

---

# Near-Optimal Density Estimation in Near-Linear Time Using Variable-Width Histograms

---

**Siu-On Chan**  
Microsoft Research  
sochan@gmail.com

**Ilias Diakonikolas**  
University of Edinburgh  
ilias.d@ed.ac.uk

**Rocco A. Servedio**  
Columbia University  
rocco@cs.columbia.edu

**Xiaorui Sun**  
Columbia University  
xiaoruisun@cs.columbia.edu

## Abstract

Let  $p$  be an unknown and arbitrary probability distribution over  $[0, 1]$ . We consider the problem of *density estimation*, in which a learning algorithm is given i.i.d. draws from  $p$  and must (with high probability) output a hypothesis distribution that is close to  $p$ . The main contribution of this paper is a highly efficient density estimation algorithm for learning using a variable-width histogram, i.e., a hypothesis distribution with a piecewise constant probability density function.

In more detail, for any  $k$  and  $\varepsilon$ , we give an algorithm that makes  $\tilde{O}(k/\varepsilon^2)$  draws from  $p$ , runs in  $\tilde{O}(k/\varepsilon^2)$  time, and outputs a hypothesis distribution  $h$  that is piecewise constant with  $O(k \log^2(1/\varepsilon))$  pieces. With high probability the hypothesis  $h$  satisfies  $d_{\text{TV}}(p, h) \leq C \cdot \text{opt}_k(p) + \varepsilon$ , where  $d_{\text{TV}}$  denotes the total variation distance (statistical distance),  $C$  is a universal constant, and  $\text{opt}_k(p)$  is the smallest total variation distance between  $p$  and any  $k$ -piecewise constant distribution. The sample size and running time of our algorithm are optimal up to logarithmic factors. The “approximation factor”  $C$  in our result is inherent in the problem, as we prove that no algorithm with sample size bounded in terms of  $k$  and  $\varepsilon$  can achieve  $C < 2$  regardless of what kind of hypothesis distribution it uses.

## 1 Introduction

Consider the following fundamental statistical task: *Given independent draws from an unknown probability distribution, what is the minimum sample size needed to obtain an accurate estimate of the distribution?* This is the question of *density estimation*, a classical problem in statistics with a rich history and an extensive literature (see e.g., [BBBB72, DG85, Sil86, Sco92, DL01]). While this broad question has mostly been studied from an information-theoretic perspective, it is an inherently algorithmic question as well, since the ultimate goal is to describe and understand algorithms that are both computationally and information-theoretically efficient. The need for computationally efficient learning algorithms is only becoming more acute with the recent flood of data across the sciences; the “gold standard” in this “big data” context is an algorithm with information-theoretically (near-) optimal sample size and running time (near-) linear in its sample size.

In this paper we consider learning scenarios in which an algorithm is given an input data set which is a sample of i.i.d. draws from an unknown probability distribution. It is natural to expect (and can be easily formalized) that, if the underlying distribution of the data is inherently “complex”, it may be hard to even approximately reconstruct the distribution. But what if the underlying distribution is “simple” or “succinct” – can we then reconstruct the distribution to high accuracy in a computationally and sample-efficient way? In this paper we answer this question in the affirmative for the

problem of learning “noisy” *histograms*, arguably one of the most basic density estimation problems in the literature.

To motivate our results, we begin by briefly recalling the role of histograms in density estimation. Histograms constitute “the oldest and most widely used method for density estimation” [Si186], first introduced by Karl Pearson in [Pea95]. Given a sample from a probability density function (pdf)  $p$ , the method partitions the domain into a number of intervals (bins)  $B_1, \dots, B_k$ , and outputs the “empirical” pdf which is constant within each bin. A  $k$ -*histogram* is a piecewise constant distribution over bins  $B_1, \dots, B_k$ , where the probability mass of each interval  $B_j, j \in [k]$ , equals the fraction of observations in the interval. Thus, the goal of the “histogram method” is to approximate an unknown pdf  $p$  by an appropriate  $k$ -histogram. It should be emphasized that the number  $k$  of bins to be used and the “width” and location of each bin are unspecified; they are parameters of the estimation problem and are typically selected in an *ad hoc* manner.

We study the following distribution learning question:

*Suppose that there exists a  $k$ -histogram that provides an accurate approximation to the unknown target distribution. Can we efficiently find such an approximation?*

In this paper, we provide a fairly complete affirmative answer to this basic question. Given a bound  $k$  on the number of intervals, we give an algorithm that uses a near-optimal sample size, runs in *near-linear time* (in its sample size), and approximates the target distribution nearly as accurately as the best  $k$ -histogram.

To formally state our main result, we will need a few definitions. We work in a standard model of learning an unknown probability distribution from samples, essentially that of [KMR<sup>+</sup>94], which is a natural analogue of Valiant’s well-known PAC model for learning Boolean functions [Val84] to the unsupervised setting of learning an unknown probability distribution.<sup>1</sup> A distribution learning problem is defined by a class  $\mathcal{C}$  of distributions over a domain  $\Omega$ . The algorithm has access to independent draws from an unknown pdf  $p$ , and its goal is to output a hypothesis distribution  $h$  that is “close” to the target distribution  $p$ . We measure the closeness between distributions using the *statistical distance* or total variation distance. In the “noiseless” setting, we are promised that  $p \in \mathcal{C}$  and the goal is to construct a hypothesis  $h$  such that (with high probability) the total variation distance  $d_{\text{TV}}(h, p)$  between  $h$  and  $p$  is at most  $\varepsilon$ , where  $\varepsilon > 0$  is the accuracy parameter.

The more challenging “noisy” or *agnostic* model captures the situation of having arbitrary (or even adversarial) noise in the data. In this setting, we do not make any assumptions about the target density  $p$  and the goal is to find a hypothesis  $h$  that is almost as accurate as the “best” approximation of  $p$  by any distribution in  $\mathcal{C}$ . Formally, given sample access to a (potentially arbitrary) target distribution  $p$  and  $\varepsilon > 0$ , the goal of an *agnostic learning algorithm for  $\mathcal{C}$*  is to compute a hypothesis distribution  $h$  such that  $d_{\text{TV}}(h, p) \leq \alpha \cdot \text{opt}_{\mathcal{C}}(p) + \varepsilon$ , where  $\text{opt}_{\mathcal{C}}(p) := \inf_{q \in \mathcal{C}} d_{\text{TV}}(q, p)$  – i.e.,  $\text{opt}_{\mathcal{C}}(p)$  is the statistical distance between  $p$  and the closest distribution to it in  $\mathcal{C}$  – and  $\alpha \geq 1$  is a constant (that may depend on the class  $\mathcal{C}$ ). We will call such a learning algorithm an  $\alpha$ -*agnostic learning algorithm for  $\mathcal{C}$* ; when  $\alpha > 1$  we sometimes refer to this as a *semi-agnostic learning algorithm*.

A distribution  $f$  over a finite interval  $I \subseteq \mathbb{R}$  is called  $k$ -*flat* if there exists a partition of  $I$  into  $k$  intervals  $I_1, \dots, I_k$  such that the pdf  $f$  is constant within each such interval. We henceforth (without loss of generality for densities with bounded support) restrict ourselves to the case  $I = [0, 1]$ . Let  $\mathcal{C}_k$  be the class of all  $k$ -*flat* distributions over  $[0, 1]$ . For a (potentially arbitrary) distribution  $p$  over  $[0, 1]$  we will denote by  $\text{opt}_k(p) := \inf_{f \in \mathcal{C}_k} d_{\text{TV}}(f, p)$ .

In this terminology, our learning problem is exactly the problem of agnostically learning the class of  $k$ -flat distributions. Our main positive result is a near-optimal algorithm for this problem, i.e., a semi-agnostic learning algorithm that has near-optimal sample size and near-linear running time. More precisely, we prove the following:

**Theorem 1 (Main).** *There is an algorithm  $A$  with the following property: Given  $k \geq 1$ ,  $\varepsilon > 0$ , and sample access to a target distribution  $p$ , algorithm  $A$  uses  $\tilde{O}(k/\varepsilon^2)$  independent draws from  $p$ , runs in time  $\tilde{O}(k/\varepsilon^2)$ , and outputs a  $O(k \log^2(1/\varepsilon))$ -flat hypothesis distribution  $h$  that satisfies  $d_{\text{TV}}(h, p) \leq O(\text{opt}_k(p)) + \varepsilon$  with probability at least  $9/10$ .*

<sup>1</sup>We remark that our model is essentially equivalent to the “minimax rate of convergence under the  $L_1$  distance” in statistics [DL01], and our results carry over to this setting as well.

Using standard techniques, the confidence probability can be boosted to  $1 - \delta$ , for any  $\delta > 0$ , with a (necessary) overhead of  $O(\log(1/\delta))$  in the sample size and the running time.

We emphasize that the difficulty of our result lies in the fact that the “optimal” piecewise constant decomposition of the domain is both *unknown* and *approximate* (in the sense that  $\text{opt}_k(p) > 0$ ); and that our algorithm is both sample-optimal and runs in (near-) *linear time*. Even in the (significantly easier) case that the target  $p \in \mathcal{C}_k$  (i.e.,  $\text{opt}_k(p) = 0$ ), and the optimal partition is explicitly given to the algorithm, it is known that a sample of size  $\Omega(k/\varepsilon^2)$  is information-theoretically necessary. (This lower bound can, e.g., be deduced from the standard fact that learning an unknown discrete distribution over a  $k$ -element set to statistical distance  $\varepsilon$  requires an  $\Omega(k/\varepsilon^2)$  size sample.) Hence, our algorithm has provably optimal sample complexity (up to a logarithmic factor), runs in essentially sample linear time, and is  $\alpha$ -agnostic for a universal constant  $\alpha > 1$ .

