

Learning k -Modal Distributions via Testing

Constantinos Daskalakis*

Ilias Diakonikolas†

Rocco A. Servedio‡

Abstract

A k -modal probability distribution over the domain $\{1, \dots, n\}$ is one whose histogram has at most k “peaks” and “valleys.” Such distributions are natural generalizations of monotone ($k = 0$) and unimodal ($k = 1$) probability distributions, which have been intensively studied in probability theory and statistics.

In this paper we consider the problem of *learning* an unknown k -modal distribution. The learning algorithm is given access to independent samples drawn from the k -modal distribution p , and must output a hypothesis distribution \hat{p} such that with high probability the total variation distance between p and \hat{p} is at most ϵ .

We give an efficient algorithm for this problem that runs in time $\text{poly}(k, \log(n), 1/\epsilon)$. For $k \leq \tilde{O}(\sqrt{\log n})$, the number of samples used by our algorithm is very close (within an $\tilde{O}(\log(1/\epsilon))$ factor) to being information-theoretically optimal. Prior to this work computationally efficient algorithms were known only for the cases $k = 0, 1$ [Bir87b, Bir97].

A novel feature of our approach is that our learning algorithm crucially uses a new *property testing* algorithm as a key subroutine. The learning algorithm uses the property tester to efficiently decompose the k -modal distribution into k (near)-monotone distributions, which are easier to learn.

1 Introduction

This paper considers a natural unsupervised learning problem involving k -modal distributions over the discrete domain $\{1, \dots, n\}$. A distribution is k -modal if the plot of its probability density function (pdf) has at most k “peaks” and “valleys” (see Section 2.1 for a precise definition). Such distributions arise both in theoretical (see e.g. [CKC83, Kem91, CT04]) and applied (see e.g. [Mur64, dTF90, FPP⁺98]) research; they naturally generalize the simpler classes of monotone ($k = 0$) and unimodal ($k = 1$) distributions that have been intensively studied in probability theory

and statistics (see the discussion of related work below).

Our main aim in this paper is to give an efficient algorithm for *learning* an unknown k -modal distribution p to total variation distance ϵ , given access only to independent samples drawn from p . As described below there is an information-theoretic lower bound of $\Omega(k \log(n/k)/\epsilon^3)$ samples for this learning problem, so an important goal for us is to obtain an algorithm whose sample complexity is as close as possible to this lower bound (and of course we want our algorithm to be computationally efficient, i.e. to run in time polynomial in the size of its input sample). Our main contribution in this paper is a computationally efficient algorithm that has nearly optimal sample complexity for small (but super-constant) values of k .

1.1 Background and related work. There is a rich body of work in the statistics and probability literatures on estimating distributions under various kinds of “shape” or “order” restrictions. In particular, many researchers have studied the risk of different estimators for monotone and unimodal distributions; see for example the works of [Rao69, Weg70, Gro85, Bir87a, Bir87b, Bir97], among many others. In the language of computational learning theory, these and related papers from the probability/statistics literature mostly deal with information-theoretic upper and lower bounds on the sample complexity of learning monotone and unimodal distributions. It should be noted that some of these works do give efficient algorithms for the cases $k = 0$ and $k = 1$; in particular we mention the results of Birgé [Bir87b, Bir97], which give computationally efficient $O(\log(n)/\epsilon^3)$ -sample algorithms for learning unknown monotone or unimodal distributions over $[n]$ respectively. (Birgé [Bir87a] also showed that this sample complexity is asymptotically optimal, as we discuss below; we describe the algorithm of [Bir87b] in more detail in Section 2.2, and indeed use it as an ingredient of our approach throughout this paper.) However, for these relatively simple $k = 0, 1$ classes of distributions the main challenge is in developing sample-efficient estimators, and the algorithmic aspects are typically rather straightforward (as is the case in [Bir87b]). In contrast, much more challenging and interesting algorithmic issues arise for the general values of k which we consider here.

1.2 Our Results. Our main result is a highly efficient algorithm for learning an unknown k -modal distribution over $[n]$:

*Electrical Engineering and Computer Science Department, MIT. Email: costis@csail.mit.edu. Research supported by NSF CAREER award CCF-0953960 and by a Sloan Foundation Fellowship.

†Computer Science Division, UC Berkeley. Email: ilias@cs.berkeley.edu. Research supported by a Simons Foundation Postdoctoral Fellowship. Some of this work was done while at Columbia University, supported by NSF grant CCF-0728736, and by an Alexander S. Onassis Foundation Fellowship.

‡Department of Computer Science, Columbia University. Email: rocco@cs.columbia.edu. Supported by NSF grants CNS-0716245, CCF-0915929, and CCF-1115703.

THEOREM 1.1. *Let p be any unknown k -modal distribution over $[n]$. There is an algorithm that uses¹*

$$\left(\frac{k \log(n/k)}{\epsilon^3} + \frac{k^3}{\epsilon^3} \cdot \log \frac{k}{\epsilon} \cdot \log \log \frac{k}{\epsilon} \right) \cdot \tilde{O}(\log(1/\delta))$$

samples from p , runs for $\text{poly}(k, \log n, 1/\epsilon, \log(1/\delta))$ bit-operations, and with probability $1 - \delta$ outputs a (succinct description of a) hypothesis distribution \hat{p} over $[n]$ such that the total variation distance between p and \hat{p} is at most ϵ .

As alluded to earlier, Birgé [Bir87a] gave a sample complexity lower bound for learning monotone distributions. The lower bound in [Bir87a] is stated for continuous distributions but the arguments are easily adapted to the discrete case; [Bir87a] shows that (for $\epsilon \geq 1/n^{\Omega(1)}$)² any algorithm for learning an unknown monotone distribution over $[n]$ to total variation distance ϵ must use $\Omega(\log(n)/\epsilon^3)$ samples. By a simple construction which concatenates k copies of the monotone lower bound construction over intervals of length n/k , using the monotone lower bound it is possible to show:

PROPOSITION 1.1. *Any algorithm for learning an unknown k -modal distribution over $[n]$ to variation distance ϵ (for $\epsilon \geq 1/n^{\Omega(1)}$) must use $\Omega(k \log(n/k)/\epsilon^3)$ samples.*

Thus our learning algorithm is nearly optimal in its sample complexity; more precisely, for $k \leq \tilde{O}(\sqrt{\log n})$ (and ϵ as bounded above), our sample complexity in Theorem 1.1 is asymptotically optimal up to a factor of $\tilde{O}(\log(1/\epsilon))$. Since each draw from a distribution over $[n]$ is a $\log(n)$ -bit string, Proposition 1.1 implies that the running time of our algorithm is optimal up to polynomial factors. We note that to the best of our knowledge, prior to this work no learning algorithm for k -modal distributions was known that simultaneously had $\text{poly}(k, \log n)$ sample complexity and $\text{poly}(n)$ running time.

1.3 Our Approach. As mentioned in Section 1.1 Birgé gave a highly efficient algorithm for learning a *monotone* distribution in [Bir87b]. Since a k -modal distribution is simply a concatenation of $k + 1$ monotone distributions (first non-increasing, then non-decreasing, then non-increasing, etc.), it is natural to try to use Birgé’s algorithm as a component of an algorithm for learning k -modal distributions, and indeed this is what we do.

The most naive way to use Birgé’s algorithm would be to guess all possible $\binom{n}{k}$ locations of the k “modes” of p . While such an approach can be shown to have good sample complexity, the resulting $\Omega(n^k)$ running time is grossly inefficient. A “moderately naive” approach, which

we analyze in Section 3.1, is to partition $[n]$ into roughly k/ϵ intervals each of weight roughly ϵ/k , and run Birgé’s algorithm separately on each such interval. Since the target distribution is k -modal, at most k of the intervals can be non-monotone; Birgé’s algorithm can be used to obtain an ϵ -accurate hypothesis on each monotone interval, and even if it fails badly on the (at most) k non-monotone intervals, the resulting total contribution towards the overall error from those failures is at most $O(\epsilon)$. This approach is much more efficient than the totally naive approach, giving running time polynomial in k , $\log n$, and $1/\epsilon$, but its sample complexity turns out to be polynomially worse than the $O(k \log(n)/\epsilon^3)$ that we are shooting for. (Roughly speaking, this is because the approach involves running Birgé’s $O(\log(n)/\epsilon^3)$ -sample algorithm k/ϵ times, so it uses at least $k \log(n)/\epsilon^4$ samples.)

Our main learning result is achieved by augmenting the “moderately naive” algorithm sketched above with a new *property testing* algorithm. We give a property testing algorithm for the following problem: given samples from a k -modal distribution p , output “yes” if p is *monotone* and “no” if p is ϵ -far from every monotone distribution. Crucially, our testing algorithm uses $O(k^2/\epsilon^2)$ samples *independent of n* for this problem. Roughly speaking, by using this algorithm $O(k/\epsilon)$ times we are able to identify $k + 1$ intervals that (i) collectively contain almost all of p ’s mass, and (ii) are each (close to) monotone and thus can be handled using Birgé’s algorithm. Thus the overall sample complexity of our approach is (roughly) $(k/\epsilon)^3$ (for the k/ϵ runs of the tester) plus $k \log(n)/\epsilon^3$ (for the k runs of Birgé’s algorithm), which gives Theorem 1.1 and is very close to optimal for k not too large.

1.4 Discussion. Our learning algorithm highlights a novel way that property testing algorithms can be useful for learning. Much research has been done on understanding the relation between property testing algorithms and learning algorithms, see e.g. [GGR98, KR00] and the lengthy survey [Ron08]. As Goldreich has noted [Gol11], an often-invoked motivation for property testing is that (inexpensive) testing algorithms can be used as a “preliminary diagnostic” to determine whether it is appropriate to run a (more expensive) learning algorithm. In contrast, in this work we are using property testing rather differently, as an inexpensive way of decomposing a “complex” object (a k -modal distribution) which we do not *a priori* know how to learn, into a collection of “simpler” objects (monotone or near-monotone distributions) which can be learned using existing techniques. We are not aware of prior learning algorithms that successfully use property testers in this way; we believe that this high-level approach to designing learning algorithms, by using property testers to decompose “complex” objects into simpler objects that can be efficiently learned, may find future applications elsewhere.

¹We write $\tilde{O}(\cdot)$ to hide factors which are polylogarithmic in the argument to $\tilde{O}(\cdot)$; thus for example $\tilde{O}(a \log b)$ denotes a quantity which is $O((a \log b) \cdot (\log(a \log b))^c)$ for some absolute constant c .

²For ϵ sufficiently small the generic upper bound of Theorem A.1, which says that any distribution over $[n]$ can be learned to variation distance ϵ using $O(n/\epsilon^2)$ samples, provides a better bound.

2 Preliminaries

2.1 Notation and Problem Statement. For $n \in \mathbb{Z}^+$, denote by $[n]$ the set $\{1, \dots, n\}$; for $i, j \in \mathbb{Z}^+, i \leq j$, denote by $[i, j]$ the set $\{i, i+1, \dots, j\}$. For $v = (v(1), \dots, v(n)) \in \mathbb{R}^n$ denote by $\|v\|_1 = \sum_{i=1}^n |v(i)|$ its L_1 -norm.

