satisfy $i \in S'$ whenever $\text{Inf}_i(f) \geq t$. Then*

$$\left| \max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[fg] - \max_{g \in \mathcal{J}_{S',k}} \mathbf{E}[fg] \right| \leq tk.$$

Proof. Let $g \in \mathcal{J}_{S,k}$ maximize $\mathbf{E}[fg]$, and let \mathcal{T} be the set of coordinates on which g depends. Define $g_{\text{avg},S'}$ by $g_{\text{avg},S'}(x) = \mathbf{E}_{\mathbf{y} \sim \text{Unif}}[g(\mathbf{y}) \mid \mathbf{y}_S = x_S]$. Then $g_{\text{avg},S'}$ is a $(S' \cap \mathcal{T})$ -Junta; it does not necessarily belong to $\mathcal{J}_{S',k}$ because it is not necessarily Boolean. However, $\max_{h \in \mathcal{J}_{S',k}} \mathbf{E}[fh] \geq \mathbf{E}[fg_{\text{avg},S'}]$ – this is because $g_{\text{avg},S'}$ can be rounded to a Boolean function without decreasing $\mathbf{E}[fg_{\text{avg},S'}]$. Therefore, it suffices to show that

$$\mathbf{E}[fg_{\text{avg},S'}] \geq \mathbf{E}[fg] - tk.$$

To do this, note that the Fourier coefficients of g and $g_{\text{avg},S'}$ are related by $\hat{g}_S = (\widehat{g_{\text{avg},S'}})_S$ whenever $S \subset S'$, and $(\widehat{g_{\text{avg},S'}})_S = 0$ otherwise. In particular,

$$\mathbf{E}[f(g - g_{\text{avg},S'})] = \sum_{\substack{S \subset \mathcal{T} \\ S \not\subset S'}} \hat{g}_S \hat{f}_S \leq \sum_{\substack{S \subset \mathcal{T} \\ S \not\subset S'}} \hat{f}_S^2 \leq \sum_{i \in \mathcal{T} \setminus S'} \text{Inf}_i(f) \leq tk,$$

where the last inequality follows because $|\mathcal{T} \setminus S'| \leq k$. \square

Thanks to Lemma 5.3, if we set $S' = \{i \in S : \text{Inf}_i(f_{\text{smooth}}) \geq \epsilon/k\}$ then $\max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[f_{\text{smooth}}g] \approx \max_{g \in \mathcal{J}_{S',k}} \mathbf{E}[f_{\text{smooth}}g]$. We do not know how to estimate this final quantity, but we can easily estimate (the larger quantity) $\max_{g \in \mathcal{J}_{S',k'}} \mathbf{E}[f_{\text{smooth}}g]$ (note that k has become k'). This is because $\mathcal{J}_{S',k'}$ consists of all boolean functions depending on the coordinates in S' . The one with maximal correlation can be found by projecting f_{smooth} onto the space of S' -Juntas and rounding the result to a boolean function: define $f_{\text{smooth},\text{avg},S'}$ by $f_{\text{smooth},\text{avg},S'}(x) = \mathbf{E}[f_{\text{smooth}}(\mathbf{y}) \mid \mathbf{y}_{S'} = x_{S'}]$. Then

$$\max_{g \in \mathcal{J}_{S',k'}} \mathbf{E}[f_{\text{smooth}}g] = \mathbf{E}[|f_{\text{smooth},\text{avg},S'}|].$$

This number (or rather, an estimate of it) will be the final output of our algorithm. Thanks to the preceding arguments, it is larger (or at least, not much smaller) than $\max_{g \in \mathcal{J}_{n,k}} \mathbf{E}[fg]$. On the other hand, it is certainly smaller than $\max_{g \in \mathcal{J}_{n,k'}} \mathbf{E}[fg]$.

In order to turn the description above into an algorithm, we need to describe how to compute all the quantities above given implicit access to the coordinates in S . In particular, we will first describe

how to simulate query access to $f_{\text{avg}, \mathcal{S}}$. Given this, it is obvious how to simulate query access to f_{smooth} . Then, we will show how to estimate \mathcal{S}' ; more accurately, we will compute some $\tilde{\mathcal{S}}'$ that contains all coordinates of influence at least $2\epsilon/2k$ and no coordinates of influence smaller than ϵ/k . Finally, we will show how to simulate query access to $f_{\text{smooth}, \text{avg}, \tilde{\mathcal{S}}'}$; we can use this query access to estimate $\mathbf{E}[|f_{\text{smooth}, \text{avg}, \tilde{\mathcal{S}}'}|]$, which is the final output of our algorithm.