It should be noted that the sample size required for our problem is well-understood; it follows from the VC theorem (Theorem 3) that  $O(k/\varepsilon^2)$  draws from  $p$  are information-theoretically sufficient. However, the theorem is non-constructive, and the “obvious” algorithm following from it has running time exponential in  $k$  and  $1/\varepsilon$ . In recent work, Chan *et al* [CDSS14] presented an approach employing an intricate combination of dynamic programming and linear programming which yields a  $\text{poly}(k/\varepsilon)$  time algorithm for the above problem. However, the running time of the [CDSS14] algorithm is  $\Omega(k^3)$  even for constant values of  $\varepsilon$ , making it impractical for applications. As discussed below our algorithmic approach is significantly different from that of [CDSS14], using neither dynamic nor linear programming.

**Applications.** Nonparametric density estimation for shape restricted classes has been a subject of study in statistics since the 1950’s (see [BBBB72] for an early book on the topic and [Gre56, Bru58, Rao69, Weg70, HP76, Gro85, Bir87] for some of the early literature), and has applications to a range of areas including reliability theory (see [Reb05] and references therein). By using the structural approximation results of Chan *et al* [CDSS13], as an immediate corollary of Theorem 1 we obtain sample optimal and *near-linear time* estimators for various well-studied classes of shape restricted densities including monotone, unimodal, and multimodal densities (with unknown mode locations), monotone hazard rate (MHR) distributions, and others (because of space constraints we do not enumerate the exact descriptions of these classes or statements of these results here, but instead refer the interested reader to [CDSS13]). Birgé [Bir87] obtained a sample optimal and linear time estimator for monotone densities, but prior to our work, no linear time and sample optimal estimator was known for any of the other classes.

Our algorithm from Theorem 1 is  $\alpha$ -agnostic for a constant  $\alpha > 1$ . It is natural to ask whether a significantly stronger accuracy guarantee is efficiently achievable; in particular, is there an agnostic algorithm with similar running time and sample complexity and  $\alpha = 1$ ? Perhaps surprisingly, we provide a negative answer to this question. Even in the simplest nontrivial case that  $k = 2$ , and the target distribution is defined over a discrete domain  $[N] = \{1, \dots, N\}$ , any  $\alpha$ -agnostic algorithm with  $\alpha < 2$  requires large sample size:

**Theorem 2** (Lower bound, Informal statement). *Any 1.99-agnostic learning algorithm for 2-flat distributions over  $[N]$  requires a sample of size  $\Omega(\sqrt{N})$ .*

See Theorem 7 in Section 4 for a precise statement. Note that there is an exact correspondence between distributions over the discrete domain  $[N]$  and pdf’s over  $[0, 1]$  which are piecewise constant on each interval of the form  $[k/N, (k+1)/N)$  for  $k \in \{0, 1, \dots, N-1\}$ . Thus, Theorem 2 implies that *no finite sample* algorithm can 1.99-agnostically learn even 2-flat distributions over  $[0, 1]$ . (See Corollary 4.1 in Section 4 for a detailed statement.)

**Related work.** A number of techniques for density estimation have been developed in the mathematical statistics literature, including kernels and variants thereof, nearest neighbor estimators, orthogonal series estimators, maximum likelihood estimators (MLE), and others (see Chapter 2 of [Sil86] for a survey of existing methods). The main focus of these methods has been on the statistical rate of convergence, as opposed to the running time of the corresponding estimators. We remark that the MLE does not exist for very simple classes of distributions (e.g., unimodal distributions with an unknown mode, see e.g. [Bir97]). We note that the notion of agnostic learning is related to the literature on model selection and oracle inequalities [MP007], however this work is of a different flavor and is not technically related to our results.

Histograms have also been studied extensively in various areas of computer science, including databases and streaming [JKM<sup>+</sup>98, GKS06, CMN98, GGI<sup>+</sup>02] under various assumptions about the input data and the precise objective. Recently, Indyk *et al* [ILR12] studied the problem of learning a  $k$ -flat distribution over  $[N]$  under the  $L_2$  norm and gave an efficient algorithm with sample complexity  $O(k^2 \log(N)/\varepsilon^4)$ . Since the  $L_1$  distance is a stronger metric, Theorem 1 implies an improved sample and time bound of  $\tilde{O}(k/\varepsilon^2)$  for their setting.

## 2 Preliminaries

Throughout the paper we assume that the underlying distributions have Lebesgue measurable densities. For a pdf  $p : [0, 1] \rightarrow \mathbb{R}_+$  and a Lebesgue measurable subset  $A \subseteq [0, 1]$ , i.e.,  $A \in \mathcal{L}([0, 1])$ , we use  $p(A)$  to denote  $\int_{z \in A} p(z)$ . The *statistical distance* or *total variation distance* between two densities  $p, q : [0, 1] \rightarrow \mathbb{R}_+$  is  $d_{\text{TV}}(p, q) := \sup_{A \in \mathcal{L}([0, 1])} |p(A) - q(A)|$ . The statistical distance satisfies the identity  $d_{\text{TV}}(p, q) = \frac{1}{2} \|p - q\|_1$  where  $\|p - q\|_1$ , the  $L_1$  distance between  $p$  and  $q$ , is  $\int_{[0, 1]} |p(x) - q(x)| dx$ ; for convenience in the rest of the paper we work with  $L_1$  distance. We refer to a nonnegative function  $p$  over an interval (which need not necessarily integrate to one over the interval) as a “sub-distribution.” Given a value  $\kappa > 0$ , we say that a (sub-)distribution  $p$  over  $[0, 1]$  is  $\kappa$ -well-behaved if  $\sup_{x \in [0, 1]} \Pr_{x \sim p}[x] \leq \kappa$ , i.e., no individual real value is assigned more than  $\kappa$  probability under  $p$ . Any probability distribution with no atoms is  $\kappa$ -well-behaved for all  $\kappa > 0$ . Our results apply for general distributions over  $[0, 1]$  which may have an atomic part as well as a non-atomic part. Given  $m$  independent draws  $s_1, \dots, s_m$  from a distribution  $p$  over  $[0, 1]$ , the *empirical distribution*  $\hat{p}_m$  over  $[0, 1]$  is the discrete distribution supported on  $\{s_1, \dots, s_m\}$  defined as follows: for all  $z \in [0, 1]$ ,  $\Pr_{x \sim \hat{p}_m}[x = z] = |\{j \in [m] \mid s_j = z\}|/m$ .

**The VC inequality.** Let  $p : [0, 1] \rightarrow \mathbb{R}$  be a Lebesgue measurable function. Given a family of subsets  $\mathcal{A} \subseteq \mathcal{L}([0, 1])$  over  $[0, 1]$ , define  $\|p\|_{\mathcal{A}} = \sup_{A \in \mathcal{A}} |p(A)|$ . The *VC dimension* of  $\mathcal{A}$  is the maximum size of a subset  $X \subseteq [0, 1]$  that is shattered by  $\mathcal{A}$  (a set  $X$  is shattered by  $\mathcal{A}$  if for every  $Y \subseteq X$ , some  $A \in \mathcal{A}$  satisfies  $A \cap X = Y$ ). If there is a shattered subset of size  $s$  for all  $s \in \mathbb{Z}_+$ , then we say that the VC dimension of  $\mathcal{A}$  is  $\infty$ . The well-known *Vapnik-Chervonenkis (VC) inequality* states the following:

**Theorem 3** (VC inequality, [DL01, p.31]). *Let  $p : I \rightarrow \mathbb{R}_+$  be a probability density function over  $I \subseteq \mathbb{R}$  and  $\hat{p}_m$  be the empirical distribution obtained after drawing  $m$  points from  $p$ . Let  $\mathcal{A} \subseteq 2^I$  be a family of subsets with VC dimension  $d$ . Then  $\mathbf{E}[\|p - \hat{p}_m\|_{\mathcal{A}}] \leq O(\sqrt{d/m})$ .*

**Partitioning into intervals of approximately equal mass.** As a basic primitive, given access to a sample drawn from a  $\kappa$ -well-behaved target distribution  $p$  over  $[0, 1]$ , we will need to partition  $[0, 1]$  into  $\Theta(1/\kappa)$  intervals each of which has probability  $\Theta(\kappa)$  under  $p$ . There is a simple algorithm, based on order statistics, which does this and has the following performance guarantee (see Appendix A.2 of [CDSS14]):

**Lemma 2.1.** *Given  $\kappa \in (0, 1)$  and access to points drawn from a  $\kappa/64$ -well-behaved distribution  $p$  over  $[0, 1]$ , the procedure `Approximately-Equal-Partition` draws  $O((1/\kappa) \log(1/\kappa))$  points from  $p$ , runs in time  $\tilde{O}(1/\kappa)$ , and with probability at least 99/100 outputs a partition of  $[0, 1]$  into  $\ell = \Theta(1/\kappa)$  intervals such that  $p(I_j) \in [\kappa/2, 3\kappa]$  for all  $1 \leq j \leq \ell$ .*

## 3 The algorithm and its analysis

In this section we prove our main algorithmic result, Theorem 1. Our approach has the following high-level structure: In Section 3.1 we give an algorithm for agnostically learning a target distribution  $p$  that is “nice” in two senses: (i)  $p$  is well-behaved (i.e., it does not have any heavy atomic elements), and (ii)  $\text{opt}_k(p)$  is bounded from above by the error parameter  $\varepsilon$ . In Section 3.2 we give a general efficient reduction showing how the second assumption can be removed, and in Section 3.3 we briefly explain how the first assumption can be removed, thus yielding Theorem 1.

