

The uniform distribution is complete with respect to testing identity to a fixed distribution

Oded Goldreich

October 3, 2019

Abstract

Inspired by Diakonikolas and Kane (2016), we reduce the class of problems consisting of testing whether an unknown distribution over $[n]$ equals a fixed distribution to the special case in which the fixed distribution is uniform over $[n]$. Our reduction preserves the parameters of the problem, which are n and the proximity parameter $\epsilon > 0$, up to a constant factor.

While this reduction yields no new bounds on the sample complexity of either problems, it provides a simple way of obtaining testers for equality to arbitrary fixed distributions from testers for the uniform distribution. The reduction first reduces the general case to the case of “grained distributions” (in which all probabilities are multiples of $\Omega(1/n)$), and then reduces this case to the case of the uniform distribution. Using grained distributions as a pivot of the exposition, we call attention to this natural class.

An early version of this work appeared as TR16-015 of *ECCC*. The original version reproduced some text from the author’s lecture notes on property testing [8], which were later used as a basis for his book [9]. The current revision is quite minimal, except for the appendix, which was significantly revised in order to provide a more transparent exposition.

1 Introduction

Inspired by Diakonikolas and Kane [5], we present, for every fixed distribution D over $[n]$, a simple reduction of the problem of testing whether an unknown distribution over $[n]$ equals D to the problem of testing whether an unknown distribution over $[n]$ equals the uniform distribution over $[n]$. Specifically, we reduce ϵ -testing of equality to D to $\epsilon/3$ -testing of equality to the uniform distribution over $[6n]$, denoted U_{6n} .

Hence, the sample (resp., time) complexity of testing equality to D , with respect to the proximity parameter ϵ , is at most the sample (resp., time) complexity of testing equality to U_{6n} with respect to the proximity parameter $\epsilon/3$. Since optimal bounds were known for both problems (cf., e.g., [10, 2, 1, 12, 4, 14]), our reduction yields no new bounds. Still, it provides a simple way of obtaining testers for equality to arbitrary fixed distributions from testers for the uniform distribution.

The setting at a glance. For any fixed distribution D over $[n]$, we consider the problem of ϵ -testing equality to D , where the tester is given samples drawn from an unknown distribution X and is required to distinguish the case that $X \equiv D$ from the case that X is ϵ -far from D , where the distance is the standard statistical distance. The sample complexity of this testing problem,

depends on D , and is viewed as a function of n and ϵ . We write $D \subseteq [n]$ to denote that D ranges over $[n]$.

Wishing to present reductions between such problems, we need to spell out what we mean by such a reduction. Confining ourselves to problems of testing equality to fixed distributions, we say that ϵ -testing equality to $D \subseteq [n]$ reduces to ϵ' -testing equality to $D' \subseteq [n']$ if there exists a randomized process F that maps $[n]$ to $[n']$ such that the distribution D is mapped to the distribution D' and any distribution that is ϵ -far from D is mapped to a distribution that is ϵ' -far from D' . We say that F maps the distribution X to the distribution Y if $Y \equiv F(X)$, where here we view the distributions as random variables. Note that the foregoing is a very stringent notion of reduction between distribution testing problems: Under this notion, a tester T for ϵ -testing equality to D is derived by invoking a ϵ' -tester T' for equality to D' and providing T' with the sample $F(i_1), \dots, F(i_s)$, where i_1, \dots, i_s is the sample provided to T . Still, our main result can be stated as follows.

Theorem 1 (completeness of testing equality to U_n): *For every distribution D over $[n]$ and every $\epsilon > 0$, it holds that ϵ -testing equality to D reduces to $\epsilon/3$ -testing equality to U_{6n} , where U_m denotes the uniform distribution over $[m]$. Furthermore, the same reduction F can be used for all $\epsilon > 0$.*

Hence, the sample complexity of ϵ -testing equality to D is upper bounded by the sample complexity of $\epsilon/3$ -testing equality to U_{6n} . We mention that in some cases, testing equality to D can be easier than testing equality to U_n ; such natural cases contain grained distributions (see below). (A general study of the dependence on D of the complexity of testing equality to D was undertaken in [14].)

The reduction at a glance. We decouple the reduction asserted in Theorem 1 into two steps. In the first step, we assume that the distribution D has a probability function q that ranges over multiples of $1/m$, for some parameter $m \in \mathbb{N}$; that is, $m \cdot q(i)$ is a non-negative integer (for every i). We call such a distribution m -grained, and reduce testing equality to any fixed m -grained distribution to testing equality to the uniform distribution over $[m]$. This reduction maps i uniformly at random to a set S_i of size $m \cdot q(i)$ such that the S_i 's are disjoint. Clearly, this reduction maps the distribution q to the uniform distribution over m fixed elements, and it can be verified that this randomized mapping preserves distances between distributions.

Since every distribution over $[n]$ is $\epsilon/2$ -close to a $O(n/\epsilon)$ -grained distribution, it stands to reason that the general case can be reduced to the grained case. This is indeed true, but the reduction is less obvious than the treatment of the grained case. Actually, we shall use a different “graining” procedure, which yields a better result (i.e., the result stated above). Specifically, we present a reduction of ϵ -testing equality to D to $\epsilon/3$ -testing equality to some $6n$ -grained distribution D' , where D' depends only on D . This reduction is described next.

Letting $q : [n] \rightarrow [0, 1]$ denote the probability function of D , the reduction maps $i \in [n]$ to itself with probability $\frac{\lfloor 6n \cdot q(i) \rfloor / 6n}{q(i)}$, and otherwise maps i to $n + 1$. This description suffices when $q(i) \geq 1/2n$ for every $i \in [n]$, since in this case $\frac{\lfloor 6n \cdot q(i) \rfloor / 6n}{q(i)} \geq \frac{2}{3}$, and in order to guarantee this condition (i.e., $q(i) \geq 1/2n$ for every $i \in [n]$) we use a preliminary reduction that maps $i \in [n]$ to itself with probability $1/2$ and maps it uniformly to $[n]$ otherwise. This preliminary reduction cuts the distance between distributions by a factor of two, and it can be shown that the main randomized mapping preserves distances between distributions up to a constant factor (of $2/3$).

History, credits, and an acknowledgement. The study of testing properties of distributions was initiated by Batu, Fortnow, Rubinfeld, Smith and White [2]. Testers of sample complexity $\text{poly}(1/\epsilon) \cdot \sqrt{n}$ for equality to U_n and for equality to an arbitrary distribution D over $[n]$ were presented by Goldreich and Ron [10] and Batu *et al.* [1], respectively, where the presentation in [10] is only implicit.¹ The tight lower and upper bound of $\Theta(\sqrt{n}/\epsilon^2)$ on the sample complexity of both problems were presented in [12, 4, 14] (see also [6, 5]). For a general survey of the area, the interested reader is referred to Canonne [3].

As stated upfront, our reductions are inspired by Diakonikolas and Kane [5], who presented a unified approach for deriving optimal testers for various properties of distributions (and pairs of distributions) via reductions to testing the equality of two unknown distributions that have small \mathcal{L}_2 -norm. We note that our reduction from testing equality to grained distributions to testing equality to the uniform distribution is implicit in [6].

Lastly, we wish to thank Ilias Diakonikolas for numerous email discussions, which were extremely helpful in many ways.

Organization. In Section 2 we recall the basic context and define the restricted notion of a reduction used in this work. The core of this work is presented in Section 3, where we prove Theorem 1. In Section 4 we briefly consider the problem of testing whether an unknown distribution is grained, leaving an open problem. The appendix addresses a side issue that arises in Section 4.

2 Preliminaries

We consider *discrete* probability distributions. Such distributions have a finite *support*, which we assume to be a subset of $[n]$ for some $n \in \mathbb{N}$, where the **support** of a distribution is the set of elements assigned positive probability mass. We represent such distributions either by random variables, like X , that are assigned values in $[n]$ (indicated by writing $X \in [n]$), or by probability functions like $p : [n] \rightarrow [0, 1]$ that satisfy $\sum_{i \in [n]} p(i) = 1$. These two representations correspond via $p(i) = \Pr[X = i]$. At times, we also refer to distributions as such, and denote them by D . (Distributions over other finite sets can be treated analogously, but in such a case we may provide the tester with a description of the set; indeed, n serves as a concise description of $[n]$.)