5.1 Averaging over irrelevant coordinates

An important primitive for us will be the ability to average over irrelevant coordinates given oracle access to the relevant ones. That is, imagine that there is a collection of coordinates $\mathcal{S} \subset [n]$. Given query access to $f : \{\pm 1\}^n \rightarrow [-1, 1]$ and a fixed point $x \in \{\pm 1\}^{\mathcal{S}}$, we would like to estimate $\mathbf{E}_{\mathbf{y} \sim \text{Unif}}[f(\mathbf{y}) \mid \mathbf{y}_{\mathcal{S}} = x]$. Were we given \mathcal{S} explicitly, this would be easy; our challenge is to do it with only oracle access to \mathcal{S} , in the sense of Definition 3.1.

Input: a function $f, x \in \{\pm 1\}^n$, an oracle \mathcal{D}, γ (accuracy parameter), δ (failure parameter)
Output: a number

- 1 Let $T = C|\mathcal{D}| \frac{\log 1/\delta}{\gamma^2}$;
- 2 Let $\mathbf{x}^{(1)} = x$;
- 3 **for** $i = 1$ to $T - 1$ **do**
- 4 **repeat**
- 5 Let \mathbf{y} be a copy of $\mathbf{x}^{(i)}$, but flip each bit independently with probability $\frac{1}{2|\mathcal{D}|}$;
- 6 **until** $g(\mathbf{y}) = g(x)$ for all $g \in \mathcal{D}$;
- 7 Let $\mathbf{x}^{(i+1)} = \mathbf{y}$;
- 8 **return** $\frac{1}{T} \sum_{i=1}^T f(\mathbf{x}^{(i)})$

Algorithm 3: Coordinate-projection

Lemma 5.4. *Coordinate-projection has the following guarantees. Given a function $f : \{\pm 1\}^n \rightarrow [-1, 1]$, a point $x \in \{\pm 1\}^n$, an oracle \mathcal{D} for the coordinates $\mathcal{S} \subset [n]$, and parameters $\delta, \gamma > 0$, the algorithm makes (in expectation) $\Theta(|\mathcal{S}| \frac{\log 1/\delta}{\gamma^2})$ queries to f and to each element of \mathcal{D} and, with probability at least $1 - \delta$, outputs a number within γ of $\mathbf{E}_{\mathbf{y} \sim \text{Unif}}[f(\mathbf{y}) \mid \mathbf{y}_{\mathcal{S}} = x]$.*

Remark 5.5. *It is not hard to see that **Coordinate-projection** succeeds even if it is given noisy query access to f . For example, if each query that **Coordinate-projection** makes to the function f can produce an error of at most γ , then (under the assumptions of Lemma 5.4), the output of **Coordinate-projection** is, with probability at least $1 - \delta$, accurate to within 2γ .*

Proof. First, observe that for every $i = 1, \dots, T$, $\mathbf{x}_{\mathcal{S}}^{(i)} = x_{\mathcal{S}}$: indeed, \mathcal{D} is an oracle for \mathcal{S} and so the test on line 6 will only pass if $\mathbf{y}_{\mathcal{S}} = x_{\mathcal{S}}$. Next, observe that conditioned on $\mathbf{x}^{(i)}$, every coordinate $j \notin \mathcal{S}$ of $\mathbf{x}^{(i+1)}$ is obtained by (independently) flipping $\mathbf{x}_j^{(i)}$ with probability $\frac{1}{2|\mathcal{S}|}$. In other words, $\mathbf{x}_{\mathcal{S}}^{(1)}, \dots, \mathbf{x}_{\mathcal{S}}^{(T)}$ is a Bonami-Beckner Markov chain with correlation parameter $\rho = 1 - \frac{1}{|\mathcal{S}|}$. If \mathbf{z} is distributed according to the stationary distribution of this Markov chain then $\mathbf{z}_{\mathcal{S}} = x_{\mathcal{S}}$ and $\mathbf{z}_{\bar{\mathcal{S}}}$ is uniformly distributed on $\{\pm 1\}^{\bar{\mathcal{S}}}$. In particular, $\mathbf{E}[f(\mathbf{z})] = \mathbf{E}_{\mathbf{y} \sim \text{Unif}}[f(\mathbf{y}) \mid \mathbf{y}_{\mathcal{S}} = x_{\mathcal{S}}]$ and so Lemma 2.5 implies that with probability at least $1 - \delta$, $\frac{1}{T} \sum_{i=1}^T f(\mathbf{x}^{(i)})$ is within γ of $\mathbf{E}_{\mathbf{y}}[f(\mathbf{y}) \mid \mathbf{y}_{\mathcal{S}} = x_{\mathcal{S}}]$.

