

A Nearly Optimal and Agnostic Algorithm for Properly Learning a Mixture of k Gaussians, for any Constant k

Jerry Li* Ludwig Schmidt†
MIT MIT
jerryzli@mit.edu ludwigs@mit.edu

June 27, 2015

Abstract

Learning a Gaussian mixture model (GMM) is a fundamental problem in machine learning, learning theory, and statistics. One notion of learning a GMM is proper learning: here, the goal is to find a mixture of k Gaussians \mathcal{M} that is close to the density f of the unknown distribution from which we draw samples. The distance between \mathcal{M} and f is typically measured in the total variation or L_1 -norm.

We give an algorithm for learning a mixture of k univariate Gaussians that is nearly optimal for any fixed k . The sample complexity of our algorithm is $\tilde{O}(\frac{k}{\epsilon^2})$ and the running time is $(k \cdot \log \frac{1}{\epsilon})^{O(k^4)} + \tilde{O}(\frac{k}{\epsilon^2})$. It is well-known that this sample complexity is optimal (up to logarithmic factors), and it was already achieved by prior work. However, the best known time complexity for proper learning a k -GMM was $\tilde{O}(\frac{1}{\epsilon^{3k-1}})$. In particular, the dependence between $\frac{1}{\epsilon}$ and k was exponential. We significantly improve this dependence by replacing the $\frac{1}{\epsilon}$ term with a $\log \frac{1}{\epsilon}$ while only increasing the exponent moderately. Hence, for any fixed k , the $\tilde{O}(\frac{k}{\epsilon^2})$ term dominates our running time, and thus our algorithm runs in time which is *nearly-linear* in the number of samples drawn. Achieving a running time of $\text{poly}(k, \frac{1}{\epsilon})$ for proper learning of k -GMMs has recently been stated as an open problem by multiple researchers, and we make progress on this open problem.

Moreover, our approach offers an agnostic learning guarantee: our algorithm returns a good GMM even if the distribution we are sampling from is not a mixture of Gaussians. To the best of our knowledge, our algorithm is the first agnostic proper learning algorithm for GMMs. Again, the closely related question of agnostic and proper learning for GMMs in the high-dimensional setting has recently been raised as an open question, and our algorithm resolves this question in the univariate setting.

We achieve these results by approaching the proper learning problem from a new direction: we start with an accurate density estimate and then fit a mixture of Gaussians to this density estimate. Hence, after the initial density estimation step, our algorithm solves an entirely deterministic optimization problem. We reduce this optimization problem to a sequence of carefully constructed systems of polynomial inequalities, which we then solve with Renegar’s algorithm. Our techniques for encoding proper learning problems as systems of polynomial inequalities are general and can be applied to properly learn further classes of distributions besides GMMs.

*Supported by NSF grant CCF-1217921 and DOE grant DE-SC0008923.

†Supported by MADALGO and a grant from the MIT-Shell Energy Initiative.

1 Introduction

Gaussian mixture models (GMMs) are one of the most popular and important statistical models, both in theory and in practice. A sample is drawn from a GMM by first selecting one of its k components according to the mixing weights, and then drawing a sample from the corresponding Gaussian, each of which has its own set of parameters. Since many phenomena encountered in practice give rise to approximately normal distributions, GMMs are often employed to model distributions composed of several distinct sub-populations. GMMs have been studied in statistics since the seminal work of Pearson [Pea94] and are now used in many fields including astronomy, biology, and machine learning. Hence the following are natural and important questions: (i) how can we efficiently “learn” a GMM when we only have access to samples from the distribution, and (ii) what rigorous guarantees can we give for our algorithms?

1.1 Notions of learning

There are several natural notions of learning a GMM, all of which have been studied in the learning theory community over the last 20 years. The known sample and time complexity bounds differ widely for these related problems, and the corresponding algorithmic techniques are also considerably different (see Table 1 for an overview and a comparison with our results). In order of decreasing hardness, these notions of learning are:

Parameter learning The goal in parameter learning is to recover the parameters of the unknown GMM (i.e., the means, variances, and mixing weights) up to some given additive error ϵ .¹

Proper learning In proper learning, our goal is to find a GMM M' such that the L_1 -distance (or equivalently, the total variation distance) between our hypothesis M' and the true unknown distribution is small.

