# Sample-Efficient Learning of Mixtures 

Hassan Ashtiani, Shai Ben-David<br>David R. Cheriton School of Computer Science<br>University of Waterloo<br>200 University Avenue West<br>Waterloo, Ontario, Canada, N2L 3G1<br>\{mhzokaei, shai\}@uwaterloo.ca

Abbas Mehrabian<br>Simons Institute for the Theory of Computing 121 Calvin Lab \#2190<br>University of California, Berkeley<br>Berkeley, CA, USA 94720-2190<br>abbasmehrabian@gmail.com


#### Abstract

We consider PAC learning of probability distributions (a.k.a. density estimation), where we are given an i.i.d. sample generated from an unknown target distribution, and want to output a distribution that is close to the target in total variation distance. Let $\mathcal{F}$ be an arbitrary class of probability distributions, and let $\mathcal{F}^{k}$ denote the class of $k$-mixtures of elements of $\mathcal{F}$. Assuming the existence of a method for learning $\mathcal{F}$ with sample complexity $m_{\mathcal{F}}(\varepsilon)$, we provide a method for learning $\mathcal{F}^{k}$ with sample complexity $O\left(k \log k \cdot m_{\mathcal{F}}(\varepsilon) / \varepsilon^{2}\right)$. Our mixture learning algorithm has the property that, if the $\mathcal{F}$ learner is proper and agnostic, then the $\mathcal{F}^{k}$-learner would be proper and agnostic as well. This general result enables us to improve the best known sample complexity upper bounds for a variety of important mixture classes. First, we show that the class of mixtures of $k$ axis-aligned Gaussians in $\mathbb{R}^{d}$ is PAC-learnable in the agnostic setting with $\widetilde{O}\left(k d / \varepsilon^{4}\right)$ samples, which is tight in $k$ and $d$ up to logarithmic factors. Second, we show that the class of mixtures of $k$ Gaussians in $\mathbb{R}^{d}$ is PAC-learnable in the agnostic setting with sample complexity $\widetilde{O}\left(k d^{2} / \varepsilon^{4}\right)$, which improves the previous known bounds of $\widetilde{O}\left(k^{3} d^{2} / \varepsilon^{4}\right)$ and $\widetilde{O}\left(k^{4} d^{4} / \varepsilon^{2}\right)$ in its dependence on $k$ and $d$. Finally, we show that the class of mixtures of $k$ log-concave distributions over $\mathbb{R}^{d}$ is PAClearnable using $\widetilde{O}\left(d^{(d+5) / 2} \varepsilon^{-(d+9) / 2} k\right)$ samples.


## 1 Introduction

Learning distributions is a fundamental problem in statistics and computer science, and has numerous applications in machine learning and signal processing. The problem can be stated as:

Given an i.i.d. sample generated from an unknown probability distribution $g$, find a distribution $\hat{g}$ that is close to $g$ in total variation distance. ${ }^{1}$

This strong notion of learning is not possible in general using a finite number of samples. However, if we assume that the target distribution belongs to or can be approximated by a family of distributions, then there is hope to acquire

[^0]algorithms with finite-sample guarantees. In this paper, we study the important family of mixture models within this framework.

Notice that we consider PAC learning of distributions (a.k.a. density estimation), which is different from parameter estimation. In the parameter estimation problem, it is assumed that the target distribution belongs to some parametric class, and the goal is to learn/identify the parameters (see, e.g., (Dasgupta 1999; Belkin and Sinha 2010; Moitra and Valiant 2010)).

As an example of our setting, assume that the target distribution is a Gaussian mixture with $k$ components in $\mathbb{R}^{d}$. Then, how many examples do we need to find a distribution that is $\varepsilon$-close to the target? This sample complexity question, as well as the corresponding computational complexity question, has received a lot of attention recently (see, e.g. (Feldman, Servedio, and O'Donnell 2006; Chan et al. 2014; Suresh et al. 2014; Diakonikolas et al. 2016; Diakonikolas, Kane, and Stewart 2017b; Acharya et al. 2017)).

In this paper, we consider a scenario in which we are given a method for learning a class of distributions (e.g., Gaussians). Then, we ask whether we can use it, as a black box, to come up with an algorithm for learning a mixture of such distributions (e.g., mixture of Gaussians). We will show that the answer to this question is affirmative.

We propose a generic method for learning mixture models. Roughly speaking, we show that by going from learning a single distribution from a class to learning a mixture of $k$ distributions from the same class, the sample complexity is multiplied by a factor of at most $\left(k \log ^{2} k\right) / \varepsilon^{2}$. This result is general, and yet it is surprisingly tight in many important cases. In this paper, we assume that the algorithm knows the number of components $k$.

As a demonstration, we show that our method provides a better sample complexity for learning mixtures of Gaussians than the state of the art. In particular, for learning mixtures of $k$ Gaussians in $\mathbb{R}^{d}$, our method requires $\widetilde{O}\left(d^{2} k / \varepsilon^{4}\right)$ samples, improving by a factor of $k^{2}$ over the $\widetilde{O}\left(d^{2} k^{3} / \varepsilon^{4}\right)$ bound of (Diakonikolas, Kane, and Stewart 2017b). Furthermore, for the special case of mixtures of axis-aligned Gaussians, we provide an upper bound of $\widetilde{O}\left(d k / \varepsilon^{4}\right)$, which is the first optimal bound with respect to $k$ and $d$ up to logarithmic fac-
tors, and improves upon the $\widetilde{O}\left(d k^{9} / \varepsilon^{4}\right)$ bound of (Suresh et al. 2014), which is only shown for the subclass of spherical Gaussians.

One merit of our approach is that it can be applied in the agnostic (a.k.a. robust) setting, where the target distribution does not necessarily belong to the mixture model of choice. To guarantee such a result, we need the black box to work in the agnostic setting. For example, an agnostic learning method for learning Gaussians can be lifted to work for Gaussian mixtures in the agnostic setting.

We would like to emphasize that our focus is on the information-theoretic aspects of learning rather than the computational ones; in particular, although our framework is algorithmic, its running time is exponential in the number of required samples. Proving sample complexity bounds is important in understanding the statistical nature of various classes of distributions (see, e.g., the recent work of (Diakonikolas, Kane, and Stewart 2017a)), and may pave the way for developing efficient algorithms.

## Our Results

Let $\mathcal{F}$ be a class of probability distributions, and let $\mathcal{F}^{k}$ denote the class of $k$-mixtures of elements of $\mathcal{F}$. In our main result, Theorem 5, assuming the existence of a method for learning $\mathcal{F}$ with sample complexity $m_{\mathcal{F}}(\varepsilon)$, we provide a method for learning $\mathcal{F}^{k}$ with sample complexity $O\left(k \log ^{2} k \cdot m_{\mathcal{F}}(\varepsilon) / \varepsilon^{2}\right)$. Our mixture learning algorithm has the property that, if the $\mathcal{F}$-learner is proper, then the $\mathcal{F}^{k}{ }_{-}$ learner would be proper as well (i.e., the learner will always output a member of $\mathcal{F}^{k}$ ). Furthermore, the algorithm works in the more general agnostic setting provided that the base learners are agnostic learners.

