

On closeness to k -wise uniformity

Ryan O'Donnell* Yu Zhao*

Abstract

A probability distribution over $\{-1, 1\}^n$ is (ϵ, k) -wise uniform if, roughly, it is ϵ -close to the uniform distribution when restricted to any k coordinates. We consider the problem of how far an (ϵ, k) -wise uniform distribution can be from any globally k -wise uniform distribution. We show that every (ϵ, k) -wise uniform distribution is $O(n^{k/2}\epsilon)$ -close to a k -wise uniform distribution in total variation distance. In addition, we show that this bound is optimal for all even k : we find an (ϵ, k) -wise uniform distribution that is $\Omega(n^{k/2}\epsilon)$ -far from any k -wise uniform distribution in total variation distance. For $k = 1$, we get a better upper bound of $O(\epsilon)$, which is also optimal.

One application of our closeness result is to the sample complexity of testing whether a distribution is k -wise uniform or δ -far from k -wise uniform. We give an upper bound of $O(n^k/\delta^2)$ (or $O(\log n/\delta^2)$ when $k = 1$) on the required samples. We show an improved upper bound of $\tilde{O}(n^{k/2}/\delta^2)$ for the special case of testing fully uniform vs. δ -far from k -wise uniform. Finally, we complement this with a matching lower bound of $\Omega(n/\delta^2)$ when $k = 2$.

Our results improve upon the best known bounds from [AAK⁺07], and have simpler proofs.

1 Introduction

1.1 k -wise uniformity and almost k -wise uniformity

We say that a probability distribution over $\{-1, 1\}^n$ is k -wise uniform if its marginal distribution on every subset of k coordinates is the uniform distribution. For Fourier analysis of the Hamming cube, it is convenient to identify the distribution with its density function $\varphi : \{-1, 1\}^n \rightarrow \mathbb{R}^{\geq 0}$ satisfying

$$\mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [\varphi(\mathbf{x})] = 1.$$

We write $\mathbf{x} \sim \varphi$ to denote that \mathbf{x} is a random variable drawn from the associated distribution with density φ :

$$\Pr_{\mathbf{x} \sim \varphi} [\mathbf{x} = x] = \frac{\varphi(x)}{2^n}$$

for any $x \in \{-1, 1\}^n$. Then a well-known fact is that a distribution is k -wise uniform if and only if the Fourier coefficient of φ is 0 on every subset $S \subseteq [n]$ of size between 1 and k :

$$\widehat{\varphi}(S) = \mathbf{E}_{\mathbf{x} \sim \varphi} \left[\prod_{i \in S} x_i \right] = 0.$$

k -wise uniformity is an essential tool in theoretical computer science. Its study dates back to work of Rao [Rao47]. He studied k -wise uniform sets, which are special cases of k -wise uniform distribution. A subset of $\{-1, 1\}^n$ is a k -wise uniform set if the uniform distribution on this subset is k -wise uniform. Rao gave constructions of a pairwise-uniform set of size $n + 1$ (when $n = 2^r - 1$ for any integer r), a 3-wise uniform set of size $2n$ (when $n = 2^r$ for any integer r), and a lower bound (reproved in [ABI86, CGH⁺85]) that a k -wise uniform set on $\{-1, 1\}^n$ requires size at least $\Omega(n^{\lfloor k/2 \rfloor})$. An alternative proof of the lower bound for even k is shown in [AGM03] using a hypercontractivity-type technique, as opposed to the linear algebra method. Coding theorists have also heavily studied k -wise uniformity, since MacWilliams and Sloane showed that

*Supported by NSF grants CCF-1618679, CCF-1717606. This material is based upon work supported by the National Science Foundation under grant numbers listed above. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation (NSF).

linear codes with dual minimum distance $k+1$ correspond to k -wise uniform sets in [MS77]. The importance in theoretical computer science of k -wise independence for derandomization arose simultaneously in many papers, with [KW85, Lub86] emphasizing derandomization via the most common pairwise-uniformity case, and [ABI86, CGH⁺85] emphasizing derandomization based on k -wise independence more generally.

A distribution is “almost k -wise uniform” if its marginal distribution on every k coordinates is very close to the uniform distribution. Typically we say two distributions φ, ψ are δ -close, if the total variation distance between φ and ψ is at most δ ; and we say they are δ -far, if the total variation distance between them is more than δ . However the precise notion of “close to uniform” has varied in previous work. Suppose ψ is the density function for the marginal distribution of φ restricted to some specific k coordinates and $\mathbf{1}$ is the density function for the uniform distribution. Several standard ways are introduced in [AGM03, AAK⁺07] to quantify closeness to uniformity, corresponding to the L_1, L_2, L_∞ norms:

- (L_1 norm): $\|\psi - \mathbf{1}\|_1 = 2d_{\text{TV}}(\psi, \mathbf{1}) \leq \epsilon$, where d_{TV} denotes total variation distance;
- (L_2 norm): $\|\psi - \mathbf{1}\|_2 = \sqrt{\chi^2(\psi, \mathbf{1})} = \sqrt{\sum_{S \neq \emptyset} \widehat{\psi}(S)^2} \leq \epsilon$, where $\chi^2(\psi, \mathbf{1})$ denotes the χ^2 -divergence of ψ from the uniform distribution;
- (L_∞ norm): $\|\psi - \mathbf{1}\|_\infty \leq \epsilon$, or in other words, for any $x \in \{-1, 1\}^n$,

$$\left| \Pr_{\mathbf{x} \sim \psi} [\mathbf{x} = x] - 2^{-k} \right| \leq 2^{-k} \epsilon.$$

Note the following: First, closeness in L_1 norm is the most natural for algorithmic derandomization purposes: it tells us that the algorithm cannot tell ψ is different from the uniform distribution up to ϵ error. Second, these definitions of closeness are in increasing order of strength. On the other hand, we also have that $\|\psi - \mathbf{1}\|_1 \leq \|\psi - \mathbf{1}\|_\infty \leq 2^k \|\psi - \mathbf{1}\|_1$; thus all of these notions are within a factor of 2^k . We generally consider k to be constant (or at worst, $O(\log n)$), so that these notions are roughly the same.

A fourth reasonable notion, proposed by Naor and Naor in [NN93], is that the distribution has a small bias over every non-empty subset of at most k coordinates. We say density function φ is (ϵ, k) -wise uniform if for every non-empty set $S \subseteq [n]$ with size at most k ,

$$|\widehat{\varphi}(S)| = \left| \Pr_{\mathbf{x} \sim \varphi} \left[\prod_{i \in S} x_i = 1 \right] - \Pr_{\mathbf{x} \sim \varphi} \left[\prod_{i \in S} x_i = -1 \right] \right| \leq \epsilon.$$

Here we also have $\epsilon = 0$ if and only if φ is exactly k -wise uniform. Clearly if the marginal density of φ over every k coordinates is ϵ -close to the uniform distribution in total variation distance, then φ is (ϵ, k) -wise uniform. On the other hand, if φ is (ϵ, k) -wise uniform, then the marginal density of φ over every k coordinates is $2^{k/2}\epsilon$ -close to uniform distribution in total variation distance. Again, if k is considered constant, this bias notion is also roughly the same as previous notions. In the rest of paper we prefer this (ϵ, k) -wise uniform notion for “almost k -wise uniform” because of its convenience for Fourier analysis.

The original paper about almost k -wise uniformity, [NN93], was concerned with derandomization; e.g., they use (ϵ, k) -wise uniformity for derandomizing the “set balancing (discrepancy)” problem. Alon et al. gave a further discussion of the relationship between almost k -wise uniformity and derandomization in [AGM03]. The key idea is the following: In many cases of randomized algorithms, the analysis only relies on the property that the random bits are k -wise uniform, as opposed to fully uniform. Since there exists an efficiently samplable k -wise uniform distribution on a set of size at most $O(n^{\lfloor k/2 \rfloor})$, one can reduce the number of random unbiased bits used in the algorithm down to $O(k \log n)$. To further reduce the number of random bits used, a natural line of thinking is to consider distributions which are “almost k -wise uniformity”. Alon et al. [AGHP92] showed that we can deterministically construct (ϵ, k) -wise uniform sets that are of size $\text{poly}(2^k, \log n, 1/\epsilon)$, much smaller than exact k -wise uniform ones (roughly $\Omega(n^{\lfloor k/2 \rfloor})$ size). Therefore we can use substantially fewer random bits by taking random strings from an almost k -wise uniform distribution.

However we need to ensure that the original analysis of the randomized algorithm still holds under the almost k -wise uniform distribution. This is to say that if the randomized algorithm behaves well on a k -wise uniform distribution, it may or may not also work as well with an (ϵ, k) -wise uniform distribution, when the parameter ϵ is small enough.

1.2 The Closeness Problem

For the analysis of derandomization, it would be very convenient if (ϵ, k) -wise uniformity – which means that “every k -local view looks close to uniform” – implies global δ -closeness to k -wise uniformity. A natural question that arises, posed in [AGM03], is the following:

How small can δ be such that the following is true? For every (ϵ, k) -wise uniform distribution φ on $\{-1, 1\}^n$, φ is δ -close to some k -wise uniform distribution.

In this paper, we will refer to this question as *the Closeness Problem*.

1.2.1 Previous work and applications

On one hand, the main message of [AGM03] is a lower bound: For every even constant $k > 4$, they gave an (ϵ, k) -wise uniform distribution with $\epsilon = O(1/n^{k/4-1})$, yet which is $\frac{1}{2}$ -far from every k -wise uniform distribution in total variation distance.

On the other hand, [AGM03] proved a very simple theorem that $\delta \leq O(n^k \epsilon)$ always holds. Despite its simplicity, this upper bound has been used many times in well known results.

One application is in circuit complexity. [AGM03]’s upper bound is used for fooling disjunctive normal formulas (DNF) [Baz09] and AC^0 [Bra10]. In these works, once the authors showed that k -wise uniformity suffices to fool DNF/ AC^0 , they deduced that $(O(1/n^k), k)$ -uniform distributions suffice, and hence $O(1/n^k)$ -biased sets sufficed trivially. [AGM03]’s upper bound is also used as a tool for the construction of two-source extractors for a similar reason in [CZ16, Li16].

Another application is for hardness of constraint satisfactory problems (CSPs). Austrin and Mossel [AM09] show that one can obtain integrality gaps and UGC-hardness for CSPs based on k -wise uniform distributions of small support size. If a predicate is k -wise uniform, Kothari et al. [KMOW17] showed that one can get SOS-hardness of refuting random instances of it when there are around $n^{(k+1)/2}$ constraints. Indeed, [KMOW17] shows that if we have a predicate that is δ -close to k -wise uniform, then with roughly $n^{(k+1)/2}$ random constraints, SOS cannot refute that a $(1 - O(\delta))$ -fraction of constraints are satisfiable. This also motivates studying δ -closeness to k -wise uniformity and how it relates to Fourier coefficients. δ -closeness to k -wise uniformity is also relevant for hardness of random CSP, as shown in [AOW15].

Alon et al. [AAK⁺07] investigated the Closeness Problem further by improving the upper bound to $\delta = O((n \log n)^{k/2} \epsilon)$. Indeed, they showed a strictly stronger fact that a distribution is $O(\sqrt{\mathbf{W}^{1\dots k}[\varphi]} \log^{k/2} n)$ -close to some k -wise uniform, where $\mathbf{W}^{1\dots k}[\varphi] = \sum_{1 \leq |S| \leq k} \widehat{\varphi}(S)^2$. Rubinfeld and Xie [RX13] generalized some of these results to non-uniform k -wise independent distributions over larger product spaces.

Let us briefly summarize the method [AAK⁺07] used to prove their upper bounds. Given an (ϵ, k) -wise uniform φ , they first try to generate a k -wise uniform “pseudo-distribution” φ' by forcing all Fourier coefficients at degree at most k to be zero. It is a “pseudo-distribution” because some points might have negative density. After this, they use a fully uniform distribution and k -wise uniform distributions with small support size to try to mend all points to be nonnegative. They bound the weight of these mending distributions to upper-bound the distance incurred by the mending process. This mending process uses the fully uniform distribution to mend the small negative weights and uses k -wise uniform distributions with small support size to correct the large negative weights point by point. By optimizing the threshold between small and large weights it introduces a factor of $(\log n)^{k/2}$.