Improper learning / density estimation Density estimation requires us to find *any* hypothesis \hat{h} such that the L_1 distance between \hat{h} and the unknown distribution is small (\hat{h} does not need to be a GMM).

Parameter learning is arguably the most desirable guarantee because it allows us to recover the unknown mixture parameters. For instance, this is important when the parameters directly correspond to physical quantities that we wish to infer. However, this power comes at a cost: recent work on parameter learning has shown that $\Omega(\frac{1}{\epsilon^{12}})$ samples are already necessary to learn the parameters of a mixture of two univariate Gaussians with accuracy ϵ [HP15] (note that this bound is tight, i.e., the paper also gives an algorithm with time and sample complexity $O(\frac{1}{\epsilon^{12}})$). Moreover, the sample complexity of parameter learning scales exponentially with the number of components: for a mixture of k univariate Gaussians, the above paper also gives a sample complexity lower bound of $\Omega(\frac{1}{\epsilon^{6k-2}})$. Hence the sample complexity of parameter learning quickly becomes prohibitive, even for a mixture of two Gaussians and reasonable choices of ϵ .

At the other end of the spectrum, improper learning has much smaller time and sample complexity. Recent work shows that it is possible to estimate the density of a univariate GMM with k components using only $\tilde{O}(\frac{k}{\epsilon^2})$ samples and time [ADLS15], which is tight up to logarithmic factors. However, the output hypothesis produced by the corresponding algorithm is only a piecewise polynomial and not a GMM. This is a disadvantage because GMMs are often desirable as a concise representation of the samples and for interpretability reasons.

Hence an attractive intermediate goal is proper learning: similar to parameter learning, we still produce a GMM as output. On the other hand, we must only satisfy the weaker L_1 -approximation guarantee between our hypothesis and the unknown GMM. While somewhat weaker than parameter learning, proper learning still offers many desirable features: for instance, the representation as a GMM requires only $3k-1$ parameters, which is significantly smaller than the at least $6k(1 + 2 \log \frac{1}{\epsilon})$ many parameters produced by the piecewise polynomial density estimate of [ADLS15] (note that the number of parameters in the piecewise polynomial

¹Since the accuracy ϵ depends on the scale of the mixture (i.e., the variance), these guarantees are often specified relative to the variance.

Problem type	Sample complexity lower bound	Sample complexity upper bound	Time complexity upper bound	Agnostic guarantee
Parameter learning				
$k = 2$	$\Theta(\frac{1}{\epsilon^{12}})$ [HP15]	$O(\frac{1}{\epsilon^{12}})$ [HP15]	$O(\frac{1}{\epsilon^{12}})$ [HP15]	no
general k	$\Omega(\frac{1}{\epsilon^{6k-2}})$ [HP15]	$O((\frac{1}{\epsilon})^{c_k})$ [MV10]	$O((\frac{1}{\epsilon})^{c_k})$ [MV10]	no
Proper learning				
$k = 2$	$\Theta(\frac{1}{\epsilon^2})$	$\tilde{O}(\frac{1}{\epsilon^2})$ [DK14]	$\tilde{O}(\frac{1}{\epsilon^5})$ [DK14]	no
general k	$\Theta(\frac{k}{\epsilon^2})$	$\tilde{O}(\frac{k}{\epsilon^2})$ [AJOS14]	$\tilde{O}(\frac{1}{\epsilon^{3k-1}})$ [DK14, AJOS14]	no
Our results				
$k = 2$		$\tilde{O}(\frac{1}{\epsilon^2})$	$\tilde{O}(\frac{1}{\epsilon^2})$	yes
general k		$\tilde{O}(\frac{k}{\epsilon^2})$	$(k \log \frac{1}{\epsilon})^{O(k^4)} + \tilde{O}(\frac{k}{\epsilon^2})$	yes
Density estimation				
general k	$\Theta(\frac{k}{\epsilon^2})$	$\tilde{O}(\frac{k}{\epsilon^2})$ [ADLS15]	$\tilde{O}(\frac{k}{\epsilon^2})$ [ADLS15]	yes

Table 1: Overview of the best known results for learning a mixture of univariate Gaussians. Our contributions (highlighted as bold) significantly improve on the previous results for proper learning: the time complexity of our algorithm is nearly optimal for any fixed k . Moreover, our algorithm gives agnostic learning guarantees. The constant c_k in the time and sample complexity of [MV10] depends only on k and is at least k . The sample complexity lower bounds for proper learning and density estimation are folklore results. The only time complexity lower bounds known are the corresponding sample complexity lower bounds, so we omit an extra column for time complexity lower bounds.