We consider discrete probability distributions over $[n]$, which are functions $p : [n] \rightarrow [0, 1]$ such that $\sum_{i=1}^n p(i) = 1$. For $S \subseteq [n]$ we write $p(S)$ to denote $\sum_{i \in S} p(i)$. For $S \subseteq [n]$, we write p_S to denote the *conditional distribution* over S that is induced by p . We use the notation P for the *cumulative distribution function (cdf)* corresponding to p , i.e. $P : [n] \rightarrow [0, 1]$ is defined by $P(j) = \sum_{i=1}^j p(i)$.

A distribution p over $[n]$ is non-increasing (resp. non-decreasing) if $p(i+1) \leq p(i)$ (resp. $p(i+1) \geq p(i)$), for all $i \in [n-1]$; p is *monotone* if it is either non-increasing or non-decreasing. We call a nonempty interval $I = [a, b] \subseteq [2, n-1]$ a *max-interval* of p if $p(i) = c$ for all $i \in I$ and $\max\{p(a-1), p(b+1)\} < c$; in this case, we say that the point a is a *left max point* of p . Analogously, a *min-interval* of p is an interval $I = [a, b] \subseteq [2, n-1]$ with $p(i) = c$ for all $i \in I$ and $\min\{p(a-1), p(b+1)\} > c$; the point a is called a *left min point* of p . If $I = [a, b]$ is either a max-interval or a min-interval (it cannot be both) we say that I is an *extreme-interval* of p , and a is called a *left extreme point* of p . Note that any distribution uniquely defines a collection of extreme-intervals (hence, left extreme points). We say that p is *k-modal* if it has at most k extreme-intervals.

Let p, q be distributions over $[n]$ with corresponding cdfs P, Q . The *total variation distance* between p and q is $d_{TV}(p, q) := \max_{S \subseteq [n]} |p(S) - q(S)| = (1/2) \cdot \|p - q\|_1$. The *Kolmogorov distance* between p and q is defined as $d_K(p, q) := \max_{j \in [n]} |P(j) - Q(j)|$. Note that $d_K(p, q) \leq d_{TV}(p, q)$.

Learning k -modal Distributions. Given independent samples from an unknown k -modal distribution $p \in \mathcal{M}_n^k$ and $\epsilon > 0$, the goal is to output a hypothesis distribution h such that with probability $1 - \delta$ we have $d_{TV}(p, h) \leq \epsilon$. We say that such an algorithm \mathcal{A} *learns p to accuracy ϵ and confidence δ* . The parameters of interest are the number of samples and the running time required by the algorithm.

2.2 Basic Tools. We will need three tools from probability theory.

Our first tool says that $O(1/\epsilon^2)$ samples suffice to learn any distribution within error ϵ with respect to the *Kolmogorov distance*. This fundamental fact is known as the *Dvoretzky-Kiefer-Wolfowitz (DKW) inequality* ([DKW56]). Given m independent samples s_1, \dots, s_m , drawn from $p : [n] \rightarrow [0, 1]$, the *empirical distribution* $\hat{p}_m : [n] \rightarrow [0, 1]$ is defined as follows: for all $i \in [n]$, $\hat{p}_m(i) = |\{j \in [m] \mid s_j = i\}|/m$. The DKW inequality states that for $m = \Omega((1/\epsilon^2) \cdot \ln(1/\delta))$, with probability $1 - \delta$ the empirical distribution \hat{p}_m will be ϵ -close to p in Kolmogorov distance. This sample bound is asymptotically optimal and

independent of the support size.

THEOREM 2.1. ([DKW56, MAS90]) *For all $\epsilon > 0$, it holds: $\Pr[d_K(p, \hat{p}_m) > \epsilon] \leq 2e^{-2m\epsilon^2}$.*

Our second tool, due to Birgé [Bir87b], provides a sample-optimal and computationally efficient algorithm to learn monotone distributions to ϵ -accuracy in total variation distance. Before we state the relevant theorem, we need a definition. We say that a distribution p is δ -close to being non-increasing (resp. non-decreasing) if there exists a non-increasing (resp. non-decreasing) distribution q such that $d_{TV}(p, q) \leq \delta$. We are now ready to state Birgé’s result:

THEOREM 2.2. ([BIR87B], THEOREM 1) *(semi-agnostic learner) There is an algorithm L^\downarrow with the following performance guarantee: Given m independent samples from a distribution p over $[n]$ which is opt-close to being non-increasing, L^\downarrow performs $\tilde{O}(m \cdot \log n + m^{1/3} \cdot (\log n)^{5/3})$ bit-operations and outputs a (succinct description of a) hypothesis distribution \tilde{p} over $[n]$ that satisfies*

$$\mathbf{E}[d_{TV}(\tilde{p}, p)] \leq 2 \cdot \text{opt} + O\left(\left(\log n / (m+1)\right)^{1/3}\right).$$

The aforementioned algorithm partitions the domain $[n]$ in $O(m^{1/3} \cdot (\log n)^{2/3})$ intervals and outputs a hypothesis distribution that is uniform within each of these intervals.

By taking $m = \Omega(\log n / \epsilon^3)$, one obtains a hypothesis such that $\mathbf{E}[d_{TV}(\tilde{p}, p)] \leq 2 \cdot \text{opt} + \epsilon$. We stress that Birgé’s algorithm for learning non-increasing distributions [Bir87b] is in fact “semi-agnostic”, in the sense that it also learns distributions that are close to being non-increasing; this robustness will be crucial for us later (since in our final algorithm we will use Birgé’s algorithm on distributions identified by our tester, that are close to monotone but not necessarily perfectly monotone). This semi-agnostic property is not explicitly stated in [Bir87b] but it can be shown to follow easily from his results. We show how the semi-agnostic property follows from Birgé’s results in Appendix A. Let L^\downarrow denote the corresponding semi-agnostic algorithm for learning non-decreasing distributions.

Our final tool is a routine to do *hypothesis testing*, i.e. to select a high-accuracy hypothesis distribution from a collection of hypothesis distributions one of which has high accuracy. The need for such a routine arises in several places; in some cases we know that a distribution is monotone, but do not know whether it is non-increasing or non-decreasing. In this case, we can run both algorithms L^\uparrow and L^\downarrow and then choose a good hypothesis using hypothesis testing. Another need for hypothesis testing is to “boost confidence” that a learning algorithm generates a high-accuracy hypothesis. Our initial version of the algorithm for Theorem 1.1 generates an ϵ -accurate hypothesis with probability at least $9/10$; by running it $O(\log(1/\delta))$ times using a hypothesis testing routine, it is possible to identify an $O(\epsilon)$ -accurate

hypothesis with probability $1 - \delta$. Routines of the sort that we require have been given in e.g. [DL01] and [DDS11]; we use the following theorem from [DDS11]:

THEOREM 2.3. *There is an algorithm `Choose-Hypothesisp`($h_1, h_2, \epsilon', \delta'$) which is given oracle access to p , two hypothesis distributions h_1, h_2 for p , an accuracy parameter ϵ' , and a confidence parameter δ' . It makes $m = O(\log(1/\delta')/\epsilon'^2)$ draws from p and returns a hypothesis $h \in \{h_1, h_2\}$. If one of h_1, h_2 has $d_{\text{TV}}(h_i, p) \leq \epsilon'$ then with probability $1 - \delta'$ the hypothesis h that `Choose-Hypothesis` returns has $d_{\text{TV}}(h, p) \leq 6\epsilon'$.*

For the sake of completeness, we describe and analyze the `Choose-Hypothesis` algorithm in Appendix B.

3 Learning k -modal Distributions

In this section, we present our main result: a nearly sample-optimal and computationally efficient algorithm to learn an unknown k -modal distribution. In Section 3.1 we present a simple learning algorithm with a suboptimal sample complexity. In Section 3.2 we present our main result which involves a property testing algorithm as a subroutine.

3.1 Warm-up: A simple learning algorithm. In this subsection, we give an algorithm that runs in time $\text{poly}(k, \log n, 1/\epsilon, \log(1/\delta))$ and learns an unknown k -modal distribution to accuracy ϵ and confidence δ . The sample complexity of the algorithm is suboptimal as a function of ϵ , by a polynomial factor.

In the following figure we give the algorithm `Learn-kmodal-simple` which produces an ϵ -accurate hypothesis with confidence $9/10$ (see Theorem 3.1). We explain how to boost the confidence to $1 - \delta$ after the proof of the theorem.

The algorithm `Learn-kmodal-simple` works as follows: We start by partitioning the domain $[n]$ into consecutive intervals of mass approximately ϵ/k . To do this, we make use of the DKW inequality, with accuracy parameter roughly ϵ^2/k . (Some care is needed in this step, since there may be “heavy” points in the support of the distribution; however, we gloss over this technical issue for the sake of this intuitive explanation.) If this step is successful, we have partitioned the domain into a set of $O(k/\epsilon)$ consecutive intervals of probability mass roughly ϵ/k . Our next step is to apply Birgé’s monotone learning algorithm to each interval.

`Learn-kmodal-simple`

Inputs: $\epsilon > 0$; sample access to k -modal distribution p over $[n]$

1. Fix $\tau := \epsilon^2/(100k)$. Draw $r = \Theta(1/\tau^2)$ samples from p and let \hat{p} denote the resulting empirical distribution.
2. Greedily partition the domain $[n]$ into ℓ atomic intervals $\mathcal{I} := \{I_i\}_{i=1}^\ell$ as follows: $I_1 := [1, j_1]$, where $j_1 := \min\{j \in [n] \mid \hat{p}([1, j]) \geq \epsilon/(10k)\}$. For $i \geq 1$, if $\cup_{j=1}^i I_j = [1, j_i]$, then $I_{i+1} := [j_i + 1, j_{i+1}]$, where j_{i+1} is defined as follows: If $\hat{p}([j_i + 1, n]) \geq \epsilon/(10k)$, then $j_{i+1} := \min\{j \in [n] \mid \hat{p}([j_i + 1, j]) \geq \epsilon/(10k)\}$, otherwise, $j_{i+1} := n$.
3. Construct a set of ℓ light intervals $\mathcal{I}' := \{I'_i\}_{i=1}^\ell$ and a set $\{b_i\}_{i=1}^\ell$ of $t \leq \ell$ heavy points as follows: For each interval $I_i = [a, b] \in \mathcal{I}$, if $\hat{p}(I) \geq \epsilon/(5k)$ define $I'_i := [a, b - 1]$ and make b a heavy point. (Note that it is possible to have $I'_i = \emptyset$.) Otherwise, define $I'_i := I_i$.
Fix $\delta' := \epsilon/(500k)$.
4. Draw $m = (k/\epsilon^4) \cdot \log(n) \cdot \tilde{\Theta}(\log(1/\delta'))$ samples $\mathbf{s} = \{s_i\}_{i=1}^m$ from p . For each light interval I'_i , $i \in [\ell]$, run both $L^{\downarrow \delta'}$ and $L^{\uparrow \delta'}$ on the conditional distribution $p_{I'_i}$ using the samples in $\mathbf{s} \cap I'_i$. Let $\tilde{p}_{I'_i}^{\downarrow}, \tilde{p}_{I'_i}^{\uparrow}$ be the corresponding conditional hypothesis distributions.
5. Draw $m' = \Theta((k/\epsilon^4) \cdot \log(1/\delta'))$ samples $\mathbf{s}' = \{s'_i\}_{i=1}^{m'}$ from p . For each light interval I'_i , $i \in [\ell]$, run `Choose-Hypothesisp`($\tilde{p}_{I'_i}^{\downarrow}, \tilde{p}_{I'_i}^{\uparrow}, \epsilon, \delta'$) using the samples in $\mathbf{s}' \cap I'_i$. Denote by $\tilde{p}_{I'_i}$ the returned conditional distribution on I'_i .
6. Output the hypothesis $h = \sum_{j=1}^\ell \hat{p}(I'_j) \cdot \tilde{p}_{I'_j} + \sum_{j=1}^t \hat{p}(b_j) \cdot \mathbf{1}_{b_j}$.