Though they did not mention it explicitly, they also give a lower bound for the Closeness Problem of $\delta \geq \Omega\left(\frac{n^{(k-1)/2}}{\log n} \epsilon\right)$ for $k > 2$ by considering the uniform distribution on a set of $O(n^k)$ random chosen strings. No previous work gave any lower bound for the most natural case of $k = 2$.

1.2.2 Our result

In this paper, we show sharper upper and lower bounds for the Closeness Problem, which are tight for k even and $k = 1$. Comparing to the result in [AAK⁺07], we get rid of the factor of $(\log n)^{k/2}$.

Theorem 1.1. Any density φ over $\{-1, 1\}^n$ is δ -close to some k -wise uniform distribution, where

$$\delta \leq e^k \sqrt{\mathbf{W}^{1\dots k}[\varphi]} = e^k \sqrt{\sum_{1 \leq |S| \leq k} \widehat{\varphi}(S)^2}.$$

Consequently, if φ is (ϵ, k) -wise uniform, i.e., $|\widehat{\varphi}(S)| \leq \epsilon$ for every non-empty set S with size at most k , then

$$\delta \leq e^k n^{k/2} \epsilon.$$

For the special case $k = 1$, the corresponding δ can be further improved to $\delta \leq \epsilon$.

Our new technique is trying to mend the original distribution to be k -wise uniform all at once. We want to show that some mixture distribution $(\varphi + w\psi)$ is k -wise uniform with small mixture weight w . The distance between the final mixture distribution and the original distribution φ is bounded by $O(w)$. Therefore we only need to show that the mending distribution ψ exists for some small weight w . Showing the existence of such a distribution ψ can be written as the feasibility of a linear program (LP). We upper bound w by bounding the dual LP, using the hypercontractivity inequality.

Our result is sharp for all even k , and is also sharp for $k = 1$. We state the matching lower bound for even k :

Theorem 1.2. For any n and even k , and small enough ϵ , there exists some (ϵ, k) -wise uniform distribution φ over $\{-1, 1\}^n$, such that φ is δ -far from every k -wise uniform distribution in total variation distance, where

$$\delta \geq \Omega\left(\frac{1}{k}\right)^k n^{k/2} \epsilon.$$

Our method for proving this lower bound is again LP duality. Our examples in the lower bound are symmetric distributions with Fourier weight only on level k . The density functions then can be written as binary Krawtchouk polynomials which behave similar to Hermite polynomials when n is large. Our dual LP bounds use various properties of Krawtchouk and Hermite polynomials.

Interestingly both our upper and lower bound utilize LP-duality, which we believe is the most natural way of looking at this problem.

We remark that we can derive a lower bound for odd k from Theorem 1.2 trivially by replacing k by $k - 1$. There exists a gap of \sqrt{n} between the resulting upper and lower bounds for odd k . We believe that the lower bound is tight, and the upper bound may be improvable by a factor of \sqrt{n} , as it is in the special case $k = 1$. We leave it as a conjecture for further work:

Conjecture 1.3. Suppose the distribution φ over $\{-1, 1\}^n$ is (ϵ, k) -wise uniform. Then φ is δ -close to some k -wise uniform distribution in total variation distance, where

$$\delta \leq O(n^{\lfloor k/2 \rfloor} \epsilon).$$

1.3 The Testing Problem

Another application of the Closeness Problem is to property testing of k -wise uniformity. Suppose we have sample access from an unknown and arbitrary distribution; we may wonder whether the distribution has a certain property. This question has received tremendous attention in the field of statistics. The main goal in the study of property testing is to design algorithms that use as few samples as possible, and to establish lower bound matching these sample-efficient algorithms. In particular, we consider the property of being k -wise uniform:

Given sample access to an unknown and arbitrary distribution φ on $\{-1, 1\}^n$, how many samples do we need to distinguish between the case that φ is k -wise uniform versus the case that φ is δ -far from every k -wise uniform distribution?

In this paper, we will refer to this question as the *Testing Problem*.

We say a testing algorithm is a δ -tester for k -wise uniformity if the algorithm outputs “Yes” with high probability when the distribution φ is k -wise uniform, and the algorithm outputs “No” with high probability when the distribution φ is δ -far from any k -wise uniform distribution (in total variation distance).

Property testing is well studied for Boolean functions and distributions. Previous work studied the testing of related properties of distribution, including uniformity [GR11, BFR⁺00, RS09] and independence [BFF⁺01, BKR04, ADK15, DK16].

The papers [AGM03, AAK⁺07, Xie12] discussed the problem of testing k -wise uniformity. [AGM03] constructed a δ -tester for k -wise uniformity with sample complexity $O(n^{2k}/\delta^2)$, and [AAK⁺07] improved it to $O(n^k \log^{k+1} n/\delta^2)$. As for lower bounds, [AAK⁺07] showed that $\Omega(n^{(k-1)/2}/\delta)$ samples are necessary, albeit only for $k > 2$. This lower bound is in particular for distinguishing the uniform distribution from δ -far-from- k -wise distributions.

We show a better upper bound for sample complexity:

Theorem 1.4. *There exists a δ -tester for k -wise uniformity of distributions on $\{-1, 1\}^n$ with sample complexity $O\left(\frac{1}{k}\right)^{k/2} \frac{n^k}{\delta^2}$. For the special case of $k = 1$, the sample complexity is $O\left(\frac{\log n}{\delta^2}\right)$.*

A natural δ -tester of k -wise uniformity is mentioned in [AAK⁺07]: Estimate all Fourier coefficients up to level k from the samples; if they are all smaller than ϵ then output “Yes”. In fact this algorithm is exactly attempting to check whether the distribution is (ϵ, k) -wise uniform. Hence the sample complexity depends on the upper bound for the Closeness Problem. Therefore we can reduce the sample complexity of this algorithm down to $O\left(\frac{n^k \log n}{\delta^2}\right)$ via our improved upper bound for the Closeness Problem. One $\log n$ factor remains because we need to union-bound over the $O(n^k)$ Fourier coefficients up to level k . To further get rid of the last $\log n$ factor, we present a new algorithm that estimates the Fourier weight up to level k , $\sum_{1 \leq |S| \leq k} \widehat{\varphi}^2(S)$, rather than estimating these Fourier coefficients one by one.

Unfortunately, a lower bound for the Closeness Problem does not imply a lower bound for the Testing Problem directly. In [AAK⁺07], they showed that a uniform distribution over a random subset of $\{-1, 1\}^n$ of size $O\left(\frac{n^{k-1}}{\delta^2}\right)$, is almost surely δ -far from any k -wise uniform distribution. On the other hand, by the Birthday Paradox, it is hard to distinguish between the fully uniform distribution on all strings of length n and a uniform distribution over a random set of such size. This gives a lower bound for the Testing Problem as $\Omega(n^{(k-1)/2}/\delta)$. Their result only holds for $k > 2$; there was no previous non-trivial lower bound for testing pairwise uniformity. We show a lower bound for the pairwise case.

Theorem 1.5. *Any δ -tester for pairwise uniformity of distributions on $\{-1, 1\}^n$ needs at least $\Omega\left(\frac{n}{\delta^2}\right)$ samples.*

For this lower bound we analyze a symmetric distribution with non-zero Fourier coefficients only on level 2. We prove that it is hard to distinguish a randomly shifted version of this distribution from the fully uniform distribution. This lower bound is also better than [AAK⁺07] in that we have a better dependence on the parameter δ ($\frac{1}{\delta^2}$ rather than $\frac{1}{\delta}$). Unfortunately we are unable to generalize our lower bound for higher k .

Notice that for our new upper and lower bounds for k -wise uniformity testing, there still remains a quadratic gap for $k \geq 2$, indicating that the upper bound might be able to be improved. Both the lower bound in our paper and that in [AAK⁺07] show that it is hard to distinguish between the fully uniform distribution and some specific sets of distributions that are far from k -wise uniform. We show that if one wants to improve the lower bound, one will need to use a distribution in the “Yes” case that is *not* fully uniform, because we give a sample-efficient algorithm for distinguishing between fully uniform and δ -far from k -wise uniform:

Theorem 1.6. *For any constant k , for testing whether a distribution is fully uniform or δ -far from every k -wise uniform distribution, there exists an algorithm with sample complexity $O(k)^k \cdot n^{k/2} \cdot \frac{1}{\delta^2} \cdot \left(\log \frac{n}{\delta}\right)^{k/2}$.*

In fact, for testing whether a distribution is αk -wise uniform or δ -far from k -wise uniform with $\alpha > 4$, there exists an algorithm with sample complexity $O(\alpha)^{k/2} \cdot n^{k/2} \cdot \frac{1}{\delta^2} \cdot \left(\frac{n^k}{\delta^4}\right)^{1/(\alpha-2)}$.

We remark that testing full uniformity can be treated as a special case of testing αk -wise uniformity approximately, by setting $\alpha = \log \frac{n}{\delta}$.

Testing full uniformity has been studied in [GR11, BFR⁺00]. Paninski [Pan08] showed that testing whether an unknown distribution on $\{-1, 1\}^n$ is $\Theta(1)$ -close to fully uniform requires $2^{n/2}$ samples. Rubinfeld and Servedio [RS09] studied testing whether an unknown monotone distribution is fully uniform or not.

The fully uniform distribution has the nice property that every pair of samples is different in $\frac{n}{2} \pm O(\sqrt{n})$ bits with high probability when the sample size is small. Our algorithm first rejects those distributions that disobey this property. We show that the remaining distributions have small Fourier weight up to level $2k$. Hence by following a similar analysis as the tester in Theorem 1.4, we can get an improved upper bound when these lower Fourier weights are small.

The lower bound remains the same as testing k -wise vs. far from k -wise. Our tester is tight up to a logarithmic factor for the pairwise case, and is tight up to a factor of $\tilde{O}(\sqrt{n})$ when $k > 2$.

We compare our results and previous best known bounds from [AAK⁺07] in Table 1. (We omit constant factors depending on k .)

	Upper bound		Lower bound	
	[AAK ⁺ 07]	Our paper	[AAK ⁺ 07]	Our paper
Closeness Problem	$O(n^{k/2}(\log n)^{k/2}\epsilon)$	$O(n^{k/2}\epsilon)$ $O(\epsilon)$ for $k = 1$	$\Omega\left(\frac{n^{(k-1)/2}}{\log n} \epsilon\right)$	$\Omega(n^{\lfloor k/2 \rfloor} \epsilon)$
Testing k -wise vs. far from k -wise	$O\left(\frac{n^k(\log n)^{k+1}}{\delta^2}\right)$	$O\left(\frac{n^k}{\delta^2}\right)$ $O\left(\frac{\log n}{\delta^2}\right)$ for $k = 1$	$\Omega\left(\frac{n^{(k-1)/2}}{\delta}\right)$ for $k > 2$	$\Omega\left(\frac{n}{\delta^2}\right)$ for $k = 2$
Testing n -wise vs. far from k -wise	$O\left(\frac{n^k(\log n)^{k+1}}{\delta^2}\right)$	$O\left(\frac{n^{k/2}}{\delta^2}(\log \frac{n}{\delta})^{k/2}\right)$ $O\left(\frac{\log n}{\delta^2}\right)$ for $k = 1$	$\Omega\left(\frac{n^{(k-1)/2}}{\delta}\right)$ for $k > 2$	$\Omega\left(\frac{n}{\delta^2}\right)$ for $k = 2$

Table 1: Comparison of our results to [AAK⁺07]

1.4 Organization

Section 2 contains definitions and notations. We will discuss upper and lower bounds for the Closeness Problem in Section 3. We will discuss the sample complexity of testing k -wise uniformity in Section 4. We present a tester for distinguishing between αk -wise uniformity (or fully uniformity) and far-from k -wise uniformity in Section 5.