also grows as the accuracy of the density estimate increases). Moreover, the representation as a GMM allows us to provide simple closed-form expressions for quantities such as the mean and the moments of the learnt distribution, which are then easy to manipulate and understand. In contrast, no such closed-form expressions exist when given a general density estimate as returned by [ADLS15]. Furthermore, producing a GMM as output hypothesis can be seen as a regularization step because the number of peaks in the density is bounded by k , and the density is guaranteed to be smooth. This is usually an advantage over improper hypotheses such as piecewise polynomials that can have many more peaks or discontinuities, which makes the hypothesis harder to interpret and process. Finally, even the most general parameter learning algorithms require assumptions on the GMMs such as identifiability, while our proper learning algorithm works for *any* GMM.

Ideally, proper learning could combine the interpretability and conciseness of parameter learning with the small time and sample complexity of density estimation. Indeed, recent work has shown that one can properly learn a mixture of k univariate Gaussians from only $\tilde{O}(\frac{k}{\epsilon^2})$ samples [DK14, AJOS14], which is tight up to logarithmic factors. However, the time complexity of proper learning is not yet well understood. This is in contrast to parameter learning and density estimation, where we have strong lower bounds and nearly-optimal algorithms, respectively. For the case of two mixture components, the algorithm of [DK14] runs in time $\tilde{O}(\frac{1}{\epsilon^5})$. However, the time complexity of this approach becomes very large for general k and scales as $O(\frac{1}{\epsilon^{3k-1}})$ [DK14, AJOS14]. Note that this time complexity is much larger than the $\tilde{O}(\frac{k}{\epsilon^2})$ required for density estimation and resembles the exponential dependence between $\frac{1}{\epsilon}$ and k in the $\Omega(\frac{1}{\epsilon^{6k-2}})$ lower bound for parameter learning. Hence the true time complexity of properly learning a GMM is an important open question. In particular, it is not known whether the exponential dependence between $\frac{1}{\epsilon}$ and k is necessary.

In our work, we answer this question and show that such an exponential dependence between $\frac{1}{\epsilon}$ and k can be avoided. We give an algorithm with the same (nearly optimal) sample complexity as previous work,

but our algorithm also runs in time which is *nearly-optimal*, i.e. nearly-linear in the number of samples, for any fixed k . It is worth noting that proper learning of k -GMMs in time $\text{poly}(k, \frac{1}{\epsilon})$ has been raised as an open problem [Moi14, BDKVDL15], and we make progress on this question. Moreover, our learning algorithm is *agnostic*, which means that the algorithm tolerates arbitrary amounts of worst-case noise in the distribution generating our samples. This is an important robustness guarantee, which we now explain further.

1.2 Robust learning guarantees

All known algorithms for properly learning or parameter learning² a GMM offer rigorous guarantees only when there is at most a very small amount of noise in the distribution generating our samples. Typically, these algorithms work in a “non-agnostic” setting: they assume that samples come from a distribution that is exactly a GMM with probability density function (pdf) f , and produce a hypothesis pdf \hat{h} such that for some given $\epsilon > 0$

$$\|f - \hat{h}\|_1 \leq \epsilon.$$

But in practice, we cannot typically expect that we draw samples from a distribution that truly corresponds to a GMM. While many natural phenomena are well approximated by a GMM, such an approximation is rarely exact. Instead, it is useful to think of such phenomena as GMMs corrupted by some amount of noise. Hence it is important to design algorithms that still provide guarantees when the true unknown distribution is far from any mixture of k Gaussians.

Agnostic learning Therefore, we focus on the problem of *agnostic* learning [KSS94], where our samples can come from *any* distribution, not necessarily a GMM. Let f be the pdf of this unknown distribution and let \mathcal{M}_k be the set of pdfs corresponding to a mixture of k Gaussians. Then we define OPT_k to be the following quantity:

$$\text{OPT}_k \stackrel{\text{def}}{=} \min_{h \in \mathcal{M}_k} \|f - h\|_1,$$

which is the error achieved by the best approximation of f with a k -GMM. Note that this is a deterministic quantity, which can also be seen as the error incurred when projecting f onto set \mathcal{M}_k . Using this definition, an agnostic learning algorithm produces a GMM with density \hat{h} such that

$$\|f - \hat{h}\|_1 \leq C \cdot \text{OPT}_k + \epsilon$$

for some given $\epsilon > 0$ and a universal constant³ C that does not depend on ϵ , k , or the unknown pdf f .