A caveat comes from the fact that not all such intervals are guaranteed to be monotone (or even close to being monotone). However, since our input distribution is assumed to be k -modal, all but (at most) k of these intervals are monotone. Call a non-monotone interval “bad”. Since all intervals have probability mass at most ϵ/k and there are at most k bad intervals, these intervals contribute at most ϵ to the total mass. So even though Birgé’s algorithm gives no guarantees for bad intervals, these intervals do not affect the error by more than ϵ .

Let us now focus on the monotone intervals. For each such interval, we do not know if it is monotone increasing or monotone decreasing. To overcome this difficulty, we run both monotone algorithms L^{\downarrow} and L^{\uparrow} for each interval and

then use hypothesis testing to choose the correct candidate distribution.

Also, note that since we have k/ϵ intervals, we need to run each instance of both the monotone learning algorithms and the hypothesis testing algorithm with confidence $1 - O(\epsilon/k)$, so that we can guarantee that the overall algorithm has confidence $9/10$. Note that Theorem 2.2 and Markov's inequality imply that if we draw $\Omega(\log n/\epsilon^3)$ samples from a non-increasing distribution p , the hypothesis \tilde{p} output by L^\downarrow satisfies $d_{TV}(\tilde{p}, p) \leq \epsilon$ with probability $9/10$. We can boost the confidence to $1 - \delta$ with an overhead of $O(\log(1/\delta) \log \log(1/\delta))$ in the sample complexity:

FACT 3.1. *Let p be a non-increasing distribution over $[n]$. There is an algorithm L^\downarrow_δ with the following performance guarantee: Given $(\log n/\epsilon^3) \cdot \tilde{O}(\log(1/\delta))$ samples from p , L^\downarrow_δ performs $\tilde{O}((\log^2 n/\epsilon^3) \cdot \log(1/\delta))$ bit-operations and outputs a (succinct description of a) hypothesis distribution \tilde{p} over $[n]$ that satisfies $d_{TV}(\tilde{p}, p) \leq \epsilon$ with probability at least $1 - \delta$.*

The algorithm L^\downarrow_δ runs L^\downarrow $O(\log(1/\delta))$ times and performs a tournament among the candidate hypotheses using `Choose-Hypothesis`. Let L^\uparrow_δ denote the corresponding algorithm for learning non-decreasing distributions with confidence δ . We postpone further details on these algorithms to Appendix C.

THEOREM 3.1. *The algorithm `Learn-kmodal-simple` uses $\frac{k \log n}{\epsilon^4} \cdot \tilde{O}(\log \frac{k}{\epsilon}) + O(\frac{k^2}{\epsilon^4})$ samples, performs $\text{poly}(k, \log n, 1/\epsilon)$ bit-operations, and learns a k -modal distribution to accuracy $O(\epsilon)$ with probability $9/10$.*

Proof. We first prove that with probability $9/10$ (over its random samples), algorithm `Learn-kmodal-simple` outputs a hypothesis h such that $d_{TV}(h, p) \leq O(\epsilon)$.

Since $r = \Theta(1/\tau^2)$ samples are drawn in Step 1, the DKW inequality implies that with probability of failure at most $1/100$, for each interval $I \subseteq [n]$ we have $|\hat{p}(I) - p(I)| \leq 2\tau$. For the rest of the analysis of `Learn-kmodal-simple` we condition on this “good” event.

Since every atomic interval $I \in \mathcal{I}$ has $\hat{p}(I) \geq \epsilon/(10k)$ (except potentially the rightmost one), it follows that the number ℓ of atomic intervals constructed in Step 2 satisfies $\ell \leq 10 \cdot (k/\epsilon)$. By the construction in Steps 2 and 3, every light interval $I' \in \mathcal{I}'$ has $\hat{p}(I') \leq \epsilon/(5k)$, which implies $p(I') \leq \epsilon/(5k) + 2\tau$. Note also that every heavy point b has $\hat{p}(b) \geq \epsilon/(10k)$ and the number of heavy points t is at most ℓ .

Since the light intervals and heavy points form a partition of $[n]$, we can write $p = \sum_{j=1}^{\ell} p(I'_j) \cdot p_{I'_j} + \sum_{j=1}^t p(b_j) \cdot 1_{b_j}$. Therefore, we can bound the variation distance as fol-

lows:

$$d_{TV}(h, p) \leq \sum_{j=1}^{\ell} |\hat{p}(I'_j) - p(I'_j)| + \sum_{j=1}^t |\hat{p}(b_j) - p(b_j)| + \sum_{j=1}^{\ell} p(I'_j) \cdot d_{TV}(\tilde{p}_{I'_j}, p_{I'_j}).$$

By the DKW inequality, each term in the first two sums is bounded from above by 2τ . Hence the contribution of these terms to the sum is at most $2\tau \cdot (\ell + t) \leq 4\tau \cdot \ell \leq 2\epsilon/5$.

We proceed to bound the contribution of the third term. Since p is k -modal, at most k of the light intervals I'_j are not monotone for p . Call these intervals “bad”. Even though we have not identified the bad intervals, we know that all such intervals are light. Therefore, their total probability mass under p is at most $k \cdot (\epsilon/(5k) + 2\tau)$. This implies that the contribution of bad intervals to the third term of the variation distance is at most $\epsilon/4$. (Note that this statement holds true independent of the samples s we draw in Step 4.) It remains to bound the contribution of monotone intervals to the third term.

Let $\ell' \leq \ell$ be the number of monotone light intervals and assume after renaming the indices that they are $\tilde{\mathcal{I}} := \{I'_i\}_{i=1}^{\ell'}$. To bound the variation distance, it suffices to show that with probability at least $19/20$ (over the samples drawn in Steps 4-5) it holds

$$(3.1) \quad \sum_{j=1}^{\ell'} p(I'_j) \cdot d_{TV}(\tilde{p}_{I'_j}, p_{I'_j}) = O(\epsilon)$$

Note first that we do not have a lower bound on the probability mass of the intervals in $\tilde{\mathcal{I}}$. We partition this set in two subsets: the subset $\tilde{\mathcal{I}}'$ containing those intervals whose probability mass under p is at most $\epsilon^2/(20k)$; and its complement $\tilde{\mathcal{I}}''$. It is clear that the contribution of $\tilde{\mathcal{I}}'$ to the above expression can be at most $\ell \cdot \epsilon/(20k) \leq \epsilon/2$. We further partition the set $\tilde{\mathcal{I}}''$ of remaining intervals into $b = \lceil \log(5/\epsilon) \rceil$ groups. For $i \in [b]$, the set $(\tilde{\mathcal{I}}'')_i$ consists of those intervals in $\tilde{\mathcal{I}}''$ that have mass under p in the range $[2^{-i} \cdot (\epsilon/5k), 2^{-i+1} \cdot (\epsilon/5k)]$. (Note that these intervals collectively cover all intervals in $\tilde{\mathcal{I}}''$, since each such interval has weight between $\epsilon^2/(20k)$ and $\epsilon/(4k)$ – recall that every light interval $I' \in \mathcal{I}'$ satisfies $p(I') \leq \epsilon/(5k) + 2\tau < \epsilon/(4k)$.) We have:

CLAIM 1. *With probability at least $19/20$ (over the sample s, s'), for each $i \in [r]$ and each monotone light interval $I'_j \in (\tilde{\mathcal{I}}'')_i$ we have $d_{TV}(\tilde{p}_{I'_j}, p_{I'_j}) = O(2^{i/3} \cdot \epsilon)$.*

Proof. Since in Step 4 we draw m samples, and each interval $I'_j \in (\tilde{\mathcal{I}}'')_i$ has $p(I'_j) \in [2^{-i} \cdot (\epsilon/5k), 2^{-i+1} \cdot (\epsilon/5k)]$, a standard coupon collector argument [NS60] tells us that with probability $99/100$, for each (i, j) pair, the interval I'_j will get at least $2^{-i} \cdot (\log(n)/\epsilon^3) \cdot \tilde{\Omega}(\log(1/\delta'))$ many samples. Let's rewrite this as $(\log(n)/(2^{i/3} \cdot \epsilon^3)) \cdot \tilde{\Omega}(\log(1/\delta'))$ samples. We condition on this event.

Fix an interval $I'_j \in (\widetilde{\mathcal{I}}'')_i$. We first show that with failure probability at most $\epsilon/(500k)$ after Step 4, either $\widetilde{p}_{I'_j}^\downarrow$ or $\widetilde{p}_{I'_j}^\uparrow$ will be $(2^{i/3} \cdot \epsilon)$ -accurate. Indeed, by Fact 3.1 and taking into account the number of samples that landed in I'_j , with probability $1 - \epsilon/(500k)$ over \mathbf{s} , $d_{\text{TV}}(\widetilde{p}_{I'_j}^{\alpha_i}, p_{I'_j}) \leq 2^{i/3}\epsilon$, where $\alpha_i = \downarrow$ if $p_{I'_j}$ is non-increasing and $\alpha_i = \uparrow$ otherwise. By a union bound over all (at most ℓ many) (i, j) pairs, it follows that with probability at least $49/50$, for each interval $I'_j \in (\widetilde{\mathcal{I}}'')_i$ one of the two candidate hypothesis distributions is $(2^{i/3}\epsilon)$ -accurate. We condition on this event.

Consider Step 5. For a fixed interval $I'_j \in (\widetilde{\mathcal{I}}'')_i$, Theorem 2.3 implies that the algorithm `Choose-Hypothesis` will output a hypothesis that is $6 \cdot (2^{i/3}\epsilon)$ -close to $p_{I'_j}$ with probability $1 - \epsilon/(500k)$. By a union bound, it follows that with probability at least $49/50$, the above condition holds for all monotone light intervals under consideration. Therefore, except with failure probability $19/20$, the statement of the Claim holds. \blacksquare

Assuming the claim, (3.1) follows by exploiting the fact that for intervals I'_j such that $p(I'_j)$ is small we can afford worse error on the variation distance. More precisely, let $w_i = |(\widetilde{\mathcal{I}}'')_i|$, the number of intervals in $(\widetilde{\mathcal{I}}'')_i$, and note that $\sum_{i=1}^b w_i \leq \ell$. Hence, we can bound the LHS of (3.1) from above by

$$\begin{aligned} & \sum_{i=1}^b w_i \cdot (\epsilon/5k) \cdot 2^{-i+1} \cdot O(2^{i/3} \cdot \epsilon) \\ & \leq O(1) \cdot (2\epsilon^2/5k) \cdot \sum_{i=1}^b w_i \cdot 2^{-2i/3}. \end{aligned}$$

Since $\sum_{i=1}^b w_i \leq \ell$, the above expression is maximized for $w_1 = \ell$ and $w_i = 0, i > 1$, and the maximum value is at most $O(1) \cdot (2\epsilon^2/5k) \cdot \ell = O(\epsilon)$. This proves (3.1).