2 Preliminaries

2.1 Fourier analysis of Boolean functions

We use $[n]$ to denote the set $\{1, \dots, n\}$. We denote the symmetric difference of two sets S and T by $S \oplus T$. For Fourier analysis we use notations consistent with [O'D14]. Every function $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ has a unique representation as a multilinear polynomial

$$f(\mathbf{x}) = \sum_{S \subseteq [n]} \widehat{f}(S) x^S \quad \text{where} \quad x^S = \prod_{i \in S} x_i.$$

We call $\widehat{f}(S)$ the Fourier coefficient of f on S . We use $\mathbf{x} \sim \{-1, 1\}^n$ to denote that \mathbf{x} is uniformly distributed on $\{-1, 1\}^n$. We can represent Fourier coefficients as

$$\widehat{f}(S) = \mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [f(\mathbf{x}) \mathbf{x}^S].$$

We define an inner product $\langle \cdot, \cdot \rangle$ on pairs of functions $f, g : \{-1, 1\}^n \rightarrow \mathbb{R}$ by

$$\langle f, g \rangle = \mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [f(\mathbf{x})g(\mathbf{x})] = \sum_{S \subseteq [n]} \widehat{f}(S)\widehat{g}(S).$$

We introduce the following p -norm notation: $\|f\|_p = (\mathbf{E}[|f(\mathbf{x})|^p])^{1/p}$, and the Fourier ℓ_p -norm is $\|\widehat{f}\|_p = \left(\sum_{S \subseteq [n]} |\widehat{f}(S)|^p\right)^{1/p}$.

We say the *degree* of a Boolean function, $\deg(f)$ is k if its Fourier polynomial is degree k . We denote $f^{=k}(x) = \sum_{|S|=k} \widehat{f}(S)x^S$, and $f^{\leq k}(x) = \sum_{|S|\leq k} \widehat{f}(S)x^S$. We denote the *Fourier weight* on level k by $\mathbf{W}^k[f] = \sum_{|S|=k} \widehat{f}(S)^2$. We denote $\mathbf{W}^{1\dots k}[\varphi] = \sum_{1 \leq |S| \leq k} \widehat{\varphi}(S)^2$.

We define the convolution $f * g$ of a pair of functions $f, g : \{-1, 1\}^n \rightarrow \mathbb{R}$ to be

$$(f * g)(x) = \mathbf{E}_{\mathbf{y} \sim \{-1, 1\}^n} [f(x)g(x \circ \mathbf{y})],$$

where \circ denotes entry-wise multiplication. The effect of convolution on Fourier coefficients is that $\widehat{f * g}(S) = \widehat{f}(S)\widehat{g}(S)$.

2.2 Densities and distances

When working with probability distribution on $\{-1, 1\}^n$, we prefer to define them via *density* function. A density function $\varphi : \{-1, 1\}^n \rightarrow \mathbb{R}^{\geq 0}$ is a nonnegative function satisfying $\widehat{\varphi}(\emptyset) = \mathbf{E}_{\mathbf{x} \sim \{-1, 1\}^n} [\varphi(\mathbf{x})] = 1$. We write $\mathbf{y} \sim \varphi$ to denote that \mathbf{y} is a random variable drawn from the distribution φ , defined by

$$\Pr_{\mathbf{y} \sim \varphi} [\mathbf{y} = y] = \frac{\varphi(y)}{2^n},$$

for all $y \in \{-1, 1\}^n$. We identify distributions with their density functions when there is no risk of confusion.

We denote $\varphi^{+t}(x) = \varphi(x \circ t)$. We denote by $\mathbf{1}_A$ the density function for the uniform distribution on support set A . The density function associated to the fully uniform distribution is the constant function $\mathbf{1}$.

The following lemma about density functions of degree at most k derives from Fourier analysis and hypercontractivity.

Lemma 2.1. *Let $\varphi : \{-1, 1\}^n \rightarrow \mathbb{R}^{\geq 0}$ be a density function of degree at most k . Then*

$$\|\widehat{\varphi}\|_2 = \sqrt{\sum_S \widehat{\varphi}(S)^2} \leq e^k.$$

Proof.

$$\|\widehat{\varphi}\|_2 = \|\varphi\|_2 \leq e^k \|\varphi\|_1 = e^k.$$

The first equality holds by Parseval's Theorem (see Section 1.4 in [O'D14]). The inequality holds by hypercontractivity (see Theorem 9.22 in [O'D14]). The last equality holds since φ is a density function. \square

A distribution φ over $\{-1, 1\}^n$ is *k-wise uniform* if and only if $\widehat{\varphi}(S) = 0$ for all $1 \leq |S| \leq k$ (see Chapter 6.1 in [O'D14]). We say that distribution φ over $\{-1, 1\}^n$ is (ϵ, k) -*wise uniform* if $|\widehat{\varphi}(S)| \leq \epsilon$ for all $1 \leq |S| \leq k$.

The most common way to measure the distance between two probability distributions is via their *total variation distance*. If the distributions have densities φ and ψ , then the total variation distance is defined to be

$$d_{\text{TV}}(\varphi, \psi) = \sup_{A \subseteq \{-1, 1\}^n} \left| \Pr_{\mathbf{x} \sim \varphi} [\mathbf{x} \in A] - \Pr_{\mathbf{x} \sim \psi} [\mathbf{x} \in A] \right| = \frac{1}{2} \mathbf{E}_{\mathbf{x}} [|\varphi(\mathbf{x}) - \psi(\mathbf{x})|] = \frac{1}{2} \|\varphi - \psi\|_1.$$

We say that φ and ψ are δ -*close* if $d_{\text{TV}}(\varphi, \psi) \leq \delta$.

Supposing H is a set of distributions, we denote

$$d_{\text{TV}}(\varphi, H) = \min_{\psi \in H} d_{\text{TV}}(\varphi, \psi).$$

In particular, we denote the set of k -wise uniform densities by kWISE . We say that density φ is δ -*close to k-wise uniform* if $d_{\text{TV}}(\varphi, \text{kWISE}) \leq \delta$, and is δ -*far* otherwise.

2.3 Krawtchouk and Hermite polynomials

Krawtchouk polynomials were introduced in [Kra29], and arise in the analysis of Boolean functions as shown in [Lev95, Kal02]. Consider the following Boolean function of degree k and input length n : $f(x) = \sum_{|S|=k} x^S$. It is symmetric and therefore only depends on the Hamming weight of x . Let t be the number of -1 's in x . Then the output of f is exactly the same as the Krawtchouk polynomial $K_k^{(n)}(t)$.

Definition 2.2. We denote by $K_k^{(n)}(t)$ the Krawtchouk polynomial:

$$K_k^{(n)}(t) = \sum_{j=0}^k (-1)^j \binom{t}{j} \binom{n-t}{k-j},$$

for $k = 0, 1, \dots, n$.

We will also use Hermite polynomials in our analysis.

Definition 2.3. We denote by $h_k(z)$ the normalized Hermite polynomial:

$$h_k(z) = \frac{1}{\sqrt{k!}} (-1)^k e^{\frac{1}{2}z^2} \frac{d^k}{dz^k} e^{-\frac{1}{2}z^2}.$$

Its explicit formula is

$$h_k(z) = \sqrt{k!} \cdot \left(\frac{z^k}{0!! \cdot k!} - \frac{z^{k-2}}{2!! \cdot (k-2)!} + \frac{z^{k-4}}{4!! \cdot (k-4)!} - \frac{z^{k-6}}{6!! \cdot (k-6)!} + \dots \right).$$

One useful fact is that the derivative of a Hermite polynomial is a scalar multiple of a Hermite polynomial (see Exercise 11.10 in [O'D14]):

Fact 2.4. For any integer $k \geq 1$, we have

$$\frac{d}{dz} h_k(z) = \sqrt{k} h_{k-1}(z).$$

The relationship between Krawtchouk and Hermite polynomials is that we can treat Hermite polynomials as a limit version of Krawtchouk polynomials when n goes to infinity (see Exercise 11.14 in [O'D14]).

Fact 2.5. For all $k \in \mathbb{N}$ and $z \in \mathbb{R}$ we have

$$\binom{n}{k}^{-1/2} \cdot K_k^{(n)} \left(\frac{n - z\sqrt{n}}{2} \right) \xrightarrow{n \rightarrow \infty} h_k(z).$$

Instead of analyzing Krawtchouk polynomials, it is easier to study Hermite polynomials when n is large because Hermite polynomials have a more explicit form. We present some basic properties of Hermite polynomials with brief proofs.

Lemma 2.6. The following are properties of $h_k(z)$:

1. $|h_k(z)| \leq h_k(k)$ for any $|z| \leq k$;
2. $h_k(z)$ is positive and increasing when $z \geq k$;
3. $h_k(Ck) \leq (Ck)^k / \sqrt{k!}$ for any constant $C \geq 1$.

Proof. We will treat the case of $k = 4i + 2$ for some integer i . The proof for the general case is similar. When $k = 4i + 2$, we can group adjacent terms into pairs:

$$h_k(z) = \sqrt{k!} \cdot \sum_{i=0}^{(k-2)/4} \frac{z^{k-4i-2}}{(4i+2)!! \cdot (k-4i)!} ((4i+2)z^2 - (k-4i)(k-4i-1)).$$

1. Notice that $|(4i+2)z^2 - (k-4i)(k-4i-1)|$ is always between $-(k-4i)(k-4i-1)$ and $(4i+2)k^2 - (k-4i)(k-4i-1)$ when $|z| \leq k$. Both the upper and lower bound have absolute value at most $(4i+2)k^2 - (k-4i)(k-4i-1)$. Therefore by the triangle inequality we have $|h_k(z)| \leq h_k(k)$.
2. It is easy to check that $((4i+2)z^2 - (k-4i)(k-4i-1))$ is positive when $z \geq k$. Then by Fact 2.4, $\frac{d}{dz}h_k(z) = \sqrt{k}h_{k-1}(z) > 0$ when $z \geq k$.
3. This is trivial from the explicit formula since each term is exactly smaller than the previous term when $z \geq k$. \square

3 The Closeness Problem

In this section, we prove the upper bound in Theorem 1.1 and the lower bound in Theorem 1.2. One interesting fact is that we use duality of linear programming (LP) in both the upper and lower bound. We think this is the proper perspective for analyzing these questions.

3.1 Upper bound

The key idea for proving the upper bound is mixture distributions. Given an (ϵ, k) -wise uniform density φ , we try to mix it with some other distribution ψ using mixture weight w , such that the mixture distribution $\frac{1}{1+w}(\varphi + w\psi)$ is k -wise uniform and is close to the original distribution. The following lemma shows that the distance between the original distribution and the mixture distribution is bounded by the weight w .

Lemma 3.1. *If $\varphi' = \frac{1}{1+w}(\varphi + w\psi)$ for some $0 \leq w \leq 1$ and density functions φ, ψ , then $d_{\text{TV}}(\varphi, \varphi') \leq w$.*

Proof. $d_{\text{TV}}(\varphi, \varphi') = \frac{1}{2}\|\varphi' - \varphi\|_1 = \frac{1}{2}\|\varphi' - ((1+w)\varphi' - w\psi)\|_1 = \frac{1}{2}w\|\varphi' - \psi\|_1 \leq w$. \square

Therefore we only need to show the existence of an appropriate ψ for some small w . The constraints on ψ can be written as an LP feasibility problem. Therefore by Farkas' Lemma we only need to show that its dual is not feasible. The variables in the dual LP can be seen as a density function of degree at most k .

Proof of Theorem 1.1 (general k case). Given density function φ , we try to find another density function ψ with constraints

$$\widehat{\psi}(S) = -\frac{1}{w}\widehat{\varphi}(S)$$

for all $1 \leq |S| \leq k$. Suppose such a density function ψ exists. Then it is trivial that $\frac{\varphi+w\psi}{1+w}$ is also a density function and is k -wise uniform. By Lemma 3.1, we conclude that $d_{\text{TV}}(\varphi, \text{kWISE}) \leq w$.