Clearly, agnostic learning guarantees are more desirable because they also apply when the distribution producing our samples does not match our model exactly (note also that agnostic learning is strictly more general than non-agnostic learning). Moreover, the agnostic learning guarantee is “stable”: when our model is close to the true distribution f , the error of the best approximation, i.e., OPT_k , is small. Hence an agnostic learning algorithm still produces a good hypothesis.

On the other hand, agnostic learning algorithms are harder to design because we cannot make any assumptions on the distribution producing our samples. To the best of our knowledge, our algorithm is the first agnostic algorithm for properly learning a mixture of k univariate Gaussians. Note that agnostic learning has recently been raised as an open question in the setting of learning high-dimensional GMMs [Vem12], and our agnostic univariate algorithm can be seen as progress on this problem. Moreover, our algorithm achieves such an agnostic guarantee without any increase in time or sample complexity compared to the non-agnostic case.

²Note that parameter learning is not well-defined if the samples do not come from a GMM. Instead, existing parameter learning algorithms are not robust in the following sense: if the unknown distribution is not a GMM, the algorithms are not guaranteed to produce a set of parameters such that the corresponding GMM is close to the unknown density.

³Strictly speaking, agnostic learning requires this constant C to be 1. However, such a tight guarantee is impossible for some learning problems such as density estimation. Hence we allow any constant C in an agnostic learning guarantee, which is sometimes also called *semi-agnostic* learning.

1.3 Our contributions

We now outline the main contributions of this paper. Similar to related work on proper learning [DK14], we restrict our attention to univariate GMMs. Many algorithms for high-dimensional GMMs work via reductions to the univariate case, so it is important to understand this case in greater detail [KMV10, MV10, HP15].

First, we state our main theorem. See Section 2 for a formal definition of the notation. The quantity OPT_k is the same as introduced above in Section 1.2.

Theorem 1. *Let f be the pdf of an arbitrary unknown distribution, let k be a positive integer, and let $\epsilon > 0$. Then there is an algorithm that draws $\tilde{O}(\frac{k}{\epsilon^2})$ samples from the unknown distribution and produces a mixture of k Gaussians such that the corresponding pdf \hat{h} satisfies*

$$\|f - \hat{h}\|_1 \leq 42 \cdot \text{OPT}_k + \epsilon.$$

Moreover, the algorithm runs in time

$$\left(k \cdot \log \frac{1}{\epsilon}\right)^{O(k^4)} + \tilde{O}\left(\frac{k}{\epsilon^2}\right).$$

We remark that we neither optimized the exponent $O(k^4)$, nor the constant in front of OPT_k . Instead, we see our result as a proof of concept that it is possible to agnostically and properly learn a mixture of Gaussians in time that is essentially fixed-parameter optimal. As mentioned above, closely related questions about efficient and agnostic learning of GMMs have recently been posed as open problems, and we make progress on these questions. In particular, our main theorem implies the following contributions:

Running time The time complexity of our algorithm is significantly better than previous work on proper learning of GMMs. For the special case of 2 mixture components studied in [DK14] and [HP15], our running time simplifies to $\tilde{O}(\frac{1}{\epsilon^2})$. This is a significant improvement over the $\tilde{O}(\frac{1}{\epsilon^5})$ bound in [DK14]. Moreover, our time complexity matches the best possible time complexity for density estimation of 2-GMMs up to logarithmic factors. This also implies that our time complexity is *optimal* up to log-factors.

For proper learning of k mixture components, prior work achieved a time complexity of $\tilde{O}(\frac{1}{\epsilon^{3k-1}})$ [DK14, AJOS14]. Compared to this result, our algorithm achieves an *exponential* improvement in the dependence between $\frac{1}{\epsilon}$ and k : our running time contains only a $(\log \frac{1}{\epsilon})$ term raised to the poly(k)-th power, not a $(\frac{1}{\epsilon})^k$. In particular, the $\tilde{O}(\frac{k}{\epsilon^2})$ term in our running time dominates for any fixed k . Hence the time complexity of our algorithm is *nearly optimal for any fixed k* .