It is clear that the algorithm has the claimed sample complexity. The running time is also easy to analyze, as it is easy to see that every step can be performed in polynomial time (in fact, nearly linear time) in the sample size. This completes the proof of Theorem 3.1. \blacksquare

To get an $O(\epsilon)$ -accurate hypothesis with probability $1 - \delta$, we can simply run `Learn-kmodal-simple` $O(\log(1/\delta))$ times and then perform a tournament using Theorem 2.3. This increases the sample complexity by a $\widetilde{O}(\log(1/\delta))$ factor. The running time increases by a factor of $O(\log^2(1/\delta))$. We postpone the details for Appendix C.

3.2 Main Result: Learning k -modal distributions using testing. Here is some intuition to motivate our k -modal distribution learning algorithm and give a high-level idea of why the dominant term in its sample complexity is $O(k \log(n/k)/\epsilon^3)$.

Let p denote the target k -modal distribution to be learned. As discussed above, optimal (in terms of time and sample complexity) algorithms are known for learning a monotone distribution over $[n]$, so if the locations of the k modes of p were known then it would be straightforward to learn p very efficiently by running the monotone distribution learner over $k + 1$ separate intervals. But it is clear that in general we cannot hope to efficiently identify the modes of p exactly (for instance it could be the case that $p(a) = p(a + 2) = 1/n$ while $p(a + 1) = 1/n + 1/2^n$). Still, it is natural to try to decompose the k -modal distribution into a collection of (nearly) monotone distributions and learn those. At a high level that is what our algorithm does, using a novel *property testing* algorithm.

More precisely, we give a distribution testing algorithm with the following performance guarantee: Let q be a k -modal distribution over $[n]$. Given an accuracy parameter τ , our tester takes $\text{poly}(k/\tau)$ samples from q and outputs “yes” with high probability if q is monotone and “no” with high probability if q is τ -far from every monotone distribution. (We stress that the assumption that q is k -modal is essential here, since an easy argument given in [BKR04] shows that $\Omega(n^{1/2})$ samples are required to test whether a general distribution over $[n]$ is monotone versus $\Theta(1)$ -far from monotone.)

With some care, by running the above-described tester $O(k/\epsilon)$ times with accuracy parameter τ , we can decompose the domain $[n]$ into

- at most $k + 1$ “superintervals,” which have the property that the conditional distribution of p over each superinterval is almost monotone (τ -close to monotone);
- at most $k + 1$ “negligible intervals”, which have the property that each one has probability mass at most $O(\epsilon/k)$ under p (so ignoring all of them incurs at most $O(\epsilon)$ total error); and
- at most $k + 1$ “heavy” points, each of which has mass at least $\Omega(\epsilon/k)$ under p .

We can ignore the negligible intervals, and the heavy points are easy to handle; however some care must be taken to learn the “almost monotone” restrictions of p over each superinterval. A naive approach, using a generic $\log(n)/\epsilon^3$ -sample monotone distribution learner that has no performance guarantees if the target distribution is not monotone, leads to an inefficient overall algorithm. Such an approach would require that τ (the closeness parameter used by the tester) be at most $1/(\text{the sample complexity of the monotone distribution learner})$, i.e. $\tau < \epsilon^3/\log(n)$. Since the sample complexity of the tester is $\text{poly}(k/\tau)$ and the tester is run k/ϵ times, this

approach would lead to an overall sample complexity that is unacceptably high.

Fortunately, instead of using a generic monotone distribution learner, we can use the semi-agnostic monotone distribution learner of Birgé (Theorem 2.2) that can handle deviations from monotonicity far more efficiently than the above naive approach. Recall that given draws from a distribution q over $[n]$ that is τ -close to monotone, this algorithm uses $O(\log(n)/\epsilon^3)$ samples and outputs a hypothesis distribution that is $(2\tau + \epsilon)$ -close to monotone. By using this algorithm we can take the accuracy parameter τ for our tester to be $\Theta(\epsilon)$ and learn the conditional distribution of p over a given superinterval to accuracy $O(\epsilon)$ using $O(\log(n)/\epsilon^3)$ samples from that superinterval. Since there are $k+1$ superintervals overall, a careful analysis shows that $O(k \log(n)/\epsilon^3)$ samples suffice to handle all the superintervals.

We note that the algorithm also requires an additional additive $\text{poly}(k/\epsilon)$ samples (independent of n) besides this dominant term (for example, to run the tester and to estimate accurate weights with which to combine the various sub-hypotheses). The overall sample complexity we achieve is stated in Theorem 3.2 below.

THEOREM 3.2. (Main) *The algorithm `Learn-kmodal` uses $O(k \log(n/k)/\epsilon^3 + (k^3/\epsilon^3) \cdot \log(k/\epsilon) \cdot \log \log(k/\epsilon))$ samples, performs $\text{poly}(k, \log n, 1/\epsilon)$ bit-operations, and learns any k -modal distribution to accuracy ϵ and confidence $9/10$.*

Theorem 1.1 follows from Theorem 3.2 by running `Learn-kmodal` $O(\log(1/\delta))$ times and using hypothesis testing to boost the confidence to $1 - \delta$. We give details in Appendix C.

Algorithm `Learn-kmodal` makes essential use of an algorithm T^\uparrow for testing whether a k -modal distribution over $[n]$ is non-decreasing. Algorithm $T^\uparrow(\epsilon, \delta)$ uses $O(\log(1/\delta)) \cdot (k/\epsilon)^2$ samples from a k -modal distribution p over $[n]$, and behaves as follows:

- (Completeness) If p is non-decreasing, then T^\uparrow outputs “yes” with probability at least $1 - \delta$;
- (Soundness) If p is ϵ -far from non-decreasing, then T^\uparrow outputs “yes” with probability at most δ .

Let T^\downarrow denote the analogous algorithm for testing whether a k -modal distribution over $[n]$ is non-increasing (we will need both algorithms). The description and proof of correctness for T^\uparrow is postponed to the following subsection (Section 3.4).

3.3 Algorithm `Learn-kmodal` and its analysis. Algorithm `Learn-kmodal` is given below with its analysis following.

`Learn-kmodal`

Inputs: $\epsilon > 0$; sample access to k -modal distribution p over $[n]$

1. Fix $\tau := \epsilon/(100k)$. Draw $r = \Theta(1/\tau^2)$ samples from p and let \hat{p} denote the empirical distribution.
2. Greedily partition the domain $[n]$ into ℓ atomic intervals $\mathcal{I} := \{I_i\}_{i=1}^\ell$ as follows: $I_1 := [1, j_1]$, where $j_1 := \min\{j \in [n] \mid \hat{p}([1, j]) \geq \epsilon/(10k)\}$. For $i \geq 1$, if $\cup_{j=1}^i I_j = [1, j_i]$, then $I_{i+1} := [j_i + 1, j_{i+1}]$, where j_{i+1} is defined as follows: If $\hat{p}([j_i + 1, n]) \geq \epsilon/(10k)$, then $j_{i+1} := \min\{j \in [n] \mid \hat{p}([j_i + 1, j]) \geq \epsilon/(10k)\}$, otherwise, $j_{i+1} := n$.
3. Set $\tau' := \epsilon/(2000k)$. Draw $r' = \Theta((k^3/\epsilon^3) \cdot \log(1/\tau') \log \log(1/\tau'))$ samples s from p to use in Steps 4-5.
4. Run both $T^\uparrow(\epsilon, \tau')$ and $T^\downarrow(\epsilon, \tau')$ over $p_{\cup_{i=1}^j I_i}$ for $j = 1, 2, \dots$, to find the leftmost atomic interval I_{j_1} such that both T^\uparrow and T^\downarrow return “no” over $p_{\cup_{i=1}^{j_1} I_i}$.

Let $I_{j_1} = [a_{j_1}, b_{j_1}]$. We consider two cases:

Case 1: If $\hat{p}[a_{j_1}, b_{j_1}] \geq 2\epsilon/(10k)$, define $I'_{j_1} := [a_{j_1}, b_{j_1} - 1]$ and b_{j_1} is a heavy point.

Case 2: If $\hat{p}[a_{j_1}, b_{j_1}] < 2\epsilon/(10k)$ then define $I'_{j_1} := I_{j_1}$.

Call I'_{j_1} a negligible interval. If $j_1 > 1$ then define the first superinterval S_1 to be $\cup_{i=1}^{j_1-1} I_i$, and set $a_1 \in \{\uparrow, \downarrow\}$ to be $a_1 = \uparrow$ if T^\uparrow returned “yes” on $p_{\cup_{i=1}^{j_1-1} I_i}$ and to be $a_1 = \downarrow$ if T^\downarrow returned “yes” on $p_{\cup_{i=1}^{j_1-1} I_i}$.

5. Repeat Step 3 starting with the next interval I_{j_1+1} , i.e. find the leftmost atomic interval I_{j_2} such that both T^\uparrow and T^\downarrow return “no” over $p_{\cup_{i=j_1+1}^{j_2} I_i}$. Continue doing this until all intervals through I_ℓ have been used.

Let S_1, \dots, S_t be the superintervals obtained through the above process and $(a_1, \dots, a_t) \in \{\uparrow, \downarrow\}^t$ be the corresponding string of bits.

6. Draw $m = \Theta(k \cdot \log(n/k)/\epsilon^3)$ samples s' from p . For each superinterval S_i , $i \in [t]$, run A^{a_i} on the conditional distribution p_{S_i} of p using the samples in $s' \cap S_i$. Let \tilde{p}_{S_i} be the hypothesis thus obtained.
7. Output the hypothesis $h = \sum_{i=1}^t \hat{p}(S_i) \cdot \tilde{p}_{S_i} + \sum_j \hat{p}(\{b_j\}) \cdot \mathbf{1}_{b_j}$.

We are now ready to prove Theorem 3.2.

Proof. [of Theorem 3.2] Before entering into the proof we

record two observations; we state them explicitly here for the sake of the exposition.

FACT 3.2. *Let $R \subseteq [n]$. If p_R is neither non-increasing nor non-decreasing, then R contains at least one left extreme point.*

FACT 3.3. *Suppose that $R \subseteq [n]$ does not contain a left extreme point. For any ϵ, τ , if $T^\uparrow(\epsilon, \tau)$ and $T^\downarrow(\epsilon, \tau)$ are both run on p_R , then the probability that both calls return “no” is at most τ .*

Proof. By Fact 3.2 p_R is either non-decreasing or non-increasing. If p_R is non-decreasing then T^\uparrow will output “no” with probability at most τ , and similarly, if p_R is non-increasing then T^\downarrow will output “no” with probability at most τ . ■

Since $r = \Theta(1/\tau^2)$ samples are drawn in the first step, the DKW inequality implies that with probability of failure at most $1/100$ each interval $I \subseteq [n]$ has $|\widehat{p}(I) - p(I)| \leq 2\tau$. For the rest of the proof we condition on this good event.

Since every atomic interval $I \in \mathcal{I}$ has $\widehat{p}(I) \geq \epsilon/(10k)$ (except potentially the rightmost one), it follows that the number ℓ of atomic intervals constructed in Step 2 satisfies $\ell \leq 10 \cdot (k/\epsilon)$. Moreover, by the DKW inequality, each atomic interval I_i has $p(I_i) \geq 8\epsilon/(100k)$.