Recall that the study of “distribution testing” refers to testing properties of distributions. That is, the object being tested is a distribution, and the property it is tested for is a property of distributions (equiv., a set of distributions). The tester itself is given samples from the distribution and is required to distinguish the case that the distribution has the property from the case that the distribution is far from having the property, where the distance between distributions is defined as the total variation distance between them (a.k.a the statistical difference). That is, X and Y are said to be ϵ -close if

$$\frac{1}{2} \cdot \sum_i |\Pr[X = i] - \Pr[Y = i]| \leq \epsilon, \tag{1}$$

and otherwise they are deemed ϵ -far. With this definition in place, we are ready to recall the standard definition of testing distributions.

¹Testing equality to U_n is implicit in a test of the distribution of the endpoint of a relatively short random walk on a bounded-degree graph.

Definition 2 (testing properties of distributions): Let $\mathcal{D} = \{\mathcal{D}_n\}_{n \in \mathbb{N}}$ be a property of distributions and $s : \mathbb{N} \times (0, 1] \rightarrow \mathbb{N}$. A tester, denoted T , of sample complexity s for the property \mathcal{D} is a probabilistic machine that, on input parameters n and ϵ , and a sequence of $s(n)$ samples drawn from an unknown distribution $X \in [n]$, satisfies the following two conditions.

1. The tester accepts distributions that belong to \mathcal{D} : If X is in \mathcal{D}_n , then

$$\Pr_{i_1, \dots, i_s \sim X}[T(n, \epsilon; i_1, \dots, i_s) = 1] \geq 2/3,$$

where $s = s(n, \epsilon)$ and i_1, \dots, i_s are drawn independently from the distribution X .

2. The tester rejects distributions that far from \mathcal{D} : If X is ϵ -far from any distribution in \mathcal{D}_n (i.e., X is ϵ -far from \mathcal{D}), then

$$\Pr_{i_1, \dots, i_s \sim X}[T(n, \epsilon; i_1, \dots, i_s) = 0] \geq 2/3,$$

where $s = s(n, \epsilon)$ and i_1, \dots, i_s are as in the previous item.

Our focus is on “singleton” properties; that is, the property is $\{D_n\}_{n \in \mathbb{N}}$, where D_n is a fixed distribution over $[n]$. Note that n fully specifies the distribution D_n , and we do not consider the complexity of obtaining an explicit description of D_n from n . For sake of simplicity, we will consider a generic n and omit it from the notation (i.e., use D rather than D_n). Furthermore, we refer to ϵ -testers derived by setting the proximity parameter to ϵ . Nevertheless, all testers discussed here are actually uniform with respect to the proximity parameter ϵ (and also with respect to n , assuming they already derived or obtained an explicit description of D_n).

Confining ourselves to problems of testing equality to distributions, we formally restate the notion of a reduction used in the introduction. In fact, we explicitly refer to the randomized mapping at the heart of the reduction, and also define a stronger (i.e., uniform over ϵ) notion of a reduction that captures the furthermore part of Theorem 1.

Definition 3 (reductions via filters): We say that a randomized process F , called a filter, reduces ϵ -testing equality to $D \subseteq [n]$ to ϵ' -testing equality to $D' \subseteq [n']$ if the following two conditions hold:

1. The filter F maps the distribution D to the distribution D' ; that is, $p'(i) = \sum_j p(j) \cdot \Pr[F(j) = i]$, where p and p' denote the probability functions of D and D' , respectively.
2. The filter F maps any distribution that is ϵ -far from D to a distribution that is ϵ' -far from D' ; that is, if q is ϵ -far from D , then $q'(i) \stackrel{\text{def}}{=} \sum_j q(j) \cdot \Pr[F(j) = i]$ is ϵ' -far from D' .

We say that F reduces testing equality to $D \subseteq [n]$ to testing equality to $D' \subseteq [n']$ if, for some constant c and every $\epsilon > 0$, it holds that F reduces ϵ -testing equality to D to ϵ/c -testing equality to D' .

Recall that we say that F maps the distribution X to the distribution Y if Y and $F(X)$ are identically distributed (i.e., $Y \equiv F(X)$), where we view the distributions as random variables. We stress that if F is invoked t times on the same i , then the t outcomes are (identically and) independently distributed. Hence, a sequence of samples drawn independently from a distribution X is mapped to a sequence of samples drawn independently from the distribution $F(X)$.

Note (added in revision): As stated in the introduction, Definition 3 captures a natural but stringent notion of a reduction. First, note that this notion extends to reducing testing any set of distributions \mathcal{D} to testing the set \mathcal{D}' (by requiring that F maps any distribution in \mathcal{D} to some distribution in \mathcal{D}' while mapping any distribution that is ϵ -far from \mathcal{D} to a distribution that is ϵ' -far from \mathcal{D}'). However, more general definitions may allow the tester of \mathcal{D} to use the sample provided to it in arbitrary ways and invoke the tester of \mathcal{D}' on an arbitrary sample as long as it distinguishes distributions in \mathcal{D} from distributions that are ϵ -far from \mathcal{D} . While such general definitions are analogous to Cook-reductions, Definition 3 seems analogous to a very restricted (i.e., “local”) notion of a Karp-reduction.

3 The reduction

Recall that testing equality to a fixed distribution D means testing the property $\{D\}$; that is, testing whether an unknown distribution equals the fixed distribution D . For any distribution D over $[n]$, we present a reduction of the task of ϵ -testing $\{D\}$ to the task of $\epsilon/3$ -testing the uniform distribution over $[6n]$.

3.1 Overview

We decouple the reduction into two steps. In the first step, we assume that the distribution D has a probability function q that ranges over multiples of $1/m$, for some parameter $m \in \mathbb{N}$; that is, $m \cdot q(i)$ is a non-negative integer (for every i). We call such a distribution *m-grained*, and reduce testing equality to any fixed *m-grained* distribution to testing uniformity (over $[m]$). Next, in the second step, we reduce testing equality to any distribution over $[n]$ to testing equality to some fixed $6n$ -grained distribution.

Definition 4 (grained distributions): *We say that a probability distribution over $[n]$ having a probability function $q : [n] \rightarrow [0, 1]$ is *m-grained* if q ranges over multiples of $1/m$; that is, if for every $i \in [n]$ there exists a non-negative integer m_i such that $q(i) = m_i/m$.*

Note that the uniform distribution over $[n]$ is n -grained, and it is the only n -grained distribution having support $[n]$. Furthermore, if a distribution D results from applying some *deterministic process* to the uniform distribution over $[m]$, then D is m -grained. On the other hand, any m -grained distribution must have support size at most m .

3.2 Testing equality to a fixed grained distribution

Fixing any m -grained distribution (represented by a probability function) $q : [n] \rightarrow \{j/m : j \in \mathbb{N} \cup \{0\}\}$, we consider a randomized transformation (or “filter”), denoted F_q , that maps the support of q to $S = \{\langle i, j \rangle : i \in [n] \wedge j \in [m_i]\}$, where $m_i = m \cdot q(i)$. (We stress that, as with any randomized process considered so far (e.g., any type of randomized algorithm including any tester), invoking the filter several times on the same input yields independently and identically distributed outcomes.) Specifically, for every i in the support of q , we map i uniformly to $S_i = \{\langle i, j \rangle : j \in [m_i]\}$; that is, $F_q(i)$ is uniformly distributed over S_i . If i is outside the support of q (i.e., $q(i) = 0$), then we map it to $\langle i, 0 \rangle$. Note that $|S| = \sum_{i \in [n]} m_i = \sum_{i \in [n]} m \cdot q(i) = m$. The key observations about this filter are:

1. *The filter F_q maps q to a uniform distribution:* If Y is distributed according to q , then $F_q(Y)$ is distributed uniformly over S ; that is, for every $\langle i, j \rangle \in S$, it holds that

$$\begin{aligned} \Pr[F_q(Y) = \langle i, j \rangle] &= \Pr[Y = i] \cdot \Pr[F_q(i) = \langle i, j \rangle] \\ &= q(i) \cdot \frac{1}{m_i} \\ &= \frac{m_i}{m} \cdot \frac{1}{m_i} \end{aligned}$$

which equals $1/m = 1/|S|$.

