2013 IEEE 54th Annual Symposium on Foundations of Computer Science

Learning Sums of Independent Integer Random Variables

Constantinos Daskalakis*, Ilias Diakonikolas†, Ryan O'Donnell‡, Rocco A. Servedio§ and Li-Yang Tan§
*MIT
†University of Edinburgh
‡CMU
§Columbia University

***Abstract*—Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be a sum of n independent integer random variables $\boldsymbol{X}_i$, where each $\boldsymbol{X}_i$ is supported on $\{0, 1, \ldots, k-1\}$ but otherwise may have an arbitrary distribution (in particular the $\boldsymbol{X}_i$'s need not be identically distributed). How many samples are required to learn the distribution $\boldsymbol{S}$ to high accuracy? In this paper we show that the answer is *completely independent of* n, and moreover we give a computationally efficient algorithm which achieves this low sample complexity. More precisely, our algorithm learns any such $\boldsymbol{S}$ to ϵ-accuracy (with respect to the total variation distance between distributions) using $\mathrm{poly}(k, 1/\epsilon)$ samples, independent of n. Its running time is $\mathrm{poly}(k, 1/\epsilon)$ in the standard word RAM model. Thus we give a broad generalization of the main result of [DDS12b] which gave a similar learning result for the special case $k = 2$ (when the distribution $\boldsymbol{S}$ is a Poisson Binomial Distribution).**

Prior to this work, no nontrivial results were known for learning these distributions even in the case $k = 3$. A key difficulty is that, in contrast to the case of $k = 2$, sums of independent $\{0, 1, 2\}$-valued random variables may behave very differently from (discretized) normal distributions, and in fact may be rather complicated — they are not log-concave, they can be $\Theta(n)$-modal, there is no relationship between Kolmogorov distance and total variation distance for the class, etc. Nevertheless, the heart of our learning result is a new limit theorem which characterizes what the sum of an arbitrary number of arbitrary independent $\{0, 1, \ldots, k-1\}$-valued random variables may look like. Previous limit theorems in this setting made strong assumptions on the "shift invariance" of the random variables $\boldsymbol{X}_i$ in order to force a discretized normal limit. We believe that our new limit theorem, as the first result for truly arbitrary sums of independent $\{0, 1, \ldots, k-1\}$-valued random variables, is of independent interest.

***Keywords*-limit theorem; discrete distribution learning; sums of independent integer random variables.**

I. Introduction

We study the problem of learning an unknown random variable given access to independent samples drawn from it. This is essentially the problem of *density estimation*, which has received significant attention in the probability and statistics literature over the course of several decades (see e.g. [DG85], [Sil86], [Sco92], [DL01] for introductory books). More recently many works in theoretical computer science have also considered problems of this sort, with an emphasis on developing computationally efficient algorithms (see e.g. [KMR+94], [Das99], [FM99], [DS00], [AK01], [VW02], [CGG02], [BGK04], [DHKS05], [MR05], [FOS05], [FOS06], [BS10], [KMV10], [MV10], [DDS12a], [DDS12b], [RSS12], [AHK12]).

In this paper we work in the following standard learning framework: the learning algorithm is given access to independent samples drawn from the unknown random variable $\boldsymbol{S}$, and it must output a hypothesis random variable $\widetilde{\boldsymbol{S}}$ such that with high probability the total variation distance $d_{\mathrm{TV}}(\boldsymbol{S}, \widetilde{\boldsymbol{S}})$ between $\boldsymbol{S}$ and $\widetilde{\boldsymbol{S}}$ is at most ϵ. This is a natural extension of the well-known PAC learning model for learning Boolean functions [Val84] to the unsupervised setting of learning an unknown random variable (i.e. probability distribution).

While density estimation has been well studied by several different communities of researchers as described above, both the number of samples and running time required to learn are not yet well understood, even for some surprisingly simple types of discrete random variables. Below we describe a simple and natural class of random variables — *sums of independent integer-valued random variables* — for which we give the first known results, both from an information-theoretic and computational perspective, characterizing the complexity of learning such random variables.

A. Sums of independent integer random variables.

Perhaps the most basic discrete distribution learning problem imaginable is learning an unknown random variable $\boldsymbol{X}$ that is supported on the k-element finite set $\{0, 1, \ldots, k-1\}$. Throughout the paper we refer to such a random variable as a *k-IRV* (for "Integer Random Variable"). Learning an unknown k-IRV is of course a well understood problem: it has long been known that a simple histogram-based algorithm can learn such a random variable to accuracy ϵ using $\Theta(k/\epsilon^2)$ samples, and that $\Omega(k/\epsilon^2)$ samples are necessary for any learning algorithm.

A natural extension of this problem is to learn a *sum* of n *independent* such random variables, i.e. to learn $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ where the $\boldsymbol{X}_i$'s are independent k-IRVs (which, we stress, need not be identically distributed and may have arbitrary distributions supported on $\{0, 1, \ldots, k-1\}$). We call such a random variable a *k-SIIRV* (for "Sum of Independent Integer Random Variables"); learning an unknown k-SIIRV is the problem we solve in this paper.

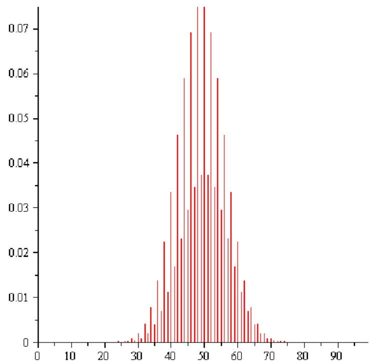

Figure 1. The probability mass function of a certain 3-SIIRV with $n = 50$.

Since every k-SIIRV is supported on $\{0, 1, \ldots, n(k-1)\}$ any such distribution can be learned using $O(nk/\epsilon^2)$ samples, but of course this simple observation does not use any of the k-SIIRV structure. On the other hand, it is clear (even when $n = 1$) that $\Omega(k/\epsilon^2)$ samples are necessary for learning k-SIIRVs. [1] A priori it is not clear how many samples (as a function of n and k) are information-theoretically sufficient to learn k-SIIRVs, even ignoring issues of computational efficiency. The $k = 2$ case of this problem (i.e., Poisson Binomial Distributions, or "PBDs") was only solved last year in [DDS12b], which gave an efficient algorithm using $\widetilde{O}(1/\epsilon^3)$ samples (independent of n) to learn any Poisson Binomial Distribution.

We stress that k-SIIRVs for general k may have a much richer structure than Poisson Binomial Distributions; even 3-SIIRVs are qualitatively very different from 2-SIIRVs. As a simple example of this more intricate structure, consider the 3-SIIRV $\boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ with $n = 50$ depicted in Figure 1, in which $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_{n-1}$ are identically distributed and uniform over $\{0, 2\}$ while $\boldsymbol{X}_n$ puts probability $2/3$ on 0 and $1/3$ on 1. It is easy to see from this simple example that even 3-SIIRVs can have significantly more daunting structure than any PBD; in particular, they can be $\Theta(1)$-far from every log-concave distribution; can be $\Theta(1)$-far from every Binomial distribution; and can have $\Theta(n)$ modes (and be $\Theta(1)$-far from every unimodal distribution). They thus dramatically fail to have all three kinds of structure (unimodality, log-concavity, and closeness to Binomial) that were exploited in the recent works [DDS12b], [CDSS13] on learning PBDs.

The main learning result. Our main learning result is that both the sample complexity (number of samples required for learning) and the computational complexity (running time) of learning k-SIIRVs is polynomial in k and $1/\epsilon$, and *completely independent of* n. [2]

Theorem I.1. *[Main Learning Result] There is a learning algorithm for k-SIIRVs with the following properties: Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be any sum of n independent (not necessarily identically distributed) random variables $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_n$ each supported on $\{0, \ldots, k-1\}$. The algorithm uses $\mathrm{poly}(k/\epsilon)$ samples from $\boldsymbol{S}$, runs in time $\mathrm{poly}(k/\epsilon)$, and with probability at least $9/10$ outputs a (succinct description of a) random variable $\widetilde{\boldsymbol{S}}$ such that $\mathrm{d}_{\mathrm{TV}}(\boldsymbol{S}, \widetilde{\boldsymbol{S}}) \le \epsilon$.*

(Note that since even learning a single k-IRV requires $\Omega(k/\epsilon^2)$ samples as noted above, this $\mathrm{poly}(k, 1/\epsilon)$ complexity is best possible up to the specific degree of the polynomial.) We give a detailed description of the "succinct description" of our hypothesis random variable $\widetilde{\boldsymbol{S}}$ in Section I-B, after we describe the new structural theorem that underlies our learning results.

B. Prior work and our techniques.

As noted above, Theorem I.1 is a broad generalization of the main learning result of [DDS12b], which established it in the special case of $k = 2$. A key ingredient in the [DDS12b] learning result is a structural theorem of Daskalakis and Papadimitriou [DP11] which states that any Poisson Binomial Distribution must be either ϵ-close to a *sparse* distribution (supported on $\mathrm{poly}(1/\epsilon)$ consecutive integers), or ϵ-close to a *translated Binomial* distribution. In our current setting of working with k-SIIRVs for general k, structural results of this sort (giving arbitrary-accuracy approximation for an arbitrary k-SIIRV) were not previously known. Our main technical contribution is proving such a structural result (see Theorem I.2 below); given this structural result, it is relatively straightforward for us to obtain our main learning result, Theorem I.1, using algorithmic ingredients for learning probability distributions from the recent works [DDS12b], [CDSS13].