Note that in Case (1) of Step 4, if $\widehat{p}[a_{j_1}, b_{j_1}] \geq 2\epsilon/(10k)$ then it must be the case that $\widehat{p}(b_{j_1}) \geq \epsilon/(10k)$ (and thus $p(b_{j_1}) \geq 8\epsilon/(100k)$). In this case, by definition of how the interval I_{j_1} was formed, we must have that $I'_{j_1} = [a_{j_1}, b_{j_1} - 1]$ satisfies $\widehat{p}(I'_{j_1}) < \epsilon/(10k)$. So both in Case 1 and Case 2, we now have that $\widehat{p}(I'_{j_1}) \leq 2\epsilon/(10k)$, and thus $p(I'_{j_1}) \leq 22\epsilon/(100k)$. Entirely similar reasoning shows that every negligible interval constructed in Steps 4 and 5 has mass at most $22\epsilon/(100k)$ under p .

In Steps 4–5 we invoke the testers T^\downarrow and T^\uparrow on the conditional distributions of (unions of contiguous) atomic intervals. Note that we need enough samples in every atomic interval, since otherwise the testers provide no guarantees. We claim that with probability at least $99/100$ over the sample s of Step 3, each atomic interval gets $b = \Omega((k/\epsilon)^2 \cdot \log(1/\tau'))$ samples. This follows by a standard coupon collector’s argument, which we now provide. As argued above, each atomic interval has probability mass $\Omega(\epsilon/k)$ under p . So, we have $\ell = O(k/\epsilon)$ bins (atomic intervals), and we want each bin to contain b balls (samples). It is well-known [NS60] that after taking $\Theta(\ell \cdot \log \ell + \ell \cdot b \cdot \log \log \ell)$ samples from p , with probability $99/100$ each bin will contain the desired number of balls. The claim now follows by our choice of parameters. Conditioning on this event, any execution of the testers $T^\uparrow(\epsilon, \tau')$ and $T^\downarrow(\epsilon, \tau')$ in Steps 4 and 5 will have the guaranteed completeness and soundness properties.

In the execution of Steps 4 and 5, there are a total of at most ℓ occasions when $T^\uparrow(\epsilon, \tau')$ and $T^\downarrow(\epsilon, \tau')$ are

both run over some union of contiguous atomic intervals. By Fact 3.3 and a union bound, the probability that (in any of these instances the interval does not contain a left extreme point and yet both calls return “no”) is at most $(10k/\epsilon)\tau' \leq 1/200$. So with failure probability at most $1/200$ for this step, each time Step 4 identifies a group of consecutive intervals I_j, \dots, I_{j+r} such that both T^\uparrow and T^\downarrow output “no”, there is a left extreme point in $\cup_{i=j}^{j+r} I_i$. Since p is k -modal, it follows that with failure probability at most $1/200$ there are at most $k + 1$ total repetitions of Step 4, and hence the number t of superintervals obtained is at most $k + 1$.

We moreover claim that with very high probability each of the t superintervals S_i is very close to non-increasing or non-decreasing (with its correct orientation given by a_i):

CLAIM 2. *With failure probability at most $1/100$, each $i \in [t]$ satisfies the following: if $a_i = \uparrow$ then p_{S_i} is ϵ -close to a non-decreasing distribution and if $a_i = \downarrow$ then p_{S_i} is ϵ -close to a non-increasing distribution.*

Proof. There are at most $2\ell \leq 20k/\epsilon$ instances when either T^\downarrow or T^\uparrow is run on a union of contiguous intervals. For any fixed execution of T^\downarrow over an interval I , the probability that T^\downarrow outputs “yes” while p_I is ϵ -far from every non-increasing distribution over I is at most τ' , and similarly for T^\uparrow . A union bound and the choice of τ' conclude the proof of the claim. ■

Thus we have established that with overall failure probability at most $5/100$, after Step 5 the interval $[n]$ has been partitioned into:

1. A set $\{S_i\}_{i=1}^t$ of $t \leq k + 1$ superintervals, with $p(S_i) \geq 8\epsilon/(100k)$ and p_{S_i} being ϵ -close to either non-increasing or non-decreasing according to the value of bit a_i .
2. A set $\{I'_i\}_{i=1}^{t'}$ of $t' \leq k + 1$ negligible intervals, such that $p(I'_i) \leq 22\epsilon/(100k)$.
3. A set $\{b_i\}_{i=1}^{t''}$ of $t'' \leq k + 1$ heavy points, each with $p(b_i) \geq 8\epsilon/(100k)$.

We condition on the above good events, and bound from above the expected total variation distance (over the sample s'). In particular, we have the following lemma:

LEMMA 3.1. *We have that $\mathbf{E}_{s'}[d_{TV}(h, p)] \leq O(\epsilon)$.*

Proof. (of Lemma 3.1) By the discussion preceding the lemma statement, the domain $[n]$ has been partitioned into a set of superintervals, a set of negligible intervals and a set of heavy points. As a consequence, we can write

$$p = \sum_{j=1}^t p(S_j) \cdot p_{S_j} + \sum_{j=1}^{t''} p(\{b_j\}) \cdot \mathbf{1}_{b_j} + \sum_{j=1}^{t'} p(I'_j) \cdot p_{I'_j}.$$

Therefore, we can bound the total variation distance as follows:

$$d_{\text{TV}}(h, p) \leq \sum_{j=1}^t |\hat{p}(S_j) - p(S_j)| + \sum_{j=1}^{t''} |\hat{p}(b_j) - p(b_j)| \\ + \sum_{j=1}^{t'} p(I'_j) + \sum_{j=1}^t p(S_j) \cdot d_{\text{TV}}(\tilde{p}_{S_j}, p_{S_j}).$$

Recall that each term in the first two sums is bounded from above by 2τ . Hence, the contribution of these terms to the RHS is at most $2\tau \cdot (2k+2) \leq \epsilon/10$. Since each negligible interval I'_j has $p(I'_j) \leq 22\epsilon/(100k)$, the contribution of the third sum is at most $t' \cdot 22\epsilon/(100k) \leq \epsilon/4$. It thus remains to bound the contribution of the last sum.

We will show that

$$\mathbf{E}_{s'} \left[\sum_{j=1}^t p(S_j) \cdot d_{\text{TV}}(\tilde{p}_{S_j}, p_{S_j}) \right] \leq O(\epsilon).$$

Denote $n_i = |S_i|$. Clearly, $\sum_{i=1}^t n_i \leq n$. Since we are conditioning on the good events (1)-(3), each superinterval is ϵ -close to monotone with a known orientation (non-increasing or non-decreasing) given by a_i . Hence we may apply Theorem 2.2 for each superinterval.

Recall that in Step 5 we draw a total of m samples. Let $m_i, i \in [t]$ be the number of samples that land in S_i ; observe that m_i is a binomially distributed random variable with $m_i \sim \text{Bin}(m, p(S_i))$. We apply Theorem 2.2 for each ϵ -monotone interval, conditioning on the value of m_i , and get

$$d_{\text{TV}}(\tilde{p}_{S_i}, p_{S_i}) \leq 2\epsilon + O\left((\log n_i / (m_i + 1))^{1/3}\right).$$

Hence, we can bound from above the desired expectation as follows

$$\sum_{j=1}^t p(S_j) \cdot \mathbf{E}_{s'} [d_{\text{TV}}(\tilde{p}_{S_j}, p_{S_j})] \leq \left(\sum_{j=1}^t 2\epsilon \cdot p(S_j) \right) + \\ O\left(\sum_{j=1}^t p(S_j) \cdot (\log n_j)^{1/3} \cdot \mathbf{E}_{s'}[(m_j + 1)^{-1/3}] \right).$$

Since $\sum_j p(S_j) \leq 1$, to prove the lemma, it suffices to show that the second term is bounded, i.e. that

$$\sum_{j=1}^t p(S_j) \cdot (\log n_j)^{1/3} \cdot \mathbf{E}_{s'}[(m_j + 1)^{-1/3}] = O(\epsilon).$$

To do this, we will first need the following claim:

CLAIM 3. *For a binomial random variable $X \sim \text{Bin}(m, q)$ it holds $\mathbf{E}[(X+1)^{-1/3}] < (mq)^{-1/3}$.*

Proof. Jensen's inequality implies that

$$\mathbf{E}[(X+1)^{-1/3}] \leq (\mathbf{E}[1/(X+1)])^{1/3}.$$

We claim that $\mathbf{E}[1/(X+1)] < 1/\mathbf{E}[X]$. This can be shown as follows: We first recall that $\mathbf{E}[X] = m \cdot q$. For the expectation of the inverse, we can write:

$$\mathbf{E}[1/(X+1)] = \\ = \sum_{j=0}^m \frac{1}{j+1} \binom{m}{j} q^j (1-q)^{m-j} \\ = \frac{1}{m+1} \cdot \sum_{j=0}^m \binom{m+1}{j+1} q^j (1-q)^{m-j} \\ = \frac{1}{q \cdot (m+1)} \cdot \sum_{i=1}^{m+1} \binom{m+1}{i} q^i (1-q)^{m+1-i} \\ = \frac{1 - (1-q)^{m+1}}{q \cdot (m+1)} < \frac{1}{m \cdot q}.$$

The claim now follows by the monotonicity of the mapping $x \rightarrow x^{1/3}$. ■

By Claim 3, applied to $m_i \sim \text{Bin}(m, p(S_i))$, we have that $\mathbf{E}_{s'}[(m_i + 1)^{-1/3}] < m^{-1/3} \cdot (p(S_i))^{-1/3}$. Therefore, our desired quantity can be bounded from above by

$$\sum_{j=1}^t \frac{p(S_j) \cdot (\log n_j)^{1/3}}{m^{1/3} \cdot (p(S_j))^{1/3}} = \\ O(\epsilon) \cdot \sum_{j=1}^t (p(S_j))^{2/3} \cdot \left(\frac{\log n_j}{k \cdot \log(n/k)} \right)^{1/3}.$$

We now claim that the second term in the RHS above is upper bounded by 2. Indeed, this follows by an application of Hölder's inequality for the vectors $(p(S_j)^{2/3})_{j=1}^t$ and $((\frac{\log n_j}{k \cdot \log(n/k)})^{1/3})_{j=1}^t$, with Hölder conjugates 3/2 and 3. That is,

$$\sum_{j=1}^t (p(S_j))^{2/3} \cdot \left(\frac{\log n_j}{k \cdot \log(n/k)} \right)^{1/3} \leq \\ \leq \left(\sum_{j=1}^t p(S_j) \right)^{2/3} \cdot \left(\sum_{j=1}^t \frac{\log n_j}{k \cdot \log(n/k)} \right)^{1/3} \\ \leq 2.$$

The first inequality is Hölder and the second uses the fact that $\sum_{j=1}^t p(S_j) \leq 1$ and $\sum_{j=1}^t \log(n_j) \leq t \cdot \log(n/t) \leq (k+1) \cdot \log(n/k)$. This last inequality is a consequence of the concavity of the logarithm and the fact that $\sum_j n_j \leq n$. This completes the proof of the Lemma. ■

By applying Markov's inequality and a union bound, we get that with probability 9/10 the algorithm `Learn-kmodal` outputs a hypothesis h that has $d_{\text{TV}}(h, p) \leq O(\epsilon)$ as required.

It is clear that the algorithm has the claimed sample complexity. The running time is also easy to analyze, as it is easy to see that every step can be performed in polynomial time (in fact, nearly linear time) in the sample size. This completes the proof of Theorem 3.2. ■

3.4 Testing whether a k -modal distribution is monotone. In this section we describe and analyze the testing algorithm T^\dagger . Given sample access to a k -modal distribution q over $[n]$ and $\tau > 0$, our tester T^\dagger uses $O(k^2/\tau^2)$ many samples from q and has the following properties:

- If q is non-decreasing, T^\dagger outputs “yes” with probability at least $2/3$.
- If q is τ -far from non-decreasing, T^\dagger outputs “no” with probability at least $2/3$.