2. *The filter preserves the variation distance between distributions:* The total variation distance between $F_q(X)$ and $F_q(X')$ equals the total variation distance between X and X' . This holds since, for $S' = S \cup \{\langle i, 0 \rangle : i \in [n]\}$, we have

$$\begin{aligned} &\sum_{\langle i, j \rangle \in S'} |\Pr[F_q(X) = \langle i, j \rangle] - \Pr[F_q(X') = \langle i, j \rangle]| \\ &= \sum_{\langle i, j \rangle \in S'} |\Pr[X = i] \cdot \Pr[F_q(i) = \langle i, j \rangle] - \Pr[X' = i] \cdot \Pr[F_q(i) = \langle i, j \rangle]| \\ &= \sum_{\langle i, j \rangle \in S'} \Pr[F_q(i) = \langle i, j \rangle] \cdot |\Pr[X = i] - \Pr[X' = i]| \\ &= \sum_{i \in [n]} |\Pr[X = i] - \Pr[X' = i]|. \end{aligned}$$

Indeed, this is a generic statement that applies to any filter that maps i to a pair $\langle i, Z_i \rangle$, where Z_i is an arbitrary distribution that only depends on i . (Equivalently, the statement holds for any filter that maps i to a random variable Z_i that only depends on i such that the supports of the different Z_i 's are disjoint.)

Noting that a knowledge of q allows to implement F_q as well as to map S to $[m]$, yields the following reduction.

Algorithm 5 (reducing testing equality to m -grained distributions to testing uniformity over $[m]$):
Let D be an m -grained distribution with probability function $q : [n] \rightarrow \{j/m : j \in \mathbb{N} \cup \{0\}\}$. On input $(n, \epsilon; i_1, \dots, i_s)$, where $i_1, \dots, i_s \in [n]$ are samples drawn according to an unknown distribution $p : [n] \rightarrow [0, 1]$, invoke an ϵ -tester for uniformity over $[m]$ by providing it with the input $(m, \epsilon; i'_1, \dots, i'_s)$ such that for every $k \in [s]$ the sample i'_k is generated as follows:

1. Generate $\langle i_k, j_k \rangle \leftarrow F_q(i_k)$.

Recall that if $m_{i_k} \stackrel{\text{def}}{=} m \cdot q(i_k) > 0$, then j_k is selected uniformly in $[m_k]$, and otherwise $j_k \leftarrow 0$. We stress that if F_q is invoked t times on the same i , then the t outcomes are (identically and) independently distributed. Hence, the s samples drawn independently from p are mapped to s samples drawn independently from p' such that $p'(\langle i, j \rangle) = p(i)/m_i$ if $j \in [m_i]$ and $p'(\langle i, 0 \rangle) = p(i)$ if $m_i = 0$.

2. If $j_k \in [m_{i_k}]$, then $\langle i_k, j_k \rangle \in S$ is mapped to its rank in S (according to a fixed order of S), where $S = \{\langle i, j \rangle : i \in [n] \wedge j \in [m_i]\}$, and otherwise $\langle i_k, j_k \rangle \notin S$ is mapped to $m + 1$.

(Alternatively, the reduction may just reject if any of the j_k 's equals 0.)²

The forgoing description presumes that the tester for uniform distributions over $[m]$ also operates well when given arbitrary distributions (which may have a support that is not a subset of $[m]$). However, any tester for uniformity can be easily extended to do so (see discussion in Section 3.4). In any case, we get

Proposition 6 (Algorithm 5 as a reduction): *The filter F_q used in Algorithm 5 reduces ϵ -testing equality to an m -grained distribution D (over $[n]$) to ϵ -testing equality to the uniform distribution over $[m]$, where the distributions tested in the latter case are over $[m + 1]$. Furthermore, if the support of D equals $[n]$ (i.e., $q(i) > 0$ for every $i \in [n]$), which may happen only if $m \geq n$, then the reduction is to testing whether a distribution over $[m]$ is uniform on $[m]$.*

Using any of the known uniformity tests that have sample complexity $O(\sqrt{n}/\epsilon^2)$,³ we obtain –

Corollary 7 (testing equality to m -grained distributions): *For any fixed m -grained distribution D , the property $\{D\}$ can be ϵ -tested in sample complexity $O(\sqrt{m}/\epsilon^2)$.*

We note that the foregoing *tester for equality to grained distributions* is of independent interest, which extends beyond its usage towards testing equality to arbitrary distributions.

3.3 From arbitrary distributions to grained ones

We now turn to the problem of testing equality to an arbitrary known distribution, represented by $q : [n] \rightarrow [0, 1]$. The basic idea is to round all probabilities to multiples of γ/n , for an error parameter γ (which will be a small constant). Of course, this rounding should be performed so that the sum of probabilities equals 1. For example, we may use a randomized filter that, on input i , outputs i with probability $\frac{m_i \gamma/n}{q(i)}$, where $m_i = \lfloor q(i) \cdot n/\gamma \rfloor$, and outputs $n + 1$ otherwise. Hence, if i is distributed according to p , then the output of this filter will be i with probability $\frac{\gamma m_i/n}{q(i)} \cdot p(i)$. This works well if $\gamma m_i/n \approx q(i)$, which is the case if $q(i) \gg \gamma/n$ (equiv., $m_i \gg 1$), but may run into trouble otherwise.

For starters, we note that if $q(i) = 0$, then $\frac{\gamma m_i/n}{q(i)}$ is undefined, and replacing it by either 0 or 1 will not do. More generally, suppose that $q(i) \in (0, \gamma/n)$ (e.g., $q(i) = 0.4\gamma/n$). In this case, setting $m_i = 0$ means that the filter is oblivious of the probability assigned to this i , and does not distinguish distributions that agree on $\{i : q(i) \geq \gamma/n\}$ but greatly differ on $\{i : q(i) < \gamma/n\}$, which means that it does not distinguish the distribution associated with q from some distributions that are 0.1γ -far from it.⁴ Hence, we modify the basic idea such to avoid this problem.

Specifically, we first use a filter that averages the input distribution p with the uniform distribution, and so guarantees that all elements occur with probability at least $1/2n$, while preserving distances between different input distributions (up to a factor of two). Only then, do we apply the

²See farther discussion in Section 3.4.

³Recall that the alternatives include the tests of [12] and [4] or the collision probability test (of [10]), per its improved analysis in [7].

⁴Consider, for example, the case that $q(i) = 0.4\gamma/n$ if $i \in [0.5n]$ and $q(i) = (2 - 0.4\gamma)/n$ otherwise, and any distribution X such that $\Pr[X=i] < \gamma/n$ if $i \in [0.5n]$ and $\Pr[X=i] = q(i)$ otherwise. Then, each of these possible X 's will be mapped by F to the same distribution, although such distributions may be 0.1γ -far from the distribution associated with q .

foregoing proposed filter (which outputs i with probability $\frac{m_i \cdot \gamma/n}{q(i)}$, where $m_i = \lfloor q(i) \cdot n/\gamma \rfloor$, and outputs $n+1$ otherwise). Details follow.

1. We first use a filter F' that, on input $i \in [n]$, outputs i with probability $1/2$, and outputs the uniform distribution (on $[n]$) otherwise. Hence, if i is distributed according to the distribution p , then $F'(i)$ is distributed according to $p' = F'(p)$ such that

$$p'(i) = \frac{1}{2} \cdot p(i) + \frac{1}{2} \cdot \frac{1}{n}. \quad (2)$$

(Indeed, we denote by $F'(p)$ the probability function of the distribution obtained by selecting i according to the probability function p and outputting $F'(i)$.)

Let $q' = F'(q)$; that is, $q'(i) = 0.5 \cdot q(i) + (1/2n) \geq 1/2n$.

2. Next, we apply a filter $F''_{q'}$, which is related to the filter F_q used in Algorithm 5. Letting $m_i = \lfloor q'(i) \cdot n/\gamma \rfloor$, on input $i \in [n]$, the filter outputs i with probability $\frac{m_i \cdot \gamma/n}{q'(i)}$, and outputs $n+1$ otherwise (i.e., with probability $1 - \frac{m_i \cdot \gamma/n}{q'(i)}$).