There is a fairly long line of research on approximate limit theorems for sums of independent integer random variables, dating back several decades (see e.g. [Pre83], [Kru86], [BHJ92]). Our main structural result employs some of the latest results in this area [CL10], [CGS11]; however, we need to extend these results beyond what is currently known. Known approximation theorems for sums of integer random variables (which are generally proved using Stein's method) typically give bounds on the variation distance between a SIIRV $\boldsymbol{S}$ and various specific types of "nice" (Gaussian-like) random variables such as translated/compound Poisson random variables or *discretized normals* (as described in Definition II.3). However, it is easy to see that in general a k-SIIRV may be very far in variation distance from

[1] It should be noted that while all our results in the paper hold for all settings of n and k, intuitively one should think of n as a "large" asymptotic parameter and $k \ll n$ as a "small" fixed parameter. If $k \ge n$ then the trivial approach described above learns using $O(nk/\epsilon^2) = O(k^2/\epsilon^2)$ samples.

[2] We work in the standard "word RAM" model in which basic arithmetic operations on $O(\log n)$-bit integers are assumed to take constant time. We give more model details in Section V.

any discretized normal distribution; see for example the 3-SIIRVs discussed in Figure 1, or the discussion following Corollary 4.5 in [BX99]. To evade this difficulty, limit theorems in the literature typically put strong restrictions on the SIIRVs they apply to so that a normal-like distribution is forced. Specifically, they bound the total variation distance between $\boldsymbol{S}$ and a "nice" distribution using an error term involving the "shift-distance" (see Definition II.4) of certain random variables closely related to $\boldsymbol{S}$; see for example Theorem 4.3 of [BX99], Theorem 7.4 of [CGS11], or Theorem 1.3 of [Fan12] for results of this sort. However it is easy to see that for general k-SIIRVs these shift-distances can be very large — large enough that no nontrivial bound is obtained. Thus previous bounds from the literature do not provide structural results that characterize general k-SIIRVs up to arbitrary accuracy.

Another approach to analyzing k-SIIRVs arises from the recent work of Valiant and Valiant [VV11]. They gave a limit theorem for sums $\vec{\boldsymbol{S}}$ of independent $\mathbb{Z}^k$-valued random variables supported on $\{0, e_1, e_2, \ldots, e_k\}$, where e_i denotes the vector $(0, \ldots, 0, 1, 0, \ldots, 0)$ with the 1 in the ith coordinate. Specifically, they bounded the total variation distance of such $\vec{\boldsymbol{S}}$ from the appropriate discretized k-dimensional normal. Note that these $\mathbb{Z}^k$-valued random sums effectively generalize k-SIIRVs, because any k-SIIRV can be obtained as the dot-product $\langle \vec{\boldsymbol{S}}, (0, 1, \ldots, k-1) \rangle$. Unfortunately we cannot use their work for two reasons. The first reason is technical: their error bound has a dependence on n, namely $\Theta(\log^{2/3} n)$, which we do not want to pay. The second reason is more conceptual; as in previous theorems their limiting distribution is a (discretized) normal, which means it cannot capture general SIIRVs. This issue manifests itself in their error term, which is large if the covariance matrix of the k-dimensional normal has a small eigenvalue. Indeed, the covariance matrix will have an $o_n(1)$ eigenvalue for k-SIIRVs of the sort illustrated in Figure 1.

Despite these difficulties, we are able to leverage prior partial results on SIIRVs to give a new structural result showing that *any* k-SIIRV can be approximated to *arbitrarily* high accuracy by a relatively "simple" random variable. More precisely, our result shows that every k-SIIRV is either close to a "sparse" random variable, or else is close to a random variable $c\boldsymbol{Z} + \boldsymbol{Y}$ which decomposes nicely into an arbitrary "local" component $\boldsymbol{Y}$ and a highly structured "global" component $c\boldsymbol{Z}$ (where as above $\boldsymbol{Z}$ is a discretized normal):

Theorem I.2. *[Main Structural Result] Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be a sum of n independent (not necessarily identically distributed) random variables $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_n$ each supported on $\{0, \ldots, k-1\}$. Then for any $\epsilon > 0$, $\boldsymbol{S}$ is either*

1) *$O(\epsilon)$-close to a random variable which is supported on at most $\frac{k^9}{\epsilon^4}$ consecutive integers; or*
2) *$O(\epsilon)$-close to a random variable of the form $c\boldsymbol{Z} + \boldsymbol{Y}$ for some $1 \le c \le k-1$, where $\boldsymbol{Y}$, $\boldsymbol{Z}$ are independent random variables such that:*
   a) *$\boldsymbol{Y}$ is a c-IRV, and*
   b) *$\boldsymbol{Z}$ is a discretized normal random variable with parameters $\frac{\mu}{c}, \frac{\sigma^2}{c^2}$ where $\mu = \mathbf{E}[\boldsymbol{S}]$ and $\sigma^2 = \mathbf{Var}[\boldsymbol{S}]$.*

An alternative statement of our main structural result is the following: for $\boldsymbol{S}$ a k-SIIRV with variance $\mathbf{Var}[\boldsymbol{S}] = \sigma^2$, there is a value $1 \le c \le k-1$ and independent random variables $\boldsymbol{Y}, \boldsymbol{Z}$ as specified in part (2) above, such that $\mathrm{d_{TV}}(\boldsymbol{S}, c\boldsymbol{Z} + \boldsymbol{Y}) \le \mathrm{poly}(k, 1/\sigma)$. (See Corollary IV.5 for a more detailed statement.) Given this detailed structural characterization of an arbitrary k-SIIRV, it is not difficult to establish our main learning result, Theorem I.1; see Section V.

We believe this approximation theorem for arbitrary k-SIIRVs should be of independent interest. One potential direction for future application comes from the field of pseudorandomness. A classic problem in this area is to find pseudorandom generators with short seed length which fool "combinatorial rectangles". A notable recent work by Gopalan et al. [GMRZ11] made new progress on this problem, as well as a generalization they described as fooling "combinatorial shapes". A combinatorial shape is nothing more than a 2-SIIRV (in which the sample space of each $\boldsymbol{X}_i$ is $[m]$ for some integer m). Indeed, much of the technical work in [GMRZ11] goes into giving a new proof methodology for 2-SIIRV limit theorems, one which is more amenable to derandomization. It seems possible that our new limit theorem for k-SIIRVs may be useful in generalizing the [GMRZ11] derandomization results from 2-SIIRVs to k-SIIRVs.

We conclude this subsection by providing the high-level idea in the proof of our structural result, as well as the structure of the hypothesis output by our learning algorithm.

The idea behind Theorem I.2. The two cases (1) and (2) of Theorem I.2 correspond to $\boldsymbol{S}$ having "small" versus "large" variance respectively. The easier case is when $\mathbf{Var}(\boldsymbol{S})$ is "small": in this case it is straightforward to show that $\boldsymbol{S}$ must have almost all its probability mass on values in a small interval, and (1) follows easily from this.

The more challenging case is when $\mathbf{Var}(\boldsymbol{S})$ is "large." Intuitively, in order for $\mathbf{Var}(\boldsymbol{S})$ to be large it must be the case that at least one of the $k-1$ values $1, 2, \ldots, k-1$ makes a "large contribution" to $\mathbf{Var}(\boldsymbol{S})$. (This is made precise by working with "0-moded" SIIRVs and analyzing the "b-weight" of the SIIRV for $b \in \{1, \ldots, k-1\}$; see Definition II.2 for details.)

It is useful to first consider the special case that all $k-1$ values $\{1, \ldots, k-1\}$ make a "large contribution" to $\mathbf{Var}(\boldsymbol{S})$; we do this in Section III. To analyze this case it is useful to view a draw of the random variable $\boldsymbol{S}$ as taking place in two stages as follows: First (stage 1) we independently choose for each $\boldsymbol{X}_i$ a value $r_i \in \{1, \ldots, k-1\}$. Then (stage 2) for each i we independently choose whether $\boldsymbol{X}_i$ will be set to 0 or to r_i. Using this perspective on $\boldsymbol{S}$, it can be shown that with high probability over the stage-1 outcomes, the resulting random variable that is sampled in Stage 2 is of the form $\sum_{j=1}^{k-1} j \cdot \boldsymbol{Y}_j$ where each $\boldsymbol{Y}_j$ is a large-variance PBD. Given this, using Theorem 7.4 of [CGS11] it is not difficult to show that the overall distribution of $\boldsymbol{S}$ is close to a discretetized normal distribution. (This is the $c = 1$ case of case (2) of Theorem I.2.)

In the general case it may be the case that some of the $k-1$ values contribute very little to $\mathbf{Var}(\boldsymbol{S})$. (This is what happened in the example illustrated in Figure 1.) Let $\mathcal{L} \subset \{1, \ldots, k-1\}$ denote the set of values that make a "small" contribution to $\mathbf{Var}(\boldsymbol{S})$ (observe that $\mathcal{L}$ is nonempty by assumption in this case, or else we are in the special case of the previous paragraph) and let $\mathcal{H} \cup \{0\}$ denote the remaining values in $\{0, 1, \ldots, k-1\}$ (observe that $\mathcal{H}$ is nonempty since otherwise $\mathbf{Var}(\boldsymbol{S})$ would be small as noted earlier). In this general case it is useful to consider a *different* decomposition of the random variable $\boldsymbol{S}$. As before we view a draw of $\boldsymbol{S}$ as taking place in stages, but now the stages are as follows: First (stage 1) for each $i \in [N]$ we independently select whether $\boldsymbol{X}_i$ will be "light" (i.e. will take a value in $\mathcal{L}$) or will be "heavy" (will take a value in $\mathcal{H} \cup \{0\}$). Then (stage 2) for each $\boldsymbol{X}_i$ that has been designated to be "light" we independently choose which particular value in $\mathcal{L}$ it will take, and similarly (stage 3) for each $\boldsymbol{X}_i$ that has been designated "heavy" we independently choose an element of $\mathcal{H} \cup \{0\}$ for it.