It remains to check how many oracle queries our procedure makes; let X_i be the number of attempts it takes to successfully generate $\mathbf{x}^{(i)}$. For each i , the probability that the test on line 6 succeeds is exactly $(1 - \frac{1}{2|\mathcal{S}|})^{|\mathcal{S}|} \geq \frac{1}{4}$. In particular, each sample takes (in expectation) at most four queries to each element of \mathcal{D} , and so the overall number of queries is at most (in expectation) $O(T)$ to each element of \mathcal{D} . \square

5.2 Estimating influences

One of the things we need to do is to estimate the influences of coordinates in \mathcal{S} . Again, if we had explicit access to these coordinates then this would be trivial. The point is to get by with only oracle access, and the difficulty is that given some $x \in \{\pm 1\}^n$, we cannot simply “flip” a bit belonging to \mathcal{S} because we don’t know which bits those are. To work around this, we introduce the notion of an influence-testing sample, which essentially is a collection of points that manage to flip each bit in \mathcal{S} :

Definition 5.6. We say that $\mathcal{X} \subset \{\pm 1\}^n$ is an influence-testing sample at $x \in \{\pm 1\}^n$ with respect to \mathcal{S} if we can enumerate $\mathcal{X} = \{y^{(i)} : i \in \mathcal{S}\}$ where $y_i^{(i)} = -x_i$ and $y_j^{(i)} = x_j$ for $j \in \mathcal{S} \setminus \{i\}$.

Note that the definition above doesn’t guarantee anything about bits not belonging to \mathcal{S} . The point of the definition above is that if the function f depends only on coordinates in \mathcal{S} , then we can use an influence-testing sample at a random point to obtain an unbiased estimator for all the influences of f .

The first important point about influence-testing samples is that we can produce them given oracle access to \mathcal{S} .

Input: $x \in \{\pm 1\}^n$, and an oracle \mathcal{D} for \mathcal{S}
Output: an influence-testing sample at x with respect to \mathcal{S}

- 1 Define (for $z \in \{\pm 1\}^n$) $I(z) = \{g \in \mathcal{D} : g(z) \neq g(x)\}$;
- 2 Initialize $\mathcal{X} = \emptyset$;
- 3 **while** $|\mathcal{X}| < |\mathcal{D}|$ **do**
- 4 Let \mathbf{y} be a copy of x , but flip each bit independently with probability $\frac{1}{2^{|\mathcal{D}|}}$;
- 5 **if** $|I(\mathbf{y})| = 1$ and $I(\mathbf{y}) \neq I(z)$ for all $z \in \mathcal{X}$ **then**
- 6 Add \mathbf{y} to \mathcal{X} ;
- 7 **return** \mathcal{X}

Algorithm 4: Influence-testing-sample

Lemma 5.7. Let \mathcal{D} be an oracle for \mathcal{S} . Given \mathcal{D} and $x \in \{\pm 1\}^n$ as input, *Influence-testing-sample* produces an influence-testing sample at x with respect to \mathcal{S} , while making (in expectation) $O(|\mathcal{S}| \log |\mathcal{S}|)$ queries to each element of \mathcal{D} .

Proof. Each element that *Influence-testing-sample* adds to \mathcal{X} differs from x on exactly one coordinate of \mathcal{S} . Moreover, the test on line 5 ensures that every coordinate of \mathcal{S} is represented by at most one element of \mathcal{X} . Hence, by the time the loop is complete, \mathcal{X} is an influence-testing sample at x with respect to \mathcal{S} .

Note that each time we sample \mathbf{y} , we have $\Pr[|I(\mathbf{y})| = 1] = (1 - 2^{|\mathcal{S}|-1})^{|\mathcal{S}|} \geq \frac{1}{4}$. Moreover, conditioned on $|I(\mathbf{y})| = 1$, $I(\mathbf{y})$ is uniformly distributed among all size-one subsets of \mathcal{D} . Hence, the number of times that we need to sample \mathbf{y} in order to see all size-one subsets at least once is distributed according to the coupon collector problem, and hence takes $O(|\mathcal{S}| \log |\mathcal{S}|)$ iterations in expectation. \square

The other important point about influence-testing samples is that if we can use them to estimate influences. The basic idea is to take the trivial algorithm for estimating influences (sample random elements, and check whether flipping the j th bit changes the value of f), but using the notion of an influence-testing sample to replace the need to bits; in *Threshold-influences*, $\mathbf{x}^{(i)}$ are the random points at which we’re testing bit-flips, and $\mathbf{y}_g^{(i)}$ is the copy of $\mathbf{x}^{(i)}$ with a bit (the one corresponding to

$g \in \mathcal{D}$) flipped.