### 3.1 The main algorithm

In this section we give our main algorithmic result, which handles well-behaved distributions  $p$  for which  $\text{opt}_k(p)$  is not too large:

**Theorem 4.** *There is an algorithm `Learn-WB-small-opt-k-histogram` that given as input  $\tilde{O}(k/\varepsilon^2)$  i.i.d. draws from a target distribution  $p$  and a parameter  $\varepsilon > 0$ , runs in time  $\tilde{O}(k/\varepsilon^2)$ , and has the following performance guarantee: If (i)  $p$  is  $\frac{\varepsilon/\log(1/\varepsilon)}{384k}$ -well-behaved, and (ii)  $\text{opt}_k(p) \leq \varepsilon$ , then with probability at least  $19/20$ , it outputs an  $O(k \cdot \log^2(1/\varepsilon))$ -flat distribution  $h$  such that  $d_{\text{TV}}(p, h) \leq 2 \cdot \text{opt}_k(p) + 3\varepsilon$ .*

We require some notation and terminology. Let  $r$  be a distribution over  $[0, 1)$ , and let  $\mathcal{P}$  be a set of disjoint intervals that are contained in  $[0, 1)$ . We say that the  $\mathcal{P}$ -flattening of  $r$ , denoted  $(r)^\mathcal{P}$ , is the sub-distribution defined as

$$r(v) = \begin{cases} r(I)/|I| & \text{if } v \in I, I \in \mathcal{P} \\ 0 & \text{if } v \text{ does not belong to any } I \in \mathcal{P} \end{cases}$$

Observe that if  $\mathcal{P}$  is a partition of  $[0, 1)$ , then (since  $r$  is a distribution)  $(r)^\mathcal{P}$  is a distribution.

We say that two intervals  $I, I'$  are *consecutive* if  $I = [a, b)$  and  $I' = [b, c)$ . Given two consecutive intervals  $I, I'$  contained in  $[0, 1)$  and a sub-distribution  $r$ , we use  $\alpha_r(I, I')$  to denote the  $L_1$  distance between  $(r)^{\{I, I'\}}$  and  $(r)^{\{I \cup I'\}}$ , i.e.,  $\alpha_r(I, I') = \int_{I \cup I'} |(r)^{\{I, I'\}}(x) - (r)^{\{I \cup I'\}}(x)| dx$ . Note here that  $\{I \cup I'\}$  is a set that contains one element, the interval  $[a, c)$ .

#### 3.1.1 Intuition for the algorithm

We begin with a high-level intuitive explanation of the `Learn-WB-small-opt-k-histogram` algorithm. It starts in Step 1 by constructing a partition of  $[0, 1)$  into  $z = \Theta(k/\varepsilon')$  intervals  $I_1, \dots, I_z$  (where  $\varepsilon' = \tilde{\Theta}(\varepsilon)$ ) such that  $p$  has weight  $\Theta(\varepsilon'/k)$  on each subinterval. In Step 2 the algorithm draws a sample of  $\tilde{O}(k/\varepsilon^2)$  points from  $p$  and uses them to define an empirical distribution  $\hat{p}_m$ . This is the only step in which points are drawn from  $p$ . For the rest of this intuitive explanation we pretend that the weight  $\hat{p}(I)$  that the empirical distribution  $\hat{p}_m$  assigns to each interval  $I$  is actually the same as the true weight  $p(I)$  (Lemma 3.1 below shows that this is not too far from the truth).

Before continuing with our explanation of the algorithm, let us digress briefly by imagining for a moment that the target distribution  $p$  actually is a  $k$ -flat distribution (i.e., that  $\text{opt}_k(p) = 0$ ). In this case there are at most  $k$  “breakpoints”, and hence at most  $k$  intervals  $I_j$  for which  $\alpha_{\hat{p}_m}(I_j, I_{j+1}) > 0$ , so computing the  $\alpha_{\hat{p}_m}(I_j, I_{j+1})$  values would be an easy way to identify the true breakpoints (and given these it is not difficult to construct a high-accuracy hypothesis).

In reality, we may of course have  $\text{opt}_k(p) > 0$ ; this means that if we try to use the  $\alpha_{\hat{p}_m}(I_j, I_{j+1})$  criterion to identify “breakpoints” of the optimal  $k$ -flat distribution that is closest to  $p$  (call this  $k$ -flat distribution  $q$ ), we may sometimes be “fooled” into thinking that  $q$  has a breakpoint in an interval  $I_j$  where it does not (but rather the value  $\alpha_{\hat{p}_m}(I_j, I_{j+1})$  is large because of the difference between  $q$  and  $p$ ). However, recall that by assumption we have  $\text{opt}_k(p) \leq \varepsilon$ ; this bound can be used to show that there cannot be too many intervals  $I_j$  for which a large value of  $\alpha_{\hat{p}_m}(I_j, I_{j+1})$  suggests a “spurious breakpoint” (see the proof of Lemma 3.3). This is helpful, but in and of itself not enough; since our partition  $I_1, \dots, I_z$  divides  $[0, 1)$  into  $k/\varepsilon'$  intervals, a naive approach based on this would result in a  $(k/\varepsilon')$ -flat hypothesis distribution, which in turn would necessitate a sample complexity of  $\tilde{O}(k/\varepsilon'^3)$ , which is unacceptably high. Instead, our algorithm performs a careful process of iteratively merging consecutive intervals for which the  $\alpha_{\hat{p}_m}(I_j, I_{j+1})$  criterion indicates that a merge will not adversely affect the final accuracy by too much. As a result of this process we end up with  $k \cdot \text{polylog}(1/\varepsilon)$  intervals for the final hypothesis, which enables us to output a  $(k \cdot \text{polylog}(1/\varepsilon'))$ -flat final hypothesis using  $\tilde{O}(k/\varepsilon'^2)$  draws from  $p$ .

In more detail, this iterative merging is carried out by the main loop of the algorithm in Step 4. Going into the  $t$ -th iteration of the loop, the algorithm has a partition  $\mathcal{P}_{t-1}$  of  $[0, 1)$  into disjoint sub-intervals, and a set  $\mathcal{F}_{t-1} \subseteq \mathcal{P}_{t-1}$  (i.e., every interval belonging to  $\mathcal{F}_{t-1}$  also belongs to  $\mathcal{P}_{t-1}$ ). Initially  $\mathcal{P}_0$  contains all the intervals  $I_1, \dots, I_z$  and  $\mathcal{F}_0$  is empty. Intuitively, the intervals in  $\mathcal{P}_{t-1} \setminus$

$\mathcal{F}_{t-1}$  are still being “processed”; such an interval may possibly be merged with a consecutive interval from  $\mathcal{P}_{t-1} \setminus \mathcal{F}_{t-1}$  if doing so would only incur a small “cost” (see condition (iii) of Step 4(b) of the algorithm). The intervals in  $\mathcal{F}_{t-1}$  have been “frozen” and will not be altered or used subsequently in the algorithm.

### 3.1.2 The algorithm

**Algorithm** Learn-WB-small-opt- $k$ -histogram:

**Input:** parameters  $k \geq 1, \varepsilon > 0$ ; access to i.i.d. draws from target distribution  $p$  over  $[0, 1)$

**Output:** If (i)  $p$  is  $\frac{\varepsilon/\log(1/\varepsilon)}{384k}$ -well-behaved and (ii)  $\text{opt}_k(p) \leq \varepsilon$ , then with probability at least 99/100 the output is a distribution  $q$  such that  $d_{\text{TV}}(p, q) \leq 2\text{opt}_k(p) + 3\varepsilon$ .

1. Let  $\varepsilon' = \varepsilon/\log(1/\varepsilon)$ . Run Algorithm Approximately-Equal-Partition on input parameter  $\frac{\varepsilon'}{6k}$  to partition  $[0, 1)$  into  $z = \Theta(k/\varepsilon')$  intervals  $I_1 = [i_0, i_1), \dots, I_z = [i_{z-1}, i_z)$ , where  $i_0 = 0$  and  $i_z = 1$ , such that with probability at least 99/100, for each  $j \in \{1, \dots, z\}$  we have  $p([i_{j-1}, i_j)) \in [\varepsilon'/12k, \varepsilon'/2k]$  (assuming  $p$  is  $\varepsilon'/(384k)$ -well-behaved).
2. Draw  $m = \tilde{O}(k/\varepsilon'^2)$  points from  $p$  and let  $\hat{p}_m$  be the resulting empirical distribution.
3. Set  $\mathcal{P}_0 = \{I_1, I_2, \dots, I_z\}$ , and  $\mathcal{F}_0 = \emptyset$ .
4. Let  $s = \log_2 \frac{1}{\varepsilon'}$ . Repeat for  $t = 1, \dots$  until  $t = s$ :
  - (a) Initialize  $\mathcal{P}_t$  to  $\emptyset$  and  $\mathcal{F}_t$  to  $\mathcal{F}_{t-1}$ .
  - (b) Without loss of generality, assume  $\mathcal{P}_{t-1} = \{I_{t-1,1}, \dots, I_{t-1,z_{t-1}}\}$  where interval  $I_{t-1,i}$  is to the left of  $I_{t-1,i+1}$  for all  $i$ . Scan left to right across the intervals in  $\mathcal{P}_{t-1}$  (i.e., iterate over  $i = 1, \dots, z_{t-1} - 1$ ). If intervals  $I_{t-1,i}, I_{t-1,i+1}$  are (i) both not in  $\mathcal{F}_{t-1}$ , and (ii)  $\alpha_{\hat{p}_m}(I_{t-1,i}, I_{t-1,i+1}) > \varepsilon'/(2k)$ , then add both  $I_{t-1,i}$  and  $I_{t-1,i+1}$  into  $\mathcal{F}_t$ .
  - (c) Initialize  $i$  to 1, and repeatedly execute one of the following four (mutually exclusive and exhaustive) cases until  $i > z_{t-1}$ :
    - [Case 1]  $i \leq z_{t-1} - 1$  and  $I_{t-1,i} = [a, b), I_{t-1,i+1} = [b, c)$  are consecutive intervals both not in  $\mathcal{F}_t$ . Add the merged interval  $I_{t-1,i} \cup I_{t-1,i+1} = [a, c)$  into  $\mathcal{P}_t$ . Set  $i \leftarrow i + 2$ .
    - [Case 2]  $i \leq z_{t-1} - 1$  and  $I_{t-1,i} \in \mathcal{F}_t$ . Set  $i \leftarrow i + 1$ .
    - [Case 3]  $i \leq z_{t-1} - 1, I_{t-1,i} \notin \mathcal{F}_t$  and  $I_{t-1,i+1} \in \mathcal{F}_t$ . Add  $I_{t-1,i}$  into  $\mathcal{F}_t$  and set  $i \leftarrow i + 2$ .
    - [Case 4]  $i = z_{t-1}$ . Add  $I_{t-1,z_{t-1}}$  into  $\mathcal{F}_t$  if  $I_{t-1,z_{t-1}}$  is not in  $\mathcal{F}_t$  and set  $i \leftarrow i + 1$ .
  - (d) Set  $\mathcal{P}_t \leftarrow \mathcal{P}_t \cup \mathcal{F}_t$ .
5. Output the  $|\mathcal{P}_s|$ -flat hypothesis distribution  $(\hat{p}_m)^{\mathcal{P}_s}$ .