The key advantage of the above decomposition is that conditioned on the stage 1 outcome, stages 2 and 3 are *independent* of each other. Using this decomposition, our analysis shows that the contribution from the "light" IRVs (stage 2) is close to a sparse random variable, and the contribution from the "heavy" random variables (stage 3) is close to a $\gcd(\mathcal{H})$-scaled discretized normal distribution (this uses the special case, sketched earlier, in which all values make a "large contribution" to the variance); here $\gcd(\mathcal{H})$ is the greatest common divisor of the absolute values of the integers in $\mathcal{H}$. This essentially gives case (2) of Theorem I.2, where as sketched above, the value "c" is $\gcd(\mathcal{H})$.

The structure of our hypotheses. The "succinct description" of the hypothesis random variable that our learning algorithm outputs naturally reflects the structure of the approximating random variable given by Theorem I.2 above. Some terminology will be useful here: we say that an IRV $\boldsymbol{A}$ is *t-flat* if $\boldsymbol{A}$ is supported on a union of $t' \leq t$ disjoint intervals $I_1 \cup \cdots \cup I_{t'}$, and for each fixed $1 \leq j \leq t'$ all points $x_1, x_2 \in I_j$ have $\mathbf{Pr}[\boldsymbol{X} = x_1] = \mathbf{Pr}[\boldsymbol{X} = x_2] = p_j$ for some $p_j > 0$ (so $\boldsymbol{A}$ is piecewise-constant across each interval I_j). An *explicit description* of a t-flat IRV $\boldsymbol{A}$ is a list of pairs $(I_1, p_1), \ldots, (I_{t'}, p_{t'})$, for some $t' \leq t$.

There are two possible forms for the output hypothesis random variable $\widetilde{\boldsymbol{S}}$ of Theorem I.1, corresponding to the two cases of Theorem I.2 above. The first possible form is simply a list of pairs $(r, p_0), \ldots, (r+\ell, p_\ell)$ where the pair (s, p) indicates that $\mathbf{Pr}[\widetilde{\boldsymbol{S}} = s] = p$ and $\sum_{j=0}^{\ell} p_j = 1$ and $\ell = k^9/\epsilon^4$. The second possible form of the hypothesis is as two lists $(I_1, p_1), \ldots, (I_\ell, p_t)$ and $(0, q_0), \ldots, (c-1, q_{c-1})$, where $I_1, \ldots, I_t$ are disjoint intervals and $\sum_{j=1}^{t} |I_j| p_j = \sum_{i=0}^{c-1} q_i = 1$. The list $(I_1, p_1), \ldots, (I_t, p_t)$ specifies a t-flat random variable $\boldsymbol{Z}'$ and the list $(0, q_0), \ldots, (c-1, q_{c-1})$ specifies a c-IRV $\boldsymbol{Y}'$. The hypothesis distribution in this case is $\widetilde{\boldsymbol{S}} = c\boldsymbol{Z}' + \boldsymbol{Y}'$.

C. Discussion: Learning independent sums of more general random variables?

It is natural to ask whether our highly efficient $\mathrm{poly}(k/\epsilon)$-sample (independent of n) learning algorithm for k-SIIRVs can be extended to n-way independent sums $\boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ of more general types of integer-valued random variables $\boldsymbol{X}_i$ than k-IRVs. Here we note that no such efficient learning results are possible for several natural generalizations of k-IRVs.

One natural generalization is to consider integer random variables which are supported on k values which need not be consecutive. Let us say that $\boldsymbol{X}$ is a *k-support IRV* if $\boldsymbol{X}$ is an IRV supported on at most k values. It turns out that sums of independent k-support IRVs can be quite difficult to learn; in particular, Theorem 3 of [DDS12b] give an information-theoretic argument showing that even for $k = 2$, any algorithm that learns a sum of n 2-support IRVs must use $\Omega(n)$ samples (even if the i-th IRV is constrained to have support $\{0, i\}$).

A different generalization of k-IRVs to consider in this context is a class that we denote as *(c,k)-moment IRVs*. A (c,k)-moment IRV is an integer-valued random variable $\boldsymbol{X}$ such that the c-th absolute moment $\mathbf{E}[|\boldsymbol{X}|^c]$ lies in $[0, (k-1)^c]$.

It is clear that any k-IRV is a (c,k)-moment IRV for all c. Moreover, it is easy to show (using Markov's inequality) that for any fixed $c > 0$, any single (c,k)-moment IRV $\boldsymbol{X}$ can be learned to accuracy ϵ using $\mathrm{poly}(k/\epsilon)$ samples. However, our sample complexity bounds for learning *sums* of n independent k-IRVs provably cannot be extended to sums of n independent (c,k)-moment IRVs, in a strong sense: any learning algorithm for such sums must (information-theoretically) use at least poly(n) samples.

Observation I.3. *Fix any integer $c \geq 1$. Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be a sum of n $(c,2)$-moment IRVs. Let L be any algorithm which, given n and access to independent samples from $\boldsymbol{S}$, with probability at least $e^{-o(n^{1/c})}$ outputs a hypothesis distribution $\widetilde{\boldsymbol{S}}$ such that $\mathrm{d_{TV}}(\boldsymbol{S}, \widetilde{\boldsymbol{S}}) < 1/41$. Then L must use at least $n^{1/c}/10$ samples.*

The argument is a simple modification of the lower bound given by Theorem 3 of [DDS12b] and is deferred to the full version.

II. Definitions and Basic Tools

In this section we give some necessary definitions and recall some useful tools from probability.

A. Definitions

We begin with a formal definition of *total variation distance*, which we specialize to the case of integer-valued random variables. For two distributions $\mathbb{P}$ and $\mathbb{Q}$ supported on $\mathbb{Z}$, their total variation distance is defined to be

$$\mathrm{d_{TV}}(\mathbb{P},\mathbb{Q}) = \sup_{A\subseteq\mathbb{Z}} |\mathbb{P}(A)-\mathbb{Q}(A)| = \frac{1}{2}\sum_{j\in\mathbb{Z}} |\mathbb{P}(\{j\})-\mathbb{Q}(\{j\})|.$$

If $\boldsymbol{X}$ and $\boldsymbol{Y}$ are integer random variables, their total variation distance, $\mathrm{d_{TV}}(\boldsymbol{X},\boldsymbol{Y})$, is defined to be the total variation distance of their distributions. Throughout the paper we are casual about the distinction between a random variable and a distribution. For example, when we say "draw a sample from random variable $\boldsymbol{X}$" we formally mean "draw a sample from the distribution of $\boldsymbol{X}$", etc.

We proceed to discuss the most basic random variables that we will be interested in, namely IRVs, k-IRVs and $\pm k$-IRVs, and sums of these random variables:

Definition II.1. An *IRV* is an integer-valued random variable. For an integer $k \geq 2$, a *k-IRV* is an IRV supported on $\{0,1,\dots,k-1\}$. (Note that a 2-IRV is the same as a Bernoulli random variable.) A *$\pm k$-IRV* is an IRV supported on $\{-k+1,-k+2,\dots,k-2,k-1\}$. We say that an IRV $\boldsymbol{X}_i$ has *mode* 0 if $\mathbf{Pr}[\boldsymbol{X}_i=0] \geq \mathbf{Pr}[\boldsymbol{X}_i=b]$ for all $b\in\mathbb{Z}$.

Definition II.2. A *SIIRV* (Sum of Independent IRVs) is any random variable $\boldsymbol{S}=\boldsymbol{X}_1+\dots+\boldsymbol{X}_n$ where the $\boldsymbol{X}_i$'s are independent IRVs. We define k-SIIRVs and $\pm k$-SIIRVs similarly; a 2-SIIRV is also called a *PBD* (Poisson Binomial Distribution). For $b\in\mathbb{Z}$ we say that the *b-weight* of the SIIRV is $\sum_{i=1}^n \mathbf{Pr}[\boldsymbol{X}_i=b]$. Finally, we say that a SIIRV is 0-*moded* if each $\boldsymbol{X}_i$ has mode 0.

As a notational convention, we will typically use $\boldsymbol{X}$ to denote an IRV, and $\boldsymbol{S}$ to denote a SIIRV.

Discretized normal distributions will play an important role in our technical results, largely because of known theorems in probability which assert that under suitable conditions sums of independent integer random variables converge in total variation distance to discretized normal distributions. We now define these distributions:

Definition II.3. Let $\mu\in\mathbb{R}$, $\sigma\in\mathbb{R}^{\geq 0}$. We let $Z(\mu,\sigma^2)$ denote the *discretized normal* distribution. The definition of $\boldsymbol{Z}\sim Z(\mu,\sigma^2)$ is that we first draw a normal $\boldsymbol{G}\sim\mathrm{N}(\mu,\sigma^2)$ and then we set $\boldsymbol{Z}=\lfloor \boldsymbol{G}\rceil$; i.e., $\boldsymbol{G}$ rounded to the nearest integer.