We provide several applications of our main result. In Theorem 11, we show that the class of mixtures of $k$ axisaligned Gaussians in $\mathbb{R}^{d}$ is PAC-learnable in the agnostic setting with sample complexity $O\left(k d \log ^{2} k / \varepsilon^{4}\right)$ (see Theorem 12). This bound is tight in terms of $k$ and $d$ up to logarithmic factors. In Theorem 14, we show that the class of mixtures of $k$ Gaussians in $\mathbb{R}^{d}$ is PAC-learnable in the agnostic setting with sample complexity $O\left(k d^{2} \log ^{2} k / \varepsilon^{4}\right)$. Finally, in Theorem 16, we prove that the class of mixtures of $k$ log-concave distributions over $\mathbb{R}^{d}$ is PAC-learnable using $\widetilde{O}\left(d^{(d+5) / 2} \varepsilon^{-(d+9) / 2} k\right)$ samples. To the best of our knowledge, this is the first upper bound on the sample complexity of learning this class.

## Related Work

PAC learning of distributions was introduced by (Kearns et al. 1994), we refer the reader to (Diakonikolas 2016) for a recent survey. A closely related line of research in statistics (in which more emphasis is on sample complexity) is density estimation, for which the book by (Devroye and Lugosi 2001) is an excellent resource.

One approach for studying the sample complexity of learning a class of distributions is bounding the VCdimension of its associated Yatracos class (see Definition 20), and applying results such as Theorem 22. (These VC-dimensions have mainly been studied for the purpose of
proving generalization bounds for neural networks with sigmoid activation functions.) In particular, the VC-dimension bound of (Anthony and Bartlett 1999, Theorem 8.14) which is based on the work of (Karpinski and Macintyre 1997) - implies a sample complexity upper bound of $O\left(\left(k^{4} d^{2}+k^{3} d^{3}\right) / \varepsilon^{2}\right)$ for PAC learning mixtures of axisaligned Gaussians, and an upper bound of $O\left(k^{4} d^{4} / \varepsilon^{2}\right)$ for PAC learning mixtures of general Gaussians (both results hold in the more general agnostic setting).

A sample complexity upper bound of $O\left(d^{2} k^{3} \log ^{2} k / \varepsilon^{4}\right)$ for learning mixtures of Gaussians in the realizable setting was proved in (Diakonikolas, Kane, and Stewart 2017b, Theorem A.1) (the running time of their algorithm is not polynomial). Our algorithm is motivated by theirs, but we have introduced several new ideas in the algorithm and in the analysis, which has resulted in improving the sample complexity bound by a factor of $k^{2}$ and an algorithm that works in the more general agnostic setting.
For mixtures of spherical Gaussians, a polynomial time algorithm for the realizable setting with sample complexity $O\left(d k^{9} \log ^{2}(d) / \varepsilon^{4}\right)$ was proposed in (Suresh et al. 2014, Theorem 11). We improve their sample complexity by a factor of $\widetilde{O}\left(k^{8}\right)$, and moreover our algorithm works in the agnostic setting, too. In the special case of $d=1$, a non-proper agnostic polynomial time algorithm with the optimal sample complexity of $\widetilde{O}\left(k / \varepsilon^{2}\right)$ was given in (Chan et al. 2014), and a proper agnostic algorithm with the same sample complexity and better running time was given in (Li and Schmidt 2017).

An important question, which we do not address in this paper, is finding polynomial time algorithms for learning distributions. To the best of our knowledge, no polynomial time algorithm for learning mixtures of general Gaussians is known. See (Diakonikolas, Kane, and Stewart 2017b) for the state-of-the-art results. Another important setting is computational complexity in the agnostic learning, see, e.g., (Diakonikolas et al. 2016) for some positive results.

A related line of research is parameter estimation for mixtures of Gaussians, see, e.g., (Dasgupta 1999; Belkin and Sinha 2010; Moitra and Valiant 2010), who gave polynomial time algorithms for this problem assuming certain separability conditions (these algorithms are polynomial in the dimension and the error tolerance but exponential in the number of components). Recall that parameter estimation is a more difficult problem and any algorithm for parameter estimation requires some separability assumptions for the target Gaussians, whereas for density estimation no such assumption is needed. E.g., consider the case that $k=2$ and the two components are identical; then there is no way to learn their mixing weights.

We finally remark that characterizing the sample complexity of learning a class of distributions in general is an open problem, even for the realizable (i.e., non-agnostic) case (see (Diakonikolas 2016, Open Problem 15.1)).

## 2 The Formal Framework

Generally speaking, a distribution learning method is an algorithm that takes a sample of i.i.d. points from distribution
$g$ as input, and outputs (a description) of a distribution $\hat{g}$ as an estimation for $g$. Furthermore, we assume that $g$ belongs to or can be approximated by class $\mathcal{F}$ of distributions, and we may require that $\hat{g}$ also belongs to this class (i.e., proper learning).

Let $f_{1}$ and $f_{2}$ be two probability distributions defined over the Borel $\sigma$-algebra $\mathcal{B}$. The total variation distance between $f_{1}$ and $f_{2}$ is defined as

$$
\left\|f_{1}-f_{2}\right\|_{T V}=\sup _{B \in \mathcal{B}}\left|f_{1}(B)-f_{2}(B)\right|=\frac{1}{2}\left\|f_{1}-f_{2}\right\|_{1}
$$

where

$$
\|f\|_{1}:=\int_{-\infty}^{+\infty}|f(x)| \mathrm{d} x
$$

is the $L_{1}$ norm of $f$. In the following definitions, $\mathcal{F}$ is a class of probability distributions, and $g$ is a distribution not necessarily in $\mathcal{F}$. Denote the set $\{1,2, \ldots, m\}$ by $[m]$. All logarithms are in the natural base. For a function $g$ and a class of distributions $\mathcal{F}$, we define

$$
\operatorname{OPT}(\mathcal{F}, g):=\inf _{f \in \mathcal{F}}\|f-g\|_{1}
$$

Definition 1 ( $\varepsilon$-approximation, $(\varepsilon, C)$-approximation). $A$ distribution $\hat{g}$ is an $\varepsilon$-approximation for $g$ if $\|\hat{g}-g\|_{1} \leq \varepsilon$. $A$ distribution $\hat{g}$ is an $(\varepsilon, C)$-approximation for $g$ with respect to $\mathcal{F}$ if

$$
\|\hat{g}-g\|_{1} \leq C \times \operatorname{OPT}(\mathcal{F}, g)+\varepsilon
$$

Definition 2 (PAC-Learning Distributions, Realizable Setting). A distribution learning method is called a (realizable) PAC-learner for $\mathcal{F}$ with sample complexity $m_{\mathcal{F}}(\varepsilon, \delta)$, if for all distribution $g \in \mathcal{F}$ and all $\varepsilon, \delta>0$, given $\varepsilon$, $\delta$, and $a$ sample of size $m_{\mathcal{F}}(\varepsilon, \delta)$, with probability at least $1-\delta$ outputs an $\varepsilon$-approximation of $g$.
Definition 3 (PAC-Learning Distributions, Agnostic Setting). For $C>0$, a distribution learning method is called a $C$-agnostic PAC-learner for $\mathcal{F}$ with sample complexity $m_{\mathcal{F}}^{C}(\varepsilon, \delta)$, if for all distributions $g$ and all $\varepsilon, \delta>0$, given $\varepsilon, \delta$, and a sample of size $m_{\mathcal{F}}^{C}(\varepsilon, \delta)$, with probability at least $1-\delta$ outputs an $(\varepsilon, C)$-approximation of $g .{ }^{2}$

Clearly, a $C$-agnostic PAC-learner (for any constant $C$ ) is also a realizable PAC-learner, with the same error parameter $\varepsilon$. Conversely a realizable PAC-learner can be thought of an $\infty$-agnostic PAC-learner.