### 3.1.3 Analysis of the algorithm and proof of Theorem 4

It is straightforward to verify the claimed running time given Lemma 2.1, which bounds the running time of Approximately-Equal-Partition. Indeed, we note that Step 2, which simply draws  $\tilde{O}(k/\varepsilon'^2)$  points and constructs the resulting empirical distribution, dominates the overall running time. In the rest of this subsection we prove correctness.

We first observe that with high probability the empirical distribution  $\hat{p}_m$  defined in Step 2 gives a high-accuracy estimate of the true probability of any union of consecutive intervals from  $I_1, \dots, I_z$ . The following lemma from [CDSS14] follows from the standard multiplicative Chernoff bound:

**Lemma 3.1** (Lemma 12, [CDSS14]). *With probability 99/100 over the sample drawn in Step 2, for every  $0 \leq a < b \leq z$  we have that  $|\hat{p}_m([i_a, i_b)) - p([i_a, i_b))| \leq \sqrt{\varepsilon'(b-a)} \cdot \varepsilon'/(10k)$ .*

We henceforth assume that this 99/100-likely event indeed takes place, so the above inequality holds for all  $0 \leq a < b \leq z$ . We use this to show that the  $\alpha_{\hat{p}_m}(I_{t-1,i}, I_{t-1,i+1})$  value that the algorithm

uses in Step 4(b) is a good proxy for the actual value  $\alpha_p(I_{t-1,i}, I_{t-1,i+1})$  (which of course is not accessible to the algorithm):

**Lemma 3.2.** *Fix  $1 \leq t \leq s$ . Then we have  $|\alpha_{\hat{p}_m}(I_{t-1,i}, I_{t-1,i+1}) - \alpha_p(I_{t-1,i}, I_{t-1,i+1})| \leq 2\varepsilon'/(5k)$ .*

Due to space constraints the proofs of all lemmas in this section are deferred to Appendix A.

For the rest of the analysis, let  $q$  denote a fixed  $k$ -flat distribution that is closest to  $p$ , so  $\|p - q\|_1 = \text{opt}_k(p)$ . (We note that while  $\text{opt}_k(p)$  is defined as  $\inf_{q \in \mathcal{C}} \|p - q\|_1$ , standard closure arguments can be used to show that the infimum is actually achieved by some  $k$ -flat distribution  $q$ .) Let  $\mathcal{Q}$  be the partition of  $[0, 1)$  corresponding to the intervals on which  $q$  is piecewise constant. We say that a *breakpoint* of  $\mathcal{Q}$  is a value in  $[0, 1]$  that is an endpoint of one of the (at most)  $k$  intervals in  $\mathcal{Q}$ .

The following important lemma bounds the number of intervals in the final partition  $\mathcal{P}_s$ :

**Lemma 3.3.**  *$\mathcal{P}_s$  contains at most  $O(k \log^2(1/\varepsilon))$  intervals.*

The following definition will be useful:

**Definition 5.** Let  $\mathcal{P}$  denote any partition of  $[0, 1)$ . We say that partition  $\mathcal{P}$  is  $\varepsilon'$ -good for  $(p, q)$  if for every breakpoint  $v$  of  $\mathcal{Q}$ , the interval  $I$  in  $\mathcal{P}$  containing  $v$  satisfies  $p(I) \leq \varepsilon'/(2k)$ .

The above definition is justified by the following lemma:

**Lemma 3.4.** *If  $\mathcal{P}$  is  $\varepsilon'$ -good for  $(p, q)$ , then  $\|p - (p)^\mathcal{P}\|_1 \leq 2\text{opt}_k(p) + \varepsilon'$ .*

We are now in a position to prove the following:

**Lemma 3.5.** *There exists a partition  $\mathcal{R}$  of  $[0, 1)$  that is  $\varepsilon'$ -good for  $(p, q)$  and satisfies*

$$\|(p)^{\mathcal{P}_s} - (p)^\mathcal{R}\|_1 \leq \varepsilon.$$

We construct the claimed  $\mathcal{R}$  based on  $\mathcal{P}_s, \mathcal{P}_{s-1}, \dots, \mathcal{P}_0$  as follows: (i) If  $I$  is an interval in  $\mathcal{P}_s$  not containing a breakpoint of  $\mathcal{Q}$ , then  $I$  is also in  $\mathcal{R}$ ; (ii) If  $I$  is an interval in  $\mathcal{P}_s$  that does contain a breakpoint of  $\mathcal{Q}$ , then we further partition  $I$  into a set of intervals  $S$  in a recursive manner using  $\mathcal{P}_{s-1}, \dots, \mathcal{P}_0$  (see Appendix A.4). Finally, by putting everything together we can prove Theorem 4:

**Proof of Theorem 4.** By Lemma 3.4 applied to  $\mathcal{R}$ , we have that  $\|p - (p)^\mathcal{R}\|_1 \leq 2\text{opt}_k(p) + \varepsilon'$ . By Lemma 3.5, we have that  $\|(p)^{\mathcal{P}_s} - (p)^\mathcal{R}\|_1 \leq \varepsilon$ ; thus the triangle inequality gives that  $\|p - (p)^{\mathcal{P}_s}\|_1 \leq 2\text{opt}_k(p) + 2\varepsilon$ . By Lemma 3.3 the partition  $\mathcal{P}_s$  contains at most  $O(k \log^2(1/\varepsilon))$  intervals, so both  $(p)^{\mathcal{P}_s}$  and  $(\hat{p}_m)^{\mathcal{P}_s}$  are  $O(k \log^2(1/\varepsilon))$ -flat distributions. Thus,  $\|(p)^{\mathcal{P}_s} - (\hat{p}_m)^{\mathcal{P}_s}\|_1 = \|(p)^{\mathcal{P}_s} - (\hat{p}_m)^{\mathcal{P}_s}\|_{\mathcal{A}_\ell}$ , where  $\ell = O(k \log^2(1/\varepsilon))$  and  $\mathcal{A}_\ell$  is the family of all subsets of  $[0, 1)$  that consist of unions of up to  $\ell$  intervals (which has VC dimension  $2\ell$ ). Consequently by the VC inequality (Theorem 3, for a suitable choice of  $m = \tilde{O}(k/\varepsilon'^2)$ ), we have that  $\mathbf{E}[\|(p)^{\mathcal{P}_s} - (\hat{p}_m)^{\mathcal{P}_s}\|_1] \leq 4\varepsilon'/100$ . Markov's inequality now gives that with probability at least  $96/100$ , we have  $\|(p)^{\mathcal{P}_s} - (\hat{p}_m)^{\mathcal{P}_s}\|_1 \leq \varepsilon'$ . Hence, with overall probability at least  $19/20$  (recall the  $1/100$  error probability incurred in Lemma 3.1), we have that  $\|p - (\hat{p}_m)^{\mathcal{P}_s}\|_1 \leq 2\text{opt}_k(p) + 3\varepsilon$ , and the theorem is proved.  $\square$

### 3.2 A general reduction to the case of small $\text{opt}$ for semi-agnostic learning

In this section we show that under mild conditions, the general problem of agnostic distribution learning for a class  $\mathcal{C}$  can be efficiently reduced to the special case when  $\text{opt}_\mathcal{C}$  is not too large compared with  $\varepsilon$ . While the reduction is simple and generic, we have not previously encountered it in the literature on density estimation, so we provide a proof in Appendix A.5. A precise statement of the reduction follows:

**Theorem 6.** *Let  $A$  be an algorithm with the following behavior:  $A$  is given as input i.i.d. points drawn from  $p$  and a parameter  $\varepsilon > 0$ .  $A$  uses  $m(\varepsilon) = \Omega(1/\varepsilon)$  draws from  $p$ , runs in time  $t(\varepsilon) = \Omega(1/\varepsilon)$ , and satisfies the following: if  $\text{opt}_\mathcal{C}(p) \leq 10\varepsilon$ , then with probability at least  $19/20$  it outputs a hypothesis distribution  $q$  such that (i)  $\|p - q\|_1 \leq \alpha \cdot \text{opt}_\mathcal{C}(p) + \varepsilon$ , where  $\alpha$  is an absolute constant, and (ii) given any  $r \in [0, 1)$ , the value  $q(r)$  of the pdf of  $q$  at  $r$  can be efficiently computed in  $T$  time steps.*

Then there is an algorithm  $A'$  with the following performance guarantee:  $A'$  is given as input i.i.d. draws from  $p$  and a parameter  $\varepsilon > 0$ .<sup>2</sup> Algorithm  $A'$  uses  $O(m(\varepsilon/10) + \log \log(1/\varepsilon)/\varepsilon^2)$  draws from  $p$ , runs in time  $O(t(\varepsilon/10)) + T \cdot \tilde{O}(1/\varepsilon^2)$ , and outputs a hypothesis distribution  $q'$  such that with probability at least  $39/40$  we have  $\|p - q'\|_1 \leq 10(\alpha + 2) \cdot \text{opt}_c(p) + \varepsilon$ .