The rest of proof is to show that such a ψ exists when $w = e^k \sqrt{\mathbf{W}^{1\dots k}[\varphi]}$. We can write the existence as an LP feasibility problem with variables $\psi(x)$ for $x \in \{-1, 1\}^n$ and constraints:

$$\begin{aligned} \widehat{\psi}(\emptyset) &= 1, \\ \widehat{\psi}(S) &= -\frac{1}{w}\widehat{\varphi}(S), & \forall 1 \leq |S| \leq k, \\ \psi(x) &\geq 0, & \forall x \in \{-1, 1\}^n, \end{aligned}$$

where $\widehat{\psi}(S) = \mathbf{E}[\psi(\mathbf{x})\mathbf{x}^S]$ is a linear combination of variables $\psi(x)$.

The dual LP has variables $\psi'(x)$ for $x \in \{-1, 1\}^n$ with constraints:

$$\begin{aligned} \widehat{\psi}'(\emptyset) &= 1, \\ \widehat{\psi}'(S) &= 0, & \forall |S| > k, \\ \psi'(x) &\geq 0, & \forall x \in \{-1, 1\}^n, \\ \frac{1}{w} \sum_{1 \leq |S| \leq k} \widehat{\varphi}(S)\widehat{\psi}'(S) &> 1. \end{aligned}$$

The original LP is feasible if and only if its dual LP is infeasible, by Farkas' Lemma. This completes the proof, since when $w = e^k \sqrt{\mathbf{W}^{1\dots k}[\varphi]}$, for any density function ψ' with degree k we have

$$\frac{1}{w} \sum_{1 \leq |S| \leq k} \widehat{\varphi}(S) \widehat{\psi}'(S) \leq \frac{1}{e^k \sqrt{\mathbf{W}^{1\dots k}[\varphi]}} \sum_{1 \leq |S| \leq k} |\widehat{\varphi}(S)| |\widehat{\psi}'(S)| \leq \frac{1}{e^k} \|\widehat{\psi}'\|_2 \leq 1,$$

where the second inequality holds by Cauchy–Schwarz, and the last inequality holds by Lemma 2.1 since ψ' has degree at most k . \square

For $k = 1$, further improvement can be achieved. We still try to use mixture distributions. Here we want to mix the distribution φ with indicator distributions on subsets of coordinates that have opposite biases to those of the original distribution.

Proof of Theorem 1.1 (case $k = 1$). By identifying each x_i with $-x_i$ if necessary, we may assume without loss of generality that $\widehat{\varphi}(\{i\}) \geq 0$ for all i . In addition, by reordering the coordinates, we may assume without loss of generality that $0 \leq \widehat{\varphi}(\{1\}) \leq \dots \leq \widehat{\varphi}(\{n\}) = \epsilon$. Define ψ_j to be the density of the distribution over $\{-1, 1\}^n$ which is uniform on coordinates x_1, \dots, x_{j-1} , and has x_i constantly fixed to be -1 for $j \leq i \leq n$. It is easy to check $\widehat{\psi}_j(\{i\}) = 0$ for $i < j$ and $\widehat{\psi}_j(\{i\}) = -1$ for $i \geq j$.

We define φ' as

$$\varphi' = \frac{1}{1 + \epsilon} \left(\varphi + \sum_{j=1}^n w_j \psi_j \right),$$

where

$$w_1 = \widehat{\varphi}(\{1\}), \quad w_j = \widehat{\varphi}(\{j\}) - \widehat{\varphi}(\{j-1\}) \quad \forall 1 < j \leq n.$$

It is easy to check that φ' is a density function and

$$\widehat{\varphi}'(\{i\}) = \frac{1}{1 + \epsilon} \left(\widehat{\varphi}(\{i\}) + \left(\sum_{j=1}^i w_j \right) (-1) \right) = 0.$$

Therefore φ' is 1-wise uniform. Then by Lemma 3.1,

$$d_{TV}(\varphi, 1\text{WISE}) \leq \frac{1}{2} \|\varphi - \varphi'\|_1 \leq \sum_{j=1}^n w_j = \epsilon. \quad \square$$

3.2 Lower bound

Interestingly, our proof of the lower bound also utilizes LP duality. We can write the Closeness Problem in the form of linear programming with variables $\varphi'(x)$ for $x \in \{-1, 1\}^n$, as follows:

$$\begin{aligned} &\text{minimize} && d_{TV}(\varphi, \varphi') = \frac{1}{2} \|\varphi - \varphi'\|_1 \\ &\text{subject to:} && \widehat{\varphi}'(\emptyset) = 1, \\ & && \widehat{\varphi}'(S) = 0, && \forall 1 \leq |S| \leq k, \\ & && \varphi'(x) \geq 0, && \forall x \in \{-1, 1\}^n. \end{aligned}$$

We ignore the factor of 1/2 in the minimization for convenience in the following analysis. The dual LP, which has variables $p(x), q(x)$ for $x \in \{-1, 1\}^n$, is the following:

$$\begin{aligned} &\text{maximize} && \langle \varphi, q \rangle - \widehat{p}(\emptyset) \\ &\text{subject to:} && p(x) - q(x) \geq 0, && \forall x \in \{-1, 1\}^n, \\ & && q(x) \leq 1, && \forall x \in \{-1, 1\}^n, \\ & && p(x) \geq -1, && \forall x \in \{-1, 1\}^n, \\ & && \deg(p) \leq k. \end{aligned}$$

Thus given a pair of Boolean functions p, q satisfying the constraints, the quantity $\langle \varphi, q \rangle - \widehat{p}(\emptyset)$ is a lower bound for our Closeness Problem. Our distribution φ achieving the lower bound is a symmetric polynomial, homogeneous of degree k (except that it has a constant term of 1, as is necessary for every density function). We can use Krawtchouk and Hermite polynomials to simplify the analysis.

Proof of Theorem 1.2. We define

$$\varphi(x) = 1 + \mu \binom{n}{k}^{-1/2} \sum_{|S|=k} x^S, \quad p(x) = \mu \binom{n}{k}^{-1/2} \sum_{|S|=k} x^S, \quad q(x) = \min(p(x), 1),$$

where μ is a small parameter to be chosen later that will ensure $\varphi(x) \geq 0$ and $p(x) \geq -1$ for all $x \in \{-1, 1\}^n$.

We have $\epsilon = \max_{1 \leq |S| \leq k} |\widehat{\varphi}(S)| = \mu \binom{n}{k}^{-1/2}$.

Since $\widehat{p}(\emptyset) = 0$, the objective function of the dual LP is

$$\begin{aligned} \langle \varphi, q \rangle &= \langle \varphi, \min(p, 1) \rangle = \langle \varphi, \mathbf{1}_{p>1} \rangle + \langle \varphi, p \mathbf{1}_{p \leq 1} \rangle = \langle \varphi, p \rangle - \langle \varphi, (p-1) \mathbf{1}_{p>1} \rangle \\ &\geq \langle \varphi, p \rangle - \sqrt{\Pr_{\mathbf{x} \sim \varphi} [p(\mathbf{x}) > 1]} \cdot \langle \varphi, (p-1)^2 \rangle, \end{aligned}$$

where the last inequality holds by Cauchy–Schwarz. It is easy to calculate the inner products $\langle \varphi, p \rangle = \mu^2$, and

$$\begin{aligned} \langle \varphi, (p-1)^2 \rangle &= \langle \varphi, p^2 \rangle - 2\langle \varphi, p \rangle + 1 \\ &= \mu^2 + \mu^3 \binom{n}{k}^{-1/2} \binom{k}{k/2} \binom{n-k}{k/2} - 2\mu^2 + 1 \\ &\leq 1 + \mu^3 \binom{k}{k/2}^{3/2} - \mu^2. \end{aligned}$$

Assuming $\mu < 2^{-\frac{3}{2}k}$, we have $\langle \varphi, (p-1)^2 \rangle < 1$.

Now we need to upper bound $\Pr_{\mathbf{x} \sim \varphi} [p(\mathbf{x}) > 1]$. Define \mathbf{z} satisfying $(n - \mathbf{z}\sqrt{n})/2 = \sum_i \mathbf{x}_i$. Then

$$\Pr_{\mathbf{x} \sim \varphi} [p(\mathbf{x}) > 1] = \Pr_{\mathbf{x} \sim \varphi} \left[\mu \binom{n}{k}^{-1/2} \cdot K_k \left(\frac{n - \mathbf{z}\sqrt{n}}{2}, n \right) > 1 \right].$$

By Fact 2.5, we know that when $z \leq k$, for sufficient large n ,

$$\binom{n}{k}^{-1/2} \cdot K_k \left(\frac{n - z\sqrt{n}}{2}, n \right) < 2h_k(z).$$

Now we set $\mu = \frac{\sqrt{k!}}{2(Ck)^k}$ with some constant $C \geq 1$. It is easy to check that $\mu < 2^{-\frac{3}{2}k}$. Using the properties in Lemma 2.6, we get

$$\begin{aligned} \Pr_{\mathbf{x} \sim \varphi} \left[\mu \binom{n}{k}^{-1/2} \cdot K_k \left(\frac{n - \mathbf{z}\sqrt{n}}{2}, n \right) > 1 \right] &\leq \Pr_{\mathbf{x} \sim \varphi} [2\mu h_k(\mathbf{z}) > 1] \\ &\leq \Pr_{\mathbf{x} \sim \varphi} [h_k(\mathbf{z}) > h_k(Ck)] \\ &= \Pr_{\mathbf{x} \sim \varphi} [|\mathbf{z}| > Ck]. \end{aligned}$$

Then using Cauchy–Schwarz again, we get

$$\begin{aligned} \Pr_{\mathbf{x} \sim \varphi} [|\mathbf{z}| > Ck] &\leq \sqrt{\mathbf{E}_{\mathbf{x} \sim \{-1,1\}^n} [\varphi(\mathbf{x})^2]} \sqrt{\Pr_{\mathbf{x} \sim \{-1,1\}^n} [|\mathbf{z}| > Ck]} \\ &\leq \sqrt{1 + \mu^2} \sqrt{2 \exp(-C^2 k^2 / 2)} \\ &\leq 2 \exp(-(Ck)^2 / 4). \end{aligned}$$

Therefore we get that the objective function is at least

$$\langle \varphi, p \rangle - \sqrt{\Pr_{\mathbf{x} \sim \varphi}[p(\mathbf{x}) > 1] \cdot \langle \varphi, (p-1)^2 \rangle} \geq \mu^2 - \sqrt{2 \exp(-(Ck)^2/4)} \geq \Omega\left(\frac{1}{k}\right)^k.$$

The last inequality holds when we choose a sufficiently large constant C .

This completes the proof, because φ is at least δ -far from k -wise uniform with $\delta = \Omega\left(\frac{1}{k}\right)^k$, and we have $\epsilon = \mu \binom{n}{k}^{-1/2} \leq \frac{n^{-k/2}}{2^{\Omega(k)}}$. Therefore we have $\delta \geq \Omega\left(\frac{1}{k}\right)^k n^{k/2} \epsilon$. \square

4 The Testing Problem

In this section, we study the problem of testing whether a distribution is k -wise uniform or δ -far from k -wise uniform. These bounds are based on new bounds for the Closeness Problem. We present a new testing algorithm for general k in Section 4.1. We give a lower bound for the pairwise case in Section 4.2.

4.1 Upper bound

Given m samples from φ , call them $\mathbf{x}_1, \dots, \mathbf{x}_m$, we will first show that

$$\Delta(\mathbf{X}) = \text{avg}_{1 \leq s < t \leq m} \left(\sum_{1 \leq |S| \leq k} \mathbf{x}_s^S \mathbf{x}_t^S \right)$$

is a natural estimator of $\mathbf{W}^{1\dots k}[\varphi]$.