Input: a function f , an oracle \mathcal{D} for \mathcal{S} , threshold parameter t , failure parameter δ

Output: an oracle \mathcal{D}'

- 1 Let $T = Ct^{-2} \log \frac{1}{\delta|\mathcal{D}|}$;
- 2 For $s = 1, \dots, T$, sample $\mathbf{x}^{(i)}$ uniformly from $\{\pm 1\}^n$;
- 3 For $s = 1, \dots, T$, let $\{\mathbf{y}_g^{(i)} : g \in \mathcal{D}\}$ be an influence testing sample at $\mathbf{x}^{(i)}$ (from Influence-testing-sample), where $\mathbf{y}_g^{(i)}$ is the element for which $g(\mathbf{y}^{(i)}) \neq g(\mathbf{x}^{(i)})$;
- 4 For $g \in \mathcal{D}$, let $\widehat{\text{Inf}}_g = \frac{1}{T} \sum_{i=1}^T (f(\mathbf{x}^{(i)}) - f(\mathbf{y}_g^{(i)}))^2$;
- 5 **return** $\mathcal{D}' = \{g \in \mathcal{D} : \widehat{\text{Inf}}_g \geq \frac{3}{2}t\}$

Algorithm 5: Threshold-influences

Lemma 5.8. *Assume that \mathcal{D} is an oracle for \mathcal{S} and that f is a \mathcal{S} -junta. With probability at least $1 - \delta$, the output of **Threshold-influences** satisfies the following: there is a set $\mathcal{S}' \subset \mathcal{S}$ such that $\text{Inf}_i(f) \geq 2t$ implies that $i \in \mathcal{S}'$, and $i \in \mathcal{S}'$ implies that $\text{Inf}_i(f) \geq t$, and the output of **Threshold-influences** is an oracle for \mathcal{S}' .*

*Moreover, **Threshold-influences** makes $O(t^{-2} \log \frac{1}{\delta|\mathcal{S}|})$ queries to f and (in expectation) $O(t^{-2}|\mathcal{S}|^2 \log \frac{1}{\delta})$ queries to each element of \mathcal{D} .*

Remark 5.9. *It is not hard to see that **Threshold-influences** succeeds even if it is given access to a slightly noisy version of f : if each evaluation of f is guaranteed to be accurate to within $t/10$ then the guarantees of Lemma 5.8 still hold (at the cost of increasing the constant C in **Threshold-influences**).*

Proof. Let us fix a single $j \in \mathcal{S}$ and its oracle $g \in \mathcal{D}$. We will show that with probability at least $1 - \frac{\delta}{|\mathcal{S}|}$, if $\text{Inf}_j(f) \leq t$ then $g \in \mathcal{D}'$ and if $\text{Inf}_j(f) \geq 2t$ then $g \notin \mathcal{D}'$ (and then the lemma will follow by a union bound over g). To prove this claim, it suffices to show that (with probability at least $1 - \frac{\delta}{|\mathcal{S}|}$) $|\text{Inf}_j(f) - \widehat{\text{Inf}}_g| \leq t/2$.

For $z \in \{\pm 1\}^n$, let \tilde{z} denote a copy of z with bit j flipped. Recall that $\text{Inf}_j(f) = \mathbf{E}_{\mathbf{z} \sim \text{Unif}}[(f(\mathbf{z}) - f(\tilde{\mathbf{z}}))^2]$; by a Chernoff bound, with probability at least $1 - \frac{\delta}{|\mathcal{S}|}$, $\text{Inf}_j(f)$ is within $t/2$ of

$$\frac{1}{T} \sum_{i=1}^T \left(f(\mathbf{x}^{(i)}) - f(\tilde{\mathbf{x}}^{(i)}) \right)^2,$$

where $\mathbf{x}^{(i)}$ are (as in **Threshold-influences**) independent and uniform in $\{\pm 1\}^n$. Finally, recall that f was assumed to be a \mathcal{S} -Junta, and recall that among coordinates of \mathcal{S} , $\mathbf{y}_g^{(i)}$ differs from $\mathbf{x}^{(i)}$ exactly in coordinate j . Hence, $f(\mathbf{y}_g^{(i)}) = f(\tilde{\mathbf{x}}^{(i)})$ for all i , and so $\widehat{\text{Inf}}_g$ coincides with the displayed quantity above. This proves the claim about the correctness of **Threshold-influences**; the claims about the number of queries follow immediately from the algorithm and Lemma 5.7. \square

5.3 The algorithm

In order to set up the algorithm, recall from Section 3 that we can begin by finding an oracle \mathcal{D} to a poly(k)-sized set of relevant coordinates \mathcal{S} . As explained before, (e.g. by applying Claim 3.8), in order to complete the task it suffices to compute $\max_{g \in \mathcal{J}_{\mathcal{S}, k'}} \mathbf{E}[fg]$ to accuracy ϵ , where $k' = k^2/\epsilon^2$.