Agnostic learning Our algorithm is the first proper learning algorithm for GMMs that is *agnostic*. Previous algorithms relied on specific properties of the normal distribution such as moments, while our techniques are more robust. Practical algorithms should offer agnostic guarantees, and we hope that our approach is a step in this direction. Moreover, it is worth noting that agnostic learning, i.e., learning under noise, is often significantly harder than non-agnostic learning. One such example is learning parity with noise, which is conjectured to be computationally hard. Hence it is an important question to understand which learning problems are tractable in the agnostic setting. While the agnostic guarantee achieved by our algorithm is certainly not optimal, our algorithm still shows that it is possible to learn a mixture of Gaussians agnostically with only a very mild dependence on $\frac{1}{\epsilon}$.

From improper to proper learning Our techniques offer a general scheme for converting *improper* learning algorithms to *proper* algorithms. In particular, our approach applies to any parametric family of distributions that are well approximated by a piecewise polynomial in which the parameters appear polynomially and the breakpoints depend polynomially (or rationally) on the parameters. As a result, we can convert purely approximation-theoretic results into proper learning algorithms for other classes of distributions, such as mixtures of Laplace or exponential distributions. Conceptually, we show how to

approach proper learning as a purely deterministic optimization problem once a good density estimate is available. Hence our approach differs from essentially all previous proper learning algorithms, which use probabilistic arguments in order to learn a mixture of Gaussians.

1.4 Techniques

At its core, our algorithm fits a mixture of Gaussians to a density estimate. In order to obtain an ϵ -accurate and agnostic density estimate, we invoke recent work that has a time and sample complexity of $\tilde{O}(\frac{k}{\epsilon^2})$ [ADLS15]. The density estimate produced by their algorithm has the form of a piecewise polynomial with $O(k)$ pieces, each of which has degree $O(\log \frac{1}{\epsilon})$. It is important to note that our algorithm does not draw any further samples after obtaining this density estimate — the process of fitting a mixture of Gaussians is entirely deterministic.

Once we have obtained a good density estimate, the task of proper learning reduces to fitting a mixture of k Gaussians to the density estimate. We achieve this via a further reduction from fitting a GMM to solving a carefully designed system of polynomial inequalities. We then solve the resulting system with Renegar’s algorithm [Ren92a, Ren92b]. This reduction to a system of polynomial inequalities is our main technical contribution and relies on the following techniques.

Shape-restricted polynomials Ideally, one could directly fit a mixture of Gaussian pdfs to the density estimate. However, this is a challenging task because the Gaussian pdf $\frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{(x-\mu)^2}{2}}$ is not convex in the parameters μ and σ . Thus fitting a mixture of Gaussians is a non-convex problem.

Instead of fitting mixtures of Gaussians directly, we instead use the notion of a *shape restricted* polynomial. We say that a polynomial is shape restricted if its coefficients are in a given semialgebraic set, i.e., a set defined by a finite number of polynomial equalities and inequalities. It is well-known in approximation theory that a single Gaussian can be approximated by a piecewise polynomial consisting of three pieces with degree at most $O(\log \frac{1}{\epsilon})$ [Tim63]. So instead of fitting a mixture of k Gaussian directly, we instead fit a mixture of k shape-restricted piecewise polynomials. By encoding that the shape-restricted polynomials must have the shape of Gaussian pdfs, we ensure that the mixture of shape-restricted piecewise polynomials found by the system of polynomial inequalities is close to a true mixture of k -Gaussians. After we have solved the system of polynomial inequalities, it is easy to convert the shape-restricted polynomials back to a proper GMM.