Lemma 4.1. *It holds that*

$$\begin{aligned} \mu &= \mathbf{E}[\Delta(\mathbf{X})] = \mathbf{W}^{1\dots k}[\varphi]; \\ \mathbf{Var}[\Delta(\mathbf{X})] &\leq \frac{4}{m^2} L_k(\varphi) + \frac{4}{m} \sqrt{L_k(\varphi)} \mu, \end{aligned} \quad (1)$$

where $L_k(\varphi) = \sum_{1 \leq |S_1|, |S_2| \leq k} \widehat{\varphi}(S_1 \oplus S_2)^2$.

Proof. We denote $F(x, y) = \sum_{1 \leq |S| \leq k} x^S y^S$. We know that

$$\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi} [\mathbf{x}^S \mathbf{y}^S] = \mathbf{E}_{\mathbf{x} \sim \varphi} [\mathbf{x}^S] \mathbf{E}_{\mathbf{y} \sim \varphi} [\mathbf{y}^S] = \widehat{\varphi}(S)^2,$$

when \mathbf{x} and \mathbf{y} are independent samples drawn from φ . Therefore by linearity of expectation, $\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi} [F(\mathbf{x}, \mathbf{y})] = \mathbf{W}^{1\dots k}[\varphi]$, and clearly by taking the average,

$$\mu = \mathbf{E}[\Delta(\mathbf{X})] = \mathbf{E}[\text{avg}_{s < t} F(\mathbf{x}_s, \mathbf{x}_t)] = \text{avg}_{s < t} \mathbf{E}[F(\mathbf{x}_s, \mathbf{x}_t)] = \mathbf{W}^{1\dots k}[\varphi].$$

We need to expand the variance:

$$\mathbf{Var} \left[\text{avg}_{s < t} (F(\mathbf{x}_s, \mathbf{x}_t)) \right] = \frac{1}{\binom{m}{2}} \sum_{\substack{s < t \\ s' < t'}} \mathbf{Cov}[F(\mathbf{x}_s, \mathbf{x}_t), F(\mathbf{x}_{s'}, \mathbf{x}_{t'})]. \quad (2)$$

We will discuss these covariances in three cases.

Case 1: $|\{s, t\} \cap \{s', t'\}| = 2$. Let $\mathbf{x}, \mathbf{y} \sim \varphi$ be independent random variables.

$$\mathbf{Cov}[F(\mathbf{x}, \mathbf{y}), F(\mathbf{x}, \mathbf{y})] = \mathbf{Var}_{\mathbf{x}, \mathbf{y} \sim \varphi} [F(\mathbf{x}, \mathbf{y})] \leq \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi} [F(\mathbf{x}, \mathbf{y})^2] = \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi} \left[\left(\sum_{1 \leq |S| \leq k} \mathbf{x}^S \mathbf{y}^S \right)^2 \right].$$

Notice here all \mathbf{x}_i 's are Rademacher variables with $\mathbf{x}_i^2 = 1$, and similarly for the \mathbf{y}_i 's. Therefore

$$\begin{aligned} \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi} \left[\left(\sum_{1 \leq |S| \leq k} \mathbf{x}^S \mathbf{y}^S \right)^2 \right] &= \sum_{1 \leq |S_1|, |S_2| \leq k} \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi} [\mathbf{x}^{S_1 \oplus S_2} \mathbf{y}^{S_1 \oplus S_2}] \\ &= \sum_{1 \leq |S_1|, |S_2| \leq k} \widehat{\varphi}(S_1 \oplus S_2)^2 = L_k(\varphi). \end{aligned}$$

Case 2: $|\{s, t\} \cap \{s', t'\}| = 1$. Let $\mathbf{x}, \mathbf{y}, \mathbf{z} \sim \varphi$ be independent random variables. Similar to Case 1, we have:

$$\begin{aligned} \mathbf{Cov}[F(\mathbf{x}, \mathbf{y}), F(\mathbf{x}, \mathbf{z})] &\leq \mathbf{E}[F(\mathbf{x}, \mathbf{y})F(\mathbf{x}, \mathbf{z})] \\ &= \mathbf{E} \left[\left(\sum_{1 \leq |S_1| \leq k} \mathbf{x}^{S_1} \mathbf{y}^{S_1} \right) \left(\sum_{1 \leq |S_2| \leq k} \mathbf{x}^{S_2} \mathbf{z}^{S_2} \right) \right] \\ &= \mathbf{E} \left[\sum_{1 \leq |S_1|, |S_2| \leq k} \mathbf{x}^{S_1 \oplus S_2} \mathbf{y}^{S_1} \mathbf{z}^{S_2} \right] \\ &= \sum_{1 \leq |S_1|, |S_2| \leq k} \widehat{\varphi}(S_1 \oplus S_2) \widehat{\varphi}(S_1) \widehat{\varphi}(S_2) \\ &\leq \sqrt{\sum_{1 \leq |S_1|, |S_2| \leq k} \widehat{\varphi}(S_1 \oplus S_2)^2} \sqrt{\sum_{1 \leq |S_1|, |S_2| \leq k} \widehat{\varphi}(S_1)^2 \widehat{\varphi}(S_2)^2} \\ &= \sqrt{L_k(\varphi)} \mu, \end{aligned}$$

where the inequality comes from Cauchy–Schwarz.

Case 3: $|\{s, t\} \cap \{s', t'\}| = 0$. Let $\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{w} \sim \varphi$ be independent random variables. Clearly $F(\mathbf{x}, \mathbf{y})$ and $F(\mathbf{z}, \mathbf{w})$ are independent and therefore $\mathbf{Cov}[F(\mathbf{x}, \mathbf{y}), F(\mathbf{z}, \mathbf{w})] = 0$.

Plugging all these cases into eq. (2), we get

$$\begin{aligned} \mathbf{Var}[\Delta(\mathbf{X})] &= \mathbf{Var} \left[\text{avg}_{s < t} (F(\mathbf{x}_s, \mathbf{x}_t)) \right] \\ &= \frac{1}{\binom{m}{2}^2} \left(\binom{m}{2} L_k(\varphi) + m(m-1)(m-2) \sqrt{L_k(\varphi)} \mu \right) \\ &\leq \frac{4}{m^2} L_k(\varphi) + \frac{4}{m} \sqrt{L_k(\varphi)} \mu. \quad \square \end{aligned}$$

Given Lemma 4.1 we can bound the samples we need for estimating $\mathbf{W}^{1\dots k}[\varphi]$.

Theorem 4.2 ($\mathbf{W}^{1\dots k}$ Estimation Test). *Let $\varphi : \{-1, 1\}^n \rightarrow \mathbb{R}^{\geq 0}$ be a density function, promised to satisfy $\mathbf{W}^i[\varphi] \leq An^{i/2}$ for all $i = 0, 1, \dots, 2k$. There is an algorithm that, given*

$$m \geq 1000 \frac{2^k \sqrt{An^{k/2}}}{\theta} \quad (3)$$

samples, distinguishes with probability at least 3/4 whether $\mathbf{W}^{1\dots k}[\varphi] \leq \frac{1}{2}\theta$ or $\mathbf{W}^{1\dots k}[\varphi] > \theta$.

Proof. The algorithm is simple: we report “ $\mu \leq \frac{1}{2}\theta$ ” if $\Delta(\mathbf{X}) \leq \frac{3}{4}\theta$ and report “ $\mu > \theta$ ” if $\Delta(\mathbf{X}) > \frac{3}{4}\theta$.

Now we need to bound $L_k(\varphi)$ to bound the variance of $\Delta(\mathbf{X})$. For a fixed subset $|S| \leq 2k$, how many pairs of $1 \leq |S_1|, |S_2| \leq k$ are there satisfying $S = S_1 \oplus S_2$? We denote $S_1 = S'_1 \cup T$, $S_2 = S'_2 \cup T$, where S'_1, S'_2, T are disjoint. Then $S = S'_1 \cup S'_2$. For a fixed set S , there are at most $2^{|S|}$ different ways to split it into two sets S'_1, S'_2 . Because $\max\{|S'_1|, |S'_2|\} \geq \lceil |S|/2 \rceil$ and $|S_1|, |S_2| \leq k$, we have $|T| \leq k - \lceil |S|/2 \rceil$. Therefore there are at most

$$\sum_{j=0}^{k - \lceil |S|/2 \rceil} \binom{n - |S|}{j} \leq \frac{2n^{k - \lceil |S|/2 \rceil}}{(k - \lceil |S|/2 \rceil)!}$$

ways to choose the set T for any fixed S'_1, S'_2 . Hence,

$$\begin{aligned}
L_k(\varphi) &= \sum_{1 \leq |S_1|, |S_2| \leq k} \widehat{\varphi}(S_1 \oplus S_2)^2 \\
&= \sum_{|S| \leq 2k} \sum_{\substack{S'_1 \cap S'_2 = \emptyset \\ S'_1 \cup S'_2 = S}} \sum_{\substack{T \cap S'_1 = \emptyset, T \cap S'_2 = \emptyset \\ |T| + \max\{|S'_1|, |S'_2|\} \leq k}} \widehat{\varphi}(S)^2 \\
&\leq \sum_{|S| \leq 2k} 2^{|S|} \frac{2n^{k - \lceil |S|/2 \rceil}}{(k - \lceil |S|/2 \rceil)!} \widehat{\varphi}(S)^2 \\
&= \sum_{i=0}^{2k} 2^i \frac{2n^{k - \lceil i/2 \rceil}}{(k - \lceil i/2 \rceil)!} \mathbf{W}^i[\varphi].
\end{aligned}$$

Plugging in $\mathbf{W}^i[\varphi] \leq An^{i/2}$, we get

$$L_k(\varphi) \leq \sum_{i=0}^{2k} 2^i \frac{2n^{k - \lceil i/2 \rceil}}{(k - \lceil i/2 \rceil)!} \mathbf{W}^i[\varphi] \leq 2^{2k+2} An^k. \quad (4)$$

By substituting eq. (4) and eq. (3) into eq. (1), we have

$$\mathbf{Var}[\Delta(\mathbf{X})] \leq \frac{4}{500^2} \theta^2 + \frac{4}{500} \theta \mu \leq \frac{1}{64} \max\{\theta^2, \mu^2\}.$$

Then we conclude our proof by Chebyshev's inequality:

$$\begin{aligned}
\Pr \left[|\Delta(\mathbf{X}) - \mu| \leq \frac{1}{4} \max\{\theta, \mu\} \right] &\geq \Pr \left[|\Delta(\mathbf{X}) - \mu| \leq 2\sqrt{\mathbf{Var}[\Delta(\mathbf{X})]} \right] \\
&\geq 1 - \left(\frac{1}{2} \right)^2 = \frac{3}{4}. \quad \square
\end{aligned}$$

This $\mathbf{W}^{1 \dots k}$ Estimation Test is just what we need for testing k -wise uniformity with the upper bound of the Closeness Problem.

Proof of Theorem 1.4. From Theorem 1.1 we know that if density φ is δ -far from k -wise uniform, then $\mathbf{W}^{1 \dots k}[\varphi] > \left(\frac{\delta}{e^k}\right)^2$. On the other hand if φ is k -wise uniform, by definition we have $\mathbf{W}^{1 \dots k}[\varphi] = 0$. Therefore distinguishing between k -wise uniform and δ -far from k -wise uniform can be reduced to distinguishing between $\mathbf{W}^{1 \dots k}[\varphi] > \left(\frac{\delta}{e^k}\right)^2$ and $\mathbf{W}^{1 \dots k}[\varphi] = 0$.

For any density function φ , $|\widehat{\varphi}(S)| = |\mathbf{E}[\varphi(\mathbf{x})\mathbf{x}^S]| \leq 1$ for any $S \subseteq [n]$. Therefore assigning $A = n^k$, we have

$$\mathbf{W}^i[\varphi] = \sum_{|S|=i} \widehat{\varphi}(S)^2 \leq n^i \leq An^{i/2}$$

for every $i = 0, 1, \dots, 2k$.

Hence we can run the $\mathbf{W}^{1 \dots k}$ Estimator Test in Theorem 4.2 with parameter $\theta = \left(\frac{\delta}{e^k}\right)^2$ and $A = n^k$, thereby we solve the Testing Problem with sample complexity $2^{O(k)} n^k / \delta^2$.