(The algorithm $T^\dagger(\tau, \delta)$ is obtained by repeating T^\dagger $O(\log(1/\delta))$ times and taking the majority vote.)

Tester $T^\dagger(\tau)$

Inputs: $\tau > 0$; sample access to k -modal distribution q over $[n]$

1. Fix $\delta := \tau/(100k)$. Draw $r = \Theta(1/\delta^2)$ samples \mathbf{s} from q and let \hat{q} be the resulting empirical distribution.
2. If there exist $a \leq b < c \in \mathbf{s} \cup \{1, n\}$ such that

$$(3.2) \quad \begin{aligned} \hat{E}(a, b, c) &:= \frac{\hat{q}([a, b])}{(b-a+1)} - \frac{\hat{q}([b+1, c])}{(c-b)} \\ &\geq \frac{(\tau/4k)}{(b-a+1)} + \frac{(\tau/4k)}{(c-b)} \end{aligned}$$

then output “no”, otherwise output “yes”.

The idea behind tester T^\dagger is simple. It is based on the observation that if q is a non-decreasing distribution, then for any two consecutive intervals $[a, b]$ and $[b+1, c]$ the average of q over $[b+1, c]$ must be at least as large as the average of q over $[a, b]$. Thus any non-decreasing distribution will pass a test that checks “all” pairs of consecutive intervals looking for a violation. Our analysis shows that in fact such a test is complete as well as sound if the distribution q is guaranteed to be k -modal. The key ingredient is the structural Lemma 9 below, which is proved using a procedure (reminiscent of Myerson ironing [Mye81]) to convert a k -modal distribution to a non-decreasing distribution.

The following theorem establishes correctness of the tester.

THEOREM 3.3. *The algorithm T^\dagger uses $O(k^2/\tau^2)$ samples from q , performs $\text{poly}(k/\tau) \cdot \log n$ bit operations and satisfies the desired completeness and soundness properties.*

Proof. The upper bound on the sample complexity is straightforward, since only Step 1 uses samples. It is also easy to see that a straightforward implementation of the algorithm runs in time $\text{poly}(k/\tau) \cdot \log n$. Below we prove that the algorithm has the claimed soundness and completeness properties.

Let us say that the sample \mathbf{s} is *good* if every interval $I \subseteq [n]$ has $|\hat{q}(I) - q(I)| \leq 2\delta$. By the DKW inequality, with probability at least $2/3$ the sample \mathbf{s} is good. Assuming that \mathbf{s} is good, we have that for any $a \leq b < c \in [n]$ the quantity

$$E(a, b, c) := \frac{q([a, b])}{(b-a+1)} - \frac{q([b+1, c])}{(c-b)}$$

differs from its empirical value $\hat{E}(a, b, c)$ (i.e. the LHS of (3.2)) by at most $\gamma(a, b, c) := \frac{2\delta}{(b-a+1)} + \frac{2\delta}{(c-b)}$. That is,

$$(3.3) \quad |E(a, b, c) - \hat{E}(a, b, c)| \leq \gamma(a, b, c).$$

We first show completeness. If q is non-decreasing the average probability value in any interval $[a, b]$ is a non-decreasing function of a . That is, for all $a \leq b < c \in [n]$ it holds $E(a, b, c) \leq 0$. Therefore, with probability at least $2/3$, it holds $\hat{E}(a, b, c) \leq \gamma(a, b, c)$ and the tester says “yes”.

For soundness, we need the following lemma:

LEMMA 3.2. *Let q be a k -modal distribution over $[n]$ that is τ -far from being non-decreasing. Then there exists a triple of points $a \leq b < c \in [n]$ such that*

$$(3.4) \quad E(a, b, c) \geq \frac{(\tau/2k)}{(b-a+1)} + \frac{(\tau/2k)}{(c-b)}.$$

We first show how the soundness follows from the lemma. For a k -modal distribution q that is τ -far from non-decreasing, we will argue that if the sample is good then there exists a triple $s_a \leq s_b < s_c \in \mathbf{s} \cup \{1, n\}$ such that $\hat{E}(s_a, s_b, s_c)$ satisfies (3.2).

By Lemma 3.2, there exists a triple $a \leq b < c \in [n]$ satisfying (3.4).

We first note that at least one sample must have landed in $[a, b]$, for otherwise the DKW inequality would give that $q([a, b]) \leq 2\delta$; this in turn would imply that $E(a, b, c) \leq 2\delta/(b-a+1)$, a contradiction, as it violates (3.4). We now define the points s_a, s_b, s_c as follows: (i) s_a is the leftmost point of the sample in $[a, b]$, (ii) s_b is the rightmost point of the sample in $[a, b]$; and (iii) s_c is either the leftmost point of the sample in $[c+1, n]$, or the rightmost point n of the interval, if $\hat{q}([c+1, n]) = 0$. We will now argue that these points satisfy (3.2). Consider the interval $[s_a, s_b]$. Then, we have that

$$(3.5) \quad \begin{aligned} \frac{\hat{q}([s_a, s_b])}{s_b - s_a + 1} &\geq \frac{\hat{q}([s_a, s_b])}{b - a + 1} \\ &= \frac{\hat{q}([a, b])}{b - a + 1} \\ &\geq \frac{q([a, b])}{b - a + 1} - \frac{2\delta}{b - a + 1} \end{aligned}$$

where the first inequality uses the fact that $[s_a, s_b] \subseteq [a, b]$, the equality uses the definition of a and b , and the final inequality follows by an application of the DKW inequality

for the interval $[a, b]$. An analogous argument can be applied for the interval $[s_b, s_c]$. Indeed, we have that

$$(3.6) \quad \begin{aligned} \frac{\widehat{q}([s_b + 1, s_c])}{s_c - s_b + 1} &\leq \frac{\widehat{q}([s_b + 1, s_c])}{c - b} \\ &= \frac{\widehat{q}([b + 1, c])}{c - b} \\ &\leq \frac{q([b + 1, c])}{c - b} + \frac{2\delta}{c - b} \end{aligned}$$

where the first inequality follows from the fact that $[s_b, s_c] \supseteq [b + 1, c]$, the equality uses the definition of b and c , and the final inequality follows by an application of the DKW inequality for the interval $[b + 1, c]$.

A combination of (3.4), (3.5), (3.6) yields the desired result. It thus remains to prove Lemma 3.2.

Proof. [Lemma 3.2] We prove the contrapositive. Let q be a k -modal distribution such that for all $a \leq b < c \in [n]$

$$(3.7) \quad E(a, b, c) \leq \frac{(\tau/2k)}{(b - a + 1)} + \frac{(\tau/2k)}{(c - b)}.$$

We will show that q is τ -close to being non-decreasing by constructing a non-decreasing distribution \tilde{q} that is τ -close to q . The construction of \tilde{q} proceeds in k stages where in each stage, we reduce the number of modes by at least one and incur error in variation distance at most τ/k . That is, we iteratively construct a sequence of distributions $\{q^{(i)}\}_{i=0}^k$, $q^{(0)} = q$ and $q^{(k)} = \tilde{q}$, such that for all $i \in [k]$ we have that $q^{(i)}$ is $(k - i)$ -modal and $d_{TV}(q^{(i-1)}, q^{(i)}) \leq \tau/k$.

Consider the graph (histogram) of the discrete density q . The x -axis represents the n points of the domain and the y -axis the corresponding probabilities. We first informally describe how to obtain $q^{(1)}$ from q . The construction of $q^{(i+1)}$ from $q^{(i)}$ is identical. Let j be the leftmost left-extreme point (mode) of q , and assume that it is a local maximum with height (probability mass) $q(j)$. (A symmetric argument works for the case that it is a local minimum.) The idea of the proof is based on the following simple process (reminiscent of Myerson's ironing process [Mye81]): We start with the horizontal line $y = q(j)$ and move it downwards until we reach a height $h_0 < q(j)$ so that the total mass "cut-off" equals the mass "missing" to the right; then make the distribution "flat" in the corresponding interval (hence, reducing the number of modes by at least one). The resulting distribution is $q^{(1)}$ and equation (3.7) implies that $d_{TV}(q^{(1)}, q) \leq \tau/k$.

We now proceed with the formal argument, assuming as above that the leftmost left-extreme point j of q is a local maximum. We say that the line $y = h$ intersects a point $i \in [n]$ in the domain of q if $q(i) \geq h$. The line $y = h$, $h \in [0, q(j)]$, intersects the graph of q at a unique interval $I(h) \subseteq [n]$ that contains j . Suppose $I(h) = [a(h), b(h)]$, where $a(h), b(h) \in [n]$ depend on h . By definition this

means that $q(a(h)) \geq h$ and $q(a(h) - 1) < h$. Recall that the distribution q is non-decreasing in the interval $[1, j]$ and that $j \geq a(h)$. The term "the mass cut-off by the line $y = h$ " means the quantity $A(h) = q(I(h)) - h \cdot (b(h) - a(h) + 1)$, i.e. the "mass of the interval $I(h)$ above the line".

The height h of the line $y = h$ defines the points $a(h), b(h) \in [n]$ as described above. We consider values of h such that q is unimodal (increasing then decreasing) over $I(h)$. In particular, let j' be the leftmost mode of q to the right of j , i.e. $j' > j$ and j' is a local minimum. We consider values of $h \in (q(j'), q(j))$. For such values, the interval $I(h)$ is indeed unimodal (as $b(h) < j'$). For $h \in (q(j'), q(j))$ we define the point $c(h) \geq j'$ as follows: It is the rightmost point of the largest interval containing j' whose probability mass does not exceed h . That is, all points in $[j', c(h)]$ have probability mass at most h and $q(c(h) + 1) > h$ (or $c(h) = n$).

Consider the interval $J(h) = [b(h) + 1, c(h)]$. This interval is non-empty, since $b(h) < j' \leq c(h)$. (Note that $J(h)$ is not necessarily a unimodal interval; it contains at least one mode j' , but it may also contain more modes.) The term "the mass missing to the right of the line $y = h$ " means the quantity $B(h) = h \cdot (c(h) - b(h)) - q(J(h))$.

Consider the function $C(h) = A(h) - B(h)$ over $[q(j'), q(j)]$. This function is continuous in its domain; moreover, we have that $C(q(j)) = A(q(j)) - B(q(j)) < 0$, as $A(q(j)) = 0$, and $C(q(j')) = A(q(j')) - B(q(j')) > 0$, as $B(q(j')) = 0$. Therefore, by the intermediate value theorem, there exists a value $h_0 \in (q(j'), q(j))$ such that $A(h_0) = B(h_0)$.

The distribution $q^{(1)}$ is constructed as follows: We move the mass $\tau' = A(h_0)$ from $I(h_0)$ to $J(h_0)$. Hence, it follows that $d_{TV}(q^{(1)}, q) \leq 2\tau'$. We also claim that $q^{(1)}$ has at least one mode less than q . Indeed, $q^{(1)}$ is non-decreasing in $[1, a(h) - 1]$ and constant in $[a(h), c(h)]$. (All the points in the latter interval have probability mass exactly h_0 .) Recalling that $q^{(1)}(a(h)) = h_0 \geq q^{(1)}(a(h) - 1) = q(a(h) - 1)$, we deduce that $q^{(1)}$ is non-decreasing in $[1, c(h)]$.