Recall that **Coordinate-projection** and **Threshold-influences** have a failure probability $\delta > 0$; we will fix take $\gamma = \text{poly}(\epsilon, 1/k)$ and $\delta = 2^{-(k/\gamma)}$, and in what follows we will guarantee to invoke these two algorithms at most $\text{poly}(k, 1/\epsilon)$ times, meaning that with high probability every single invocation will succeed.

Here is the final algorithm, together with a justification of the sample complexity:

- $f_{\text{avg},S} : \{\pm 1\} \rightarrow [-1, 1]$ is defined by $f_{\text{avg},S}(x) = \mathbf{E}_{\mathbf{y} \sim \text{Unif}}[f(\mathbf{y}) \mid \mathbf{y}_S = x_S]$. According to Lemma 5.4, for every $x \in \{\pm 1\}^n$ we can approximate $f_{\text{avg},S}(x)$ to within γ using $\text{poly}(k, 1/\epsilon)$ queries. Moreover (recalling our choice of δ), as long as we repeat this process at most $\text{poly}(k, 1/\epsilon)$ times, with high probability we will never fail to obtain accuracy γ . Also, $f_{\text{avg},S}$ is an S -Junta.
- f_{smooth} is defined by $f_{\text{smooth}} = T_{1-\epsilon/(2k)} f_{\text{avg},S}$. By naïve sampling and a Chernoff bound, for every x in $\{\pm 1\}^n$ we can (with probability at least $1 - \delta$) approximate $f_{\text{smooth}}(x)$ to within 2γ using $\text{poly}(1/\gamma, \log(1/\delta))$ queries (each with accuracy γ) to $f_{\text{avg},S}$. Thanks to the previous point, this can be done using $\text{poly}(k, 1/\epsilon)$ queries to f , and thanks to the choice of δ we can repeat this $\text{poly}(k, 1/\epsilon)$ times without failure.
- Thanks to Lemma 5.8 and Remark 5.9 (with $t = \epsilon/k$), we can compute an oracle \tilde{D}' to some set \tilde{S}' such that $i \in \tilde{S}'$ whenever $\text{Inf}_i(f_{\text{smooth}}) \geq 2\epsilon/k$, and $\text{Inf}_i(f_{\text{smooth}}) < \epsilon/k$ for all $i \in \tilde{S}'$. This can be done using $\text{poly}(k, 1/\epsilon)$ queries to f_{smooth} , which (thanks to the previous point) can be done using $\text{poly}(k, 1/\epsilon)$ queries to f . By Lemma 5.2, $|\tilde{S}'| \leq k^2/\epsilon^2$.
- Applying Lemma 5.4 again (this time, also applying Remark 5.5), for every $x \in \{\pm 1\}^n$ we can approximate $f_{\text{smooth,avg},\tilde{S}'}(x) := \mathbf{E}_{\mathbf{y}}[f_{\text{smooth}}(\mathbf{y}) \mid \mathbf{y}_{\tilde{S}'} = x_{\tilde{S}'}]$ to within 3γ using $\text{poly}(k, 1/\epsilon)$ queries to f_{smooth} (which in turn requires $\text{poly}(k, 1/\epsilon)$ queries to f).
- Finally, use $\text{poly}(k, 1/\epsilon)$ independent samples to estimate $\mathbf{E}_{\mathbf{y}}[|f_{\text{smooth,avg},\tilde{S}'}|]$ to within accuracy 4γ (which we can assume is at most ϵ). This estimate is the output of the algorithm.

To complete the proof that this algorithm is correct (i.e. it fulfills the claims made in Theorem 1.5), let us combine our previous bounds to show that the final quantity being estimated in our algorithm (namely, $\mathbf{E}[|f_{\text{smooth,avg},\tilde{S}'}|]$) is approximately between $\max_{g \in \mathcal{J}_{S,k'}} \mathbf{E}[fg]$ and $\max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[fg]$.