Note that $\frac{m_i \cdot \gamma/n}{q'(i)} \leq 1$, since $m_i \leq q'(i) \cdot n/\gamma$. On the other hand, recalling that $q'(i) \geq 1/2n$ and observing that $m_i \cdot \gamma/n > ((q'(i) \cdot n/\gamma) - 1) \cdot \gamma/n = q'(i) \cdot n - \gamma$, it follows that $\frac{m_i \cdot \gamma/n}{q'(i)} > 1 - 2\gamma$.

Now, if i is distributed according to the distribution p' , then $F''_{q'}(i)$ is distributed according to $p'' : [n+1] \rightarrow [0, 1]$ such that, for every $i \in [n]$, it holds that

$$p''(i) = p'(i) \cdot \frac{m_i \cdot \gamma/n}{q'(i)} \quad (3)$$

and $p''(n+1) = 1 - \sum_{i \in [n]} p''(i)$.

Let q'' denote the probability function related (by $F''_{q'}$) to q' . Then, for every $i \in [n]$, it holds that $q''(i) = q'(i) \cdot \frac{m_i \cdot \gamma/n}{q'(i)} = m_i \cdot \gamma/n \in \{j \cdot \gamma/n : j \in \mathbb{N} \cup \{0\}\}$ and $q''(n+1) = 1 - \sum_{i \in [n]} m_i \cdot \gamma/n < \gamma$, since $m \stackrel{\text{def}}{=} \sum_{i \in [n]} m_i > \sum_{i \in [n]} ((n/\gamma) \cdot q'(i) - 1) = (n/\gamma) - n$. Note that if n/γ is an integer, then q'' is n/γ -grained, since in this case $q''(n+1) = 1 - m \cdot \gamma/n = (n/\gamma - m) \cdot \gamma/n$. Furthermore, if $m = n/\gamma$, which happens if and only if $q'(i) = m_i \cdot \gamma/n$ for every $i \in [n]$, then q'' has support $[n]$, and otherwise it has support $[n+1]$.

Combining these two filters, we obtain the desired reduction.

Algorithm 8 (reducing testing equality to a general distribution to testing equality to a $O(n)$ -grained distributions): *Let D be an arbitrary distribution with probability function $q : [n] \rightarrow [0, 1]$, and T be an ϵ' -tester for m -grained distributions having sample complexity $s(m, \epsilon')$. On input $(n, \epsilon; i_1, \dots, i_s)$, where $i_1, \dots, i_s \in [n]$ are $s = s(O(n), \epsilon/3)$ samples drawn according to an unknown distribution p , the tester proceeds as follows:*

1. *It produces a s -long sequence (i''_1, \dots, i''_s) by applying $F''_{F'(q)} \circ F'$ to (i_1, \dots, i_s) , where F' and $F''_{q'}$ are as in Eq. (2)&(3); that is, for every $k \in [s]$, it produces $i'_k \leftarrow F'(i_k)$ and $i''_k \leftarrow F''_{F'(q)}(i'_k)$.*

(Recall that $F''_{q'}$ depends on a universal constant γ , which we shall set to $1/6$.)

2. It invokes the $\epsilon/3$ -tester T for q'' providing it with the sequence (i''_1, \dots, i''_s) . Note that this is a sequence over $[n+1]$.

Using the notations as in Eq. (2)&(3), we first observe that the total variation distance between $p' = F'(p)$ and $q' = F'(q)$ is half the total variation distance between p and q (since $p'(i) = 0.5 \cdot p(i) + (1/2n)$ and ditto for q'). Next, we observe that the total variation distance between $p'' = F''_{q'}(p')$ and $q'' = F''_{q'}(q')$ is lower bounded by a constant fraction of the total variation distance between p' and q' . To see this, let X and Y be distributed according to p' and q' , respectively, and observe that

$$\begin{aligned} \sum_{i \in [n]} |\Pr[F''_{q'}(X) = i] - \Pr[F''_{q'}(Y) = i]| &= \sum_{i \in [n]} \left| p'(i) \cdot \frac{m_i \gamma / n}{q'(i)} - q'(i) \cdot \frac{m_i \gamma / n}{q'(i)} \right| \\ &= \sum_{i \in [n]} \frac{m_i \gamma / n}{q'(i)} \cdot |p'(i) - q'(i)| \\ &\geq \min_{i \in [n]} \left\{ \frac{m_i \gamma / n}{q'(i)} \right\} \cdot \sum_{i \in [n]} |p'(i) - q'(i)|. \end{aligned}$$

As stated above, recalling that $q'(i) \geq 1/2n$ and $m_i = \lfloor (n/\gamma) \cdot q'(i) \rfloor > (n/\gamma) \cdot q'(i) - 1$, it follows that

$$\frac{m_i \gamma / n}{q'(i)} > \frac{((n/\gamma) \cdot q'(i) - 1) \cdot \gamma / n}{q'(i)} = 1 - \frac{\gamma / n}{q'(i)} \geq 1 - \frac{\gamma / n}{1/2n} = 1 - 2\gamma.$$

Hence, if p is ϵ -far from q , then p' is $\epsilon/2$ -far from q' , and p'' is $\epsilon/3$ -far from q'' , where we use $\gamma \leq 1/6$. On the other hand, if $p = q$, then $p'' = q''$. Noting that q'' is an n/γ -grained distribution, provided that n/γ is an integer (as is the case for $\gamma = 1/6$), we complete the analysis of the reduction. Hence,

Proposition 9 (Algorithm 8 as a reduction): *The filter $F''_{F'(q)} \circ F'$ used in Algorithm 8 reduces ϵ -testing equality to any fixed distribution D (over $[n]$) to ϵ -testing equality to an $6n$ -grained distribution over $[n']$, where $n' \in \{n, n+1\}$ depends on q .⁵ Furthermore, the support of $F''_{F'(q)} \circ F'(q)$ equals $[n']$.*

Hence, the sample complexity of ϵ -testing equality to arbitrary distributions over $[n]$ equals the sample complexity of $\epsilon/3$ -testing equality to $6n$ -grained distributions (which is essentially a special case).

Digest. One difference between the filter underlying Algorithm 5 and the one underlying Algorithm 8 is that the former preserves the exact distance between distributions, whereas the later only preserves them up to a constant factor. The difference is reflected in the fact that the first filter maps the different i 's to distributions of disjoint support, whereas the second filter (which is composed of the filters of Eq. (2)&(3)) maps different i 's to distributions of non-disjoint support. (Specifically, the filter of Eq. (2) maps every $i \in [n]$ to a distribution that assigns each $i' \in [n]$ probability at least $1/2n$, whereas the filter of Eq. (3) typically maps each $i \in [n]$ to a distribution with a support that contains the element $n+1$.)

⁵Typically, $n' = n+1$. Recall that $n' = n$ if and only if D itself is $6n$ -grained, in which case the reduction is not needed anyhow.

3.4 From arbitrary distributions to the uniform one

Combining the reductions stated in Propositions 6 and 9, we obtain a proof of Theorem 1.

Theorem 10 (Theorem 1, restated) *For every probability function $q : [n] \rightarrow [0, 1]$ the filter $F_{q''} \circ F''_{F'(q)} \circ F'$, where $q'' = F''_{F'(q)} \circ F'(q)$ is as in Algorithm 8 and $F_{q''}$ is as in Algorithm 5, reduces ϵ -testing equality to q to $\epsilon/3$ -testing equality to the uniform distribution over $[6n]$.*

Proof: First, setting $\gamma = 1/6$ and using the filter $F''_{F'(q)} \circ F'$, we reduce the problem of ϵ -testing equality to q to the problem of $\epsilon/3$ -testing equality to the $6n$ -grained distribution q'' , while noting that the distribution q'' has support $[n']$, where $n' \in \{n, n+1\}$ (depending on q). Note that the latter assertion relies on the furthermore part of Proposition 9. Next, using the furthermore part of Proposition 6, we note that $F_{q''}$ reduces $\epsilon/3$ -testing equality to q'' to $\epsilon/3$ -testing equality to the uniform distribution over $[6n]$. ■

Observe that the proof of Theorem 10 avoids the problem discussed right after the presentation of Algorithm 5, which refers to the fact that testing equality to an m -grained distribution $q : [n] \rightarrow [0, 1]$ is reduced to testing whether distributions over $[n']$ are uniform over $[m]$, where in some cases $n' \in [n, n+m]$ rather than $n' = m$. These bad cases arise when the support of q is a strict subset of $[n]$, and it was avoided since we applied the filter of Algorithm 5 to distributions $q'' : [n'] \rightarrow [0, 1]$ that have support $[n']$. Nevertheless, it is nice to have a reduction from the general case of “testing uniformity” to the special case, where the general case refers to testing whether distributions over $[n]$ are uniform over $[m]$, for any n and m , and the special case mandates that $m = n$. Such a reduction is provided next.