We note that in the "large-variance" regime that we shall be concerned with, discretized normals are known to be close in variation distance to other types of distributions such as Binomial distributions and Translated Poisson distributions (see e.g. [RÖ7], [RR12]); however we shall not need to work with these other distributions.

Some of our arguments will use the following notion of *shift-distance* of a random variable:

Definition II.4. For $\boldsymbol{X}$ a random variable we define its *shift-distance* to be $\mathrm{d_{shift}}(\boldsymbol{X})=\mathrm{d_{TV}}(\boldsymbol{X},\boldsymbol{X}+1)$.

Finally, for completeness we record the following:

Definition II.5. Given a sequence or set C of nonzero integers we define $\gcd(C)$ to be the greatest common divisor of the absolute values of the integers in C. We adopt the convention that $\gcd(\emptyset)=0$.

B. Basic results from probability

Our proofs use various basic results from probability; these include bounds on total variation distance, results on normal and discretized normal distributions, bounds on shift-distance, and uniform convergence bounds.

Proposition II.6. *Let $\boldsymbol{A},\boldsymbol{A}',\boldsymbol{B},\boldsymbol{B}'$ be integer random variables such that $(\boldsymbol{A},\boldsymbol{A}')$ is independent of $(\boldsymbol{B},\boldsymbol{B}')$. Then* $\mathrm{d_{TV}}(\boldsymbol{A}+\boldsymbol{B},\boldsymbol{A}'+\boldsymbol{B}')\leq \mathrm{d_{TV}}(\boldsymbol{A},\boldsymbol{A}')+\mathrm{d_{TV}}(\boldsymbol{B},\boldsymbol{B}')$.

Proposition II.7. *Let $\boldsymbol{X}$ and $\boldsymbol{Y}$ be integer random variables which are independent conditioned on the outcome of a third discrete random variable $\boldsymbol{L}$. Further, let $\boldsymbol{Z}$ be an integer random variable independent of $\boldsymbol{X}$, $\boldsymbol{Y}$, and $\boldsymbol{L}$. Finally, let G be a set of "good" outcomes for $\boldsymbol{L}$ such that:*

- $\mathbf{Pr}[\boldsymbol{L}\notin G]\leq\eta$*;*
- *for each $L\in G$ we have* $\mathrm{d_{TV}}((\boldsymbol{Y}\mid\boldsymbol{L}=L),\boldsymbol{Z})\leq\epsilon$.

Then $\mathrm{d_{TV}}(\boldsymbol{X}+\boldsymbol{Y},\boldsymbol{X}\cdot\mathbf{1}_{\boldsymbol{L}\in G}+\boldsymbol{Z})\leq(1-\eta)\epsilon+\eta\leq\epsilon+\eta$.

Proposition II.8. *Let $\mu_1,\mu_2\in\mathbb{R}$ and $0<\sigma_1\leq\sigma_2$. Then* $\mathrm{d_{TV}}(\mathrm{N}(\mu_1,\sigma_1^2),\mathrm{N}(\mu_2,\sigma_2^2))\leq\frac{1}{2}\left(\frac{|\mu_1-\mu_2|}{\sigma_1}+\frac{\sigma_2^2-\sigma_1^2}{\sigma_1^2}\right)$.

Proposition II.9. *Let $\boldsymbol{G}\sim\mathrm{N}(\mu,\sigma^2)$ and $\lambda\in\mathbb{R}$. Then* $\mathrm{d_{TV}}(\lfloor\boldsymbol{G}+\lambda\rceil,\lfloor\boldsymbol{G}\rceil+\lfloor\lambda\rceil)\leq\frac{1}{2\sigma}$. *Consequently, for $\lambda\in\mathbb{Z}$, if $\boldsymbol{Z}\sim Z(\mu,\sigma^2)$ and $\boldsymbol{Z}'\sim Z(\mu-\lambda,\sigma^2)$ then* $\mathrm{d_{TV}}(\boldsymbol{Z},\lambda+\boldsymbol{Z}')\leq\frac{1}{2\sigma}$.

Proposition II.10. *Let $\boldsymbol{G}\sim\mathrm{N}(\mu,\sigma^2)$. Then* $\mathrm{d_{shift}}(\boldsymbol{G})\leq\frac{1}{2\sigma}$. *Consequently, if $\boldsymbol{Z}\sim Z(\mu,\sigma^2)$ then* $\mathrm{d_{shift}}(\boldsymbol{Z})\leq\frac{1}{2\sigma}$.

III. A useful special case: each $b\in\{1,\dots,k-1\}$ has large weight

In this section we state a useful special case of our our desired structural theorem for k-SIIRVs; its proof is deferred to the full version of this paper. In later sections we will use this special case to prove the general result.

Recall that for $b\in\{0,1,\dots,k-1\}$ the *b-weight* of a k-SIIRV $\boldsymbol{S}=\boldsymbol{X}_1+\dots+\boldsymbol{X}_n$ is $\sum_{i=1}^n\mathbf{Pr}[\boldsymbol{X}_i=b]$. The special case we consider in this section is that every $b \in \{1, \dots, k-1\}$ has large b-weight. The result we prove in this special case is that $\boldsymbol{S}$ is close to a discretized normal distribution:

Theorem III.1. *Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be a 0-moded $\pm k$-SIIRV, and assume no $\boldsymbol{X}_i$ is constantly 0. For each nonzero integer c with $|c| < k$, let M_c denote the c-weight of $\boldsymbol{S}$ and let $C = \{c \in \mathbb{Z} : c \neq 0, |c| < k, \text{ and } M_c > 0\} \neq \emptyset$. Further assume $\gcd(C) = 1$ and $M_c \geq M$ for all $c \in C$ where $M = \omega(k \log k)$. Let $\boldsymbol{Z} \sim Z(\mu, \sigma^2)$, where $\mu = \mathbf{E}[\boldsymbol{S}]$ and $\sigma^2 = \mathbf{Var}[\boldsymbol{S}]$. Then $\mathrm{d_{TV}}(\boldsymbol{S}, \boldsymbol{Z}) \leq O(k^{3.5}/\sqrt{M})$ and $\sigma^2 \geq M/8k$.*

Observe that Theorem III.1 corresponds to Case (2) of Theorem I.2 with $c = 1$ (so $\boldsymbol{Y}$ is a 1-IRV, i.e. the constant-0 random variable).

IV. Proof of Main Structural Result

A. Intuition and Preparatory Work

Throughout this section we let δ be an error parameter, and $M = M(k, 1/\delta)$ a large enough polynomial in k and $1/\delta$ to be determined later. With respect to M we make the following definition:

Definition IV.1 (Light integers and heavy integers)**.** Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be a 0-moded $\pm k$-SIIRV. We say that a non-zero integer $|b| < k$ is *M-light* if the b-weight of $\boldsymbol{S}$ is at most M; otherwise we call it *M-heavy*. We denote by $\mathcal{L}$ the set of M-light integers, and by $\mathcal{H}$ the set of M-heavy integers.

The bulk of the work towards proving Theorem I.2 is showing that any 0-moded $\pm k$-SIIRV $\boldsymbol{S}$ is close to the sum of a sparse distribution (supported on some $\mathrm{poly}(k/\delta)$ consecutive integers) and a discretized normal random variable scaled by $\gcd(\mathcal{H})$. To see why this may be true, we can distinguish the following cases:

- If $\mathcal{H} = \emptyset$ then $\boldsymbol{S}$ should be close to a sparse random variable, by Markov's inequality and $\mathbf{E}[|\boldsymbol{S}|] \leq \mathbf{E}[\sum_i |\boldsymbol{X}_i|] \leq \sum_i k \, \mathbf{Pr}[\boldsymbol{X}_i \in \mathcal{L}] = \sum_i k \sum_{j \in \mathcal{L}} \mathbf{Pr}[\boldsymbol{X}_i = j] \leq 2k^2 M$, where the last inequality holds because there are at most $2k$ integers in $\mathcal{L}$ and each of them is M-light.
- On the other hand, if $\mathcal{L} = \emptyset$, then Theorem III.1 is readily applicable, showing that $\boldsymbol{S}$ is close to a discretized normal random variable with the same mean and variance as $\boldsymbol{S}$.
- The remaining possibility is that $\mathcal{L}$, $\mathcal{H} \neq \emptyset$. If we condition on the event $\boldsymbol{X}_i \notin \mathcal{L}$ for all i, then the conditional distribution of $\boldsymbol{S}$ is still, by Theorem III.1, close to a discretized normal random variable, except that this discretized normal random variable is now scaled by $\gcd(\mathcal{H})$. (Indeed the conditioning only boosts the b-weight of integers $b \in \mathcal{H}$.) But "$\boldsymbol{X}_i \notin \mathcal{L}$ for all i" may be a rare event. Regardless, a typical sample from $\boldsymbol{S}$ shouldn't have a large set of indices $\boldsymbol{L} := \{i \mid \boldsymbol{X}_i \in \mathcal{L}\}$, because $\mathbf{E}[|\boldsymbol{L}|] \leq 2kM$. Indeed, one would expect that, conditioning on a typical $\boldsymbol{L}$, the b-weight of $\sum_{i \notin \boldsymbol{L}} \boldsymbol{X}_i$ for $b \in \mathcal{H}$ is still very large. Hence, conditioning on typical $\boldsymbol{L}$'s, $\sum_{i \notin \boldsymbol{L}} \boldsymbol{X}_i$ should still be close to a discretized normal (scaled by $\gcd(\mathcal{H})$). Moreover, we may hope that the normals arising by conditioning on different typical $\boldsymbol{L}$'s are close to a fixed "typical" discretized normal. Indeed, the fluctuations in the mean and variance of $\sum_{i \notin \boldsymbol{L}} \boldsymbol{X}_i$, conditioned on typical $\boldsymbol{L}$'s, should not be severe, since $\boldsymbol{L}$ is small. These considerations suggest that $\boldsymbol{S}$ is close to the sum of a sparse random variable, "the contribution of $\mathcal{L}$", and a $\gcd(\mathcal{H})$-scaled discretized normal, "the contribution of $\mathcal{H} \cup \{0\}$".