In fact by more precise calculation we can further improve the constant factor involving k to $O\left(\frac{1}{k}\right)^{k/2}$, but we will omit the proof here for the sake of brevity. \square

4.2 Lower bound for the pairwise case

An upper bound for the Closeness Problem implies an upper bound for the Testing Problem. But a lower bound for Closeness does not obviously yield a lower bound for the Testing Problem. The function used to show the lower bound for the Closeness Problem is far from k -wise uniform, but it is not sufficient to say that it is hard to distinguish between it and some k -wise uniform distribution. In [AAK⁺07], they show that

it is hard to distinguish between the fully uniform distribution and the uniform distribution on a random set of size around $O(n^{k-1}/\delta^2)$; this latter distribution is far from k -wise uniform with high probability for $k > 2$.

We show that the density function φ we used for the lower bound for the Closeness Problem is a useful density to use for a testing lower bound in the pairwise case. However it is not hard to distinguish between the fully uniform distribution and φ . Our trick is shifting φ by a random “center”. We remind the reader that we denote by $\varphi^{+t}(x) = \varphi(x \circ t)$ the distribution φ shifted by vector t . We claim that with $m = o(n/\delta^2)$ samples, it is hard to distinguish the fully uniform distribution from φ^{+t} with a uniformly randomly chosen t .

Lemma 4.3. *Let φ be the density function defined by $\varphi(x) = 1 + \frac{\delta}{n} \sum_{i < j} x_i x_j$. Assume $m < n/\delta^2$. Let $\Phi : (\{-1, 1\}^n)^m \rightarrow \mathbb{R}^{\geq 0}$ be the density associated to the distribution on m -tuples of strings defined as follows: First, choose \mathbf{t} in $\{-1, 1\}^n$ uniformly; then choose m strings independently from $\varphi^{+\mathbf{t}}$. Let $\mathbf{1}$ denote the constantly 1 function on $(\{-1, 1\}^n)^m$, the density associated to the uniform distribution. Then the χ^2 -divergence between Φ and $\mathbf{1}$, $\|\Phi - \mathbf{1}\|_2^2$, is bounded by*

$$\|\Phi - \mathbf{1}\|_2^2 \leq O\left(\frac{m\delta^2}{n}\right).$$

Proof. We need to show that $\mathbf{E}[(\Phi - \mathbf{1})^2] = \mathbf{E}[\Phi^2] - \mathbf{1} \leq O(m\delta^2/n)$. For uniform and independent $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}$,

$$\begin{aligned} \mathbf{E}[\Phi(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)})^2] &= \mathbf{E}_{\mathbf{x}} \left[\left(\mathbf{E}_{\mathbf{t}} \left[\prod_{i=1}^m \varphi^{+\mathbf{t}}(\mathbf{x}^{(i)}) \right] \right)^2 \right] \\ &= \mathbf{E}_{\mathbf{x}, \mathbf{t}, \mathbf{t}'} \left[\prod_{i=1}^m \varphi^{+\mathbf{t}}(\mathbf{x}^{(i)}) \varphi^{+\mathbf{t}'}(\mathbf{x}^{(i)}) \right] \\ &= \mathbf{E}_{\mathbf{t}, \mathbf{t}'} [\langle \varphi^{+\mathbf{t}}, \varphi^{+\mathbf{t}'} \rangle^m]. \end{aligned}$$

It is a trivial fact that $\langle \varphi^{+\mathbf{t}}, \varphi^{+\mathbf{t}'} \rangle = \varphi * \varphi(\mathbf{t} + \mathbf{t}')$. Therefore

$$\mathbf{E}[\Phi(\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)})^2] = \mathbf{E}[(\varphi * \varphi)^m].$$

We know that $\widehat{\varphi * \varphi}(S) = \varphi(S)^2$. Therefore

$$\varphi * \varphi = 1 + \frac{\delta^2}{n^2} \sum_{i < j} x_i x_j.$$

To compute $\mathbf{E}[(\varphi * \varphi)^m]$, we just need to calculate the constant term of $(1 + \frac{\delta^2}{n^2} \sum_{i < j} x_i x_j)^m$ since $x_i^2 = 1$. Suppose that when expanding this out, we take l terms of $x_i x_j$; we think these as l (possibly parallel) edges in the complete graph on n vertices. Then if these l terms “cancel out”, the associated edges form a collection of cycles, since each vertex has even degree. There are at most n^l collections of cycles with l edges. Considering choosing those l terms (edges) in order, we get an upper bound of $(mn)^l$ for the number of ways of choosing l terms of $x_i x_j$ to get canceled. Therefore we have

$$\mathbf{E} \left[\left(1 + \frac{\delta^2}{n^2} \sum_{i \neq j} x_i x_j \right)^m \right] \leq \sum_{l=0}^m (mn)^l \left(\frac{\delta^2}{n^2} \right)^l \leq \sum_{l=0}^m \left(\frac{m\delta^2}{n} \right)^l \leq 1 + O\left(\frac{m\delta^2}{n}\right),$$

which completes the proof. \square

Now we are ready to give the lower bound for sample complexity of testing fully uniform vs. far-from-pairwise uniform.

Proof of Theorem 1.5. If $m = o(n/\delta^2)$, by Lemma 4.3 we have $\|\Phi - \mathbf{1}\|_2^2 \leq o(1)$. Then any tester cannot distinguish, with more than $o(1)$ advantage, whether those m samples are fully uniform or they are drawn from φ^{+t} for some random t .

On the other hand, the proof of Theorem 1.2 shows that φ is $\Omega(\delta)$ -far from pairwise uniform, and from the Fourier characterization, we have that φ^{+t} is pairwise uniform whenever φ is. We can conclude that testing fully uniform versus δ -far-from-pairwise-uniform needs sample complexity at least $\Omega(n/\delta^2)$. \square

Unfortunately, we do not see an obvious way to generalize this lower bound to $k > 2$.

5 Testing αk -wise/fully uniform vs. far from k -wise uniform

5.1 The algorithm

In this section we show a sample-efficient algorithm for testing whether a distribution is αk -wise/fully uniform or δ -far from k -wise uniform. As a reminder, Theorem 4.2 indicates that the sample complexity of estimating $\mathbf{W}^{1\dots k}[\varphi]$ is bounded by the Fourier weight up to level $2k$. This suggests using a filter test to try to “kick out” those distributions with noticeable Fourier weight up to degree $2k$.

Filter Test. Draw m_1 samples from φ . If there exists a pair of samples \mathbf{x}, \mathbf{y} such that $|\sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i| > t\sqrt{n}$, output “Reject”; otherwise, output “Accept”.

The Overall Algorithm is combining the Filter Test and the $\mathbf{W}^{1\dots k}$ Estimation Test.

Overall Algorithm. Do the Filter Test with m_1 samples and parameter t . If it rejects, say “No”. Otherwise, do the $\mathbf{W}^{1\dots k}$ Estimation Test with m_2 samples and $\theta = (\delta/e^k)^2$. Say “No” if it outputs “ $\mathbf{W}^{1\dots k}[\varphi] > \theta$ ” and say “Yes” otherwise.

Here “Yes” means φ is αk -wise/fully uniform, and “No” means φ is δ -far from k -wise uniform. We will decide the parameters m_1, t, m_2 in the Overall Algorithm later.

For simplicity, we denote $\bar{k} = \alpha k$. We will focus on testing αk -wise uniform vs. far from k -wise uniform in the analysis. For fully uniformity, the analysis is almost the same, and we will discuss it at the end of this subsection.

First of all, we will prove that if φ is \bar{k} -wise uniform, it will pass the Filter Test with high probability, provided we choose m_1 and t properly.

Lemma 5.1. *If φ is \bar{k} -wise uniform (assuming \bar{k} is even), the Filter Test will accept with probability at least .9 when $m_1^2 \leq \frac{t^{\bar{k}}}{5\bar{k}^{\bar{k}/2}}$.*

Proof. If φ is \bar{k} -wise uniform with \bar{k} even, then by Markov’s inequality on the \bar{k} -th moment, we have

$$\Pr_{\substack{\mathbf{x}, \mathbf{y} \sim \varphi \\ \text{independent}}} \left[\left| \sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i \right| > t\sqrt{n} \right] = \Pr_{\mathbf{x}, \mathbf{y} \sim \varphi} \left[\left(\sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i \right)^{\bar{k}} > (t\sqrt{n})^{\bar{k}} \right] \leq \frac{\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi} \left[\left(\sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i \right)^{\bar{k}} \right]}{t^{\bar{k}} n^{\bar{k}/2}}.$$

When we expand $(\sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i)^{\bar{k}}$, each term is at most degree \bar{k} in x or y . Because \mathbf{x} and \mathbf{y} are independent random variables chosen from \bar{k} -wise uniform distribution φ , the whole polynomial behaves the same as if \mathbf{x} and \mathbf{y} were chosen from the fully uniform distribution:

$$\begin{aligned} \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi} \left[\left(\sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i \right)^{\bar{k}} \right] &= \mathbf{E}_{\mathbf{z} \sim \{-1, 1\}^n} \left[\left(\sum_{i=1}^n z_i \right)^{\bar{k}} \right] \\ &\leq \bar{k}^{\bar{k}/2} \left(\mathbf{E}_{\mathbf{z} \sim \{-1, 1\}^n} \left[\left(\sum_{i=1}^n z_i \right)^2 \right]^{\bar{k}/2} \right) \\ &= \bar{k}^{\bar{k}/2} n^{\bar{k}/2}. \end{aligned}$$

The inequality uses hypercontractivity; see Theorem 9.21 in [O'D14]. Hence we have

$$\Pr_{\mathbf{x}, \mathbf{y} \sim \varphi} \left[\left| \sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i \right| > t\sqrt{n} \right] \leq \frac{\bar{k}^{\bar{k}/2}}{t^{\bar{k}}}.$$

When drawing m_1 examples, there are at most $\binom{m_1}{2} \leq \frac{1}{2}m_1^2$ pairs. Hence by the union bound, the probability of φ getting rejected is at most $\frac{m_1^2 \bar{k}^{\bar{k}/2}}{2t^{\bar{k}}} \leq \frac{1}{10}$. \square

Secondly, we claim that for any distribution φ that does not get rejected by the Filter Test, it is close to a distribution φ' with upper bounds on the Fourier weights of each of its levels.

Lemma 5.2. *Any distribution φ either gets rejected by the Filter Test with probability at least .9, or there exists some distribution φ' such that:*

1. φ' and φ are $\frac{8}{m_1}$ -close in total variation distance;
2. $\mathbf{W}^i[\varphi'] \leq \frac{10^7}{m_1^2} n^i + t^i n^{i/2}$ for all $i = 1, \dots, n$.

We will present the proof of Lemma 5.2 in the next subsection.

If φ is not rejected by the Filter Test, Lemma 5.2 tells us that it is close to some distribution φ' with bounded Fourier weights on each of its levels. Even though we are drawing samples from φ , we can “pretend” that we are drawing samples from φ' since they are close:

Claim 5.3. *Let $m_2 \leq \frac{m_1}{200}$, and let $A(\mathbf{X}^{(m_2)})$ be any event related to m_2 samples in $\{-1, 1\}^n$, $\mathbf{X}^{(m_2)} = \{x_1, \dots, x_{m_2}\}$. Then we have*

$$\left| \Pr_{\mathbf{X}^{(m_2)} \sim \varphi} [A(\mathbf{X}^{(m_2)})] - \Pr_{\mathbf{X}^{(m_2)} \sim \varphi'} [A(\mathbf{X}^{(m_2)})] \right| \leq .08,$$

when φ and φ' are $\frac{8}{m_1}$ -close.