## 3 Learning Mixture Models

Let $\Delta_{n}$ denote the $n$-dimensional simplex:

$$
\Delta_{n}=\left\{\left(w_{1}, \ldots, w_{n}\right): w_{i} \geq 0, \sum_{i=1}^{k} w_{i}=1\right\}
$$

Definition 4. Let $\mathcal{F}$ be a class of probability distributions. Then the class of $k$-mixtures of $\mathcal{F}$, written $\mathcal{F}^{k}$, is defined as $\underline{\mathcal{F}^{k}}:=\left\{\sum_{i=1}^{k} w_{i} f_{i}:\left(w_{1}, \ldots, w_{k}\right) \in \Delta_{k}, f_{1}, \ldots, f_{k} \in \mathcal{F}\right\}$.

[^1]Assume that we have a method to PAC-learn $\mathcal{F}$. Does this mean that we can PAC-learn $\mathcal{F}^{k}$ ? And if so, what is the sample complexity of this task? Our main theorem gives an affirmative answer to the first question, and provides a bound for sample complexity of learning $\mathcal{F}^{k}$.
Theorem 5. Assume that $\mathcal{F}$ has a $C$-agnostic PAC-learner with sample complexity $m_{\mathcal{F}}^{C}(\varepsilon, \delta)=\lambda(\mathcal{F}, \delta) / \varepsilon^{\alpha}$ for some $C>0, \alpha \geq 1$ and some function $\lambda(\mathcal{F}, \delta)=\Omega(\log (1 / \delta))$. Then there exists a 3C-agnostic PAC-learner for the class $\mathcal{F}^{k}$ requiring $m_{\mathcal{F}^{k}}^{3 C}(\varepsilon, \delta)=$

$$
O\left(\frac{\lambda\left(\mathcal{F}, \frac{\delta}{3 k}\right) k \log k}{\varepsilon^{\alpha+2}}\right)=O\left(\frac{k \log k \cdot m_{\mathcal{F}}\left(\varepsilon, \frac{\delta}{3 k}\right)}{\varepsilon^{2}}\right)
$$

samples.
Since a realizable PAC-learner is an $\infty$-agnostic PAClearner, we immediately obtain the following corollary.
Corollary 6. Assume that $\mathcal{F}$ has a realizable PAC-learner with sample complexity $m_{\mathcal{F}}(\varepsilon, \delta)=\lambda(\mathcal{F}, \delta) / \varepsilon^{\alpha}$ for some $\alpha \geq 1$ and some function $\lambda(\mathcal{F}, \delta)=\Omega(\log (1 / \delta))$. Then there exists a realizable PAC-learner for the class $\mathcal{F}^{k}$ requiring $m_{\mathcal{F}^{k}}(\varepsilon, \delta)=$

$$
O\left(\frac{\lambda\left(\mathcal{F}, \frac{\delta}{3 k}\right) k \log k}{\varepsilon^{\alpha+2}}\right)=O\left(\frac{k \log k \cdot m_{\mathcal{F}}\left(\varepsilon, \frac{\delta}{3 k}\right)}{\varepsilon^{2}}\right)
$$

samples.
Some remarks:

1. Our mixture learning algorithm has the property that, if the $\mathcal{F}$-learner is proper, then the $\mathcal{F}^{k}$-learner is proper as well.
2. The computational complexity of the resulting algorithm is exponential in the number of required samples.
3. The condition $\lambda(\mathcal{F}, \delta)=\Omega(\log (1 / \delta))$ is a technical condition that holds for all interesting classes $\mathcal{F}$.
4. One may wonder about tightness of this theorem. In Theorem 2 in (Suresh et al. 2014), it is shown that if $\mathcal{F}$ is the class of spherical Gaussians, we have $m_{\mathcal{F}^{k}}^{O(1)}(\varepsilon, \delta)=$ $\Omega\left(k m_{\mathcal{F}}(\varepsilon, \delta / k)\right)$, therefore, the factor of $k$ is necessary in general. However, it is not clear whether the additional factor of $\log k / \varepsilon^{2}$ in the theorem is tight.
5. The constant 3 (in the $3 C$-agnostic result) comes from (Devroye and Lugosi 2001, Theorem 6.3) (see Theorem 8), and it is not clear whether it is necessary. If we allow for randomized algorithms (which produce a random distribution whose expected distance to the target is bounded by $\varepsilon$ ), then the constant can be improved to 2 , see (Mahalanabis and Stefankovic 2008, Theorem 22).
In the rest of this section we prove Theorem 5. Let $g$ be the true data generating distribution, and let

$$
\begin{equation*}
g^{*}=\underset{f \in \mathcal{F}^{k}}{\operatorname{argmin}}\|g-f\|_{1} \text { and } \rho=\left\|g^{*}-g\right\|_{1}=\operatorname{OPT}\left(\mathcal{F}^{k}, g\right) \tag{1}
\end{equation*}
$$

Note that although $g^{*} \in \mathcal{F}^{k}, g$ itself is not necessarily in the form of a mixture. Since our algorithm works for mixtures, we would first like to write $g$ in the form of a mixture of $k$ distributions, such that each distribution is $\rho$-close to being in $\mathcal{F}$. This is done via the following lemma.
Lemma 7. Suppose that $g$ is a distribution with $\operatorname{OPT}\left(\mathcal{F}^{k}, g\right)=\rho$. Then we may write $g=\sum_{i \in[k]} w_{i} G_{i}$, such that each $G_{i}$ is a distribution, and for each $i$ we have $\operatorname{OPT}\left(\mathcal{F}, G_{i}\right) \leq \rho$.
Proof. Since $\operatorname{OPT}\left(\mathcal{F}^{k}, g\right)=\rho$, we can write

$$
\begin{equation*}
g=\sum_{i=1}^{k} w_{i} f_{i}+h=\sum_{i=1}^{k} w_{i}\left(f_{i}+h\right) \tag{2}
\end{equation*}
$$

with $\left(w_{1}, \ldots, w_{k}\right) \in \Delta_{k}$ and each $f_{i} \in \mathcal{F}$, and $\|h\|_{1}=\rho$. Note that $f_{i}+h$ is not necessarily a density function. Let $\mathcal{D}$ denote the set of density functions, that is, the set of nonnegative functions with unit $L_{1}$ norm. Note that this is a convex set. Since projection is a linear operator, by projecting both sides of (2) onto $\mathcal{D}$ we find

$$
g=\sum_{i=1}^{k} w_{i} G_{i}
$$

where $G_{i}$ is the $L_{1}$ projection of $f_{i}+h$ onto $\mathcal{D}$ (since $g \in$ $\mathcal{D}$, the projection of $g$ onto $\mathcal{D}$ is itself). Also, since $f_{i} \in$ $\mathcal{F} \cap \mathcal{D}$ and projection onto a convex set does not increases distances, we have
$\operatorname{OPT}\left(\mathcal{F}, G_{i}\right) \leq\left\|f_{i}-G_{i}\right\|_{1} \leq\left\|f_{i}-\left(f_{i}+h\right)\right\|_{1}=\|h\|_{1}=\rho$, as required.