Lemma 5.10. *With $f_{\text{smooth,avg},\tilde{S}'}$ defined as above,*

$$\max_{g \in \mathcal{J}_{S,k'}} \mathbf{E}[fg] \geq \mathbf{E}[|f_{\text{smooth,avg},\tilde{S}'}|] = \max_{g \in \mathcal{J}_{n,k'}} \mathbf{E}[f_{\text{smooth,avg},\tilde{S}'}g] \geq \max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[fg] - \frac{3}{2}\epsilon.$$

Proof. The first inequality follows immediately from Jensen's inequality and the fact that $f_{\text{smooth,avg},\tilde{S}'}$ is defined by repeatedly averaging f in various senses; we will focus on the last inequality. As we discussed previously,

$$\max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[fg] = \max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[f_{\text{avg},S}g].$$

By Lemma 5.1 applied to $f_{\text{avg},S}$,

$$\max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[f_{\text{smooth}}g] \geq \max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[f_{\text{avg},S}g] - \epsilon/2.$$

By Lemma 5.3 applied with $S' = \tilde{S}'$ and $t = 2\epsilon/k$,

$$\max_{g \in \mathcal{J}_{S',k}} \mathbf{E}[f_{\text{smooth}}g] \geq \max_{g \in \mathcal{J}_{S,k}} \mathbf{E}[f_{\text{smooth}}g] - \epsilon.$$

Finally (and for the same reason as the first point),

$$\mathbf{E}[|f_{\text{smooth,avg},\tilde{S}'}|] = \max_{g \in \mathcal{J}_{S',k'}} \mathbf{E}[f_{\text{smooth,avg},\tilde{S}'}g] = \max_{g \in \mathcal{J}_{S',k'}} \mathbf{E}[f_{\text{smooth}}g].$$

□

References

- [1] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability-towards tight results. *SIAM Journal on Computing*, 27(3):804–915, 1998.
- [2] E. Blais. Improved bounds for testing juntas. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 317–330. Springer, 2008.
- [3] E. Blais. Testing juntas nearly optimally. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 151–158. ACM, 2009.
- [4] E. Blais, C. Canonne, T. Eden, A. Levi, and D. Ron. Tolerant junta testing and the connection to submodular optimization and function isomorphism. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2113–2132. Society for Industrial and Applied Mathematics, 2018.
- [5] E. Blais and D. Kane. Tight bounds for testing k-linearity. In *APPROX RANDOM*, pages 435–446. Springer, 2012.
- [6] A. Blum. Relevant examples and relevant features: Thoughts from computational learning theory. in *AAAI Fall Symposium on ‘Relevance’*, 1994.
- [7] A. Blum and P. Langley. Selection of relevant features and examples in machine learning. *Artificial Intelligence*, 97(1-2):245–271, 1997.
- [8] S. Chakraborty, E. Fischer, D. García-Soriano, and A. Matsliah. Junto-symmetric functions, hypergraph isomorphism and crunching. In *27th Annual Conference on Computational Complexity (CCC)*, pages 148–158. IEEE, 2012.
- [9] X. Chen, Z. Liu, Rocco A. Servedio, Y. Sheng, and J. Xie. Distribution free junta testing. In *Proceedings of the ACM STOC 2018*, 2018.
- [10] Xi Chen, Rocco A. Servedio, Li-Yang Tan, Erik Waingarten, and Jinyu Xie. Settling the query complexity of non-adaptive junta testing. In *Proceedings of the 32Nd Computational Complexity Conference*, pages 26:1–26:19, 2017.
- [11] A. De. Boolean Function Analysis Meets Stochastic Optimization: An Approximation Scheme for Stochastic Knapsack. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA ’18*, pages 1286–1305, 2018.
- [12] Anindya De, Elchanan Mossel, and Joe Neeman. Is your function low-dimensional? 2018. at <https://arxiv.org/abs/1806.10057>.
- [13] 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.
- [14] E. Fischer, G. Kindler, D. Ron, S. Safra, and A. Samorodnitsky. Testing juntas. *J. Computer & System Sciences*, 68(4):753–787, 2004.
- [15] M. Furst, J. Saxe, and M. Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984.
- [16] A. Garg, Y.-T. Lee, Z. Song, and N. Srivastava. A matrix expander Chernoff bound. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1102–1114, 2018.

- [17] D. Gillman. A chernoff bound for random walks on expander graphs. *SIAM Journal on Computing*, 27(4):1203–1220, 1998.
- [18] J. Håstad. *Computational Limitations for Small Depth Circuits*. MIT Press, Cambridge, MA, 1986.
- [19] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001.
- [20] A. Healy. Randomness-Efficient Sampling within NC. *Computational Complexity*, 17(1):3–37, 2008.
- [21] D. Levin and Y. Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- [22] P. Lezaud. Chernoff-type bound for finite Markov chains. *Annals of Applied Probability*, pages 849–867, 1998.
- [23] R. O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, Cambridge, 2014.
- [24] M. Parnas, D. Ron, and R. Rubinfeld. Tolerant property testing and distance approximation. *Journal of Computer and System Sciences*, 72(6):1012–1042, 2006.
- [25] M. Parnas, D. Ron, and A. Samorodnitsky. Testing Basic Boolean Formulae. *SIAM J. Disc. Math.*, 16:20–46, 2002.
- [26] M. Sağlam. Near Log-Convexity of Measured Heat in (Discrete) Time and Consequences. In *59th IEEE Annual Symposium on Foundations of Computer Science*, pages 967–978, 2018.
- [27] R. Servedio, L-Y. Tan, and J. Wright. Adaptivity helps for testing juntas. In *Proceedings of CCC*, volume 33. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- [28] S. Vadhan. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012.