The last case is clearly the hardest and is handled by Theorem IV.3 in the next section. Before proceeding, let us make our intuition a bit more precise. First, let us formally disentangle $\boldsymbol{S}$ into the contributions of $\mathcal{L}$ and $\mathcal{H} \cup \{0\}$, by means of the following alternate sampling procedure for $\boldsymbol{S}$.

Definition IV.2. ["The Light-Heavy Experiment".] Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be a 0-moded $\pm k$-SIIRV. We define here an alternate experiment for sampling the random variable $\boldsymbol{S}$, called the "Light-Heavy Experiment". There are three stages:

1) [Stage 1]: We first sample a random subset $\boldsymbol{L} \subseteq [n]$, by independently including each i into $\boldsymbol{L}$ with probability $\mathbf{Pr}[\boldsymbol{X}_i \in \mathcal{L}]$.
2) [Stage 2]: Independently we sample for each $i \in [n]$ a random variable $\underline{\boldsymbol{X}}_i \in \mathcal{L}$ as follows:
$$\underline{\boldsymbol{X}}_i = b, \text{with probability } \frac{\mathbf{Pr}[\boldsymbol{X}_i = b]}{\mathbf{Pr}[\boldsymbol{X}_i \in \mathcal{L}]};$$
i.e. $\underline{\boldsymbol{X}}_i$ is distributed according to the conditional distribution of $\boldsymbol{X}_i$, conditioning on $\boldsymbol{X}_i \in \mathcal{L}$. In the exceptional case that $\mathbf{Pr}[\boldsymbol{X}_i \in \mathcal{L}] = 0$, we define $\underline{\boldsymbol{X}}_i = 0$ with probability 1.
3) [Stage 3]: Independently we sample for each $i \in [n]$ a random variable $\overline{\boldsymbol{X}}_i \in \mathcal{H} \cup \{0\}$ as follows:
$$\overline{\boldsymbol{X}}_i = b, \text{with probability } \frac{\mathbf{Pr}[\boldsymbol{X}_i = b]}{\mathbf{Pr}[\boldsymbol{X}_i \notin \mathcal{L}]};$$
i.e. $\overline{\boldsymbol{X}}_i$ is distributed according to the conditional distribution of $\boldsymbol{X}_i$, conditioning on $\boldsymbol{X}_i \notin \mathcal{L}$.

After these three stages we output $\sum_{i \in \boldsymbol{L}} \underline{\boldsymbol{X}}_i + \sum_{i \notin \boldsymbol{L}} \overline{\boldsymbol{X}}_i$ as a sample from $\boldsymbol{S}$, where $\sum_{i \in \boldsymbol{L}} \underline{\boldsymbol{X}}_i$ represents "the contribution of $\mathcal{L}$" and $\boldsymbol{S}_{\boldsymbol{L}} := \sum_{i \notin \boldsymbol{L}} \overline{\boldsymbol{X}}_i$ "the contribution of $\mathcal{H} \cup \{0\}$." We note that the two contributions are not independent, but they are independent conditioned on the outcome of $\boldsymbol{L}$. This concludes the definition of the Light-Heavy Experiment.

Coming back to our proof strategy, we aim to argue that:

(i) The contribution of $\mathcal{L}$ is close to a sparse random variable. This is clear from the definition of $\mathcal{L}$, since $\mathbf{E}[|\sum_{i\in \boldsymbol{L}} \underline{\boldsymbol{X}}_i|] \le \mathbf{E}[\sum_{i\in \boldsymbol{L}} |\underline{\boldsymbol{X}}_i|] \le k\sum_{i=1}^n \mathbf{Pr}[\boldsymbol{X}_i \in \mathcal{L}] = k\sum_{j\in\mathcal{L}}\sum_{i=1}^n \mathbf{Pr}[\boldsymbol{X}_i = j] \le 2k^2M$.

(ii) With probability close to 1 (with respect to $\boldsymbol{L}$), the contribution of $\mathcal{H}\cup\{0\}$ is close to a fixed, $\gcd(\mathcal{H})$-scaled discretized normal random variable $\boldsymbol{Z}$, which is independent of $\boldsymbol{S}$. Showing this is the heart of our argument in the proof of Theorem IV.3, in the next section.

Given (i) and (ii), we can readily finish the proof of Theorem IV.3 using Proposition II.7: Indeed, if we set $\boldsymbol{X}$ to be the contribution of $\mathcal{L}$ and $\boldsymbol{Y}$ to be the contribution of $\mathcal{H}\cup\{0\}$, we get that $\boldsymbol{S}$ is close to the sum of $\boldsymbol{X}$ times the indicator that $\boldsymbol{L}$ is typical (which is close to a sparse random variable) and a discretized normal independent of $\boldsymbol{X}$, scaled by $\gcd(\mathcal{H})$.

B. The Structural Result

We make the intuition of the previous section precise, by providing the proof of the following.

Theorem IV.3. *Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be a 0-moded $\pm k$-SIIRV with mean μ and variance $\sigma^2 \ge 15k^4\log(1/\delta)\cdot M$, where $1 \le M = \omega(k\log k)$ and $\delta \in (0, \frac{1}{10})$ are parameters. Let also $c = \gcd(\mathcal{H})$, where $\mathcal{L}$, $\mathcal{H}$ are defined in terms of M and k as in Definition IV.1. Then there are independent random variables $\boldsymbol{Y}$ and $\boldsymbol{Z}$ such that:*

- *$\boldsymbol{Y}$ is a $\pm(kM')$-IRV, where $M' = 4k\log(1/\delta)\cdot M$;*
- *$\boldsymbol{Z} \sim Z(\frac{\mu}{c}, \frac{\sigma^2}{c^2})$;*
- *$\mathrm{d}_{\mathrm{TV}}(\boldsymbol{S}, \boldsymbol{Y} + c\boldsymbol{Z}) \le \delta + 2k\exp(-M/8) + O(k^{3.5}/\sqrt{M}) + O(k^2\log(1/\delta)M/\sigma)$.*

In particular, taking $M = k^7/\delta^2$ the total variation bound becomes $O(\delta + (k^9/\delta^2)\log(1/\delta)/\sigma)$.

Proof: Throughout we will assume that $\boldsymbol{S}$ is drawn according to the Light-Heavy Experiment from Definition IV.2. We use that definition's notation: $\boldsymbol{L}$, $\underline{\boldsymbol{X}}_i$, and $\overline{\boldsymbol{X}}_i$. For each outcome L of $\boldsymbol{L}$, we introduce the notation

$$\boldsymbol{S}_L = \sum_{i\notin L} \overline{\boldsymbol{X}}_i.$$

Note that each random variable $\boldsymbol{S}_L$ is a $\pm k$-SIIRV. Finally, for each $i \in [n]$ we introduce the shorthand $\ell_i = \mathbf{Pr}[\boldsymbol{X}_i \in \mathcal{L}]$. Note that $\sum_{i=1}^n \ell_i \le 2kM$ and $\ell_i < 1 - \frac{1}{2k}$ (since $\boldsymbol{X}_i$ has mode 0).

Understanding Typical $\boldsymbol{L}$'s.: We study typical outcomes of $\boldsymbol{L}$. First, we argue that typical $\boldsymbol{L}$'s have small cardinality. Indeed, let us define the following event:

$$\mathit{BAD}_0 = \{\text{outcomes } L \text{ for } \boldsymbol{L} \text{ having } |L| \ge M'\}.$$

Since $\mathbf{E}[|\boldsymbol{L}|] = \sum_i \ell_i \le 2kM$, our choice of M' and a multiplicative Chernoff bound imply that $\mathbf{Pr}[\mathit{BAD}_0] \le \delta$.

Next, we argue that, conditioning on typical outcomes L for $\boldsymbol{L}$, the random variable $\boldsymbol{S}_L$ has b-weight at least $M/2$ for each $b \in \mathcal{H}$. In particular, for each $b \in \mathcal{H}$ define the event BAD_b to be the outcomes L for $\boldsymbol{L}$ in which the b-weight of the $\pm k$-SIIRV $\boldsymbol{S}_{\boldsymbol{L}}$ is less than $M/2$.