By Lemma 7, we have

$$
g=\sum_{i \in[k]} w_{i} G_{i}
$$

where each $G_{i}$ is a probability distribution. Let $\rho_{i}:=$ $\operatorname{OPT}\left(\mathcal{F}, G_{i}\right)$, and by the lemma we have

$$
\begin{equation*}
\sum_{i \in[k]} w_{i} \rho_{i} \leq \rho \tag{3}
\end{equation*}
$$

The idea now is to learn each of the $G_{i}$ 's separately using the agnostic learner for $\mathcal{F}$. We will view $g$ as a mixture of $k$ distributions $G_{1}, G_{2}, \ldots, G_{k}$.

For proving Theorem 5, we will use the following theorem on learning finite classes of distributions, which immediately follows from (Devroye and Lugosi 2001, Theorem 6.3) and a standard Chernoff bound.
Theorem 8. Suppose we are given $M$ candidate distributions $f_{1}, \ldots, f_{M}$ and we have access to i.i.d. samples from an unknown distribution $g$. Then there exists an algorithm that given the $f_{i}$ 's and $\varepsilon>0$, takes $\log \left(3 M^{2} / \delta\right) / 2 \varepsilon^{2}$ samples from $g$, and with probability $\geq 1-\delta / 3$ outputs an index $j \in[M]$ such that

$$
\left\|f_{j}-g\right\|_{1} \leq 3 \min _{i \in[M]}\left\|f_{i}-g\right\|_{1}+4 \varepsilon
$$

Input: $k, \varepsilon, \delta$ and an i.i.d. sample $S$
0 . Let $\widehat{W}$ be an $(\varepsilon / k)$-cover for $\Delta_{k}$ in $\ell_{\infty}$ distance. 1. $\mathcal{C}=\emptyset$ (set of candidate distributions)
2. For each $\left(\widehat{w}_{1}, \ldots, \widehat{w}_{k}\right) \in \widehat{W}$ do:
3. For each possible partition of $S$ into $A_{1}, A_{2}, \ldots, A_{k}$ :
4. Provide $A_{i}$ to the $\mathcal{F}$-learner, and let $\widehat{G}_{i}$ be its output.
5. Add the candidate distribution $\sum_{i \in[k]} \widehat{w}_{i} \widehat{G}_{i}$ to $\mathcal{C}$.
6. Apply the algorithm for finite classes (Theorem 8) to $\mathcal{C}$ and output its result.

Figure 1: Algorithm for learning the mixture class $\mathcal{F}^{k}$

We now describe an algorithm that with probability $\geq$ $1-\delta$ outputs a distribution with $L_{1}$ distance $13 \varepsilon+3 C \rho$ to $g$ (the error parameter is $13 \varepsilon$ instead of $\varepsilon$ just for convenience of the proof; it is clear that this does not change the order of magnitude of sample complexity). The algorithm, whose pseudocode is shown in Figure 1, has two main steps. In the first step we generate a set of candidate distributions, such that at least one of them is $(3 \varepsilon+\rho)$-close to $g$ in $L_{1}$ distance. These candidates are of the form $\sum_{i=1}^{k} \widehat{w}_{i} \widehat{G}_{i}$, where the $\widehat{G}_{i}$ 's are extracted from samples and are estimates for the real components $G_{i}$, and the $\widehat{w}$ 's come from a fixed discretization of $\Delta_{k}$, and are estimates for the real mixing weights $w_{i}$. In the second step, we use Theorem 8 to obtain a distribution that is $(13 \varepsilon+3 C \rho)$-close to $g$.

We start with describing the first step. We take

$$
\begin{equation*}
s=\max \left\{\frac{2 k \lambda(\mathcal{F}, \delta / 3 k)}{\varepsilon^{\alpha}}, \frac{16 k \log (3 k / \delta)}{\varepsilon}\right\} \tag{4}
\end{equation*}
$$

i.i.d. samples from $g$. Let $S$ denote the set of generated points. Note that $\lambda(\mathcal{F}, \delta)=\Omega(\log (1 / \delta))$ implies

$$
s=O\left(k \lambda(\mathcal{F}, \delta / 3 k) \times \varepsilon^{-\alpha}\right)
$$

Let $\widehat{W}$ be an $\varepsilon / k$-cover for $\Delta_{k}$ in $\ell_{\infty}$ distance of cardinality $(k / \varepsilon+1)^{k}$. That is, for any $x \in \Delta_{k}$ there exists $w \in \widehat{W}$ such that $\|w-x\|_{\infty} \leq \varepsilon / k$. This can be obtained from a grid in $[0,1]^{k}$ of side length $\varepsilon / k$, which is an $\varepsilon / k$-cover for $[0,1]^{k}$, and projecting each of its points onto $\Delta_{k}$.

By an assignment, we mean a function $A: S \rightarrow[k]$. The role of an assignment is to "guess" each sample point is coming from which component, by mapping them to a component index. For each pair $\left(A,\left(\widehat{w}_{1}, \ldots, \widehat{w}_{k}\right)\right)$, where $A$ is an assignment and $\left(\widehat{w}_{1}, \ldots, \widehat{w}_{k}\right) \in \widehat{W}$, we generate a candidate distribution as follows: let $A^{-1}(i) \subseteq S$ be those sample points that are assigned to component $i$. For each $i \in[k]$, we provide the set $A^{-1}(i)$ of samples to our $\mathcal{F}$ learner, and the learner provides us with a distribution $\widehat{G}_{i}$. We add the distribution $\sum_{i \in[k]} \widehat{w}_{i} \widehat{G}_{i}$ to the set of candidate distributions.

Lemma 9. With probability $\geq 1-2 \delta / 3$, at least one of the generated candidate distributions is $(3 \varepsilon+C \rho)$-close to $g$.

Before proving the lemma, we show that it implies our main result, Theorem 5. By the lemma, we obtain a set of candidates such that at least one of them is $(3 \varepsilon+C \rho)$-close to $g$ (with failure probability $\leq 2 \delta / 3$ ). This step takes $s=$ $O\left(k \lambda(\mathcal{F}, \delta / 3 k) \times \varepsilon^{-\alpha}\right)$ many samples. Then, we apply Theorem 8 to output one of those candidates that is $(13 \varepsilon+3 C \rho)$ close to $g$ (with failure probability $\leq \delta / 3$ ), therefore using $\log \left(3 M^{2} / \delta\right) / 2 \varepsilon^{2}$ additional samples. Note that the number of generated candidate distributions is $M=k^{s} \times(1+k / \varepsilon)^{k}$. Hence, in the second step of our algorithm, we take

$$
\begin{aligned}
\log \left(3 M^{2} / \delta\right) / 2 \varepsilon^{2} & =O\left(\frac{\lambda(\mathcal{F}, \delta / 3 k) k \log k}{\varepsilon^{\alpha+2}}\right) \\
& =O\left(\frac{m_{\mathcal{F}}(\varepsilon, \delta / 3 k) k \log k}{\varepsilon^{2}}\right)
\end{aligned}
$$

additional samples. The proof is completed noting the total failure probability is at most $\delta$ by the union bound.