A Description of the algorithm Find-best-fit

In this section, we describe the algorithm Find-best-fit (from Lemma 4.1) and give its proof of correctness. We begin by describing the algorithm.

Input: f (target function), k (arity of Junta), ϵ (distance parameter)
 \mathcal{S} (oracles to dictator functions), δ (confidence parameter)
Output: $\widehat{\text{Corr}}_{f,\mathcal{S},k} \in [0, 1]$ and $h_{\mathcal{S},k} : \{-1, 1\}^k \rightarrow \{-1, 1\}$
 // Set parameters for the algorithm
 1 Set $N = 2^k \cdot \frac{1}{\epsilon^2} \cdot 4 \cdot (|\mathcal{S}| + \log(1/\delta) + k^2)$ and $M = N \cdot 2^{-k}$;
 // Evaluation of the oracles
 2 Let $N' = \text{Poi}(N)$ and sample $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N')} \in \{-1, 1\}^n$;
 3 Evaluate f on $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N')}$;
 4 For all $\mathcal{S} \ni \text{Dict}_j(\cdot)$, evaluate Dict_j on $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N')}$;
 // Hypothesis testing
 5 **for** $\mathcal{T} \subseteq \mathcal{S}$ and $|\mathcal{T}| = k$ **do**
 6 Let $\mathcal{T} = \{\text{Dict}_{i_1}, \dots, \text{Dict}_{i_k}\}$;
 7 **for** $y \in \{-1, 1\}^k$ **do**
 8 Define $\mathcal{A}_{\mathcal{T},y} = \{\mathbf{x}^{(s)} : \text{for } 1 \leq s \leq N', 1 \leq j \leq k \text{ and } \text{Dict}_{i_j}(\mathbf{x}^{(s)}) = y_j\}$;
 9 Define $\widehat{\text{Corr}}_{\mathcal{T},y} = \frac{\sum_{\mathbf{x}^{(s)} \in \mathcal{A}_{\mathcal{T},y}} f(\mathbf{x}^{(s)})}{M}$
 10 Define $\widehat{\text{Corr}}_{\mathcal{T}} = \frac{\sum_y |\widehat{\text{Corr}}_{\mathcal{T},y}|}{2^k}$ and $h_{\mathcal{T}} : \{-1, 1\}^k \rightarrow \{-1, 1\}$ as $h_{\mathcal{T}}(y) = \text{sign}(\widehat{\text{Corr}}_{\mathcal{T},y})$.
 11 Define $\mathcal{T}^* = \arg \max_{\mathcal{T} \subseteq \mathcal{S}: |\mathcal{T}|=k} \widehat{\text{Corr}}_{\mathcal{T}}$;
 12 Output $\widehat{\text{Corr}}_{f,\mathcal{S},k} \leftarrow \widehat{\text{Corr}}_{\mathcal{T}^*}$ and $h_{\mathcal{S},k} \leftarrow h_{\mathcal{T}^*} : \{-1, 1\}^k \rightarrow \{-1, 1\}$.

Algorithm 6: Find-best-fit

For convenience of the reader, we restate Lemma 4.1 and then give the proof.

Lemma 4.1. *There is an algorithm Find-best-fit with the following guarantee:*

1. The algorithm gets as input oracle access to a function $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ as well as a set of oracles $\mathcal{S} \subseteq \{\text{Dict}_1, \dots, \text{Dict}_n\}$. We clarify that algorithm is only given the oracle Dict_i (for $\text{Dict}_i \in \mathcal{S}$) but not i .
2. The algorithm gets as input error parameter $\epsilon > 0$, arity parameter k and confidence parameter $\delta > 0$.
3. The algorithm makes $N(k, |\mathcal{S}|, \epsilon, \delta) = O(2^k/\epsilon^2 \cdot |\mathcal{S}| \cdot (\log(1/\delta) + k^2 + |\mathcal{S}|))$ queries with probability $1 - \delta$.
4. Each query point (to either f or oracle in \mathcal{S}) is distributed as a uniformly random element of $\{-1, 1\}^n$.
5. With probability $1 - \delta$, the algorithm outputs a number $\widehat{\text{Corr}}_{f,\mathcal{S},k}$ and $h_{\mathcal{S},k} : \{-1, 1\}^k \rightarrow \{-1, 1\}$ such that there exists oracles $\text{Dict}_{i_1}, \dots, \text{Dict}_{i_k} \in \mathcal{S}$

$$\mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot h_{\mathcal{S},k}(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k})] \geq \max_{\ell \in \mathcal{J}_{\mathcal{S},k}} \mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot \ell(\mathbf{x})] - \epsilon \text{ and } |\widehat{\text{Corr}}_{f,\mathcal{S},k} - \max_{\ell \in \mathcal{J}_{\mathcal{S},k}} \mathbf{E}_{\mathbf{x}}[f(\mathbf{x}) \cdot \ell(\mathbf{x})]| \leq \epsilon.$$