We will now argue that $\tau' \leq \tau/(2k)$ which completes the proof of the lemma. To this end we use our starting assumption, equation (3.7). Recall that we have $A(h_0) = B(h_0) = \tau'$, which can be written as

$$\begin{aligned} q([a(h), b(h)]) - h_0 \cdot (b(h) - a(h) + 1) \\ = h_0 \cdot (c(h) - b(h)) - q([b(h) + 1, c(h)]) \\ = \tau'. \end{aligned}$$

From this, we get

$$\begin{aligned} \frac{q([a(h), b(h)])}{(b(h) - a(h) + 1)} - \frac{q([b(h) + 1, c(h)])}{(c(h) - b(h))} \\ = \frac{\tau'}{(b(h) - a(h) + 1)} + \frac{\tau'}{(c(h) - b(h))}. \end{aligned}$$

Combining with (3.7) proves Lemma 3.2. ■

This completes the proof of Theorem 3.3. ■

4 Conclusions and future work

At the level of techniques, this work illustrates the viability of a new general strategy for developing efficient learning algorithms, namely by using “inexpensive” property testers to decompose a complex object (for us these objects are k -modal distributions) into simpler objects (for us these are monotone distributions) that can be more easily learned. It would be interesting to apply this paradigm in other contexts such as learning Boolean functions.

At the level of the specific problem we consider – learning k -modal distributions – our results show that k -modality is a useful type of structure which can be strongly exploited by sample-efficient and computationally efficient learning algorithms. Our results motivate the study of computationally efficient learning algorithms for distributions that satisfy other kinds of “shape restrictions.” Possible directions here include multivariate k -modal distributions, log-concave distributions, monotone hazard rate distributions and more.

Finally, at a technical level, any improvement in the sample complexity of our property testing algorithm of Section 3.4 would directly improve the “extraneous” additive $\tilde{O}((k/\epsilon)^3)$ term in the sample complexity of our algorithm. We suspect that it may be possible to improve our testing algorithm (although we note that it is easy to give an $\Omega(\sqrt{k})$ lower bound using standard constructions).

References

- [Bir87a] L. Birgé. Estimating a density under order restrictions: Nonasymptotic minimax risk. *Annals of Statistics*, 15(3):995–1012, 1987.
- [Bir87b] L. Birgé. On the risk of histograms for estimating decreasing densities. *Annals of Statistics*, 15(3):1013–1022, 1987.
- [Bir97] L. Birgé. Estimation of unimodal densities without smoothness assumptions. *Annals of Statistics*, 25(3):970–981, 1997.
- [BKR04] Tugkan Batu, Ravi Kumar, and Ronitt Rubinfeld. Sublinear algorithms for testing monotone and unimodal distributions. In *ACM Symposium on Theory of Computing*, pages 381–390, 2004.
- [CKC83] L. Cobb, P. Koppstein, and N.H. Chen. Estimation and moment recursion relations for multimodal distributions of the exponential family. *J. American Statistical Association*, 78(381):124–130, 1983.
- [CT04] K.S. Chan and H. Tong. Testing for multimodality with dependent data. *Biometrika*, 91(1):113–123, 2004.
- [DDS11] C. Daskalakis, I. Diakonikolas, and R.A. Servedio. Learning poisson binomial distributions. Arxiv report, 2011, 2011.
- [DKW56] A. Dvoretzky, J. Kiefer, and J. Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Ann. Mathematical Statistics*, 27(3):642–669, 1956.
- [DL96a] L. Devroye and G. Lugosi. Nonasymptotic universal smoothing factors, kernel complexity and Yatracos classes. *Annals of Statistics*, 25:2626–2637, 1996.
- [DL96b] L. Devroye and G. Lugosi. A universally acceptable smoothing factor for kernel density estimation. *Annals of Statistics*, 24:2499–2512, 1996.
- [DL01] L. Devroye and G. Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics, Springer, 2001.
- [dTf90] G.A. de Toledo and J.M. Fernandez. Patch-clamp measurements reveal multimodal distribution of granule sizes in rat mast cells. *Journal of Cell Biology*, 110(4):1033–1039, 1990.
- [FPP⁺98] F.R. Ferraro, B. Paltrinieri, F.F. Pecci, R.T. Rood, and B. Dorman. Multimodal distributions along the horizontal branch. *The Astrophysical Journal*, 500:311–319, 1998.
- [GGR98] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45:653–750, 1998.
- [Gol11] O. Goldreich. Highlights of the Bertinoro workshop on Sublinear Algorithms (unpublished comments). Posted at <http://www.wisdom.weizmann.ac.il/oded/MC/072.html>, accessed June 17, 2011, 2011.
- [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.
- [Kem91] J.H.B. Kemperman. Mixtures with a limited number of modal intervals. *Annals of Statistics*, 19(4):2120–2144, 1991.
- [KR00] M. Kearns and D. Ron. Testing problems with sub-learning sample complexity. *J. Comp. Sys. Sci.*, 61:428–456, 2000.
- [Mas90] P. Massart. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *Annals of Probability*, 18(3):1269–1283, 1990.
- [Mur64] E.A. Murphy. One cause? many causes?: The argument from the bimodal distribution. *J. Chronic Diseases*, 17(4):301–324, 1964.
- [Mye81] R.B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [NS60] D. J. Newman and L. Shepp. The double dixie cup problem. *The American Mathematical Monthly*, 67(1):pp. 58–61, 1960.
- [Rao69] B.L.S. Prakasa Rao. Estimation of a unimodal density. *Sankhya Ser. A*, 31:23–36, 1969.
- [Ron08] D. Ron. Property Testing: A Learning Theory Perspective. *Foundations and Trends in Machine Learning*, 1(3):307–402, 2008.
- [Weg70] E.J. Wegman. Maximum likelihood estimation of a unimodal density. I. and II. *Ann. Math. Statist.*, 41:457–471, 2169–2174, 1970.
- [Yat85] Y. G. Yatracos. Rates of convergence of minimum distance estimators and Kolmogorov’s entropy. *Annals of Statistics*, 13:768–774, 1985.

A Birgé’s algorithm as a semi-agnostic learner

In this section we briefly explain why Birgé’s algorithm [Bir87b] also works in the semi-agnostic setting. To do this, we need to explain his approach. For this, we will need the following theorem, which gives a tight bound on the number of samples required to learn an arbitrary distribution with respect to *total variation distance*.

THEOREM A.1. (FOLKLORE) *Let p be any distribution over $[n]$. We have: $\mathbf{E}[d_{\text{TV}}(p, \hat{p}_m)] \leq 2\sqrt{n/m}$.*

Let p be a non-increasing distribution over $[n]$. (The analysis for the non-decreasing case is identical.) Conceptually, we view algorithm L^\perp as working in three steps:

- In the first step, it partitions the set $[n]$ into a carefully chosen set I_1, \dots, I_ℓ of consecutive intervals, with $\ell = O(m^{1/3} \cdot (\log n)^{2/3})$. Consider the *flattened* distribution p_f over $[n]$ obtained from p by averaging the weight that p assigns to each interval over the entire interval. That is, for $j \in [\ell]$ and $i \in I_j$, $p_f(i) = \sum_{t \in I_j} p(t)/|I_j|$. Then a simple argument given in [Bir87b] gives that $d_{\text{TV}}(p_f, p) = O((\log n/(m+1))^{1/3})$.
- Let p_r be the *reduced* distribution corresponding to p and the partition I_1, \dots, I_ℓ . That is, p_r is a distribution over $[\ell]$ with $p_r(i) = p(I_i)$ for $i \in [\ell]$. In the second step, the algorithm uses the m samples to learn p_r . (Note that p_r is not necessarily monotone.) After m samples, one obtains a hypothesis \hat{p}_r such that $\mathbf{E}[d_{\text{TV}}(p_r, \hat{p}_r)] = O(\sqrt{\ell/m}) = O((\log n/(m+1))^{1/3})$. The first equality follows from Theorem A.1 (since p_r is distribution over ℓ elements) and the second inequality follows from the choice of ℓ .
- Finally, the algorithm outputs the flattened hypothesis $(\hat{p}_r)_f$ over $[n]$ corresponding to \hat{p}_r , i.e. obtained by \hat{p}_r by subdividing the mass of each interval uniformly within the interval. It follows from the above two steps that $\mathbf{E}[d_{\text{TV}}((\hat{p}_r)_f, p_f)] = O((\log n/(m+1))^{1/3})$.
- The combination of the first and third steps yields that $\mathbf{E}[d_{\text{TV}}((\hat{p}_r)_f, p)] = O((\log n/(m+1))^{1/3})$.

The above arguments are entirely due to Birgé [Bir87b]. We now explain how his analysis can be extended to show that his algorithm is in fact a semi-agnostic learner as claimed in Theorem 2.2. To avoid clutter in the expressions below let us fix $\delta := O((\log n/(m+1))^{1/3})$.

The second and third steps in the algorithm description above are used to learn the distribution p_f to variation distance δ . Note that these steps do not use the assumption that p is non-increasing. The following claim, which generalizes Step 1 above, says that if p is τ -close to non-increasing, the

flattened distribution p_f (defined as above) is $(2\tau + \delta)$ -close to p . Therefore, it follows that, for such a distribution p , algorithm L^\perp succeeds with expected (total variation distance) error $(2\tau + \delta) + \delta$.

We have:

CLAIM 4. *Let p be a distribution over $[n]$ that is τ -close to non-increasing. Then, the flattened distribution p_f (obtained from p by averaging its weight on every interval I_j) satisfies $d_{\text{TV}}(p_f, p) \leq (2\tau + \delta)$.*

Proof. Let p^\perp be the non-increasing distribution that is τ -close to p . Let τ_j denote the L_1 -distance between p and p^\perp in the interval I_j . Then, we have that

$$(A.1) \quad \sum_{j=1}^{\ell} \tau_j \leq \tau.$$

By Birgé’s arguments, it follows that the flattened distribution $(p^\perp)_f$ corresponding to p^\perp is δ -close to p^\perp , hence $(\tau + \delta)$ -close to p . That is,

$$(A.2) \quad d_{\text{TV}}((p^\perp)_f, p) \leq \tau + \delta.$$

We want to show that

$$(A.3) \quad d_{\text{TV}}((p^\perp)_f, p_f) \leq \tau.$$

Assuming (A.3) holds, we can conclude by the triangle inequality that

$$d_{\text{TV}}(p, p_f) \leq 2\tau + \delta$$

as desired.

Observe that, by assumption, p and p^\perp have L_1 -distance at most τ_j in each I_j interval. In particular, this implies that, for all $j \in [\ell]$, it holds

$$|p(I_j) - p^\perp(I_j)| \leq \tau_j.$$

Now note that, within each interval I_j , p_f and $(p^\perp)_f$ are both uniform. Hence, the contribution of I_j to the variation distance between p_f and $(p^\perp)_f$ is at most $|p(I_j) - p^\perp(I_j)|$.