Proof. We denote by Φ (respectively, Φ') the joint distribution of m_2 samples from φ (respectively, φ'). Then by a union bound we know that Φ and Φ' are $.04$ -close, since $m_2 \frac{8}{m_1} \leq .04$. We denote $\mathbf{1}[A(\mathbf{X}^{(m_2)})]$ as the indicator function of event A happening on $\mathbf{X}^{(m_2)}$. Then we have

$$\begin{aligned} \left| \Pr_{\mathbf{X}^{(m_2)} \sim \varphi} [A(\mathbf{X}^{(m_2)})] - \Pr_{\mathbf{X}^{(m_2)} \sim \varphi'} [A(\mathbf{X}^{(m_2)})] \right| &= \left| \sum_{\mathbf{X}^{(m_2)}} \mathbf{1}[A(\mathbf{X}^{(m_2)})] \left(\Phi(\mathbf{X}^{(m_2)}) - \Phi'(\mathbf{X}^{(m_2)}) \right) \right| \\ &\leq \sum_{\mathbf{X}^{(m_2)}} \left| \Phi(\mathbf{X}^{(m_2)}) - \Phi'(\mathbf{X}^{(m_2)}) \right| \\ &= 2d_{\text{TV}}(\Phi, \Phi') \leq .08 \end{aligned}$$

which completes the proof. \square

Now we are ready to analyze the Overall Algorithm.

Proof of Theorem 1.6. We discuss distinguishing between \bar{k} -wise uniform and δ -far from k -wise uniform first. In the Overall Algorithm, we set the parameters $t = \left(10^{11} (4e^4)^k \bar{k}^{\bar{k}/2} \frac{n^k}{\delta^4} \right)^{\frac{1}{\bar{k}-2k}}$ and $m_1 = \sqrt{\frac{t^{\bar{k}}}{5\bar{k}^{\bar{k}/2}}}$ in the Filter Test; and, we set $m_2 = \frac{1}{200}m_1$ and $\theta = \left(\frac{\delta}{e^k} \right)^2$ in the $\mathbf{W}^{1\dots k}$ Estimation test.

In total we use $m_1 + m_2 = O\left(\sqrt{\frac{t^{\bar{k}}}{\bar{k}^{\bar{k}/2}}}\right)$ samples in the Overall Algorithm. By plugging in the definition of t and $\bar{k} = \alpha k$, we can simplify the sample complexity to $O(\alpha)^{k/2} \cdot n^{k/2} \cdot \frac{1}{\delta^2} \cdot \left(\frac{n^k}{\delta^4}\right)^{1/(\alpha-2)}$.

The rest of the proof is to show the correctness of this algorithm. We discuss the two cases.

“Yes” case: Suppose φ is \bar{k} -wise uniform. By Lemma 5.1 we know that φ will pass the Filter Test with probability at least .9 since $m_1^2 = \frac{t^{\bar{k}}}{5\bar{k}^{k/2}}$.

Now φ is \bar{k} -wise uniform with $\bar{k} > 2k$, which means $\widehat{\varphi}(S) = 0$ for any $1 \leq |S| \leq 2k$. Therefore by setting $\delta = \left(\frac{\theta}{e^k}\right)^2$ and $A = 1$, Theorem 4.2 tells us that m_2 samples are large enough for $\mathbf{W}^{1\dots k}$ Estimation Test to output “ $\mathbf{W}^{1\dots k}[\varphi] \leq \frac{1}{2}\theta$ ” with probability $3/4$.

The overall probability of the Overall Algorithm saying “Yes” is therefore at least $.9 \times \frac{3}{4} > \frac{2}{3}$.

“No” case: Suppose φ is δ -far from k -wise uniform. Either φ gets rejected by the Filter Test with probability .9, or according to Lemma 5.2, we know that there exists some distribution φ' which is $\frac{8}{m_1}$ -close to φ and $\mathbf{W}^i[\varphi'] \leq \frac{10^7}{m_1^2}n^i + t^i n^{i/2}$ for all $i = 1, \dots, n$.

The second stage is slightly tricky. As described in Claim 5.3, at the expense of losing .08 probability, we may pretend we are drawing samples from φ' rather than φ . Notice that $m_1^2 = \frac{t^{\bar{k}}}{5\bar{k}^{k/2}} = \omega(n^k)$. We have

$$\mathbf{W}^i[\varphi'] \leq \frac{10^7}{m_1^2}n^i + t^i n^{i/2} = (1 + o(1))t^i n^{i/2} \leq An^{i/2}$$

for $i = 0, \dots, 2k$ with parameter $A = 1.01t^{2k}$. Then plugging $A = 1.01t^{2k}$ and $\theta = \left(\frac{\delta}{e^k}\right)^2$ into Theorem 4.2, we know that the $\mathbf{W}^{1\dots k}$ Estimation Test will say “ $\mathbf{W}^{1\dots k}[\varphi] > \theta$ ” with probability at least $\frac{3}{4}$ when φ' is δ -far from k -wise uniform, provided we have at least $1005 \frac{(2e^2)^k t^k n^{k/2}}{\delta^2}$ samples. It is easy to check $m_2 = \frac{1}{200} \sqrt{\frac{t^{\bar{k}}}{5\bar{k}^{k/2}}}$ is sufficient.

However, in the real algorithm we are drawing samples from φ rather than φ' . From Claim 5.3, we know that the estimator will accept with probability at least $\frac{3}{4} - .08 > \frac{2}{3}$ when φ' is δ -far from k -wise uniform. Notice that φ and φ' are $\frac{8}{m_1}$ -close, where $\frac{8}{m_1} = o\left(\frac{\delta^4}{n^k}\right)$. Hence if φ is δ -far from k -wise uniform, φ' is also δ -far from k -wise uniform, which completes the proof.

Finally, for distinguishing between a distribution being fully uniform and a distribution being δ -far from k -wise uniform, the modification we need is that in Lemma 5.1 we use Hoeffding’s inequality to get

$$\Pr_{\mathbf{x}, \mathbf{y} \sim \varphi} \left[\left| \sum_{i=1}^n \mathbf{x}_i \mathbf{y}_i \right| > t\sqrt{n} \right] \leq 2e^{-t^2/2},$$

and then we have the constraint $m_1^2 \leq \frac{1}{10}e^{t^2/2}$. Following exactly the same analysis, we get the same algorithm with sample complexity $O(k)^k \cdot n^{k/2} \cdot \frac{1}{\delta^2} \cdot \left(\log \frac{n}{\delta}\right)^{k/2}$. \square

5.2 Proof of Lemma 5.2

The rest of this section is devoted to proving Lemma 5.2. We will use the following definition in the analysis.

Definition 5.4. For $x, y \in \{-1, 1\}^n$, we say (x, y) is *skewed* if $|\sum_{i=1}^n x_i y_i| > t\sqrt{n}$. We say that x is β -bad for distribution φ if $\Pr_{\mathbf{y} \sim \varphi}[(x, \mathbf{y}) \text{ is skewed}] > \beta$.

Claim 5.5. If $\Pr_{\mathbf{x} \sim \varphi} \left[\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad for } \varphi \right] > \frac{8}{m_1}$, then φ will be rejected by the Filter Test with probability at least .9.

Proof. Suppose $\Pr_{\mathbf{x} \sim \varphi} \left[\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad for } \varphi \right] > \frac{8}{m_1}$. We will divide the samples we draw for the Filter Test into two sets with size $m_1/2$ each. Then the probability of choosing an $\frac{8}{m_1}$ -bad x among the first $m_1/2$ samples is at least

$$\Pr_{\mathbf{x}_1, \dots, \mathbf{x}_{m_1/2} \sim \varphi} \left[\exists \mathbf{x} \frac{8}{m_1}\text{-bad for } \varphi \text{ among } \mathbf{x}_1, \dots, \mathbf{x}_{m_1/2} \right] > 1 - \left(1 - \frac{8}{m_1}\right)^{m_1/2} \geq 1 - e^{-4}.$$

Now if we have such an $\frac{8}{m_1}$ -bad x among the first $m_1/2$ samples, each (x, \mathbf{x}_t) will be skewed with probability at least $\frac{8}{m_1}$ for any $t = m_1/2 + 1, \dots, m$. Therefore

$$\Pr_{\mathbf{x}_{m/2+1}, \dots, \mathbf{x}_m} [(x, \mathbf{x}_t) \text{ is skewed for some } t = \frac{m}{2} + 1, \dots, m] \geq 1 - \left(1 - \frac{8}{m_1}\right)^{m_1/2} \geq 1 - e^{-4}.$$

Combining the two inequalities together, we know that the probability of at least one pair being skewed is at least $(1 - e^{-4})^2 \geq .9$. \square

Now we only need to consider the case when the probability of drawing a bad x from φ is very small. We want to show a stronger claim that even the probability of drawing a skewed pair from φ is small. However this might not be true for φ itself. Thus we look at another distribution φ' , which is defined to be φ conditioned on outcomes being not bad. Define φ' as

$$\varphi'(x) = \varphi(x) \frac{\mathbf{1}[x \text{ not } \frac{8}{m_1}\text{-bad for } \varphi]}{1 - \Pr_{\mathbf{x} \sim \varphi} [\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad for } \varphi]}. \quad (5)$$

We show that φ' is close to φ and that φ' has no bad samples:

Claim 5.6. *Suppose φ satisfies $\Pr_{\mathbf{x} \sim \varphi} [\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad for } \varphi] \leq \frac{8}{m_1}$ and $m_1 \geq 16$. Let φ' be defined as in eq. (5). Then:*

1. φ and φ' are $\frac{8}{m_1}$ -close;
2. $\varphi'(x) = 0$ for any x that is $\frac{16}{m_1}$ -bad for φ' .

Proof. 1. Notice that $\varphi'(x) = 0 \leq \varphi(x)$ when x is $\frac{8}{m_1}$ -bad for φ , and $\varphi'(x) \geq \varphi(x)$ otherwise. Hence,

$$\begin{aligned} d_{\text{TV}}(\varphi, \varphi') &= \frac{1}{2} \mathbf{E}_{\mathbf{x}} [|\varphi(\mathbf{x}) - \varphi'(\mathbf{x})|] \\ &= \frac{1}{2^n} \sum_{\varphi'(x) < \varphi(x)} (\varphi(x) - \varphi'(x)) \\ &\leq \Pr_{\mathbf{x} \sim \varphi} \left[\mathbf{x} \frac{8}{m_1}\text{-bad on } \varphi \right] \leq \frac{8}{m_1}. \end{aligned}$$

2. $\varphi'(x)$ is either 0 or at most $(1 + \frac{16}{m_1})\varphi(x)$ given $\Pr_{\mathbf{x} \sim \varphi} [\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad for } \varphi] \leq \frac{8}{m_1}$ and $m_1 \geq 16$. Therefore if $\varphi'(x) > 0$, x is not $\frac{8}{m_1}$ -bad for φ . Hence,

$$\begin{aligned} \Pr_{\mathbf{y} \sim \varphi'} [(x, \mathbf{y}) \text{ is skewed}] &\leq \left(1 + \frac{16}{m_1}\right) \Pr_{\mathbf{y} \sim \varphi} [(x, \mathbf{y}) \text{ is skewed}] \\ &\leq \left(1 + \frac{16}{m_1}\right) \frac{8}{m_1} \leq \frac{16}{m_1}. \quad \square \end{aligned}$$

Claim 5.7. *Suppose distribution φ satisfies $\Pr_{\mathbf{x} \sim \varphi} [\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad for } \varphi] \leq \frac{8}{m_1}$. Let φ' be defined as in eq. (5). If $\Pr_{\mathbf{x}, \mathbf{y} \sim \varphi'} [(x, \mathbf{y}) \text{ is skewed}] > \frac{10^7}{m_1^2}$, then with probability at least .9, φ will be rejected by the Filter Test.*

We want to clarify that the constraint is about φ' , but we are drawing samples from φ in the Filter Test.

Proof. We only consider the first $m'_1 = \frac{m_1}{200}$ samples. From Claim 5.6 we know that φ and φ' are $\frac{8}{m_1}$ -close. Therefore, we only need to show that if the samples are drawn from φ' , the probability of appearing a skewed pair among these m'_1 samples is at least .98. Then φ will be rejected by the Filter Test with probability at least $.98 - .08 \geq .9$ according to Claim 5.3.