Theorem 11 (testing uniform distributions, a reduction between two versions): *There exists a simple filter that maps U_m to U_{2m} , while mapping any distribution X that is ϵ -far from U_m to a distribution over $[2m]$ that is $\epsilon/2$ -far from U_{2m} . We stress that X is not necessarily distributed over $[m]$ and remind the reader that U_n denotes the uniform distribution over $[n]$.*

Thus, this filter reduces ϵ -testing whether distributions over $[n]$ are uniform over $[m]$ to $\epsilon/2$ -testing whether distributions over $[2m]$ are uniform over $[2m]$.

Proof: The filter, denoted F , maps $i \in [m]$ uniformly at random to an element in $\{i, m+i\}$, while mapping any $i \notin [m]$ uniformly at random to an element in $[m]$. Observe that any distribution over $[n]$ is mapped to a distribution over $[2m]$ and that $F(U_m) \equiv U_{2m}$. Note that F does not necessarily preserve distances between arbitrary distributions over $[n]$ (e.g., both the uniform distribution over $[2m]$ and the uniform distribution over $[m] \cup [2m+1, 3m]$ are mapped to the same distribution), but (as shown next) F preserves distances to the relevant uniform distributions up to a constant factor. Specifically, note that

$$\sum_{i \in [m+1, 2m]} |\Pr[F(X)=i] - \Pr[U_{2m}=i]| = \frac{1}{2} \cdot \sum_{i \in [m]} |\Pr[X=i] - \Pr[U_m=i]|$$

and

$$\sum_{i \in [m]} |\Pr[F(X)=i] - \Pr[U_{2m}=i]| \geq \Pr[F(X) \in [m]] - \Pr[U_{2m} \in [m]]$$

$$\begin{aligned}
&= \left(\frac{1}{2} \cdot \Pr[X \in [m]] + \Pr[X \notin [m]] \right) - \frac{1}{2} \\
&= \frac{1}{2} \cdot \sum_{i \notin [m]} |\Pr[X=i] - \Pr[U_m=i]|,
\end{aligned}$$

since $0.5 \cdot \Pr[X \in [m]] + \Pr[X \notin [m]] \geq 0.5$. Hence, the total variation distance between $F(X)$ and U_{2m} is at least half the total variation distance between X and U_m . ■

4 On the complexity of testing whether a distribution is grained

A natural question that arises from the interest in grained distributions refers to the complexity of testing whether an unknown distribution is grained. Specifically, given n and m (and a proximity parameter ϵ), how many samples are required in order to determine whether an unknown distribution X over $[n]$ is m -grained or ϵ -far from any m -grained distribution. This question can be partially answered by invoking the results of Valiant and Valiant [13]. Specifically, for an upper bound we use their “learning up to relabelling” algorithm, which may be viewed as a learner of histograms (which is what it actually does). Recall that the histogram of the probability function p is defined as the multiset $\{p(i) : i \in [n]\}$ (equiv., as the set of pairs $\{(v, m) : m = |\{i \in [n] : p(i) = v\}| > 0\}$).

Theorem 12 (learning the histogram [13, Thm. 1]):⁶ *There exists an $O(\epsilon^{-2} \cdot n/\log n)$ time algorithm that, on input n, ϵ and $O(\epsilon^{-2} \cdot n/\log n)$ samples drawn from an unknown distribution $p : [n] \rightarrow [0, 1]$, outputs, with probability $1 - 1/\text{poly}(n)$, a histogram of a distribution that is ϵ -close to p .*

The implication of this result on testing any label-invariant property of distributions is immediate, where a property of distribution \mathcal{D} is called **label-invariant** if for every distribution $p : [n] \rightarrow [0, 1]$ in \mathcal{D} and every permutation $\pi : [n] \rightarrow [n]$ it holds that $p \circ \pi$ is in \mathcal{D} . In our case, the tester consists of employing the algorithm of Theorem 12 with proximity parameter $\epsilon/2$ and accepting if and only if the output fits a histogram of a distribution that is $\epsilon/2$ -close to being m -grained. The same holds with respect to estimating the distance from the set of m -grained distributions (which can be captured as a special case of label-invariant properties). Hence, we get

Corollary 13 (testing whether a distribution is grained): *For every $n, m \in \mathbb{N}$, the set of m -grained distributions over $[n]$ has a tester of sample complexity $O(\epsilon^{-2} \cdot n/\log n)$. Furthermore, the distance of an unknown distribution to the set of m -grained distributions over $[n]$ can be approximated up to an additive error of ϵ using the same number of samples.*

We comment that it seems that, using the techniques of [13], one can reduce the complexity to $O(\epsilon^{-2} \cdot n'/\log n')$, where $n' = \min(n, m)$. (For the case of testing, this is shown in the appendix, using a reduction.) On the other hand, for $m \in [\Omega(n), O(n)]$, the above distance approximator is optimal, whereas it makes no sense to consider $m > n/\epsilon$ (since any distribution over $[n]$ is ϵ -close to being n/ϵ -grained). The negative result follows from the corresponding result of Valiant and Valiant [13].

⁶Valiant and Valiant [13] stated this result for the “relative earthmover distance” (REMD) and commented that the total variation distance up to relabelling is upper bounded by REMD. This claim appears as a special case of [15, Fact 1] (using $\tau = 0$), and a detailed proof appears in [11].

Theorem 14 (optimality of Theorem 12, [13, Thm. 2]):⁷ For every sufficiently small $\epsilon > 0$, there exist two distributions $p_1, p_2 : [n] \rightarrow [0, 1]$ that are indistinguishable by any label-invariant algorithm that takes $O(\epsilon^{-1}n/\log n)$ samples although p_1 is ϵ -close to the uniform distribution over $[n]$ and p_2 is ϵ -close to the uniform distribution over some set of $n/2$ elements.

Let us spell out that, in the current context, an algorithm A is called **label-invariant** if for every permutation $\pi : [n] \rightarrow [n]$ and every sample i_1, \dots, i_s , it holds that $A(n, \epsilon; i_1, \dots, i_s) \equiv A(n, \epsilon; \pi(i_1), \dots, \pi(i_s))$. Indeed, when estimating the distance to a label-invariant property, we may assume (w.l.o.g.) that the algorithm is label-invariant. Combining Theorem 14 with the latter fact, we get –

Corollary 15 (optimality of Corollary 13): For any $m \in [\Omega(n), O(n)]$, estimating the distance to the set of m -grained distributions over $[n]$ up to a sufficiently small additive constant requires $\Omega(n/\log n)$ samples.

Similarly, tolerant testing in the sense of distinguishing distributions that are ϵ_1 -close to being m -grained from distributions that are ϵ_2 -far from being m -grained requires $\Omega(n/\log n)$ samples, for any constant $\epsilon_2 \in (0, 1/(2 \cdot \lfloor 2m/n \rfloor))$ and $\epsilon_1 \in (0, \epsilon_2)$.

Proof: The case of $m = n/2$ follows by invoking Theorem 14, while observing that p_1 is ϵ -close to being m -grained, whereas p_2 is ϵ -close to a distribution that is $(0.499 - \epsilon)$ -far from being m -grained,⁸ where p_1, p_2 and ϵ are as in Theorem 14. Hence, distinguishing the distributions p_2 and p_1 (in a label-invariant manner) is reducible to $(0.499 - 2\epsilon)$ -testing the set of distributions that are ϵ -close to being m -grained, which implies that the latter task has sample complexity $\Omega(n/\log n)$.

For $m < n/2$, we invoke Theorem 14 while resetting n to $2m$, which means that we consider distributions over $[n]$ with a support that is a subset of $[2m]$. (So the lower bound is $\Omega(m/\log m) = \Omega(n/\log n)$, where the inequality uses $m = \Omega(n)$.)