Therefore, by (A.1) we deduce

$$d_{\text{TV}}(p_f, (p^\perp)_f) \leq \tau$$

which completes the proof of the claim. \blacksquare

B Hypothesis Testing

Our hypothesis testing routine `Choose-Hypothesisp` runs a simple “competition” to choose a winner between two candidate hypothesis distributions h_1 and h_2 over $[n]$ that it is given in the input either explicitly, or in some succinct way. We show that if at least one of the two candidate hypotheses is close to the target distribution p , then with high probability over the samples drawn from p the routine selects as winner a candidate that is close to p . This basic

approach of running a competition between candidate hypotheses is quite similar to the ‘‘Scheffé estimate’’ proposed by Devroye and Lugosi (see [DL96b, DL96a] and Chapter 6 of [DL01]), which in turn built closely on the work of [Yat85], but there are some small differences between our approach and theirs; the [DL01] approach uses a notion of the ‘‘competition’’ between two hypotheses which is not symmetric under swapping the two competing hypotheses, whereas our competition is symmetric.

We now prove Theorem 2.3.

Proof. [of Theorem 2.3] Let \mathcal{W} be the support of p . To set up the competition between h_1 and h_2 , we define the following subset of \mathcal{W} :

$$(B.4) \quad \mathcal{W}_1 = \mathcal{W}_1(h_1, h_2) := \{w \in \mathcal{W} \mid h_1(w) > h_2(w)\}.$$

Let then $p_1 = h_1(\mathcal{W}_1)$ and $q_1 = h_2(\mathcal{W}_1)$. Clearly, $p_1 > q_1$ and $d_{TV}(h_1, h_2) = p_1 - q_1$.

The competition between h_1 and h_2 is carried out as follows:

1. If $p_1 - q_1 \leq 5\epsilon'$, declare a draw and return either h_i . Otherwise:
2. Draw $m = O\left(\frac{\log(1/\delta')}{\epsilon'^2}\right)$ samples s_1, \dots, s_m from p , and let $\tau = \frac{1}{m}|\{i \mid s_i \in \mathcal{W}_1\}|$ be the fraction of samples that fall inside \mathcal{W}_1 .
3. If $\tau > p_1 - \frac{3}{2}\epsilon'$, declare h_1 as winner and return h_1 ; otherwise,
4. if $\tau < q_1 + \frac{3}{2}\epsilon'$, declare h_2 as winner and return h_2 ; otherwise,
5. declare a draw and return either h_i .

It is not hard to check that the outcome of the competition does not depend on the ordering of the pair of distributions provided in the input; that is, on inputs (h_1, h_2) and (h_2, h_1) the competition outputs the same result for a fixed sequence of samples s_1, \dots, s_m drawn from p .

The correctness of Choose-Hypothesis is an immediate consequence of the following lemma.

LEMMA B.1. *Suppose that $d_{TV}(p, h_1) \leq \epsilon'$. Then:*

- (i) *If $d_{TV}(p, h_2) > 6\epsilon'$, then the probability that the competition between h_1 and h_2 does not declare h_1 as the winner is at most $e^{-m\epsilon'^2/2}$. (Intuitively, if h_2 is very bad then it is very likely that h_1 will be declared winner.)*
- (ii) *If $d_{TV}(p, h_2) > 4\epsilon'$, the probability that the competition between h_1 and h_2 declares h_2 as the winner is at most $e^{-m\epsilon'^2/2}$. (Intuitively, if h_2 is only moderately bad then a draw is possible but it is very unlikely that h_2 will be declared winner.)*

Proof. Let $r = p(\mathcal{W}_1)$. The definition of the total variation distance implies that $|r - p_1| \leq \epsilon'$. Let us define the 0/1 (indicator) random variables $\{Z_j\}_{j=1}^m$ as $Z_j = 1$ iff $s_j \in \mathcal{W}_1$. Clearly, $\tau = \frac{1}{m} \sum_{j=1}^m Z_j$ and $\mathbb{E}[\tau] = \mathbb{E}[Z_j] = r$. Since the Z_j 's are mutually independent, it follows from the Chernoff bound that $\Pr[\tau \leq r - \epsilon'/2] \leq e^{-m\epsilon'^2/2}$. Using $|r - p_1| \leq \epsilon'$ we get that $\Pr[\tau \leq p_1 - 3\epsilon'/2] \leq e^{-m\epsilon'^2/2}$.

- For part (i): If $d_{TV}(p, h_2) > 6\epsilon'$, from the triangle inequality we get that $p_1 - q_1 = d_{TV}(h_1, h_2) > 5\epsilon'$. Hence, the algorithm will go beyond Step 1, and with probability at least $1 - e^{-m\epsilon'^2/2}$, it will stop at Step 3, declaring h_1 as the winner of the competition between h_1 and h_2 .
- For part (ii): If $p_1 - q_1 \leq 5\epsilon'$ then the competition declares a draw, hence h_2 is not the winner. Otherwise we have $p_1 - q_1 > 5\epsilon'$ and the above arguments imply that the competition between h_1 and h_2 will declare h_2 as the winner with probability at most $e^{-m\epsilon'^2/2}$.

This concludes the proof of Lemma B.1. ■

The proof of the theorem is now complete. ■

C Using the Hypothesis Tester

In this section, we explain in detail how we use the hypothesis testing algorithm Choose-Hypothesis throughout this paper. In particular, the algorithm Choose-Hypothesis is used in the following places:

- In Step 4 of algorithm Learn-kmodal-simple we need an algorithm $L^\downarrow_{\delta'}$ (resp. $L^\uparrow_{\delta'}$) that learns a non-increasing (resp. non-increasing) distribution within total variation distance ϵ and confidence δ' . Note that the corresponding algorithms L^\downarrow and L^\uparrow provided by Theorem 2.2 have confidence $9/10$. To boost the confidence of L^\downarrow (resp. L^\uparrow) we run the algorithm $O(\log(1/\delta'))$ times and use Choose-Hypothesis in an appropriate tournament procedure to select among the candidate hypothesis distributions.
- In Step 5 of algorithm Learn-kmodal-simple we need to select among two candidate hypothesis distributions (with the promise that at least one of them is close to the true conditional distribution). In this case, we run Choose-Hypothesis once to select between the two candidates.
- Also note that both algorithms Learn-kmodal-simple and Learn-kmodal generate an ϵ -accurate hypothesis with probability $9/10$. We would like to boost the probability of success to $1 - \delta$. To achieve this we again run the corresponding algorithm $O(\log(1/\delta))$ times and use Choose-Hypothesis in an appropriate tournament to select among the candidate hypothesis distributions.

We now formally describe the “tournament” algorithm to boost the confidence to $1 - \delta$.

LEMMA C.1. *Let p be any distribution over a finite set \mathcal{W} . Suppose that \mathcal{D}_ϵ is a collection of N distributions over \mathcal{W} such that there exists $q \in \mathcal{D}_\epsilon$ with $d_{TV}(p, q) \leq \epsilon$. Then there is an algorithm that uses $O(\epsilon^{-2} \log N \log(1/\delta))$ samples from p and with probability $1 - \delta$ outputs a distribution $p' \in \mathcal{D}_\epsilon$ that satisfies $d_{TV}(p, p') \leq 6\epsilon$.*

Devroye and Lugosi (Chapter 7 of [DL01]) prove a similar result by having all pairs of distributions in the cover compete against each other using their notion of a competition, but again there are some small differences: their approach chooses a distribution in the cover which wins the maximum number of competitions, whereas our algorithm chooses a distribution that is never defeated (i.e. won or achieved a draw against all other distributions in the cover). Instead we follow the approach from [DDS11].

Proof. The algorithm performs a tournament by running the competition `Choose-Hypothesisp($h_i, h_j, \epsilon, \delta/(2N)$)` for every pair of distinct distributions h_i, h_j in the collection \mathcal{D}_ϵ . It outputs a distribution $q^* \in \mathcal{D}_\epsilon$ that was never a loser (i.e. won or achieved a draw in all its competitions). If no such distribution exists in \mathcal{D}_ϵ then the algorithm outputs “failure.”

By definition, there exists some $q \in \mathcal{D}_\epsilon$ such that $d_{TV}(p, q) \leq \epsilon$. We first argue that with high probability this distribution q never loses a competition against any other $q' \in \mathcal{D}_\epsilon$ (so the algorithm does not output “failure”). Consider any $q' \in \mathcal{D}_\epsilon$. If $d_{TV}(p, q') > 4\epsilon$, by Lemma B.1(ii) the probability that q loses to q' is at most $2e^{-m\epsilon^2/2} = O(1/N)$. On the other hand, if $d_{TV}(p, q') \leq 4\delta$, the triangle inequality gives that $d_{TV}(q, q') \leq 5\epsilon$ and thus q draws against q' . A union bound over all N distributions in \mathcal{D}_ϵ shows that with probability $1 - \delta/2$, the distribution q never loses a competition.

We next argue that with probability at least $1 - \delta/2$, every distribution $q' \in \mathcal{D}_\epsilon$ that never loses has small variation distance from p . Fix a distribution q' such that $d_{TV}(q', p) > 6\epsilon$; Lemma B.1(i) implies that q' loses to q with probability $1 - 2e^{-m\epsilon^2/2} \geq 1 - \delta/(2N)$. A union bound gives that with probability $1 - \delta/2$, every distribution q' that has $d_{TV}(q', p) > 6\epsilon$ loses some competition.

Thus, with overall probability at least $1 - \delta$, the tournament does not output “failure” and outputs some distribution q^* such that $d_{TV}(p, q^*)$ is at most 6ϵ . This proves the lemma. \blacksquare

We now explain how the above lemma is used in our context: Suppose we perform $O(\log(1/\delta))$ runs of a learning algorithm that constructs an ϵ -accurate hypothesis with probability at least $9/10$. Then, with failure probability at most $\delta/2$, at least one of the hypotheses gen-

erated is ϵ -close to the true distribution in variation distance. Conditioning on this good event, we have a collection of distributions with cardinality $O(\log(1/\delta))$ that satisfies the assumption of the lemma. Hence, using $O((1/\epsilon^2) \cdot \log \log(1/\delta) \cdot \log(1/\delta))$ samples we can learn to accuracy 6ϵ and confidence $1 - \delta/2$. The overall sample complexity is $O(\log(1/\delta))$ times the sample complexity of the (learning algorithm with confidence $9/10$) plus this additional $O((1/\epsilon^2) \cdot \log \log(1/\delta) \cdot \log(1/\delta))$ term.

In terms of running time, we make the following easily verifiable remarks: When the hypothesis testing algorithm `Choose-Hypothesis` is run on a pair of distributions that are produced by Birgé’s algorithm, its running time is polynomial in the succinct description of these distributions, i.e. in $\log^2(n)/\epsilon$. Similarly, when `Choose-Hypothesis` is run on a pair of outputs of `Learn-kmodal-simple` or `Learn-kmodal`, its running time is polynomial in the succinct description of these distributions. More specifically, in the former case, the succinct description has bit complexity $O(k \cdot \log^2(n)/\epsilon^2)$ (since the output consists of $O(k/\epsilon)$ monotone intervals, and the conditional distribution on each interval is the output of Birgé’s algorithm for that interval). In the latter case, the succinct description has bit complexity $O(k \cdot \log^2(n)/\epsilon)$, since the algorithm `Learn-kmodal` constructs only k monotone intervals. Hence, in both cases, each execution of the testing algorithm performs $\text{poly}(k, \log n, 1/\epsilon)$ bit operations. Since the tournament invokes the algorithm `Choose-Hypothesis` $O(\log^2(1/\delta))$ times (for every pair of distributions in our pool of $O(\log(1/\delta))$ candidates) the upper bound on the running time follows.