We now prove Lemma 9. We will use the following concentration inequality, which holds for any binomial random variable $X$ (see (Mitzenmacher and Upfal 2005, Theorem 4.5(2))):

$$
\begin{equation*}
\operatorname{Pr}\{X<\mathbf{E} X / 2\} \leq \exp (-\mathbf{E} X / 8) \tag{5}
\end{equation*}
$$

Say a component $i$ is negligible if

$$
w_{i} \leq \frac{8 \log (3 k / \delta)}{s}
$$

Let $L \subseteq[k]$ denote the set of negligible components. Let $i$ be a non-negligible component. Note that, the number of points coming from component $i$ is binomial with parameters $s$ and $w_{i}$ and thus has mean $s w_{i}$, so (5) implies that, with probability at least $1-\delta / 3 k, S$ contains at least $w_{i} s / 2$ points from $i$. Since we have $k$ components in total, the union bound implies that, with probability at least $1-\delta / 3$, uniformly for all $i \notin L, S$ contains at least $w_{i} s / 2$ points from component $i$.

Now consider the pair $\left(A,\left(\widehat{w}_{1}, \ldots, \widehat{w}_{k}\right)\right)$ such that $A$ assigns samples to their correct indices, and has the property that $\left|\widehat{w}_{i}-w_{i}\right| \leq \varepsilon / k$ for all $i \in[k]$. We claim that the resulting candidate distribution is $(3 \varepsilon+C \rho)$-close to $g$.

Let $\widehat{G}_{1}, \ldots, \widehat{G}_{k}$ be the distributions provided by the learner. For each $i \in[k]$ define

$$
\varepsilon_{i}:=\left(\frac{2 \lambda(\mathcal{F}, \delta / 3 k)}{w_{i} s}\right)^{1 / \alpha}
$$

For any $i \notin L$, since there exists at least $w_{i} s / 2$ samples for component $i$, and since

$$
w_{i} s / 2=\lambda(\mathcal{F}, \delta / 3 k) \varepsilon_{i}^{-\alpha}=m_{\mathcal{F}}\left(\varepsilon_{i}, \delta / 3 k\right),
$$

we are guaranteed that $\left\|\widehat{G}_{i}-G_{i}\right\|_{1} \leq C \rho_{i}+\varepsilon_{i}$ with probability $1-\delta / 3 k$ (recall that each $G_{i}$ is $\rho_{i}$-close to the class $\mathcal{F})$. Therefore, $\left\|\widehat{G}_{i}-G_{i}\right\|_{1} \leq C \rho_{i}+\varepsilon_{i}$ holds uniformly over all $i \notin L$, with probability $\geq 1-\delta / 3$. Note that since
$\alpha \geq 1$, the function $w_{i}^{1-1 / \alpha}$ is concave in $w_{i}$, so by Jensen's inequality we have

$$
\sum_{i \in[k]} w_{i}^{1-1 / \alpha} \leq k\left(\left(\sum_{i \in[k]} w_{i} / k\right)^{1-1 / \alpha}\right)=k^{1 / \alpha}
$$

hence

$$
\begin{aligned}
\sum_{i \notin L} w_{i} \varepsilon_{i} & =\left(\frac{2 \lambda(\mathcal{F}, \delta / 3 k)}{s}\right)^{1 / \alpha} \sum_{i \notin L} w_{i}^{1-1 / \alpha} \\
& \leq\left(\frac{2 k \lambda(\mathcal{F}, \delta / 3 k)}{s}\right)^{1 / \alpha}
\end{aligned}
$$

Also recall from (3) that $\sum_{i \in[k]} w_{i} \rho_{i} \leq \rho$. Proving the lemma is now a matter of careful applications of the triangle inequality:

$$
\begin{aligned}
& \left\|\sum_{i \in[k]} \widehat{w}_{i} \widehat{G}_{i}-g\right\|_{1}=\left\|\sum_{i \in[k]} \widehat{w}_{i} \widehat{G}_{i}-\sum_{i \in[k]} w_{i} G_{i}\right\|_{1} \\
& \leq\left\|\sum_{i \in[k]} w_{i}\left(\widehat{G}_{i}-G_{i}\right)\right\|_{1}+\left\|\sum_{i \in[k]}\left(\widehat{w}_{i}-w_{i}\right) \widehat{G}_{i}\right\|_{1} \\
& \leq\left\|\sum_{i \in L} w_{i}\left(\widehat{G}_{i}-G_{i}\right)\right\|_{1}+\left\|\sum_{i \notin L} w_{i}\left(\widehat{G}_{i}-G_{i}\right)\right\|_{1} \\
& \quad+\sum_{i \in[k]}\left|\widehat{w}_{i}-w_{i}\right|\left\|\widehat{G}_{i}\right\|_{1} \\
& \leq 2 \sum_{i \in L} w_{i}+\sum_{i \notin L} w_{i}\left(\varepsilon_{i}+C \rho_{i}\right)+\sum_{i \in[k]} \varepsilon / k \times 1 \\
& \leq 2 k \times \frac{8 \log (3 k / \delta)}{s}+\left(\frac{2 k \lambda(\mathcal{F}, \delta / 3 k)}{s}\right)^{1 / \alpha}+C \rho+\varepsilon \\
& \leq \varepsilon+\varepsilon+\varepsilon+C \rho,
\end{aligned}
$$

where for the last inequality we used the definition of $s$ in (4). This completes the proof of Lemma 9.

## 4 Learning Mixtures of Gaussians

Gaussian Mixture Models (GMMs) are probably the most widely studied mixture classes with numerous applications; yet, the sample complexity of learning this class is not fully understood, especially when the number of dimensions is large. In this section, we will show that our method for learning mixtures can improve the state of the art for learning GMMs in terms of sample complexity. In the following, $\mathcal{N}_{d}(\mu, \Sigma)$ denotes a Gaussian density function defined over $\mathbb{R}^{d}$, with mean $\mu$ and covariance matrix $\Sigma$.

## Mixtures of Axis-Aligned Gaussians

A Gaussian is called axis-aligned if its covariance matrix $\Sigma$ is diagonal. The class of axis-aligned Gaussian Mixtures is an important special case of GMMs that is thoroughly studied in the literature (e.g. (Feldman, Servedio, and O'Donnell 2006)).

Theorem 10. Let $\mathcal{F}$ denote the class of d-dimensional axisaligned Gaussians. Then $\mathcal{F}$ is 3-agnostic PAC-learnable with $m_{\mathcal{F}}^{3}(\varepsilon, \delta)=O\left((d+\log (1 / \delta)) / \varepsilon^{2}\right)$.

We defer the proof of this result to Section 7. Combining this theorem with Theorem 5 we obtain the following result:
Theorem 11. The class $\mathcal{F}^{k}$ of mixtures of $k$ axis-aligned Gaussians in $\mathbb{R}^{d}$ is 9-agnostic PAC-learnable with sample complexity $m_{\mathcal{F}^{k}}^{9}(\varepsilon, \delta)=O\left(k d \log k \log (k / \delta) / \varepsilon^{4}\right)$. Accordingly, it is also PAC-learnable in the realizable case with the same number of samples.