\mathcal{A}_K -distance The system of polynomial inequalities we use for finding a good mixture of piecewise polynomials must encode that the mixture should be close to the density estimate. In our final guarantee for proper learning, we are interested in an approximation guarantee in the L_1 -norm. However, directly encoding the L_1 -norm in the system of polynomial inequalities is challenging because it requires knowledge of the intersections between the density estimate and the mixture of piecewise polynomials in order to compute the integral of their difference accurately. Instead of directly minimizing the L_1 -norm, we instead minimize the closely related \mathcal{A}_K -norm from VC (Vapnik–Chervonenkis) theory [DL01]. For functions with at most $K - 1$ sign changes, the \mathcal{A}_K -norm exactly matches the L_1 -norm. Since two mixtures of k Gaussians have at most $O(k)$ intersections, we have a good bound on the order of the \mathcal{A}_K -norm we use to replace the L_1 -norm. In contrast to the L_1 -norm, we can encode the \mathcal{A}_K -norm without increasing the size of our system of polynomial inequalities significantly — directly using the L_1 -norm would lead to an exponential dependence on $\log \frac{1}{\epsilon}$ in our system of polynomial inequalities.

Adaptively rescaling the density estimate In order to use Renegar’s algorithm for solving our system of polynomial inequalities, we require a bound on the accuracy necessary to find a good set of mixture components. While Renegar’s algorithm has a good dependence on the accuracy parameter, our goal is to give an algorithm for proper learning without *any* assumptions on the GMM. Therefore, we must be able to produce good GMMs even if the parameters of the unknown GMM are, e.g., doubly-exponential in $\frac{1}{\epsilon}$ or even larger. Note that this issue arises in spite of the fact that our algorithm works in the real-RAM model:

since different mixture parameters can have widely varying scales, specifying a single accuracy parameter for Renegar’s algorithm is not sufficient.

We overcome this technical challenge by adaptively rescaling the parametrization used in our system of polynomial inequalities based on the lengths of the intervals I_1, \dots, I_s that define the piecewise polynomial density estimate p_{dens} . Since p_{dens} can only be large on intervals I_i of small length, the best Gaussian fit to p_{dens} can only have large parameters near such intervals. Hence, this serves as a simple way of identifying where we require more accuracy when computing the mixture parameters.

Putting things together Combining the ideas outlined above, we can fit a mixture of k Gaussians with a carefully designed system of polynomial inequalities. A crucial aspect of the system of polynomial inequalities is that the number of variables is $O(k)$, that the number of inequalities is $k^{O(k)}$, and the degree of the polynomials is bounded by $O(\log \frac{1}{\epsilon})$. These bounds on the size of the system of polynomial inequalities then lead to the running time stated in Theorem 1. In particular, the size of the system of polynomial inequalities is almost independent of the number of samples, and hence the running time required to solve the system scales only *poly-logarithmically* with $\frac{1}{\epsilon}$.

1.5 Related work

Due to space constraints, it is impossible to summarize the entire body of work on learning GMMs here. Therefore, we limit our attention to results with provable guarantees corresponding to the notions of learning outlined in Subsection 1.1. Note that this is only one part of the picture: for instance, the well-known Expectation-Maximization (EM) algorithm is still the subject of current research (see [BWY14] and references therein).

For parameter learning, the seminal work of Dasgupta [Das99] started a long line of research in the theoretical computer science community, e.g., [SK01, VW04, AM05, KSV08, BV08, KMV10]. We refer the reader to [MV10] for a discussion of these and related results. The papers [MV10] and [BS10] were the first to give polynomial time algorithms (polynomial in ϵ and the dimension of the mixture) with provably minimal assumptions for k -GMMs. More recently, Hardt and Price gave tight bounds for learning the parameters of a mixture of 2 univariate Gaussians [HP15]: $\Theta(\frac{1}{\epsilon^{12}})$ samples are necessary and sufficient, and the time complexity is linear in the number of samples. Moreover, Hardt and Price give a strong lower bound of $\Omega(\frac{1}{\epsilon^{6k-2}})$ for the sample complexity of parameter learning a k -GMM. While our proper learning algorithm offers a weaker guarantee than these parameter learning approaches, our time complexity does not have an exponential dependence between $\frac{1}{\epsilon}$ and k . Moreover, proper learning retains many of the attractive features of parameter learning (see Subsection 1.1).

Interestingly, parameter learning becomes more tractable as the number of dimensions increases. A recent line of work investigates this phenomenon under a variety of assumptions (e.g., non-degeneracy or smoothed analysis) [HK13, BCMV14, ABG⁺14, GHK15]. However, all of these algorithms require a lower bound on the dimension d such as $d \geq \Omega(k)$ or $d \geq \Omega(k^2)$. Since we focus on the one-dimensional case, our results are not directly comparable. Moreover, to the best of our knowledge, none of the parameter learning algorithms (in any dimension) provide proper learning guarantees in the agnostic setting.