Notice that the b-weight of $\boldsymbol{S}_{\boldsymbol{L}}$ is the sum of independent random variables $\boldsymbol{W}_i$, where $\boldsymbol{W}_i = 0$ with probability ℓ_i and $\boldsymbol{W}_i = \mathbf{Pr}[\boldsymbol{X}_i = b]/(1-\ell_i) \le 1$ with probability $1-\ell_i$. Thus the expectation (over $\boldsymbol{L}$) of the b-weight of $\boldsymbol{S}_{\boldsymbol{L}}$ is simply the b-weight of $\boldsymbol{S}$; since this is at least M, a multiplicative Chernoff bound implies that $\mathbf{Pr}[\mathit{BAD}_b] \le \exp(-M/8)$. Defining *BAD* to be the union of all the bad events, we conclude that

$$\mathbf{Pr}[\mathit{BAD}] \le \delta + 2k\exp(-M/8). \qquad (1)$$

Concluding the Proof.: Let $\boldsymbol{Z} \sim Z(\frac{\mu}{c}, \frac{\sigma^2}{c^2})$ be independent of all other random variables, as in the statement of the theorem. The remainder of the proof will be devoted to showing that for every $L \notin \mathit{BAD}$

$$\mathrm{d}_{\mathrm{TV}}(\boldsymbol{S}_L, c\boldsymbol{Z}) \le O(k^{3.5}/\sqrt{M}) + O(k^2\log(1/\delta)M/\sigma). \qquad (2)$$

Given (2) we can conclude the proof by applying Proposition II.7, with $\sum_{i\in\boldsymbol{L}} \underline{\boldsymbol{X}}_i$ playing the role of "$\boldsymbol{X}$", $\boldsymbol{S}_L$ playing the role of "$\boldsymbol{Y}$", $c\boldsymbol{Z}$ playing the role of "$\boldsymbol{Z}$", and "G" being the complement of *BAD*. Note that $(\sum_{i\in\boldsymbol{L}} \underline{\boldsymbol{X}}_i)\cdot \mathbf{1}_{\boldsymbol{L}\notin \mathit{BAD}}$ is indeed a $\pm kM'$-IRV.

Establishing (2). Notice first that if $\mathcal{H} = \emptyset$, then (2) is trivially true as then $\boldsymbol{Z} = 0$. Otherwise, $\boldsymbol{Z} \sim Z(\frac{\mu}{c}, \frac{\sigma^2}{c^2})$ and let us fix an arbitrary outcome $L \notin \mathit{BAD}$. Write $\mu_L = \mathbf{E}[\boldsymbol{S}_L]$, $\sigma_L^2 = \mathbf{Var}[\boldsymbol{S}_L]$, and define $\boldsymbol{S}' = \frac{1}{c}\boldsymbol{S}_L$. (Also, delete any identically zero summands from $\boldsymbol{S}'$.) By virtue of $L \notin \mathit{BAD}_b$ for all $b \in \mathcal{H}$ we are in a position to apply Theorem III.1 to $\boldsymbol{S}'$ (except with $M/2$ in place of M). We deduce that for $\boldsymbol{Z}' \sim Z(\frac{\mu_L}{c}, \frac{\sigma_L^2}{c^2})$ we have

$$\mathrm{d}_{\mathrm{TV}}(\boldsymbol{S}', \boldsymbol{Z}') \le O(k^{3.5}/\sqrt{M}).$$

If we can furthermore show

$$\mathrm{d}_{\mathrm{TV}}(\boldsymbol{Z}', \boldsymbol{Z}) \le O(k^2\log(1/\delta)M/\sigma) \qquad (3)$$

then we will have established (2).

It therefore remains to show (3); i.e., to show that

$$\mathrm{d}_{\mathrm{TV}}\left(Z\left(\frac{\mu_L}{c}, \frac{\sigma_L^2}{c^2}\right), Z\left(\frac{\mu}{c}, \frac{\sigma^2}{c^2}\right)\right) \le O(k^2\log(1/\delta)M/\sigma).$$

This is in turn implied by the claim

$$\mathrm{d}_{\mathrm{TV}}\left(\mathrm{N}(\mu_L, \sigma_L^2), \mathrm{N}(\mu, \sigma^2)\right) \le O(k^2\log(1/\delta)M/\sigma). \qquad (4)$$

To establish (4) we claim the following:

Claim IV.4. *For $L \notin \mathit{BAD}$,*

$$|\mu - \mu_L| \le 4k^2(\log(1/\delta)+1)\cdot M; \qquad (5)$$

$$|\sigma^2 - \sigma_L^2| \le 14k^4\log(1/\delta)\cdot M. \qquad (6)$$

Proof of Claim IV.4.: (Bounding the Mean Difference.) We have:

$$|\mu - \mu_L| = \left|\sum_{i\in L}\mathbf{E}[\boldsymbol{X}_i] + \sum_{i\notin L}(\mathbf{E}[\boldsymbol{X}_i] - \mathbf{E}[\overline{\boldsymbol{X}}_i])\right| \le \sum_{i\in L}\mathbf{E}[|\boldsymbol{X}_i|] + \sum_{i\notin L}|\,\mathbf{E}[\boldsymbol{X}_i] - \mathbf{E}[\overline{\boldsymbol{X}}_i]|.$$

Notice that $\sum_{i\in L}\mathbf{E}[|\boldsymbol{X}_i|] \le k|L| \le kM' \le 4k^2\log(1/\delta)\cdot M$. Next we bound each difference $|\,\mathbf{E}[\boldsymbol{X}_i] - \mathbf{E}[\overline{\boldsymbol{X}}_i]|$ separately, using the law of total expectation. Namely, if $\boldsymbol{I}$ is the indicator for $\boldsymbol{X}_i \in \mathcal{L}$, we have

$$\mathbf{E}[\boldsymbol{X}_i] = \mathbf{E}[\mathbf{E}[\boldsymbol{X}_i \mid \boldsymbol{I}]] = \mathbf{Pr}[\boldsymbol{I}]E[\underline{\boldsymbol{X}}_i] + (1-\mathbf{Pr}[\boldsymbol{I}])E[\overline{\boldsymbol{X}}_i].$$

Hence,

$$|\,\mathbf{E}[\boldsymbol{X}_i] - \mathbf{E}[\overline{\boldsymbol{X}}_i]| = \mathbf{Pr}[\boldsymbol{I}]|E[\underline{\boldsymbol{X}}_i] - E[\overline{\boldsymbol{X}}_i]| \le \ell_i 2k$$

and consequently $\sum_{i\notin L}|\,\mathbf{E}[\boldsymbol{X}_i] - \mathbf{E}[\overline{\boldsymbol{X}}_i]| \le 2k\sum_i \ell_i \le 4k^2M$. (5) follows.

(Bounding the Variance Difference.) We have

$$|\sigma^2 - \sigma_L^2| = \left|\sum_{i\in L}\mathbf{Var}[\boldsymbol{X}_i] + \sum_{i\notin L}(\mathbf{Var}[\boldsymbol{X}_i] - \mathbf{Var}[\overline{\boldsymbol{X}}_i])\right| \le \sum_{i\in L}\mathbf{Var}[\boldsymbol{X}_i] + \sum_{i\notin L}|\,\mathbf{Var}[\boldsymbol{X}_i] - \mathbf{Var}[\overline{\boldsymbol{X}}_i]|. \quad (7)$$

The first term is at most $k^2\cdot|L| \le k^2\cdot M' = 4k^3\log(1/\delta)\cdot M$. To bound the second term we bound each $|\,\mathbf{Var}[\boldsymbol{X}_i] - \mathbf{Var}[\overline{\boldsymbol{X}}_i]|$ using the law of total variance. Letting $\boldsymbol{I}$ be the indicator for $\boldsymbol{X}_i \in \mathcal{L}$ we have

$$\begin{aligned}\mathbf{Var}[\boldsymbol{X}_i] &= \mathbf{E}[\mathbf{Var}[\boldsymbol{X}_i \mid \boldsymbol{I}]] + \mathbf{Var}[\mathbf{E}[\boldsymbol{X}_i \mid \boldsymbol{I}]] \\ &= \mathbf{Pr}[\boldsymbol{I}]\,\mathbf{Var}[\underline{\boldsymbol{X}}_i] + (1-\mathbf{Pr}[\boldsymbol{I}])\,\mathbf{Var}[\overline{\boldsymbol{X}}_i] \\ &\quad + \mathbf{Pr}[\boldsymbol{I}](1-\mathbf{Pr}[\boldsymbol{I}])(\mathbf{E}[\underline{\boldsymbol{X}}_i] - \mathbf{E}[\overline{\boldsymbol{X}}_i])^2. \quad (8)\end{aligned}$$

From this we get:

$$\mathbf{Var}[\boldsymbol{X}_i] \le \ell_i\cdot k^2 + \mathbf{Var}[\overline{\boldsymbol{X}}_i] + \ell_i\cdot 4k^2$$
$$\Longrightarrow \mathbf{Var}[\boldsymbol{X}_i] - \mathbf{Var}[\overline{\boldsymbol{X}}_i] \le \ell_i\cdot 5k^2.$$

For a lower bound, we get from (8):

$$\begin{aligned}&(1-\mathbf{Pr}[\boldsymbol{I}])(\mathbf{Var}[\boldsymbol{X}_i] - \mathbf{Var}[\overline{\boldsymbol{X}}_i]) \\ &= \mathbf{Pr}[\boldsymbol{I}](\mathbf{Var}[\underline{\boldsymbol{X}}_i] - \mathbf{Var}[\boldsymbol{X}_i]) \\ &\quad + \mathbf{Pr}[\boldsymbol{I}](1-\mathbf{Pr}[\boldsymbol{I}])(\mathbf{E}[\underline{\boldsymbol{X}}_i] - \mathbf{E}[\overline{\boldsymbol{X}}_i])^2 \\ &\ge -\mathbf{Pr}[\boldsymbol{I}]\,\mathbf{Var}[\boldsymbol{X}_i].\end{aligned}$$

Hence, $\mathbf{Var}[\boldsymbol{X}_i] - \mathbf{Var}[\overline{\boldsymbol{X}}_i] \ge -\frac{\ell_i}{1-\ell_i}\mathbf{Var}[\boldsymbol{X}_i] \ge -\frac{\ell_i}{1-\ell_i}k^2 \ge -2k^3\ell_i$.