This theorem improves the upper bound of $O\left(d k^{9} \log ^{2}(d / \delta) / \varepsilon^{4}\right)$ proved in (Suresh et al. 2014, Theorem 11) for spherical Gaussians in the realizable setting. Spherical Gaussians are special cases of axis-aligned Gaussians in which all eigenvalues of the covariance matrix are equal, i.e., $\Sigma$ is a multiple of the identity matrix. The following minimax lower bound (i.e., worst-case on all instances) on the sample complexity of learning mixtures of spherical Gaussians is proved in the same paper.
Theorem 12 (Theorem 2 in (Suresh et al. 2014)). The class $\mathcal{F}^{k}$ of mixtures of $k$ axis-aligned Gaussians in $\mathbb{R}^{d}$ in the realizable setting has $m_{\mathcal{F}^{k}}(\varepsilon, 1 / 2)=\Omega\left(d k / \varepsilon^{2}\right)$.

Therefore, our upper bound of Theorem 11 is optimal in terms of dependence on $d$ and $k$ (up to logarithmic factors) for axis-aligned Gaussians.

## Mixtures of General Gaussians

For general Gaussians, we have the following result.
Theorem 13. Let $\mathcal{F}$ denote the class of d-dimensional Gaussians. Then, $\mathcal{F}$ is 3-agnostic PAC-learnable with $m_{\mathcal{F}}^{3}(\varepsilon, \delta)=O\left(\left(d^{2}+\log (1 / \delta)\right) / \varepsilon^{2}\right)$.

We defer the proof of this result to Section 7. Combining this theorem with Theorem 5, we obtain the following result: Theorem 14. The class $\mathcal{F}^{k}$ of mixtures of $k$ Gaussians in $\mathbb{R}^{d}$ is 9-agnostic PAC-learnable with sample complexity $m_{\mathcal{F}^{k}}^{9}(\varepsilon, \delta)=O\left(k d^{2} \log k \log (k / \delta) / \varepsilon^{4}\right)$. Accordingly, it is also PAC-learnable in the realizable case with the same number of samples.

This improves by a factor of $k^{2}$ the upper bound of $O\left(k^{3} d^{2} \log k / \varepsilon^{4}\right)$ in the realizable setting, proved in (Diakonikolas, Kane, and Stewart 2017b, Theorem A.1).

Note that Theorem 12 gives a lower bound of $\Omega\left(k d / \varepsilon^{2}\right)$ for $m_{\mathcal{F}^{k}}(\varepsilon, \delta)$, hence the dependence of Theorem 14 on $k$ is optimal (up to logarithmic factors). However, there is a factor of $d / \varepsilon^{2}$ between the upper and lower bounds.

## 5 Mixtures of Log-Concave Distributions

A probability density function over $\mathbb{R}^{d}$ is log-concave if its logarithm is a concave function. The following result about the sample complexity of learning log-concave distributions is the direct consequence of the recent work of (Diakonikolas, Kane, and Stewart 2017a).
Theorem 15. Let $\mathcal{F}$ be the class of distributions corresponding to the set of all log-concave densities over $\mathbb{R}^{d}$. Then $\mathcal{F}$ is 3-agnostic PAC learnable using $m^{3}(\varepsilon, \delta)=$ $O\left((d / \varepsilon)^{(d+5) / 2} \log ^{2}(1 / \varepsilon)\right)$ samples.

Using Theorem 5, we come up with the first result about the sample complexity of learning mixtures of log-concave distributions.

Theorem 16. The class of mixtures of $k$ log-concave distributions over $\mathbb{R}^{d}$ is 9-agnostic PAC-learnable using $\widetilde{O}\left(d^{(d+5) / 2} \varepsilon^{-(d+9) / 2} k\right)$ samples.

## 6 Conclusions

We studied PAC learning of classes of distributions that are in the form of mixture models, and proposed a generic approach for learning such classes in the cases where we have access to a black box method for learning a singlecomponent distribution. We showed that by going from one component to a mixture model with $k$ components, the sample complexity is multiplied by a factor of at most $\left(k \log ^{2} k\right) / \varepsilon^{2}$.

Furthermore, as a corollary of this general result, we provided upper bounds for the sample complexity of learning GMMs and axis-aligned GMMs- $O\left(k d^{2} \log ^{2} k / \varepsilon^{4}\right)$ and $O\left(k d \log ^{2} k / \varepsilon^{4}\right)$ respectively. Both of these results improve upon the state of the art in terms of dependence on $k$ and $d$.

It is worthwhile to note that for the case of GMMs, the dependence of our bound is $1 / \varepsilon^{4}$. Therefore, proving an upper bound of $k d^{2} / \varepsilon^{2}$ remains open.

Also, note that our result can be readily applied to the general case of mixtures of the exponential family. Let $\mathcal{F}_{d}$ denote the $d$-parameter exponential family. Then the VCdimension of the corresponding Yatracos class (see Definition 20) is $O(d)$ (see Theorem 8.1 in (Devroye and Lugosi 2001)) and therefore by Theorem 22, the sample complexity of PAC learning $\mathcal{F}_{d}$ is $O\left(d / \varepsilon^{2}\right)$. Finally, applying Theorem 5 gives a sample complexity upper bound of $\widetilde{O}\left(k d / \varepsilon^{4}\right)$ for learning $\mathcal{F}_{d}^{k}$.

## 7 Proofs of Theorems 10 and 13

We follow the general methodology of (Devroye and Lugosi 2001) to prove upper bounds on the sample complexity of learning Gaussian distributions. The idea is to first connect distribution learning to the VC-dimension of a class of a related set system (called the Yatracos class of the corresponding distribution family), and then provide upper bounds on VC-dimension of this system. Our Theorem 22 gives an upper bound for the sample complexity of agnostic learning, given an upper bound for the VC-dimension of the Yatracos class. We remark that a variant of this result, without explicit dependence on the failure probability, is proved implicitly in (Chan et al. 2014) and also appears explicitly in (Diakonikolas, Kane, and Stewart 2017a, Lemma 6).
Definition $17\left(\mathcal{A}\right.$-Distance). Let $\mathcal{A} \subset 2^{X}$ be a class of subsets of domain $X$. Let $p$ and $q$ be two probability distributions over $X$. Then the $\mathcal{A}$-distance between $p$ and $q$ is defined as

$$
\|p-q\|_{\mathcal{A}}:=\sup _{A \in \mathcal{A}}|p(A)-q(A)|
$$

Definition 18 (Empirical Distribution). Let $S=\left\{x_{i}\right\}_{i=1}^{m}$ be a sequence of members of $X$. The empirical distribu-
tion corresponding to this sequence is defined by $\hat{p}_{S}(x)=$ $\sum_{i=1}^{m} \frac{\mathbb{1}\left\{x=x_{i}\right\}}{m}$.

The following lemma is a well known refinement of the uniform convergence theorem, see, e.g., (Anthony and Bartlett 1999, Theorem 4.9).