### 3.3 Dealing with distributions that are not well behaved

The assumption that the target distribution  $p$  is  $\tilde{\Theta}(\varepsilon/k)$ -well-behaved can be straightforwardly removed by following the approach in Section 3.6 of [CDSS14]. That paper presents a simple linear-time sampling-based procedure, using  $\tilde{O}(k/\varepsilon)$  samples, that with high probability identifies all the “heavy” elements (atoms which cause  $p$  to not be well-behaved, if any such points exist).

Our overall algorithm first runs this procedure to find the set  $S$  of “heavy” elements, and then runs the algorithm presented above (which succeeds for well-behaved distributions, i.e., distributions that have no “heavy” elements) using as its target distribution the conditional distribution of  $p$  over  $[0, 1] \setminus S$  (let us denote this conditional distribution by  $p'$ ). A straightforward analysis given in [CDSS14] shows that (i)  $\text{opt}_k(p) \geq \text{opt}_k(p')$ , and moreover (ii)  $d_{\text{TV}}(p, p') \leq \text{opt}_k(p)$ . Thus, by the triangle inequality, any hypothesis  $h$  satisfying  $d_{\text{TV}}(h, p') \leq C \text{opt}_k(p') + \varepsilon$  will also satisfy  $d_{\text{TV}}(h, p) \leq (C + 1) \text{opt}_k(p) + \varepsilon$  as desired.

## 4 Lower bounds on agnostic learning

In this section we establish that  $\alpha$ -agnostic learning with  $\alpha < 2$  is information theoretically impossible, thus establishing Theorem 2.

Fix any  $0 < t < 1/2$ . We define a probability distribution  $\mathcal{D}_t$  over a finite set of discrete distributions over the domain  $[2N] = \{1, \dots, 2N\}$  as follows. (We assume without loss of generality below that  $t$  is rational and that  $tN$  is an integer.) A draw of  $p_{S_1, S_2, t}$  from  $\mathcal{D}_t$  is obtained as follows.

1. A set  $S_1 \subset [N]$  is chosen uniformly at random from all subsets of  $[N]$  that contain precisely  $tN$  elements. For  $i \in [N]$ , the distribution  $p_{S_1, S_2, t}$  assigns probability weight as follows:

$$p_{S_1, S_2, t}(i) = \frac{1}{4N} \text{ if } i \in S_1, \quad p_{S_1, S_2, t}(i) = \frac{1}{2N} \left( 1 + \frac{t}{2(1-t)} \right) \text{ if } i \in [N] \setminus S_1.$$

2. A set  $S_2 \subset [N + 1, \dots, 2N]$  is chosen uniformly at random from all subsets of  $[N + 1, \dots, 2N]$  that contain precisely  $tN$  elements. For  $i \in [N + 1, \dots, 2N]$ , the distribution  $p_{S_1, S_2, t}$  assigns probability weight as follows:

$$p_{S_1, S_2, t}(i) = \frac{3}{4N} \text{ if } i \in S_2, \quad \frac{1}{2N} \left( 1 - \frac{t}{2(1-t)} \right) \text{ if } i \in [N] \setminus S_1.$$

Using a birthday paradox type argument, we show that no  $o(\sqrt{N})$ -sample algorithm can successfully distinguish between a distribution  $p_{S_1, S_2, t} \sim \mathcal{D}_t$  and the uniform distribution over  $[2N]$ . We then leverage this indistinguishability to show that any  $(2 - \delta)$ -semi-agnostic learning algorithm, even for 2-flat distributions, must use a sample of size  $\Omega(\sqrt{N})$  (see Appendix B for these proofs):

**Theorem 7.** *Fix any  $\delta > 0$  and any function  $f(\cdot)$ . There is no algorithm  $A$  with the following property: given  $\varepsilon > 0$  and access to independent points drawn from an unknown distribution  $p$  over  $[2N]$ , algorithm  $A$  makes  $o(\sqrt{N}) \cdot f(\varepsilon)$  draws from  $p$  and with probability at least  $51/100$  outputs a hypothesis distribution  $h$  over  $[2N]$  satisfying  $\|h - p\|_1 \leq (2 - \delta) \text{opt}_2(p) + \varepsilon$ .*

As described in the Introduction, via the obvious correspondence that maps distributions over  $[N]$  to distributions over  $[0, 1]$ , we get the following:

**Corollary 4.1.** *Fix any  $\delta > 0$  and any function  $f(\cdot)$ . There is no algorithm  $A$  with the following property: given  $\varepsilon > 0$  and access to independent draws from an unknown distribution  $p$  over  $[0, 1]$ , algorithm  $A$  makes  $f(\varepsilon)$  draws from  $p$  and with probability at least  $51/100$  outputs a hypothesis distribution  $h$  over  $[0, 1]$  satisfying  $\|h - p\|_1 \leq (2 - \delta) \text{opt}_2(p) + \varepsilon$ .*

<sup>2</sup> Note that now there is no guarantee that  $\text{opt}_c(p) \leq \varepsilon$ ; indeed, the point here is that  $\text{opt}_c(p)$  may be arbitrary.

## References

- [AJOS14] J. Acharya, A. Jafarpour, A. Orlicsky, and A.T. Suresh. Near-optimal-sample estimators for spherical gaussian mixtures. Technical Report <http://arxiv.org/abs/1402.4746>, 19 Feb 2014. [A.5](#)
- [BBBB72] R.E. Barlow, D.J. Bartholomew, J.M. Bremner, and H.D. Brunk. *Statistical Inference under Order Restrictions*. Wiley, New York, 1972. [1](#), [1](#)
- [Bir87] L. Birgé. Estimating a density under order restrictions: Nonasymptotic minimax risk. *Annals of Statistics*, 15(3):995–1012, 1987. [1](#)
- [Bir97] L. Birgé. Estimation of unimodal densities without smoothness assumptions. *Annals of Statistics*, 25(3):970–981, 1997. [1](#)
- [Bru58] H. D. Brunk. On the estimation of parameters restricted by inequalities. *Ann. Math. Statist.*, 29(2):pp. 437–454, 1958. [1](#)
- [CDSS13] S. Chan, I. Diakonikolas, R. Servedio, and X. Sun. Learning mixtures of structured distributions over discrete domains. In *SODA*, pages 1380–1394, 2013. [1](#)
- [CDSS14] S. Chan, I. Diakonikolas, R. Servedio, and X. Sun. Efficient density estimation via piecewise polynomial approximation. Technical Report <http://arxiv.org/abs/1305.3207>, conference version in *STOC*, pages 604–613, 2014. [1](#), [2](#), [3.1.3](#), [3.1](#), [3.3](#), [A.2](#)
- [CMN98] S. Chaudhuri, R. Motwani, and V. Narasayya. Random sampling for histogram construction: How much is enough? In *SIGMOD Conference*, pages 436–447, 1998. [1](#)
- [DDS12] A. De, I. Diakonikolas, and R. Servedio. Inverse problems in approximate uniform generation. Available at <http://arxiv.org/pdf/1211.1722v1.pdf>, 2012. [A.5](#)
- [DG85] L. Devroye and L. Györfi. *Nonparametric Density Estimation: The  $L_1$  View*. John Wiley & Sons, 1985. [1](#)
- [DK14] C. Daskalakis and G. Kamath. Faster and sample near-optimal algorithms for proper learning mixtures of gaussians. In *COLT*, pages 1183–1213, 2014. [A.5](#)
- [DL01] L. Devroye and G. Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics, Springer, 2001. [1](#), [1](#), [3](#), [A.5](#)
- [GGI<sup>+</sup>02] A. Gilbert, S. Guha, P. Indyk, Y. Kotidis, S. Muthukrishnan, and M. Strauss. Fast, small-space algorithms for approximate histogram maintenance. In *STOC*, pages 389–398, 2002. [1](#)
- [GKS06] S. Guha, N. Koudas, and K. Shim. Approximation and streaming algorithms for histogram construction problems. *ACM Trans. Database Syst.*, 31(1):396–438, 2006. [1](#)
- [Gre56] U. Grenander. On the theory of mortality measurement. *Skand. Aktuarietidskr.*, 39:125–153, 1956. [1](#)
- [Gro85] P. Groeneboom. Estimating a monotone density. In *Proc. of the Berkeley Conference in Honor of Jerzy Neyman and Jack Kiefer*, pages 539–555, 1985. [1](#)
- [HP76] D. L. Hanson and G. Pledger. Consistency in concave regression. *The Annals of Statistics*, 4(6):pp. 1038–1050, 1976. [1](#)
- [ILR12] P. Indyk, R. Levi, and R. Rubinfeld. Approximating and Testing  $k$ -Histogram Distributions in Sub-linear Time. In *PODS*, pages 15–22, 2012. [1](#)
- [JKM<sup>+</sup>98] H. V. Jagadish, N. Koudas, S. Muthukrishnan, V. Poosala, K. Sevcik, and T. Suel. Optimal histograms with quality guarantees. In *VLDB*, pages 275–286, 1998. [1](#)
- [KMR<sup>+</sup>94] M. Kearns, Y. Mansour, D. Ron, R. Rubinfeld, R. Schapire, and L. Sellie. On the learnability of discrete distributions. In *Proc. 26th STOC*, pages 273–282, 1994. [1](#)
- [MP007] Concentration inequalities and model selection. Lecture Notes in Mathematics, 33, 2003, Saint-Flour, Cantal, 2007. Massart, P. and Picard, J., Springer. [1](#)
- [Pea95] K. Pearson. Contributions to the mathematical theory of evolution. ii. skew variation in homogeneous material. *Philosophical Trans. of the Royal Society of London*, 186:343–414, 1895. [1](#)
- [Rao69] B.L.S. Prakasa Rao. Estimation of a unimodal density. *Sankhya Ser. A*, 31:23–36, 1969. [1](#)
- [Reb05] L. Réboul. Estimation of a function under shape restrictions. Applications to reliability. *Ann. Statist.*, 33(3):1330–1356, 2005. [1](#)
- [Sco92] D.W. Scott. *Multivariate Density Estimation: Theory, Practice and Visualization*. Wiley, New York, 1992. [1](#)
- [Sil86] B. W. Silverman. *Density Estimation*. Chapman and Hall, London, 1986. [1](#), [1](#)
- [Val84] L. G. Valiant. A theory of the learnable. In *Proc. 16th Annual ACM Symposium on Theory of Computing (STOC)*, pages 436–445. ACM Press, 1984. [1](#)
- [Weg70] E.J. Wegman. Maximum likelihood estimation of a unimodal density. I. and II. *Ann. Math. Statist.*, 41:457–471, 2169–2174, 1970. [1](#)