Proof. The proof of Item 4 is immediate from Step 2 of the algorithm. Similarly, note that the total number of queries made is $|\mathcal{S}| \cdot \text{Poi}(N)$ where $N = O(2^k/\epsilon^2 \cdot (\log(1/\delta) + k^2 + |\mathcal{S}|))$. Item 3 now follows from tail bounds on the Poisson distribution.

Thus, it just remains to prove Item 5. To prove this, consider a fixed \mathcal{T} (from Step 5 of the algorithm Find-best-fit). Define $\mathcal{B}_{\mathcal{T},y} = \{x \in \{-1, 1\}^n : \text{for } 1 \leq j \leq k, \text{Dict}_{i_j}(x) = y_j\}$. Observe that $\mathcal{A}_{\mathcal{T},y}$ can

be seen as a sampling of $\text{Poi}(M)$ elements from $\mathcal{B}_{\mathcal{T},y}$. Let us define $\text{Corr}_{\mathcal{T},y} = \mathbf{E}_{\mathbf{x} \in \mathcal{B}_{\mathcal{T},y}}[f(\mathbf{x})]$. Then, observe that $\mathbf{E}_{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}}[\widehat{\text{Corr}}_{\mathcal{T},y}] = \text{Corr}_{\mathcal{T},y}$. In fact, by Chernoff-Hoeffding bounds, it follows that

$$\Pr_{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}}[|\text{Corr}_{\mathcal{T},y} - \widehat{\text{Corr}}_{\mathcal{T},y}| > \epsilon/2] \leq \frac{2^{-k} \cdot \delta}{10 \cdot 2^{|\mathcal{S}|}}.$$

By a union bound, this implies that

$$\text{For all } y \in \{-1, 1\}^k, \Pr_{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}}[|\text{Corr}_{\mathcal{T},y} - \widehat{\text{Corr}}_{\mathcal{T},y}| > \epsilon/2] \leq \frac{\delta}{10 \cdot 2^{|\mathcal{S}|}}. \quad (1)$$

This implies that for any subset \mathcal{T} ,

$$\Pr_{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}}[|\text{Corr}_{\mathcal{T},y} - \widehat{\text{Corr}}_{\mathcal{T},y}| > \epsilon/2] \leq \frac{\delta}{10 \cdot 2^{|\mathcal{S}|}}. \quad (2)$$

Finally, note that for any subset \mathcal{T} ,

$$\begin{aligned} \mathbf{E}_{\mathbf{x} \in \{-1, 1\}^n} [h_{\mathcal{T}}(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k}) \cdot f(\mathbf{x})] &= \mathbf{E}_{\mathbf{y} \in \{-1, 1\}^k} \mathbf{E}_{\mathbf{x} \in \mathcal{A}_{\mathcal{T},y}} [f(\mathbf{x}) \cdot \text{sign}(\widehat{\text{Corr}}_{\mathcal{T},y})] \\ &= \mathbf{E}_{\mathbf{y} \in \{-1, 1\}^k} [\text{sign}(\widehat{\text{Corr}}_{\mathcal{T},y}) \cdot \text{Corr}_{\mathcal{T},y}] \\ &\geq \mathbf{E}_{\mathbf{y} \in \{-1, 1\}^k} [\text{Corr}_{\mathcal{T},y}] - 2 \max_{\mathbf{y} \in \{-1, 1\}^k} [|\text{Corr}_{\mathcal{T},y} - \widehat{\text{Corr}}_{\mathcal{T},y}|] \end{aligned}$$

Using (1), we have that

$$\Pr_{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}} [\mathbf{E}_{\mathbf{x} \in \{-1, 1\}^n} [h_{\mathcal{T}}(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k}) \cdot f(\mathbf{x})] \geq \text{Corr}_{\mathcal{T}} - \epsilon] \geq 1 - \frac{\delta}{10 \cdot 2^{|\mathcal{S}|}}. \quad (3)$$

A union bound over all subsets $\mathcal{T} \subseteq \mathcal{S}$ on (2) and (3) yields Item 5. \square