For $m > n/2$ (equiv., $n < 2m$), we show a reduction of the distinguishing task underlying Theorem 14 to the testing problem at hand. Specifically, let $t = \lfloor 2m/n \rfloor$, and assume that $t \in [2, O(1)]$ divides m (otherwise use $\lfloor m/t \rfloor$ instead of m/t , and reduce testing $(t \cdot \lfloor m/t \rfloor)$ -grained distributions to testing m -grained distributions).⁹ Consider a randomized filter, denoted $F_{m,t}$, that maps $i \in [m/t]$ to $(m/t) + i$ with probability $1/t$ and otherwise maps it to itself, but always maps $i \in [n] \setminus [m/t]$ to $i - (m/t)$. Then:

- $F_{m,t}$ maps the uniform distribution over $[m/t]$ to an m -grained distribution, since $q'_2(j) \stackrel{\text{def}}{=} \Pr[F_{m,t}(U_{m/t}) = j]$ equals $\Pr[U_{m/t} = j] \cdot (t - 1)/t = (t - 1)/m$ if $j \in [m/t]$ and equals $\Pr[U_{m/t} = j - (m/t)] \cdot 1/t = 1/m$ if $j \in [(m/t) + 1, 2m/t]$.
- $F_{m,t}$ maps the uniform distribution over $[2m/t]$ to a distribution that is $(0.999/2t)$ -far from being m -grained, since $q'_2(j) \stackrel{\text{def}}{=} \Pr[F_{m,t}(U_{2m/t}) = j]$ equals $\Pr[U_{2m/t} = j + (m/t)] + \Pr[U_{2m/t} = j] \cdot (t - 1)/t = (2t - 1)/2m$ if $j \in [m/t]$ and equals $\Pr[U_{2m/t} = j - (m/t)] \cdot 1/t = 1/2m$ if $j \in [(m/t) + 1, 2m/t]$.

⁷Like in Footnote 6, we note that Valiant and Valiant [13] stated this result for the “relative earthmover distance” (REMD) and commented that the total variation distance up to relabelling is upper bounded by REMD. This claim appears as a special case of [15, Fact 1] (using $\tau = 0$), and a detailed proof appears in [11].

⁸The constant 0.499 stands for an arbitrary large constant that is smaller than 0.5. Recall that the definition of δ -far mandates that the relevant distance be *greater* than δ .

⁹For example, the reduction may use a filter that maps $i \in [n]$ to itself with probability $t \cdot \lfloor m/t \rfloor / m$ and maps it to n otherwise.

Applying the filter $F_{m,t}$ to the distributions p_1 and p_2 of Theorem 14 (while setting $n = 2 \cdot m/t$), we obtain distributions p'_2 and p'_1 such that p'_2 is ϵ -close to being m -grained, whereas p'_1 is $((0.999/2t) - \epsilon)$ -far from being m -grained, since filters can only decrease the distance between distributions.¹⁰ Hence, distinguishing the distributions p_2 and p_1 (over $[2m/t]$) is reducible to $((0.999/2t) - 2\epsilon)$ -testing the set of distributions that are ϵ -close to being m -grained, which implies that the latter task has sample complexity $\Omega((2m/t)/\log(2m/t))$. (The claim follows by recalling that $1/t = \Omega(1)$, since $m = O(n)$.) ■

Open Problems. Note that Corollary 15 does not refer to testing, but rather to distance approximation, and there are natural cases in which the complexity of testing a property of distributions is significantly lower than the corresponding distance approximation task (cf. [10] versus [13]). Hence, we ask –

Open Problem 16 (the sample complexity of testing whether a distribution is grained): *For any m and n , what is the sample complexity of testing the property that consists of all m -grained distributions over $[n]$.*

This question can be generalized to properties that allow m to reside in some predetermined set M , where the most natural case is that M is an interval, say of the form $[m', 2m']$.

Open Problem 17 (Problem 16, generalized): *For any finite set $M \subset \mathbb{N}$ and $n \in \mathbb{N}$, what is the sample complexity of testing the property that consists of all distributions over $[n]$ that are each m -grained for some $m \in M$.*

Appendix: Reducing testing m -grained distributions (over $[n]$) to the case of $n = O(m)$

Recall that Corollary 13 asserts that *for every $n, m \in \mathbb{N}$, the set of m -grained distributions over $[n]$ has a tester of sample complexity $O(\epsilon^{-2} \cdot n/\log n)$* . As commented in the main text, we believe that using the techniques of [13] one can reduce the complexity to $O(\epsilon^{-2} \cdot n'/\log n')$, where $n' = \min(n, m)$. Here we show an alternative proof of this result. Specifically, we shall reduce ϵ -testing m -grained distributions over $[n]$ to $\Omega(\epsilon)$ -testing m -grained distributions over $[O(m)]$, and apply Corollary 13.

The reduction will consist of using a deterministic filter $f : [n] \rightarrow [k]$, where $k = O(m)$, which will be selected uniformly at random among all such filters. We stress that this is fundamentally different from the randomized filters F used in the main text. Specifically, when applying F several times to the same input, we obtained outcomes that are independently and identically distributed, whereas when we apply a function f (which is selected at random) several times to the same input we obtain the same output.

Note that applying any function $f : [n] \rightarrow [k]$ to any m -grained distribution yields an m -grained distribution. Our main result is that, for any distribution X over $[n]$ that is ϵ -far from being m -grained, for almost all functions $f : [n] \rightarrow [O(m)]$, the distribution $f(X)$ is $\Omega(\epsilon)$ -far from being m -grained.

¹⁰That is, p'_2 is ϵ -close to q'_2 , which is m -grained, since p_2 is ϵ -close to the uniform distribution over $[m/t]$. Likewise, p'_1 is ϵ -close to q'_1 , which is $(0.999/2t)$ -far from being grained, since p_1 is ϵ -close to the uniform distribution over $[2m/t]$.

Lemma 18 (relative preservation of distance from m -grained distributions): *For all sufficiently small $c > 0$ and all sufficiently large n and m , the following holds. If a distribution X over $[n]$ is ϵ -far from being m -grained, then, with probability at least $1 - 36c$ over the choice of a function $f : [n] \rightarrow [m/c]$, the distribution $f(X)$ is $0.05 \cdot \epsilon$ -far from being m -grained.*

Hence, we obtain a randomized reduction of the general problem of testing m -grained distributions (over $[n]$) to the special case of $n = O(m)$, where the reduction consists of selecting at random a function $f : [n] \rightarrow [m/c]$ and using it as a (deterministic) filter for reducing the general problem to its special case.

Proof: Let $k = m/c$ and let $p : [n] \rightarrow [0, 1]$ denote the probability function that describes X . Define $r : [n] \rightarrow [0, 1/m]$ such that $r(i) = p(i) - \lfloor m \cdot p(i) \rfloor / m$. Denoting by $\Delta_G(p)$ the statistical distance between p and the set of m -grained distributions (i.e., half the norm-1 distance), we have

$$2 \cdot \Delta_G(p) \geq \sum_{i \in [n]} \min(r(i), (1/m) - r(i)) \quad (4)$$

$$2 \cdot \Delta_G(p) \leq 2 \cdot \sum_{i \in [n]} \min(r(i), (1/m) - r(i)) \quad (5)$$

where Eq. (4) is due to the need to transform each $p(i)$ to a multiple of $1/m$ and Eq. (5) is justified by a two-step correction process in which we first round each $p(i)$ to the closest multiple of $1/m$, and then we correct the resulting function so that it sums up to 1 (while keeping its values as multiples of $1/m$).¹¹ Hence, using Eq. (5), the lemma's hypothesis implies that $\sum_{i \in [n]} \min(r(i), (1/m) - r(i)) > \epsilon$. We shall prove the lemma by lower-bounding (w.h.p.) the corresponding sum that refers to the distribution $f(X)$, when f is selected at random. Specifically, for $p'(j) = \sum_{i: f(i)=j} p(i)$, we shall lower-bound the probability that $\sum_{j \in [k]} \min(r'(j), (1/m) - r'(j)) > 0.1 \cdot \epsilon$, where $r'(j) = p'(j) - \lfloor m \cdot p'(j) \rfloor / m$, and then apply Eq. (4).