## Appendix

### A Omitted Proofs from Section 3

#### A.1 Proof of Lemma 3.2

Observe that in iteration  $t$ , two consecutive intervals  $I_{t-1,i}$  and  $I_{t-1,i+1}$  correspond to two unions of consecutive intervals  $I_a \cup \dots \cup I_b$  and  $I_{b+1} \cup \dots \cup I_c$  respectively from the original partition  $\mathcal{P}_0$ . Moreover, since each interval in  $\mathcal{P}_{t-1} \setminus \mathcal{F}_{t-1}$ ,  $t > 1$ , is formed by merging two consecutive intervals from  $\mathcal{P}_{t-2} \setminus \mathcal{F}_{t-2}$ , it must be the case that  $b - a + 1, c - b + 1 \leq 2^{t-1} < 2^{s-1} \leq 1/(2\varepsilon')$ . Hence, by Lemma 3.1, we have

$$|p(I_{t-1,i}) - \widehat{p}_m(I_{t-1,i})| \leq \sqrt{\varepsilon' \cdot 2^{s-1}} \cdot \frac{\varepsilon'}{10k} \leq \frac{\varepsilon'}{10\sqrt{2}k}$$

and similarly,

$$|p(I_{t-1,i+1}) - \widehat{p}_m(I_{t-1,i+1})| \leq \frac{\varepsilon'}{10\sqrt{2}k}.$$

To simplify notation, let  $I = I_{t-1,i}$  and  $J = I_{t-1,i+1}$ . By definition of  $\alpha$ ,

$$\begin{aligned} \alpha_p(I, J) &= \left| \frac{p(I)}{|I|} - \frac{p(I) + p(J)}{|I| + |J|} \right| |I| + \left| \frac{p(J)}{|J|} - \frac{p(I) + p(J)}{|I| + |J|} \right| |J| \\ &= \frac{2}{|I| + |J|} |p(I)|J| - p(J)|I||. \end{aligned} \quad (1)$$

A straightforward calculation now gives that

$$\begin{aligned} |\alpha_p(I, J) - \alpha_{\widehat{p}_m}(I, J)| &= \frac{2}{|I| + |J|} \left| |p(I)|J| - p(J)|I|| - |\widehat{p}_m(I)|J| - \widehat{p}_m(J)|I|| \right| \\ &\leq \frac{2}{|I| + |J|} \left( |p(I) - \widehat{p}_m(I)||J| + |p(J) - \widehat{p}_m(J)||I| \right) \\ &\leq 2\varepsilon'/(5k). \end{aligned}$$

#### A.2 Proof of Lemma 3.3

We start by recording a basic fact that will be useful in the proof of the lemma. Let  $p$  be a distribution over an interval  $I$  and let  $q$  be any sub-distribution over  $I$ . Perhaps contrary to initial intuition, the optimal scaling  $c \cdot q$ ,  $c > 0$ , of  $q$  to approximate  $p$  (with respect to the  $L_1$ -distance) is not necessarily obtained by scaling  $q$  so that  $c \cdot q$  is a distribution over  $I$ . However, a simple argument (see e.g., Appendix A.1 of [CDSS14]) shows that scaling so that  $c \cdot q$  is a distribution cannot result in  $L_1$ -error more than twice that of the optimal scaling:

**Claim A.1.** *Let  $p, g : I \rightarrow \mathbb{R}^{\geq 0}$  be probability distributions over  $I$  (so  $\int_I p(x)dx = \int_I g(x)dx = 1$ ). Then, writing  $\|f\|_1$  to denote  $\int_I |f(x)|dx$ , for every  $a > 0$  we have that  $\|p - g\|_1 \leq 2\|p - ag\|_1$ .*

We now proceed with the proof of Lemma 3.3.

We first show that a total of at most  $O(k \log(1/\varepsilon'))$  intervals are ever added into  $\mathcal{F}_t$  across all executions of Step 4(b).

Suppose that intervals  $I_{t-1,i}, I_{t-1,i+1}$  are added into  $\mathcal{F}_t$  in some execution of Step 4(b). We consider the following two cases:

**Case 1:**  $I_{t-1,i} \cup I_{t-1,i+1}$  contains at least one breakpoint of  $\mathcal{Q}$ . Since  $\mathcal{Q}$  has at most  $k$  breakpoints, this can happen at most  $k$  times in total.

**Case 2:**  $I_{t-1,i} \cup I_{t-1,i+1}$  does not contain any breakpoint of  $\mathcal{Q}$ . Then  $I_{t-1,i} \cup I_{t-1,i+1}$  is a subset of an interval in  $\mathcal{Q}$ . Recalling that intervals  $I_{t-1,i}, I_{t-1,i+1}$  were added into  $\mathcal{F}_t$  in an execution of Step 4(b), we have that  $\alpha_{\widehat{p}_m}(I_{t-1,i}, I_{t-1,i+1}) > \varepsilon'/(2k)$ , and hence by Lemma 3.2, we have that  $\alpha_p(I_{t-1,i}, I_{t-1,i+1}) \geq \frac{1}{5} \cdot \frac{\varepsilon'}{k}$ . Claim A.1 now implies that the contribution to the

$L_1$  distance between  $p$  and  $q$  from  $I_{t-1,i} \cup I_{t-1,i+1}$ , i.e.,  $\int_{I_{t-1,i} \cup I_{t-1,i+1}} |p(x) - q(x)| dx$ , is at least  $\frac{1}{10} \frac{\varepsilon'}{k}$ .

Since  $\|p - q\|_1 = \text{opt}_k(p)$ , there can be at most

$$k + O\left(\frac{\text{opt}_k(p) \cdot k}{\varepsilon'}\right) = O\left(k \cdot \log \frac{1}{\varepsilon}\right)$$

intervals ever added into  $\mathcal{F}_t$  across all executions of Step 4(b) (note that for the last equality we have used the assumption that  $\text{opt}_k(p) \leq \varepsilon$ ).

Next, we argue that each  $\mathcal{F}_t$  satisfies  $|\mathcal{F}_t| \leq O(k \log^2(1/\varepsilon))$ . We have bounded the number of intervals added into  $\mathcal{F}_t$  in Step 4(b) by  $O(k \log(1/\varepsilon'))$ , so it remains to bound the number of intervals added in Step 4(c)(Case 3) and 4(c)(Case 4). It is clear that a total of at most  $O(\log(1/\varepsilon'))$  intervals are ever added in 4(c)(Case 4). Inspection of Step 4(c)(Case 3) shows that for a given value of  $t$ , the number of intervals that this step adds to  $\mathcal{F}_t$  is at most the number of “blocks” of consecutive  $\mathcal{F}_t$ -intervals. Since each interval added in Step 4(c)(Case 3) extends some blocks of consecutive  $\mathcal{F}_t$ -intervals but does not create a new one (and hence does not increase their number), across the  $s = \log(1/\varepsilon')$  stages, the total number of intervals that can be added in executions of Step 4(c)(Case 3) is at most  $O(k \log^2(1/\varepsilon'))$ . It follows that we have  $|\mathcal{F}_s| = O(k \log^2(1/\varepsilon))$  as claimed.

To bound  $|\mathcal{P}_t \setminus \mathcal{F}_t|$ , we observe that by inspection of the algorithm, for each  $t$  we have  $|\mathcal{P}_t \setminus \mathcal{F}_t| \leq \frac{1}{2} |\mathcal{P}_{t-1} \setminus \mathcal{F}_{t-1}|$ . Since  $|\mathcal{P}_0| = \Theta(k/\varepsilon')$ , it follows that  $|\mathcal{P}_s \setminus \mathcal{F}_s| = O(k)$ , and the lemma is proved.