Define random variable $U_{s,t}$ to be the indicator associated with the event that $(\mathbf{x}_s, \mathbf{x}_t)$ is skewed, and $U = \sum_{1 \leq s < t \leq m'_1} U_{s,t}$. We need to prove that $\Pr[U = 0] \leq .02$. (From now on, all probabilities and expectations are based on choosing samples from distribution φ' .) By Chebyshev's inequality, we know that $\Pr[U = 0] \leq \frac{\text{Var}[U]}{\mathbf{E}[U]^2}$, so we need to calculate $\text{Var}[U]$ and $\mathbf{E}[U]$.

Denote $\mu = \Pr_{\mathbf{x}, \mathbf{y} \sim \varphi'}[(\mathbf{x}, \mathbf{y}) \text{ is skewed}]$. Then $\mathbf{E}[U_{s,t}] = \mu$ for any $s < t$ and hence we have

$$\mathbf{E}[U] = \sum_{s < t} \mathbf{E}[U_{s,t}] = \binom{m'_1}{2} \mu.$$

It remains to calculate $\mathbf{E}[U^2]$. We can expand it as

$$\mathbf{E}[U^2] = \mathbf{E} \left[\left(\sum_{s < t} U_{s,t} \right)^2 \right] = \sum_{\substack{s < t \\ s' < t'}} \mathbf{E}[U_{s,t} U_{s',t'}].$$

Similar to the proof of Lemma 4.1, we discuss these expectations in three cases.

Case 1: $|\{s, t\} \cap \{s', t'\}| = 2$. Since $U_{s,t}$ is a Bernoulli random variable, we know that

$$\mathbf{E}[U_{s,t}^2] = \mathbf{E}[U_{s,t}] = \mu.$$

Case 2: $|\{s, t\} \cap \{s', t'\}| = 1$. Without loss of generality we assume $s = s'$. We consider drawing \mathbf{x}_s first. For any fixed x_s with $\varphi'(x_s) > 0$,

$$\mathbf{E}_{\mathbf{x}_{t'}}[U_{s,t'}] = \Pr_{\mathbf{x}_{t'}}[(x_s, \mathbf{x}_{t'}) \text{ get skewed}] \leq \frac{16}{m_1} = \frac{2}{25m'_1},$$

where the inequality comes from Claim 5.6. Therefore,

$$\mathbf{E}[U_{s,t} U_{s,t'}] = \mathbf{E}_{\mathbf{x}_s, \mathbf{x}_t} [U_{s,t} \mathbf{E}_{\mathbf{x}_{t'}}[U_{s,t'}]] \leq \frac{2\mu}{25m'_1}.$$

Case 3: $|\{s, t\} \cap \{s', t'\}| = 0$. Since s, t, s', t' are all distinct, we have

$$\mathbf{E}[U_{s,t} U_{s',t'}] = \mathbf{E}[U_{s,t}] \mathbf{E}[U_{s',t'}] = \mu^2.$$

Combining these cases together, we get

$$\mathbf{E}[U^2] = \binom{m'_1}{2} \mu + m'_1(m'_1 - 1)(m'_1 - 2) \frac{2\mu}{25m'_1} + \binom{m'_1}{2} \binom{m'_1 - 2}{2} \mu^2.$$

Then we have

$$\frac{\text{Var}[U]}{\mathbf{E}[U]^2} = \frac{\mathbf{E}[U^2]}{\mathbf{E}[U]^2} - 1 \leq \frac{58}{25m_1^2 \mu}.$$

By substituting $\mu \geq \frac{10^7}{m_1^2} = \frac{10^3}{4m_1^2}$, we conclude $\Pr[U = 0] = \frac{\text{Var}[U]}{\mathbf{E}[U]^2} \leq .02$, which completes the proof. \square

Now we only need to consider those distributions φ where their corresponding φ' satisfies that $\Pr_{\mathbf{x}, \mathbf{y} \sim \varphi'}[(\mathbf{x}, \mathbf{y}) \text{ is skewed}] \leq \frac{10^7}{m_1^2}$. This gives us an upper bound on the Fourier weight on all levels of φ' .

Claim 5.8. *If $\Pr_{\mathbf{x}, \mathbf{y} \sim \varphi'}[(\mathbf{x}, \mathbf{y}) \text{ is skewed}] \leq \frac{10^7}{m_1^2}$, then*

$$\mathbf{W}^i[\varphi'] \leq \frac{10^7}{m_1^2} n^i + t^i n^{i/2}$$

for $i = 1, \dots, n$.

Proof. We will first show that $\mathbf{W}^i[\varphi'] \leq \mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi'}[(\sum_{j=1}^n \mathbf{x}_j \mathbf{y}_j)^i]$. Since $(\sum_{j=1}^n x_j y_j)^i$ is a symmetric function, we can expand it as

$$\left(\sum_{j=1}^n x_j y_j \right)^i = \sum_{\substack{0 \leq k \leq i \\ i-k \text{ even}}} \alpha_k \left(\sum_{|S|=k} x^S y^S \right),$$

with positive integer coefficients α_k . Notice that

$$\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi'} \left[\sum_{|S|=k} x^S y^S \right] = \mathbf{W}^k[\varphi'].$$

Therefore

$$\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi'} \left[\left(\sum_{j=1}^n \mathbf{x}_j \mathbf{y}_j \right)^i \right] = \sum_{\substack{0 \leq k \leq i \\ i-k \text{ even}}} \alpha_k \mathbf{W}^k[\varphi'] \geq \mathbf{W}^i[\varphi'].$$

The last inequality holds because the α_k 's are positive integers and each $\mathbf{W}^k[\varphi']$ is non-negative.

The rest of the proof is devoted to bounding $\mathbf{E}_{\mathbf{x}, \mathbf{y} \sim \varphi'}[(\sum_{j=1}^n \mathbf{x}_j \mathbf{y}_j)^i]$. When (x, y) is not skewed, $\sum_j x_j y_j$ is at most n ; otherwise by the definition of “being skewed”, $\sum_j x_j y_j$ is at most $t\sqrt{n}$. Therefore,

$$\mathbf{E} \left[\left(\sum_{j=1}^n \mathbf{x}_j \mathbf{y}_j \right)^i \right] \leq \frac{10^7}{m_1^2} n^i + t^i n^{i/2}$$

for all $i = 1, \dots, n$. □

Combining the above discussion, we get the proof of Lemma 5.2.

Proof of Lemma 5.2. We consider three cases for φ .

Case 1: If $\Pr_{\mathbf{x} \sim \varphi}[\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad on } \varphi] > \frac{8}{m_1}$, Claim 5.5 tells us that φ is rejected by the Filter Test with probability at least .9.

For the remaining two cases we know that $\Pr_{\mathbf{x} \sim \varphi}[\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad on } \varphi] \leq \frac{8}{m_1}$. We construct φ' as in eq. (5).

Case 2: If $\Pr_{\mathbf{x} \sim \varphi}[\mathbf{x} \text{ is } \frac{8}{m_1}\text{-bad on } \varphi] \leq \frac{8}{m_1}$ but $\Pr_{\mathbf{x}, \mathbf{y} \sim \varphi'}[(\mathbf{x}, \mathbf{y}) \text{ is skewed}] > \frac{10^7}{m_1^2}$, Claim 5.7 tells us that φ also gets rejected with probability at least .9.

Case 3: If $\Pr_{\mathbf{x}, \mathbf{y} \sim \varphi'}[(\mathbf{x}, \mathbf{y}) \text{ is skewed}] > \frac{10^7}{m_1^2}$, then according to Claim 5.8, $\mathbf{W}^i[\varphi'] \leq \frac{10^7}{m_1^2} n^i + t^i n^{i/2}$ for all $i = 1, \dots, n$. Also by Claim 5.6 we know that φ and φ' are $\frac{8}{m_1}$ -close. □

References

- [AAK⁺07] Noga Alon, Alexandr Andoni, Tali Kaufman, Kevin Matulef, Ronitt Rubinfeld, and Ning Xie. Testing k -wise and almost k -wise independence. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 496–505, 2007.
- [ABI86] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4):567–583, 1986.
- [ADK15] Jayadev Acharya, Constantinos Daskalakis, and Gautam Kamath. Optimal testing for properties of distributions. In *Advances in Neural Information Processing Systems*, pages 3591–3599, 2015.
- [AGHP92] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k -wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.

- [AGM03] Noga Alon, Oded Goldreich, and Yishay Mansour. Almost k -wise independence versus k -wise independence. *Information Processing Letters*, 88(3):107–110, 2003.
- [AM09] Per Austrin and Elchanan Mossel. Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18(2):249–271, 2009.
- [AOW15] Sarah R. Allen, Ryan O’Donnell, and David Witmer. How to refute a random CSP. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, pages 689–708, 2015.
- [Baz09] Louay M. J. Bazzi. Polylogarithmic independence can fool DNF formulas. *SIAM Journal on Computing*, 38(6):2220–2272, 2009.
- [BFF⁺01] Tuğkan Batu, Eldar Fischer, Lance Fortnow, Ravi Kumar, Ronitt Rubinfeld, and Patrick White. Testing random variables for independence and identity. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science*, pages 442–451, 2001.
- [BFR⁺00] Tuğkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D. Smith, and Patrick White. Testing that distributions are close. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 259–269, 2000.
- [BKR04] Tuğkan Batu, Ravi Kumar, and Ronitt Rubinfeld. Sublinear algorithms for testing monotone and unimodal distributions. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 381–390, 2004.
- [Bra10] Mark Braverman. Polylogarithmic independence fools AC^0 circuits. *Journal of the ACM*, 57(5):28:1–28:10, 2010.
- [CGH⁺85] Benny Chor, Oded Goldreich, Johan Håstad, Joel Friedman, Steven Rudich, and Roman Smolensky. The bit extraction problem of t -resilient functions. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science*, pages 396–407, 1985.
- [CZ16] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 670–683, 2016.
- [DK16] Ilias Diakonikolas and Daniel M Kane. A new approach for testing properties of discrete distributions. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, pages 685–694. IEEE, 2016.
- [GR11] Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 68–75. Springer, 2011.
- [Kal02] Gil Kalai. A Fourier-theoretic perspective on the Condorcet paradox and Arrow’s theorem. *Advances in Applied Mathematics*, 29(3):412–426, 2002.
- [KMOW17] Pravesh K. Kothari, Ryuhei Mori, Ryan O’Donnell, and David Witmer. Sum of squares lower bounds for refuting any CSP. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 132–145, 2017.
- [Kra29] Mikhail Krawtchouk. Sur une généralisation des polynômes d’hermite. *Comptes Rendus de l’Académie des sciences*, 189:620–622, 1929.
- [KW85] Richard M. Karp and Avi Wigderson. A fast parallel algorithm for the maximal independent set problem. *Journal of the ACM*, 32(4):762–773, 1985.
- [Lev95] Vladimir I. Levenshtein. Krawtchouk polynomials and universal bounds for codes and designs in hamming spaces. *IEEE Transactions on Information Theory*, 41(5):1303–1321, 1995.

- [Li16] Xin Li. Improved two-source extractors, and affine extractors for polylogarithmic entropy. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, pages 168–177. IEEE, 2016.
- [Lub86] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.
- [MS77] Florence Jessie MacWilliams and Neil James Alexander Sloane. *The theory of error-correcting codes*. Elsevier, 1977.
- [NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993.
- [O’D14] Ryan O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, 2014.
- [Pan08] Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.
- [Rao47] Calyampudi Radhakrishna Rao. Factorial experiments derivable from combinatorial arrangements of arrays. *Journal of the Royal Statistical Society*, 9(1):128–139, 1947.
- [RS09] Ronitt Rubinfeld and Rocco A. Servedio. Testing monotone high-dimensional distributions. *Random Structures & Algorithms*, 34(1):24–44, 2009.
- [RX13] Ronitt Rubinfeld and Ning Xie. Robust characterizations of k -wise independence over product spaces and related testing results. *Random Structures & Algorithms*, 43(3):265–312, 2013.
- [Xie12] Ning Xie. *Testing k -wise independent distributions*. PhD thesis, Massachusetts Institute of Technology, 2012.