Lemma 19. Let p be a probability distribution over $X$. Let $\mathcal{A} \subseteq 2^{X}$ and let $v$ be the $V C$-dimension of $\mathcal{A}$. Then, there exist universal positive constants $c_{1}, c_{2}, c_{3}$ such that

$$
\operatorname{Pr}_{S \sim p^{m}}\left\{\left\|p-\hat{p}_{S}\right\|_{\mathcal{A}} \geq \varepsilon\right\} \leq \exp \left(c_{1}+c_{2} v-c_{3} m \varepsilon^{2}\right)
$$

Definition 20 (Yatracos class). For a class $\mathcal{F}$ of functions from $X$ to $\mathbb{R}$, their Yatracos class is the family of subsets of $X$ defined as
$\mathcal{Y}(\mathcal{F}):=\left\{\left\{x \in X: f_{1}(x) \geq f_{2}(x)\right\}\right.$ for some $\left.f_{1}, f_{2} \in \mathcal{F}\right\}$
Observe that if $f, g \in \mathcal{F}$ then $\|f-g\|_{T V}=\|f-g\|_{\mathcal{Y}(F)}$.
Definition 21 (Empirical Yatracos Minimizer). Let $\mathcal{F}$ be a class of distributions over domain $X$. The empirical Yatra$\cos$ minimizer is defined as $L^{\mathcal{F}}: \cup_{m=1}^{\infty} X^{m} \rightarrow \mathcal{F}$ satisfying

$$
L^{\mathcal{F}}(S)=\arg \min _{q \in \mathcal{F}}\left\|q-\hat{p}_{S}\right\|_{\mathcal{Y}(\mathcal{F})}
$$

Theorem 22 (PAC Learning Families of Distributions). Let $\mathcal{F}$ be a class of probability distributions, and let $S \sim p^{m}$ be an i.i.d. sample of size $m$ generated from an arbitrary probability distribution $p$, which is not necessarily in $\mathcal{F}$. Then with probability at least $1-\delta$ we have

$$
\left\|p-L^{\mathcal{F}}(S)\right\|_{T V} \leq 3 \operatorname{OPT}(\mathcal{F}, p)+\alpha \sqrt{\frac{v+\log \frac{1}{\delta}}{m}}
$$

where $v$ is VC-dimension of $\mathcal{Y}(\mathcal{F})$, and $\operatorname{OPT}(\mathcal{F}, p)=$ $\inf _{q^{*} \in \mathcal{F}}\left\|q^{*}-p\right\|_{T V}$, and $\alpha$ is a universal constant. In particular, in the realizable setting $p \in \mathcal{F}$, we have

$$
\left\|p-L^{\mathcal{F}}(S)\right\|_{T V} \leq \alpha \sqrt{\frac{v+\log \frac{1}{\delta}}{m}}
$$

Remark 23. The $L_{1}$ distance is precisely twice the total variation distance.

Proof. Let $q^{*}=\arg \min _{q \in \mathcal{F}}\|p-q\|_{T V}$, so $\left\|q^{*}-p\right\|_{\mathcal{Y}(\mathcal{F})} \leq$ $\left\|q^{*}-p\right\|_{T V}=\operatorname{OPT}(\mathcal{F}, p)$. Since $L^{\mathcal{F}}(S), q^{*} \in \mathcal{F}$ we have $\left\|L^{\mathcal{F}}(S)-q^{*}\right\|_{T V}=\left\|L^{\mathcal{F}}(S)-q^{*}\right\|_{\mathcal{Y}_{(\mathcal{F})}}$. By Lemma 19, with probability $\geq 1-\delta$ we have $\| p-$ $\hat{p}_{S} \|_{\mathcal{A}} \leq \alpha \sqrt{\left(v+\log \frac{1}{\delta}\right) / m}$ for some universal constant $\alpha$. Also, since $L^{\mathcal{F}}(S)$ is the empirical minimizer of the $\mathcal{Y}(\mathcal{F})$ distance, we have $\left\|L^{\mathcal{F}}(S)-\hat{p}_{S}\right\|_{\mathcal{Y}(\mathcal{F})} \leq\left\|q^{*}-\hat{p}_{S}\right\|_{\mathcal{Y}(\mathcal{F})}$. The proof follows from these facts combined with multiple
applications of the triangle inequality:

$$
\begin{aligned}
& \left\|p-L^{\mathcal{F}}(S)\right\|_{T V} \leq\left\|L^{\mathcal{F}}(S)-q^{*}\right\|_{T V}+\left\|q^{*}-p\right\|_{T V} \\
& \quad=\left\|L^{\mathcal{F}}(S)-q^{*}\right\|_{\mathcal{Y}(\mathcal{F})}+\operatorname{OPT}(\mathcal{F}, p) \\
& \quad \leq\left\|L^{\mathcal{F}}(S)-\hat{p}_{S}\right\|_{\mathcal{Y}(\mathcal{F})}+\left\|\hat{p}_{S}-q^{*}\right\|_{\mathcal{Y}(\mathcal{F})}+\operatorname{OPT}(\mathcal{F}, p) \\
& \quad \leq\left\|q^{*}-\hat{p}_{S}\right\|_{\mathcal{Y}(\mathcal{F})}+\left(\left\|\hat{p}_{S}-p\right\|_{\mathcal{A}}+\left\|p-q^{*}\right\|_{\mathcal{Y}(\mathcal{F})}\right)+ \\
& \operatorname{OPT}(\mathcal{F}, p) \leq\left(\left\|q^{*}-p\right\|_{\mathcal{Y}(\mathcal{F})}+\left\|p-\hat{p}_{S}\right\|_{\mathcal{A}}\right)+ \\
& \quad\left\|p-\hat{p}_{S}\right\|_{\mathcal{Y}(\mathcal{F})}+2 \operatorname{OPT}(\mathcal{F}, p) \\
& \quad \leq\left\|q^{*}-p\right\|_{T V}+2\left\|p-\hat{p}_{S}\right\|_{\mathcal{Y}(\mathcal{F})}+2 \operatorname{OPT}(\mathcal{F}, p) \\
& \quad \leq 2 \alpha \sqrt{\frac{v+\log \frac{1}{\delta}}{m}}+3 \operatorname{OPT}(\mathcal{F}, p) .
\end{aligned}
$$

Theorem 22 provides a tool for proving upper bounds on the sample complexity of distribution learning. To prove Theorems 13 and 10, it remains to show upper bounds on the VC dimensions of the Yatracos class of (axis-aligned) Gaussian densities.

For classes $\mathcal{F}$ and $\mathcal{G}$ of functions, let

$$
\mathrm{NN}(\mathcal{G}):=\{\{x: f(x) \geq 0\} \text { for some } f \in \mathcal{G}\}
$$

and

$$
\Delta \mathcal{F}:=\left\{f_{1}-f_{2}: f_{1}, f_{2} \in \mathcal{F}\right\}
$$

and notice that

$$
\mathcal{Y}(\mathcal{F})=\mathrm{NN}(\Delta \mathcal{F})
$$

We upper bound the VC-dimension of $\mathrm{NN}(\Delta \mathcal{F})$ via the following well known result in statistical learning theory, see, e.g., (Devroye and Lugosi 2001, Lemma 4.2).

Theorem 24 (Dudley). Let $\mathcal{G}$ be an $n$-dimensional vector space of real-valued functions. Then $V C(\mathrm{NN}(\mathcal{G})) \leq n$.