### A.3 Proof of Lemma 3.4

Fix an interval  $I$  in  $\mathcal{P}$ . If there does not exist an interval  $J$  in  $\mathcal{Q}$  such that  $I \subseteq J$ , then  $I$  must contain a breakpoint of  $\mathcal{Q}$ , and hence since  $\mathcal{P}$  is  $\varepsilon'$ -good for  $(p, q)$ , we have  $p(I) \leq \varepsilon'/(2k)$ . This implies that the contribution to  $\|(p)^\mathcal{P} - q\|_1$  that comes from  $I$ , namely  $\int_I |(p)^\mathcal{P}(x) - q(x)| dx$ , satisfies

$$\begin{aligned} \int_I |(p)^\mathcal{P}(x) - q(x)| dx &\leq \int_I |(p)^\mathcal{P}(x) - p(x)| dx + \int_I |p(x) - q(x)| dx \\ &\leq \int_I |p(x) - q(x)| dx + 2p(I) \\ &\leq \int_I |p(x) - q(x)| dx + \frac{\varepsilon'}{k}. \end{aligned}$$

The other possibility is that there exists an interval  $J$  in  $\mathcal{Q}$  such that  $I \subseteq J$ . In this case, we have that

$$\int_I |(p)^\mathcal{P}(x) - q(x)| dx \leq \int_I |p(x) - q(x)| dx.$$

Since there are at most  $k$  intervals in  $\mathcal{P}$  containing breakpoints of  $\mathcal{Q}$ , summing the above inequalities over all intervals  $I$  in  $\mathcal{P}$ , we get that

$$\|(p)^\mathcal{P} - q\|_1 \leq \|p - q\|_1 + \varepsilon' = \text{opt}_k(p) + \varepsilon',$$

and hence

$$\|(p)^\mathcal{P} - p\|_1 \leq \|(p)^\mathcal{P} - q\|_1 + \|p - q\|_1 \leq 2\text{opt}_k(p) + \varepsilon'.$$

### A.4 Proof of Lemma 3.5

We construct the claimed  $\mathcal{R}$  based on  $\mathcal{P}_s, \mathcal{P}_{s-1}, \dots, \mathcal{P}_0$  as follows:

- (i) If  $I$  is an interval in  $\mathcal{P}_s$  not containing a breakpoint of  $\mathcal{Q}$ , then  $I$  is also in  $\mathcal{R}$ .
- (ii) If  $I$  is an interval in  $\mathcal{P}_s$  that does contain a breakpoint of  $\mathcal{Q}$ , then we further partition  $I$  into a set of intervals  $S$  by calling procedure `Refine-partition(s, I)`. This recursive procedure exploits the local structure of the earlier, finer partitions  $\mathcal{P}_{s-1}, \mathcal{P}_{s-2}, \dots$  as described below.

**Procedure** Refine-partition:

**Input:** Integer  $t$ , Interval  $J$

**Output:**  $S$ , a partition of interval  $J$

1. If  $t = 0$ , then output  $\{J\}$ .
2. If  $J$  is an interval in  $\mathcal{P}_t$ , then
  - (a) If  $J$  contains a breakpoint of  $\mathcal{Q}$ , then output  $\text{Refine-partition}(t - 1, J)$ .
  - (b) Otherwise output  $\{J\}$ .
3. Otherwise,  $J$  is a union of two intervals in  $\mathcal{P}_t$ . Let  $J_1$  and  $J_2$  denote the two intervals in  $\mathcal{P}_t$  such that  $J_1 \cup J_2 = J$ . Output  $\text{Refine-partition}(t, J_1) \cup \text{Refine-partition}(t, J_2)$ .

We claim that  $|\mathcal{R}|$  (the number of intervals in  $\mathcal{R}$ ) is at most  $|\mathcal{P}_s| + O(k \cdot \log \frac{1}{\varepsilon})$ . To see this, note that each interval  $I \in \mathcal{P}_s$  not containing a breakpoint of  $\mathcal{Q}$  (corresponding to (i) above) translates directly to a single interval of  $\mathcal{R}$ . For each interval of type (ii) in  $\mathcal{P}_s$ , inspection of the `Refine-Partition` procedure shows that these intervals are partitioned into at most  $O(k \log(1/\varepsilon))$  intervals in  $\mathcal{R}$ .

In the rest of the proof, we show that for any interval  $J$  in  $\mathcal{P}_s$  containing at least one breakpoint of  $\mathcal{Q}$ , the contribution to the  $L_1$  distance between  $(p)^{\mathcal{P}_s}$  and  $(p)^{\mathcal{R}}$  coming from interval  $J$  is at most  $|b_J| \cdot \frac{\varepsilon' \log \frac{1}{\varepsilon}}{k}$ , where  $b_J$  is the set of breakpoints of  $\mathcal{Q}$  in  $J$ .

Consider a fixed breakpoint  $v$  of  $\mathcal{Q}$ . Let  $I_{t,v}$  denote the interval containing  $v$  in the partition  $\mathcal{P}_t$ . If  $I_{t,v}$  merges with another interval in  $\mathcal{P}_t$  in Case 1 of Step 4(c), we denote that other interval as  $I'_{t,v}$ . Since  $I_{t,v}$  merges with  $I'_{t,v}$  in Case 1 of Step 4(c), these intervals are both not in  $\mathcal{F}_t$  and hence were both not in  $\mathcal{F}_{t-1}$  in Step 4(b). Consequently when  $t > 1$  it must be the case that condition (ii) of Step 4(b) does not hold for these intervals, i.e.  $\alpha_{\widehat{p}_m}(I_{t,v}, I'_{t,v}) \leq \varepsilon'/(2k)$ . It follows that by Lemma 3.2, we have that  $\alpha_p(I_{t,v}, I'_{t,v})$  is at most  $\frac{4\varepsilon'}{5k}$ . When  $t = 1$ , we have a similar bound  $\alpha_p(I_{t,v}, I'_{t,v}) \leq \varepsilon'/k$ , by using (1) and the fact that  $p(I_{t,v}), p(I'_{t,v}) \leq \varepsilon'/2k$  when  $I_{t,v}, I'_{t,v} \in \mathcal{P}_0$ .

On the other hand, inspection of the procedure `Refine-Partition` gives that if two intervals in  $\mathcal{P}_t$  are unions of some intervals in  $\text{Refine-partition}(s, I)$ , and their union is an interval in  $\mathcal{P}_{t+1}$ , then there exists  $v$  which is a breakpoint of  $\mathcal{Q}$  such that the two intervals are  $I_{t,v}$  and  $I'_{t,v}$ .

Thus, the contribution to the  $L_1$  distance between  $(p)^{\mathcal{P}_s}$  and  $(p)^{\mathcal{R}}$  coming from interval  $J$  is at most  $\frac{\varepsilon'}{k} \cdot \log \frac{1}{\varepsilon'} \cdot |b_J|$ . Summing over all intervals  $J$  that contain at least one breakpoint and recalling that the total number of breakpoints is at most  $k$ , we get that the overall  $L_1$  distance between  $(p)^{\mathcal{P}_s}$  and  $(p)^{\mathcal{R}}$  is at most  $\varepsilon$ .

## A.5 Proof of Theorem 6

*Proof.* The algorithm  $A'$  works in two stages, which we describe and analyze below.

In the first stage,  $A'$  iterates over  $\lceil \log(20/\varepsilon) \rceil$  “guesses” for the value of  $\text{opt}_{\mathcal{C}}(p)$ , where the  $i$ -th guess  $g_i$  is  $\frac{\varepsilon}{10} \cdot 2^{i-1}$  (so  $g_1 = \frac{\varepsilon}{10}$  and  $g_{\lceil \log(20/\varepsilon) \rceil} \geq 1$ ). For each value of  $g_i$ , it performs  $r = O(1)$  runs of Algorithm  $A$  (using a fresh sample from  $p$  for each run) using parameter  $g_i$  as the “ $\varepsilon$ ” parameter for each run; let  $h_{1,i}, \dots, h_{r,i}$  be the  $r$  hypotheses thus obtained for the  $i$ -th guess. It is clear that this stage uses  $O(m(\varepsilon/10) + m(2\varepsilon/10) + \dots) = O(m(\varepsilon))$  draws from  $p$ , and similarly that it runs in time  $O(t(\varepsilon))$ . If  $\text{opt}_{\mathcal{C}}(p) \leq \varepsilon$ , then (for a suitable choice of  $r = O(1)$ ) we get that with probability at least  $39/40$ , some hypothesis  $h_{1,\ell}$  satisfies  $\|p - h_{1,\ell}\| \leq \alpha \cdot \text{opt}_{\mathcal{C}}(p) + \varepsilon/10$ . Otherwise, there must be some  $i \in \{2, \dots, \lceil \log(20/\varepsilon) \rceil\}$  such that  $g_i/2 < \text{opt}_{\mathcal{C}}(p) \leq g_i$ ; in this case, for a suitable choice of  $r = O(1)$  we get that with probability at least  $39/40$ , there is some hypothesis  $h_{i,\ell}$  that satisfies  $\|p - h_{i,\ell}\|_1 \leq \alpha \cdot \text{opt}_{\mathcal{C}}(p) + g_i \leq (\alpha + 2) \cdot \text{opt}_{\mathcal{C}}(p)$ . Thus in either event, with probability at least  $39/40$  some  $h_{i,\ell}$  satisfies  $\|p - h_{i,\ell}\|_1 \leq (\alpha + 2) \cdot \text{opt}_{\mathcal{C}}(p) + \varepsilon/10$ .