Before doing so, we introduce a few notations. Firstly, we let $s(i) = \min(r(i), (1/m) - r(i))$, and let $\delta = \sum_{i \in [n]} s(i)$, which is greater than ϵ by the hypothesis. Next, we let $H = \{i \in [n] : p(i) \geq 1/3m\}$ denote the set of ‘‘heavy’’ elements in X . We observe that $|H| \leq 3m$ and that for every $i \in \overline{H} \stackrel{\text{def}}{=} [n] \setminus H$ it holds that $s(i) = r(i) = p(i)$, since $p(i) < 1/2m$ holds for every $i \in \overline{H}$. We consider two cases, according to whether or not the sum $\sum_{i \in \overline{H}} p(i)$ is smaller than $0.5 \cdot \delta$.

Claim 18.1 (the first case): *Suppose that $\sum_{i \in \overline{H}} p(i) < 0.5 \cdot \delta$. Then, with probability at least $1 - 16c$ over the choice of f , it holds that $f(X)$ is 0.1ϵ -far from being m -grained.*

Proof: In this case $\sum_{i \in H} s(i) > 0.5 \cdot \delta$, and we shall focus on the contribution of $f(H)$ to the distance of $f(X)$ from being m -grained. We shall show that, for almost all functions f , much of this weight is mapped (by f) in a one-to-one manner, and that the elements in \overline{H} do not change by much the weight mapped by f to $f(H)$. Specifically, we consider a uniformly selected function $f : [n] \rightarrow [k]$, and the following two good events defined on this probability space.

1. The first (good) event is that the function f maps at least 0.2δ of the $s(i)$ -mass of the i 's in H to distinct images. Intuitively, this is very likely given that the total $s(i)$ -mass of i 's in H

¹¹Specifically, let $q : [n] \rightarrow [0, 2)$ be the function resulting from the first step (i.e., $q(i) = \lfloor m \cdot p(i) \rfloor / m$ if $r(i) \leq 1/2m$ and $q(i) = \lceil m \cdot p(i) \rceil / m$ otherwise). Then, $\delta \stackrel{\text{def}}{=} \sum_{i \in [n]} |q(i) - p(i)| = \sum_{i \in [n]} \min(r(i), (1/m) - r(i))$ and $|1 - \sum_{i \in [n]} q(i)| \leq \delta$.

is greater than 0.5δ and that $|H| \ll k$. Formally, denoting by H_f the (random variable that represents the) set of $i \in H$ that satisfy $f(i) \notin f(H \setminus \{i\})$ (i.e., for every $i \in H_f$ it holds that $f^{-1}(f(i)) \cap H = \{i\}$), we claim that $\Pr_f \left[\sum_{i \in H_f} s(i) > 0.2\delta \right] > 1 - c$.

To see this, we first note that, for every $i \in H$, conditioned on the values assigned to $H \setminus \{i\}$, the probability that $f(i) \notin f(H \setminus \{i\})$ is at least $\frac{k - (|H| - 1)}{k} > 1 - |H|/k \geq 0.9$, where the inequality is due to $|H| \leq 3m \leq 0.1k$. Hence, each $i \in H$ contributes $s(i) \leq 1/2m$ to the sum (of $s(i)$'s with $i \in H_f$) with probability at least 0.9, also when conditioned on all other values assigned by f . It follows that $\Pr_f \left[\sum_{i \in H_f} s(i) > 0.2\delta \right] > 1 - c$, where the (typical) case of $\delta = \omega(1/m)$ is straightforward.¹²

2. The second (good) event is that the function f does not map much $p(i)$ -mass of i 's in \overline{H} to the images occupied by H . Again, this is very likely given that $|H| \ll k$. Specifically, observe that $\mathbb{E}_f \left[\sum_{i \in \overline{H}: f(i) \in f(H)} p(i) \right] \leq \frac{|H|}{k} \cdot \sum_{i \in \overline{H}} p(i) < 3c \cdot \delta/2$, since $p(i) = s(i)$ for every $i \in \overline{H}$ (and $|H| \leq 3m$ and $k = m/c$). Letting $S_f = \sum_{i \in \overline{H}: f(i) \in f(H)} p(i)$, we get $\Pr_f[S_f < 0.1\delta] > 1 - \frac{3c\delta/2}{0.1\delta} = 1 - 15c$.

Assuming that the two good events occur (which happens with probability at least $1 - 6c$), it follows that at least 0.2δ of the $s(\cdot)$ -mass of H is mapped by f to distinct images and at most 0.1δ of the mass of \overline{H} is mapped to these images. Hence, $f(X)$ corresponds to a probability function p' such that $r'(i) = p'(i) - \lfloor m \cdot p'(i) \rfloor / m$ satisfies

$$\begin{aligned} \sum_{i \in H_f} \min(r'(i), (1/m) - r'(i)) &\geq \sum_{i \in H_f} s(i) - \sum_{i \in \overline{H}: f(i) \in f(H)} p(i) \\ &> 0.2\delta - 0.1\delta, \end{aligned}$$

where $H_f = \{i \in H : f^{-1}(f(i)) \cap H = \{i\}\}$ (as above). Hence, recalling that $\delta > \epsilon$, with probability at least $1 - 16c$ over the choice of f , it holds that $f(X)$ is 0.1ϵ -far from being m -grained. ■

Claim 18.2 (the second case): *Suppose that $\delta' \stackrel{\text{def}}{=} \sum_{i \in \overline{H}} p(i) \geq 0.5 \cdot \delta$. Then, with probability at least $1 - 36c$ over the choice of f , it holds that $f(X)$ is 0.05ϵ -far from being m -grained.*

Proof: In this case $\sum_{i \in \overline{H}} s(i) > 0.5 \cdot \delta$, and we shall focus on the contribution of $f(\overline{H})$ to the distance of $f(X)$ from being m -grained. We shall show that, for almost all functions f , much of this weight is mapped (by f) to $[k] \setminus H$ and that the mass of the elements of \overline{H} is distributed almost

¹²Specifically, letting $\zeta_i = \zeta_i(f)$ denote the contribution of $i \in H$ to $\sum_{i \in H_f} s(i)$, we have $\mathbb{E}[\zeta_i] \geq 0.9 \cdot s(i)$ and $\mathbb{V}[\zeta_i] \leq \mathbb{E}[\zeta_i^2] \leq s(i)/2m$. Hence, $\Pr \left[\sum_{i \in H} \zeta_i \leq 0.2\delta \right] < \frac{\delta/2m}{(0.25\delta)^2}$, since $\mathbb{E} \left[\sum_{i \in H} \zeta_i \right] > 0.45\delta$. This suffices for $\delta = \omega(1/m)$. Actually, the same argument holds if $\sum_{i \in H} s(i)^2 = o(\delta^2)$. In general (esp., if $\sum_{i \in H} s(i)^2 = \Omega(\delta^2)$), for a sufficiently small $c' > 0$, we define $H' \stackrel{\text{def}}{=} \{i \in H : s(i) \geq c' \cdot \delta\}$, and consider two cases.

- (a) If $\sum_{i \in H \setminus H'} s(i) > 0.3 \cdot \delta$, then we use $H \setminus H'$ instead of H , while noting that $\Pr \left[\sum_{i \in H \setminus H'} \zeta_i \leq 0.2\delta \right] < \frac{c' \cdot \delta^2}{(0.07\delta)^2} < c$, since $\mathbb{E} \left[\sum_{i \in H \setminus H'} \zeta_i \right] > 0.27\delta$ and $\mathbb{V}[\sum_{i \in H \setminus H'} \zeta_i] \leq c' \delta \cdot \delta$.
- (b) If $\sum_{i \in H'} s(i) > 0.2 \cdot \delta$, then we use H' instead of H , while noting that the probability that $|f(H')| < |H'|$ is at most $\binom{1/c'}{2} \cdot (1/k) < c$.

(We proceed with H replaced by either H' or $H \setminus H'$.)

uniformly. Specifically, we first show that more than half of the probability mass of \overline{H} is mapped disjointly of H . That is,

$$\Pr_{f:[n] \rightarrow [k]} \left[\sum_{i \in \overline{H}: f(i) \notin f(H)} p(i) > 0.5 \cdot \delta' \right] \geq 1 - 6c \quad (6)$$

where the probability is taken uniformly over all possible choices of f . The proof is similar to the analysis of the second event in the proof of Claim 18.2. Specifically, we consider random variables ζ_i 's such that $\zeta_i = p(i)$ if $f(i) \notin f(H)$ and $\zeta_i = 0$ otherwise, and observe that $\mathbb{E}[\zeta_i] \geq \frac{k-|H|}{k} \cdot p(i) \geq (1-3c) \cdot p(i)$ (since $|H| \leq 3m$ and $m = ck$). Thus, $\mathbb{E}[\sum_{i \in \overline{H}} \zeta_i] \geq (1-3c) \cdot \delta'$ and Eq. (6) follows by Markov Inequality while using $\sum_{i \in \overline{H}} \zeta_i \leq \sum_{i \in \overline{H}} p(i) = \delta'$. This holds also if we fix the values of f on H and condition on it, which is what we do from this point on. Hence, we fix an arbitrary sequence of value for $f(H)$, and consider the uniform distribution of f conditioned on this fixing as well as on the event in Eq. (6).