Now let $h$ be an indicator function for an arbitrary element in $\mathrm{NN}\left(f_{1}-f_{2}\right)$, where $f_{1}, f_{2}$ are densities of (axis-aligned) Gaussians. Then $h$ is a $\{0,1\}$-valued function and we have:

$$
\begin{aligned}
h(x)= & \mathbb{1}\left\{\mathcal{N}\left(\mu_{1}, \Sigma_{1}\right)>\mathcal{N}\left(\mu_{2}, \Sigma_{2}\right)\right\} \\
= & \mathbb{1}\left\{\alpha_{1} \exp \left(\frac{-1}{2}\left(x-\mu_{1}\right)^{T} \Sigma_{1}^{-1}\left(x-\mu_{1}\right)\right)>\right. \\
& \left.\alpha_{2} \exp \left(\frac{-1}{2}\left(x-\mu_{2}\right)^{T} \Sigma_{2}^{-1}\left(x-\mu_{2}\right)\right)\right\} \\
= & \mathbb{1}\left\{\left(x-\mu_{1}\right)^{T} \Sigma_{1}^{-1}\left(x-\mu_{1}\right)\right. \\
& \left.\quad-\left(x-\mu_{2}\right)^{T} \Sigma_{2}^{-1}\left(x-\mu_{2}\right)-\log \frac{\alpha_{2}}{\alpha_{1}}>0\right\} .
\end{aligned}
$$

The inner expression is a quadratic form, and the linear dimension of all quadratic functions is $O\left(d^{2}\right)$. Furthermore, for axis-aligned Gaussians, $\Sigma_{1}$ and $\Sigma_{2}$ are diagonal, and therefore, the inner function lies in an $O(d)$-dimensional space of functions spanned by $\left\{1, x_{1}, \ldots, x_{d}, x_{1}^{2}, \ldots, x_{d}^{2}\right\}$. Hence, by Dudley's theorem, we have the required upper bound ( $d$ or $d^{2}$ ) on the VC-dimension of the Yatracos classes. Finally, Theorems 13 and 10 follow from applying Theorem 22 to the class of (axis-aligned) Gaussian distributions.

## Acknowledgments

Abbas Mehrabian was supported by an NSERC Postdoctoral Fellowship and a Simons-Berkeley Research Fellowship. The authors thank the reviewers of ALT 2017 conference for their comments and pointing out a mistake in an earlier version.

## Addendum

Very recently, the authors have improved Theorem 11 and showed that the class of mixtures of $k$ axis-aligned Gaussians in $\mathbb{R}^{d}$ is 3 -agnostic PAC-learnable with sample complexity $\widetilde{O}\left(k d / \varepsilon^{2}\right)$. The proof uses novel techniques, see (Ashtiani, Ben-David, and Mehrabian 2017).

## References

Acharya, J.; Diakonikolas, I.; Li, J.; and Schmidt, L. 2017. Sample-optimal density estimation in nearly-linear time. In Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '17, 1278-1289. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.
Anthony, M., and Bartlett, P. 1999. Neural network learning: theoretical foundations. Cambridge University Press.
Ashtiani, H.; Ben-David, S.; and Mehrabian, A. 2017. Agnostic distribution learning via compression. arXiv preprint arXiv:1710.05209.
Belkin, M., and Sinha, K. 2010. Polynomial learning of distribution families. In Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, FOCS '10, 103-112. Washington, DC, USA: IEEE Computer Society.
Chan, S.-O.; Diakonikolas, I.; Servedio, R. A.; and Sun, X. 2014. Efficient density estimation via piecewise polynomial approximation. In Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing, STOC '14, 604613. New York, NY, USA: ACM.

Dasgupta, S. 1999. Learning mixtures of Gaussians. In Foundations of Computer Science, 1999. 40th Annual Symposium on, 634-644. IEEE.

Devroye, L., and Lugosi, G. 2001. Combinatorial methods in density estimation. Springer Series in Statistics. SpringerVerlag, New York.
Diakonikolas, I.; Kamath, G.; Kane, D. M.; Li, J.; Moitra, A.; and Stewart, A. 2016. Robust estimators in high dimensions without the computational intractability. In 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), 655-664.
Diakonikolas, I.; Kane, D. M.; and Stewart, A. 2017a. Learning multivariate log-concave distributions. In Proceedings of Machine Learning Research, volume 65 of COLT'17, 1-17.
Diakonikolas, I.; Kane, D. M.; and Stewart, A. 2017b. Statistical query lower bounds for robust estimation of high-dimensional Gaussians and Gaussian mixtures. arXiv
preprint arXiv:1611.03473v2 [cs.LG]. To appear in Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS '17).
Diakonikolas, I. 2016. Learning Structured Distributions. In Peter Bühlmann; Petros Drineas; Michael Kane; and Mark van der Laan., eds., Handbook of Big Data. Chapman and Hall/CRC. chapter 15, 267-283.
Feldman, J.; Servedio, R. A.; and O’Donnell, R. 2006. PAC learning axis-aligned mixtures of Gaussians with no separation assumption. In Proceedings of the 19th Annual Conference on Learning Theory, COLT'06, 20-34. Berlin, Heidelberg: Springer-Verlag.
Karpinski, M., and Macintyre, A. 1997. Polynomial bounds for VC dimension of sigmoidal and general pfaffian neural networks. Journal of Computer and System Sciences 54(1):169-176.
Kearns, M.; Mansour, Y.; Ron, D.; Rubinfeld, R.; Schapire, R. E.; and Sellie, L. 1994. On the learnability of discrete distributions. In Proceedings of the Twenty-sixth Annual ACM Symposium on Theory of Computing, STOC '94, 273-282. New York, NY, USA: ACM.
Li, J., and Schmidt, L. 2017. Robust and proper learning for mixtures of gaussians via systems of polynomial inequalities. In Kale, S., and Shamir, O., eds., Proceedings of the 2017 Conference on Learning Theory, volume 65 of Proceedings of Machine Learning Research, 1302-1382. Amsterdam, Netherlands: PMLR.
Mahalanabis, S., and Stefankovic, D. 2008. Density estimation in linear time. In Servedio, R. A., and Zhang, T., eds., 21 st Annual Conference on Learning Theory - COLT 2008, Helsinki, Finland, July 9-12, 2008, 503-512. Omnipress.
Mitzenmacher, M., and Upfal, E. 2005. Probability and Computing: Randomized Algorithms and Probabilistic Analysis. New York, NY, USA: Cambridge University Press.
Moitra, A., and Valiant, G. 2010. Settling the polynomial learnability of mixtures of Gaussians. In Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, FOCS '10, 93-102. Washington, DC, USA: IEEE Computer Society.
Suresh, A. T.; Orlitsky, A.; Acharya, J.; and Jafarpour, A. 2014. Near-optimal-sample estimators for spherical Gaussian mixtures. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N. D.; and Weinberger, K. Q., eds., Advances in Neural Information Processing Systems 27. Curran Associates, Inc. 1395-1403.


[^0]:    Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.
    ${ }^{1}$ Total variation distance is a prominent distance measure between distributions. For a discussion on this and other choices see (Devroye and Lugosi 2001, Chapter 5).

[^1]:    ${ }^{2}$ Note that in some papers, only the case $C \leq 1$ is called agnostic learning, and the case $C>1$ is called semi-agnostic learning.