Thus $|\,\mathbf{Var}[\boldsymbol{X}_i] - \mathbf{Var}[\overline{\boldsymbol{X}}_i]| \le 5k^3\ell_i$ and so the second sum in (7) is at most $5k^3\sum_i \ell_i \le 10k^4\cdot M$.

We conclude

$$|\sigma^2 - \sigma_L^2| \le 14k^4\log(1/\delta)\cdot M.$$

■

Given Claim IV.4, (4) follows from (5), (6), Proposition II.8, and our assumption that $\sigma^2 \ge 15k^4\log(1/\delta)\cdot M$. This completes the proof of Theorem IV.3. ■

Corollary IV.5. *Let $\boldsymbol{S}$ be a k-SIIRV with mean μ and variance σ^2. Moreover, let $0 < \delta < \frac{1}{10}$ and assume $\sigma^2 \ge 15(k^{18}/\delta^6)\log^2(1/\delta)$. Then, for some integer c with $1 \le c \le k-1$, we have that $\mathrm{d_{TV}}(\boldsymbol{S}, \boldsymbol{Y}+c\boldsymbol{Z}) \le O(\delta)$, where $\boldsymbol{Y}$ and $\boldsymbol{Z}$ are independent, $\boldsymbol{Y}$ is a c-IRV and $\boldsymbol{Z} \sim Z(\frac{\mu}{c}, \frac{\sigma^2}{c^2})$.*

Proof: The claim is trivial for $k = 1$ so we assume that $k \ge 2$. By subtracting an appropriate integer constant from each component $\boldsymbol{X}_i$ of the k-SIIRV $\boldsymbol{S}$, we can obtain a 0-moded $\pm k$-SIIRV $\boldsymbol{S}'$ such that $\boldsymbol{S} = \boldsymbol{S}' + m$ for some $m \in \mathbb{Z}$. Note that $\boldsymbol{S}'$ has mean $\mu - m$ and variance σ^2. Now apply Theorem IV.3 to $\boldsymbol{S}'$ with $M = k^7/\delta^2$, calling the obtained random variables $\boldsymbol{Y}'$ and $\boldsymbol{Z}'$. (We leave it to the reader to verify, with the aid of Proposition IV.7 below, that the lower bound on σ means there is at least one M-heavy integer and hence the obtained c is nonzero.) $\boldsymbol{Y}'$ and $\boldsymbol{Z}'$ are independent, $\boldsymbol{Z}' \sim Z(\frac{\mu-m}{c}, \frac{\sigma^2}{c^2})$, and $\boldsymbol{Y}'$ is a $\pm M''$-IRV, where $M'' = 4(k^9/\delta^2)\log(1/\delta)$. Moreover,

$$\mathrm{d_{TV}}(\boldsymbol{S}', \boldsymbol{Y}' + c\boldsymbol{Z}') \le O(\delta + (k^9/\delta^2)\log(1/\delta)/\sigma) \le O(\delta), \quad (9)$$

where the second inequality is by our assumed lower bound on σ.

Next, write $m = qc + r$ for some integers q, r with $|r| \le c/2 \le k$. Defining $\boldsymbol{Y}'' = \boldsymbol{Y}' + r$, clearly $\boldsymbol{Y}''$ is a $\pm(M''+k)$-IRV. Moreover, it follows from Proposition II.9 that $\mathrm{d_{TV}}(\boldsymbol{Z}'+q, \boldsymbol{Z}) \le 1/\sigma$. So assuming $\boldsymbol{Z}$ is independent of $\boldsymbol{Y}''$, using the triangle inequality, Proposition II.6, and (9), we obtain:

$$\mathrm{d_{TV}}(\boldsymbol{S}, \boldsymbol{Y}'' + c\boldsymbol{Z}) \le O(\delta). \quad (10)$$

Finally, define two dependent random variables $\boldsymbol{Y} = \boldsymbol{Y}(\boldsymbol{Y}'')$ and $\boldsymbol{Q} = \boldsymbol{Q}(\boldsymbol{Y}'')$ such that $\boldsymbol{Y} = \boldsymbol{Y}''$ mod c, which is a c-IRV, and $\boldsymbol{Y}'' = c\boldsymbol{Q} + \boldsymbol{Y}$, so that $\boldsymbol{Q}$ is a $\pm\lfloor\frac{M''+k}{c}\rfloor$-IRV. With this definition, we have $\boldsymbol{Y}'' + c\boldsymbol{Z} = \boldsymbol{Y} + c(\boldsymbol{Z} + \boldsymbol{Q})$. The proof is concluded by noting the following:

Claim IV.6. $\mathrm{d_{TV}}(\boldsymbol{Y} + c(\boldsymbol{Z}+\boldsymbol{Q}), \boldsymbol{Y} + c\boldsymbol{Z}) \le \frac{\lfloor\frac{M''+k}{c}\rfloor}{2\sigma}$.

Proof of Claim IV.6: First, by iterating Proposition II.10 and using the triangle inequality, we have that for any integer λ, $\mathrm{d_{TV}}(\boldsymbol{Z}, \boldsymbol{Z}+\lambda) \le \frac{\lambda}{2\sigma}$.

For the following derivation, whenever $\boldsymbol{X}$ is a random variable, we write $f_{\boldsymbol{X}}$ for the probability density function of $\boldsymbol{X}$. We have:

$$
\begin{aligned}
&\mathrm{d_{TV}}(\boldsymbol{Y}+c(\boldsymbol{Z}+\boldsymbol{Q}),\boldsymbol{Y}+c\boldsymbol{Z}) \\
&= \frac{1}{2}\int_{-\infty}^{+\infty} |f_{\boldsymbol{Y}+c(\boldsymbol{Z}+\boldsymbol{Q})}(x) - f_{\boldsymbol{Y}+c\boldsymbol{Z}}(x)|dx \\
&= \frac{1}{2}\int_{-\infty}^{+\infty} |f_{\boldsymbol{Y}+c(\boldsymbol{Z}+\boldsymbol{Q})}(x) - f_{\boldsymbol{Y}+c\boldsymbol{Z}}(x)|dx \\
&= \frac{1}{2}\int_{-\infty}^{+\infty} \Bigg| \int_{-\infty}^{+\infty} f_{\boldsymbol{Y}(y'')+c(\boldsymbol{Z}+\boldsymbol{Q}(y''))}(x) f_{\boldsymbol{Y}''}(y'')dy'' \\
&\qquad\qquad - f_{\boldsymbol{Y}(y'')+c\boldsymbol{Z}}(x) f_{\boldsymbol{Y}''}(y'')dy'' \Bigg| dx \\
&\leq \int_{-\infty}^{+\infty} \frac{1}{2}\int_{-\infty}^{+\infty} \Bigg| f_{\boldsymbol{Y}(y'')+c(\boldsymbol{Z}+\boldsymbol{Q}(y''))}(x) \\
&\qquad\qquad - f_{\boldsymbol{Y}(y'')+c\boldsymbol{Z}}(x) \Bigg| dx f_{\boldsymbol{Y}''}(y'')dy'' \\
&\leq \int_{-\infty}^{+\infty} \mathrm{d_{TV}}(\boldsymbol{Y}(y'') \\
&\qquad\qquad + c(\boldsymbol{Z}+\boldsymbol{Q}(y'')), \boldsymbol{Y}(y'')+c\boldsymbol{Z}) f_{\boldsymbol{Y}''}(y'')dy'' \\
&\leq \int_{-\infty}^{+\infty} \mathrm{d_{TV}}(\boldsymbol{Z}+\boldsymbol{Q}(y''), \boldsymbol{Z}) f_{\boldsymbol{Y}''}(y'')dy'' \\
&\leq \frac{\lfloor \frac{M''+k}{c} \rfloor}{2\sigma}.
\end{aligned}
$$

■

Using the triangle inequality, Claim IV.6 and (10), we obtain: $\mathrm{d_{TV}}(\boldsymbol{S}, \boldsymbol{Y}+c\boldsymbol{Z}) \leq O(\delta) + \frac{\lfloor \frac{M''+k}{c} \rfloor}{2\sigma} = O(\delta)$, by our lower bound on σ. ■

Proposition IV.7. *Let $\boldsymbol{X}$ be a $\pm k$-IRV with mode 0. Let $w = \mathbf{Pr}[\boldsymbol{X} \neq 0]$. Then $\frac{1}{8}w \leq \mathbf{Var}[\boldsymbol{X}] \leq k^2 w$.*

Proof: For the upper bound we have $\mathbf{Var}[\boldsymbol{X}] \leq \mathbf{E}[\boldsymbol{X}^2] \leq \mathbf{Pr}[\boldsymbol{X} \neq 0]k^2 = k^2 w$. As for the lower bound, let $\mu = \mathbf{E}[\boldsymbol{X}]$ and write $m = \lfloor \mu \rceil$. If $m = 0$ then whenever $\boldsymbol{X} \neq 0$ we have $|\boldsymbol{X} - \mu| \geq \frac{1}{2}$; hence $\mathbf{Var}[\boldsymbol{X}] = \mathbf{E}[(\boldsymbol{X}-\mu)^2] \geq (\frac{1}{2})^2 w \geq \frac{1}{8}w$. If $m \neq 0$ then we have $\mathbf{Pr}[\boldsymbol{X} \neq m] \geq \frac{1}{2}$ (else m would be the mode of $\boldsymbol{X}$); hence $\mathbf{E}[(\boldsymbol{X}-\mu)^2] \geq \frac{1}{2}(\frac{1}{2})^2 = \frac{1}{8} \geq \frac{1}{8}w$. ■

We conclude this section with the following corollary, which is another way of stating Theorem I.2.