Actually, we decompose $f : [n] \rightarrow [k]$ into three parts, denoted f', f'' and f''' , that represents its restriction to the three-way partition of $[n]$ into (H, B, G) such that $B = \{i \in \overline{H} : f(i) \in f(H)\}$ (and $G = \{i \in \overline{H} : f(i) \notin f(H)\}$); indeed, $f' : H \rightarrow [k]$ is the restriction of f to H , whereas $f'' : B \rightarrow f(H)$ and $f''' : G \rightarrow [k] \setminus f(H)$ are its restrictions to the two parts of \overline{H} . We fix arbitrary $f' : H \rightarrow [k]$ and $f'' : B \rightarrow f(H)$, where $B = \{i \in \overline{H} : f''(i) \in f(H)\}$, such that $\sum_{i \in B} p(i) < 0.5\delta'$, while bearing in mind that such fixing (of f' and f'') arise from the choice of a random f with probability at least $1 - 6c$. Our aim will be to show that, with high probability over the choice of $f''' : G \rightarrow [k] \setminus f(H)$, it holds that

$$\sum_{i \in G: f'''(i) \in J(f''')} p(i) > 0.4\delta'. \quad (7)$$

where $J(f''') \stackrel{\text{def}}{=} \{j \in [k] : \sum_{i \in G: f'''(i)=j} p(i) \leq 0.8/m\}$. (Recall that for any $i \in G \subseteq \overline{H}$ it holds that $p(i) = r(i) = s(i) < 1/3m$.) This would imply that, with high probability, the distance of $f(X)$ from being m -grained is at least

$$\begin{aligned} \sum_{j \in J(f''')} \min \left(\sum_{i: f'''(i)=j} p(i), \frac{1}{m} - \frac{0.8}{m} \right) &\geq \sum_{j \in J(f''')} 0.25 \cdot \sum_{i: f'''(i)=j} p(i) \\ &> 0.25 \cdot 0.4\delta' \geq 0.05\delta, \end{aligned}$$

where the first inequality is due to the fact that $p'(j) \stackrel{\text{def}}{=} \Pr[f(X)=j] = \sum_{i \in G: f'''(i)=j} p(i) \leq 0.8/m$ for every $j \in J(f''')$ and so $\min(p'(j), 0.2/m) \geq p'(j)/4$. So all that remains is to show that Eq. (7) holds with high probability over the choice of f''' .

Letting $K' \stackrel{\text{def}}{=} [k] \setminus f(H)$, we start by observing that, for every $i \in G$, it holds that

$$\begin{aligned} \Pr_{f''': G \rightarrow K'} [f'''(i) \notin J(f''')] &\leq \Pr_{f''': G \rightarrow K'} \left[\sum_{\ell \in G \setminus \{i\}: f'''(\ell)=f'''(i)} p(\ell) > \frac{0.8}{m} - \frac{1}{3m} \right] \\ &= \Pr_{f''': G \rightarrow K'} \left[f'''(i) \in \left\{ j \in K' : \sum_{\ell \in G \setminus \{i\}: f'''(\ell)=j} p(\ell) > \frac{1.4}{3m} \right\} \right] \end{aligned}$$

$$\begin{aligned} &\leq \frac{3m}{1.4} \cdot \frac{1}{|K'|} \\ &\leq \frac{3ck}{1.4} \cdot \frac{1}{k - 3ck} \end{aligned}$$

where the equality can be seen by first fixing f''' -values for all elements in $G \setminus \{i\}$ and then selecting $f'''(i)$ uniformly in K' . Assuming that $c \leq 0.4/4.2$, we get $\Pr_{f''' : G \rightarrow K'}[f'''(i) \notin J(f''')] \leq 3c$, and it follows that

$$\begin{aligned} \mathbb{E}_{f''' : G \rightarrow K'} \left[\sum_{i \in G : f'''(i) \notin J(f''')} p(i) \right] &= \sum_{i \in G} \Pr_{f''' : G \rightarrow K'}[f'''(i) \notin J(f''')] \cdot p(i) \\ &= \sum_{i \in G} 3c \cdot p(i) \\ &\leq 3c \cdot \delta' \end{aligned}$$

since $\sum_{i \in G} p(i) \leq \sum_{i \in \overline{H}} p(i) = \delta'$. It follows that $\Pr_{f''' : G \rightarrow K'} \left[\sum_{i \in G : f'''(i) \notin J(f''')} p(i) \geq 0.1\delta' \right]$ is upper-bounded by $3c/0.1 = 30c$. Recalling that $\sum_{i \in B} p(i) < 0.5\delta'$, which implies $\sum_{i \in G} p(i) > 0.5\delta'$, this implies that Eq. (7) holds with probability at least $1 - 30c$ (over the choice of f'''). Lastly, recall that $\sum_{i \in B} p(i) < 0.5\delta'$, where $B = \{i \in \overline{H} : f''(i) \in f'(H)\}$, holds with probability at least $1 - 6c$ (over the choice of f' and f''). The claim follows, since (as argued above) Eq. (7) implies that $f(X)$ is 0.05δ -far from being m -grained (and $\delta > \epsilon$). ■

Combining Claims 18.1 and 18.2, the lemma follows. ■

References

- [1] T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *42nd FOCS*, pages 442–451, 2001.
- [2] T. Batu, L. Fortnow, R. Rubinfeld, W.D. Smith, and P. White. Testing that Distributions are Close. In *41st FOCS*, pages 259–269, 2000.
- [3] C.L. Canonne. A Survey on Distribution Testing: Your Data is Big. But is it Blue? *ECCC*, TR015-063, 2015.
- [4] S. Chan, I. Diakonikolas, P. Valiant, and G. Valiant. Optimal Algorithms for Testing Closeness of Discrete Distributions. In *25th ACM-SIAM Symposium on Discrete Algorithms*, pages 1193–1203, 2014.
- [5] I. Diakonikolas and D. Kane. A New Approach for Testing Properties of Discrete Distributions. [arXiv:1601.05557 \[cs.DS\]](https://arxiv.org/abs/1601.05557), 2016.
- [6] I. Diakonikolas, D. Kane, V. Nikishkin. Testing Identity of Structured Distributions. In *26th ACM-SIAM Symposium on Discrete Algorithms*, pages 1841–1854, 2015.
- [7] I. Diakonikolas, T. Gouleakis, J. Peebles, and E. Price. In preparation.
- [8] O. Goldreich. *Introduction to Property Testing: Lecture Notes*. Superseded by [9]. Drafts are available from <http://www.wisdom.weizmann.ac.il/~oded/pt-ln.html>

- [9] O. Goldreich. *Introduction to Property Testing*. Cambridge University Press, 2017.
- [10] O. Goldreich and D. Ron. On Testing Expansion in Bounded-Degree Graphs. *ECCC*, TR00-020, March 2000.
- [11] O. Goldreich and D. Ron. On the relation between the relative earth mover distance and the variation distance (an exposition). Available from http://www.wisdom.weizmann.ac.il/~oded/p_remd.html
- [12] L. Paninski. A coincidence-based test for uniformity given very sparsely-sampled discrete data. *IEEE Transactions on Information Theory*, Vol. 54, pages 4750–4755, 2008.
- [13] G. Valiant and P. Valiant. Estimating the unseen: an $n/\log(n)$ -sample estimator for entropy and support size, shown optimal via new CLTs. In *43rd ACM Symposium on the Theory of Computing*, pages 685-694, 2011.
- [14] G. Valiant and P. Valiant. Instance-by-instance optimal identity testing. *ECCC*, TR13-111, 2013.
- [15] G. Valiant and P. Valiant. Instance Optimal Learning. CoRR abs/1504.05321, 2015.