In the second stage,  $A'$  runs a hypothesis selection procedure to choose one of the candidate hypotheses  $h_{i,\ell}$ . A number of such procedures are known (see e.g. Section 6.6 of [DL01] or

[DDS12, DK14, AJOS14]); all of them work by running some sort of “tournament” over the hypotheses, and all have the guarantee that with high probability they will output a hypothesis from the pool of candidates which has  $L_1$  error (with respect to the target distribution  $p$ ) not much worse than that of the best candidate in the pool. We use the classic Scheffé algorithm (see [DL01]) as described and analyzed in [AJOS14] (see Algorithm SCHEFFE\* in Appendix B of that paper). Adapted to our context, this algorithm has the following performance guarantee:

**Proposition A.2.** *Let  $p$  be a target distribution over  $[0, 1)$  and let  $\mathcal{D}_\tau = \{p_j\}_{j=1}^N$  be a collection of  $N$  distributions over  $[0, 1)$  with the property that there exists  $i \in [N]$  such that  $\|p - p_i\|_1 \leq \tau$ . There is a procedure SCHEFFE which is given as input a parameter  $\varepsilon > 0$  and a confidence parameter  $\delta > 0$ , and is provided with access to*

- (i) *i.i.d. draws from  $p$  and from  $p_i$  for all  $i \in [N]$ , and*
- (ii) *an evaluation oracle  $eval_{p_i}$  for each  $i \in [N]$ . This is a procedure which, on input  $r \in [0, 1)$ , outputs the value  $p_i(r)$  of the pdf of  $p_i$  at the point  $r$ .*

*The procedure SCHEFFE has the following behavior: It makes  $s = O((1/\varepsilon^2) \cdot (\log N + \log(1/\delta)))$  draws from  $p$  and from each  $p_i$ ,  $i \in [N]$ , and  $O(s)$  calls to each oracle  $eval_{p_i}$ ,  $i \in [N]$ , and performs  $O(sN^2)$  arithmetic operations. With probability at least  $1 - \delta$  it outputs an index  $i^* \in [N]$  that satisfies  $\|p - p_{i^*}\|_1 \leq 10 \max\{\tau, \varepsilon\}$ .*

The algorithm  $A'$  runs the procedure SCHEFFE using the  $N = O(\log(1/\varepsilon))$  hypotheses  $h_{i,\ell}$ , with its “ $\varepsilon$ ” parameter set to  $\frac{1}{10}$  (the input parameter  $\varepsilon$  that is given to  $A'$ ) and its “ $\delta$ ” parameter set to  $1/40$ . By Proposition A.2, with overall probability at least  $19/20$  the output is a hypothesis  $h_{i,\ell}$  satisfying  $\|p - h_{i,\ell}\|_1 \leq 10(\alpha + 2)\text{opt}_C(p) + \varepsilon$ . The overall running time and sample complexity are easily seen to be as claimed, and the theorem is proved.  $\square$

## B Proof of Theorem 7

We write  $\mathcal{U}_{2N}$  to denote the uniform distribution over  $[2N]$ . The following proposition shows that  $\mathcal{U}_{2N}$  has  $L_1$  distance from  $p_{S_1, S_2, t}$  almost twice that of the optimal 2-flat distribution:

**Proposition B.1.** *Fix any  $0 < t < 1/2$ .*

1. *For any distribution  $p_{S_1, S_2, t}$  in the support of  $\mathcal{D}_t$ , we have*

$$\|\mathcal{U}_{2N} - p_{S_1, S_2, t}\|_1 = t.$$

2. *For any distribution  $p_{S_1, S_2, t}$  in the support of  $\mathcal{D}_t$ , we have*

$$\text{opt}_2(p_{S_1, S_2, t}) \leq \frac{t}{2} \left(1 + \frac{t}{1-t}\right).$$

*Proof.* Part (1.) is a simple calculation. For part (2.), consider the 2-flat distribution

$$q(i) = \begin{cases} \frac{1}{2N} \left(1 + \frac{t}{2(1-t)}\right) & \text{if } i \in [N] \\ \frac{1}{2N} \left(1 - \frac{t}{2(1-t)}\right) & \text{if } i \in [N+1, \dots, 2N] \end{cases}$$

It is straightforward to verify that  $\|p_{S_1, S_2, t} - q\|_1 = \frac{t}{2} \left(1 + \frac{t}{1-t}\right)$  as claimed.  $\square$

For a distribution  $p$  we write  $A^p$  to indicate that algorithm  $A$  is given access to i.i.d. points drawn from  $p$ .

The following simple proposition states that no algorithm can successfully distinguish between a distribution  $p_{S_1, S_2, t} \sim \mathcal{D}_t$  and  $\mathcal{U}_{2N}$  using fewer than (essentially)  $\sqrt{N}$  draws:

**Proposition B.2.** *There is an absolute constant  $c > 0$  such that the following holds: Fix any  $0 < t < 1/2$ , and let  $B$  be any “distinguishing algorithm” which receives  $c\sqrt{N}$  i.i.d. draws from a distribution over  $[2N]$  and outputs either “uniform” or “non-uniform”. Then*

$$\left| \Pr[B^{\mathcal{U}_{[2N]}} \text{ outputs “uniform”}] - \Pr_{p_{S_1, S_2, t} \sim \mathcal{D}_t}[B^{p_{S_1, S_2, t}} \text{ outputs “uniform”}] \right| \leq 0.01. \quad (2)$$

The proof is an easy consequence of the fact that in both cases (the distribution is  $\mathcal{U}_{[2N]}$ , or the distribution is  $p_{S_1, S_2, t} \sim \mathcal{D}_t$ ), with probability at least 0.99 the  $c\sqrt{N}$  draws received by  $A$  are a uniform random set of  $c\sqrt{N}$  distinct elements from  $[2N]$  (this can be shown straightforwardly using a birthday paradox type argument).

Now we use Proposition B.2 to show that any  $(2 - \delta)$ -semi-agnostic learning algorithm even for 2-flat distributions must use a sample of size  $\Omega(\sqrt{N})$ , and thereby prove Theorem 7:

**Theorem 7.** *Fix any  $\delta > 0$  and any function  $f(\cdot)$ . There is no algorithm  $A$  with the following property: given  $\varepsilon > 0$  and access to independent points drawn from an unknown distribution  $p$  over  $[2N]$ , algorithm  $A$  makes  $o(\sqrt{N}) \cdot f(\varepsilon)$  draws from  $p$  and with probability at least  $51/100$  outputs a hypothesis distribution  $h$  over  $[2N]$  satisfying  $\|h - p\|_1 \leq (2 - \delta)\text{opt}_2(p) + \varepsilon$ .*

*Proof.* Fix a value of  $\delta > 0$  and suppose, for the sake of contradiction, that there exists such an algorithm  $A$ . We describe how the existence of such an algorithm  $A$  yields a distinguishing algorithm  $B$  that violates Proposition B.2.

The algorithm  $B$  works as follows, given access to i.i.d. draws from an unknown distribution  $p$ . It first runs algorithm  $A$  with its “ $\varepsilon$ ” parameter set to  $\varepsilon := \frac{\delta^3}{12(2+\delta)}$ , obtaining (with probability at least  $51/100$ ) a hypothesis distribution  $h$  over  $[2N]$  such that  $\|h - p\|_1 \leq (2 - \delta)\text{opt}_2(p) + \varepsilon$ . It then computes the value  $\|h - \mathcal{U}_{2N}\|_1$  of the  $L_1$ -distance between  $h$  and the uniform distribution (note that this step uses no draws from the distribution). If  $\|h - \mathcal{U}_{2N}\|_1 < 3\varepsilon/2$  then it outputs “uniform” and otherwise it outputs “non-uniform.”

Since  $\delta$  (and hence  $\varepsilon$ ) is independent of  $N$ , the algorithm  $B$  makes fewer than  $c\sqrt{N}$  draws from  $p$  (for  $N$  sufficiently large). To see that the above-described algorithm  $B$  violates (2), consider first the case that  $p$  is  $\mathcal{U}_{[2N]}$ . In this case  $\text{opt}_2(p) = 0$  and so with probability at least  $51/100$  the hypothesis  $h$  satisfies  $\|h - \mathcal{U}_{2N}\|_1 \leq \varepsilon$ , and hence algorithm  $B$  outputs “uniform” with probability at least  $51/100$ .

On the other hand, suppose that  $p = p_{S_1, S_2, t}$  is drawn from  $\mathcal{D}_t$ , where  $t = \frac{\delta}{2+\delta}$ . In this case, with probability at least  $51/100$  the hypothesis  $h$  satisfies

$$\|h - p_{S_1, S_2, t}\|_1 \leq (2 - \delta)\text{opt}_2(p_{S_1, S_2, t}) + \varepsilon \leq (2 - \delta) \cdot \frac{t}{2} \cdot \left(1 + \frac{t}{1-t}\right) + \varepsilon,$$

by part (2.) of Proposition B.1. Since by part (1.) of Proposition B.1 we have  $\|\mathcal{U}_{2N} - p_{S_1, S_2, t}\|_1 = t$ , the triangle inequality gives that

$$\|h - \mathcal{U}_{2N}\|_1 \geq t - (2 - \delta) \cdot \frac{t}{2} \cdot \left(1 + \frac{t}{1-t}\right) - \varepsilon = 2\varepsilon,$$

where to obtain the final equality we recalled the settings  $\varepsilon = \frac{\delta^3}{12(2+\delta)}$ ,  $t = \frac{\delta}{2+\delta}$ . Hence algorithm  $B$  outputs “uniform” with probability at most  $49/100$ . Thus we have

$$\left| \Pr[B^{\mathcal{U}_{[2N]}} \text{ outputs “uniform”}] - \Pr_{p_{S_1, S_2, t} \sim \mathcal{D}_t}[B^{p_{S_1, S_2, t}} \text{ outputs “uniform”}] \right| \geq 0.02$$

which contradicts (2) and proves the theorem.  $\square$