Corollary IV.8. *Let $\boldsymbol{S} = \boldsymbol{X}_1 + \cdots + \boldsymbol{X}_n$ be a k-SIIRV for some positive integer k. Let μ and σ^2 be respectively the mean and variance of $\boldsymbol{S}$. Then, for all $\epsilon > 0$, the distribution of $\boldsymbol{S}$ is $O(\epsilon)$-close in total variation distance to one of the following:*

1) *a random variable supported on $\frac{k^9}{\epsilon^4}$ consecutive integers; or*
2) *the sum of two independent random variables $\boldsymbol{S}_1 + c\boldsymbol{S}_2$, where c is some positive integer $1 \leq c \leq k-1$, $\boldsymbol{S}_2$ is distributed according to $Z(\mu, \sigma^2)$, and $\boldsymbol{S}_1$ is a c-IRV; in this case, $\sigma^2 = \Omega\left(\frac{k^{18}}{\epsilon^6}\log^2(1/\epsilon)\right)$.*

Proof: Assume $\epsilon < 1/10$. We distinguish two cases depending on whether σ^2 is $<$ or $\geq 15(k^{18}/\epsilon^6)\log^2(1/\epsilon)$. In the former case, we have by Chebyshev's inequality that $\boldsymbol{S}$ is ϵ-close to a random variable supported on $O(\frac{k^9}{\epsilon^4})$ consecutive integers, as in the first case of the statement. In the latter case, we can apply Corollary IV.5 to get that $\boldsymbol{S}$ is close to $\boldsymbol{S}_1 + c\boldsymbol{S}_2$ as in the second case of the statement. ■

V. Learning Sums of Independent Integer Random Variables

Using ideas and tools from previous work on learning discrete distributions [DDS12b], [CDSS13], we apply our main structural result, Corollary IV.8, to prove our main learning result, Theorem I.1. The algorithm of Theorem I.1 works by first running two different learning procedures, corresponding to the "small variance" and "large variance" cases of Corollary IV.8 respectively. It then does hypothesis testing to select a final hypothesis from the hypotheses thus obtained. Due to space limitations the proof is deferred to the full version of this paper.

Acknowledgment

Part of this work was done while I.D. was at UC Berkeley supported by a Simons Postdoctoral Fellowship. R.O. is supported by NSF grants CCF-0747250 and CCF-1116594, a Sloan fellowship, and a grant from the MSR–CMU Center for Computational Thinking. R.S. is supported by NSF grant CCF-1115703. L.-Y.T. is supported by NSF grant CCF-1115703; part of this research was completed while visiting CMU.

References

[AHK12] A. Anandkumar, D. Hsu, and S. Kakade. A method of moments for mixture models and Hidden Markov Models. *Journal of Machine Learning Research - Proceedings Track*, 23:33.1–33.34, 2012.

[AK01] S. Arora and R. Kannan. Learning mixtures of arbitrary Gaussians. In *Proceedings of the 33rd Symposium on Theory of Computing*, pages 247–257, 2001.

[BGK04] T. Batu, S. Guha, and S. Kannan. Inferring mixtures of Markov chains. In *COLT*, pages 186–199, 2004.

[BHJ92] A.D. Barbour, L. Holst, and S. Janson. *Poisson Approximation*. Oxford University Press, New York, NY, 1992.

[BS10] Mikhail Belkin and Kaushik Sinha. Polynomial learning of distribution families. In *FOCS*, pages 103–112, 2010.

[BX99] Andrew Barbour and Aihua Xia. Poisson perturbations. *European Series in Applied and Industrial Mathematics. Probability and Statistics*, 3:131–150, 1999.

[CDSS13] S. Chan, I. Diakonikolas, R. Servedio, and X. Sun. Learning mixtures of structured distributions over discrete domains. In *SODA*, 2013.

[CGG02] M. Cryan, L. Goldberg, and P. Goldberg. Evolutionary trees can be learned in polynomial time in the two state general Markov model. *SIAM Journal on Computing*, 31(2):375–397, 2002.

[CGS11] Louis Chen, Larry Goldstein, and Qi-Man Shao. *Normal Approximation by Stein's Method*. Springer, 2011.

[CL10] Louis H. Y. Chen and Yu K. Leong. From zero-bias to discretized normal approximation. 2010.

[Das99] S. Dasgupta. Learning mixtures of Gaussians. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, pages 634–644, 1999.

[DDS12a] C. Daskalakis, I. Diakonikolas, and R.A. Servedio. Learning k-modal distributions via testing. In *SODA*, pages 1371–1385, 2012.

[DDS12b] C. Daskalakis, I. Diakonikolas, and R.A. Servedio. Learning Poisson Binomial Distributions. In *STOC*, pages 709–728, 2012.

[DG85] L. Devroye and L. Györfi. *Nonparametric Density Estimation: The L_1 View*. John Wiley & Sons, 1985.

[DHKS05] A. Dasgupta, J. Hopcroft, J. Kleinberg, and M. Sandler. On learning mixtures of heavy-tailed distributions. In *FOCS*, pages 491–500, 2005.

[DL01] L. Devroye and G. Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics, Springer, 2001.

[DP11] C. Daskalakis and C. Papadimitriou. On Oblivious PTAS's for Nash Equilibrium. *STOC* 2009, pp. 75–84. Full version available as ArXiV report, 2011.

[DS00] S. Dasgupta and L. Schulman. A two-round variant of EM for Gaussian mixtures. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, pages 143–151, 2000.

[Fan12] X. Fang. Discretized normal approximation by Steins method. arXiv:1111.3162v3, 2012.

[FM99] Y. Freund and Y. Mansour. Estimating a mixture of two product distributions. In *Proceedings of the Twelfth Annual Conference on Computational Learning Theory*, pages 183–192, 1999.

[FOS05] J. Feldman, R. O'Donnell, and R. Servedio. Learning mixtures of product distributions over discrete domains. In *Proc. 46th Symposium on Foundations of Computer Science (FOCS)*, pages 501–510, 2005.

[FOS06] J. Feldman, R. O'Donnell, and R. Servedio. PAC learning mixtures of Gaussians with no separation assumption. In *Proc. 19th Annual Conference on Learning Theory (COLT)*, pages 20–34, 2006.

[GMRZ11] Parikshit Gopalan, Raghu Meka, Omer Reingold, and David Zuckerman. Pseudorandom generators for combinatorial shapes. In *STOC*, pages 253–262, 2011.

[KMR$^+$94] M. Kearns, Y. Mansour, D. Ron, R. Rubinfeld, R. Schapire, and L. Sellie. On the learnability of discrete distributions. In *Proceedings of the 26th Symposium on Theory of Computing*, pages 273–282, 1994.

[KMV10] Adam Tauman Kalai, Ankur Moitra, and Gregory Valiant. Efficiently learning mixtures of two Gaussians. In *STOC*, pages 553–562, 2010.

[Kru86] J. Kruopis. Precision of approximation of the generalized binomial distribution by convolutions of poisson measures. *Lithuanian Mathematical Journal*, 26(1):37–49, 1986.

[MR05] E. Mossel and S. Roch. Learning nonsingular phylogenies and Hidden Markov Models. In *To appear in Proceedings of the 37th Annual Symposium on Theory of Computing (STOC)*, 2005.

[MV10] Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of Gaussians. In *FOCS*, pages 93–102, 2010.

[Pre83] E.L. Presman. Approximation of binomial distributions by infinitely divisible ones. *Theory Probab. Appl.*, 28:393–403, 1983.

[RÖ7] A. Röllin. Translated Poisson Approximation Using Exchangeable Pair Couplings. *Annals of Applied Probability*, 17(5/6):1596–1614, 2007.

[RR12] A. Röllin and N. Ross. Local limit theorems via Landau-Kolmogorov inequalities. arXiv:1011.3100v2, 2012.

[RSS12] Y. Rabani, L. Schulman, and C. Swamy. Learning mixtures of arbitrary distributions over large discrete domains. at http://arxiv.org/abs/1212.1527, 2012.

[Sco92] D.W. Scott. *Multivariate Density Estimation: Theory, Practice and Visualization*. Wiley, New York, 1992.

[Sil86] B. W. Silverman. *Density Estimation*. Chapman and Hall, London, 1986.

[Val84] L. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

[VV11] Gregory Valiant and Paul Valiant. Estimating the unseen: an $n/\log(n)$-sample estimator for entropy and support size, shown optimal via new CLTs. In *STOC*, pages 685–694, 2011.

[VW02] S. Vempala and G. Wang. A spectral algorithm for learning mixtures of distributions. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science*, pages 113–122, 2002.