The first work to consider proper learning of k -GMMs without separation assumptions on the components was [FSO06]. Their algorithm takes $\text{poly}(d, 1/\epsilon, L)$ samples and returns a mixture whose KL-divergence to the unknown mixture is at most ϵ . Unfortunately, their algorithm has a pseudo-polynomial dependence on L , which is a bound on the means and the variances of the underlying components. Note that such an assumption is not necessary a priori, and our algorithm works without any such requirements. Moreover, their sample complexity is exponential in the number of components k .

The work closest to ours are the papers [DK14] and [AJOS14], who also consider the problem of properly learning a k -GMM. Their algorithms are based on constructing a set of candidate GMMs that are then compared via an improved version of the Scheffé-estimate. While this approach leads to a nearly-optimal sample complexity of $\tilde{O}(\frac{k}{\epsilon^2})$, their algorithm constructs a large number of candidate hypothesis. This leads to a time complexity of $O(\frac{1}{\epsilon^{3k-1}})$. As pointed out in Subsection 1.1, our algorithm significantly improves the

dependence between $\frac{1}{\epsilon}$ and k . Moreover, none of their algorithms are agnostic.

Another related paper on learning GMMs is [BSZ15]. Their approach reduces the learning problem to finding a sparse solution to a non-negative linear system. Conceptually, this approach is somewhat similar to ours in that they also fit a mixture of Gaussians to a set of density estimates. However, their algorithm does not give a proper learning guarantee: instead of k mixture components, the GMM returned by their algorithm contains $O(\frac{k}{\epsilon^3})$ components. Note that this number of components is significantly larger than the k components returned by our algorithm. Moreover, their number of components increases as the accuracy parameter ϵ improves. In the univariate case, the time and sample complexity of their algorithm is $O(\frac{k}{\epsilon^6})$. Note that their sample complexity is not optimal and roughly $\frac{1}{\epsilon^4}$ worse than our approach. For any fixed k , our running time is also better by roughly $\frac{1}{\epsilon^4}$. Furthermore, the authors do not give an agnostic learning guarantee for their algorithm.

For density estimation, there is a recent line of work on improperly learning structured distributions [CDSS13, CDSS14, ADLS15]. While the most recent paper from this line achieves a nearly-optimal time and sample complexity for density estimation of k -GMMs, the hypothesis produced by their algorithm is a piecewise polynomial. As mentioned in Subsection 1.1, GMMs have several advantages as output hypothesis.

1.6 Outline of our paper

In Section 2, we introduce basic notation and important known results that we utilize in our algorithm. Section 3 describes our learning algorithm for the special case of *well-behaved* density estimates. This assumption allows us to introduce two of our main tools (shape-restricted polynomials and the \mathcal{A}_K -distance as a proxy for L_1) without the technical details of adaptively reparametrizing the shape-restricted polynomials. Section 4 then removes this assumption and gives an algorithm that works for agnostically learning *any* mixture of Gaussians. In Section 5, we show how our techniques can be extended to properly learn further classes of distributions.

2 Preliminaries

Before we construct our learning algorithm for GMMs, we introduce basic notation and the necessary tools from density estimation, systems of polynomial inequalities, and approximation theory.

2.1 Basic notation and definitions

For a positive integer k , we write $[k]$ for the set $\{1, \dots, k\}$. Let $I = [\alpha, \beta]$ be an interval. Then we denote the length of I with $|I| = \beta - \alpha$. For a measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$, the L_1 -norm of f is $\|f\|_1 = \int f(x) dx$. All functions in this paper are measurable.

Since we work with systems of polynomial inequalities, it will be convenient for us to parametrize the normal distribution with the *precision*, i.e., one over the standard deviation, instead of the variance. Thus, throughout the paper we let

$$\mathcal{N}_{\mu, \tau}(x) \stackrel{\text{def}}{=} \frac{\tau}{\sqrt{2\pi}} e^{-\tau^2(x-\mu)^2/2}$$

denote the pdf of a normal distribution with mean μ and precision τ . A k -GMM is a distribution with pdf of the form $\sum_{i=1}^k w_i \cdot \mathcal{N}_{\mu_i, \tau_i}(x)$, where we call the w_i *mixing weights* and require that the w_i satisfy $w_i \geq 0$ and $\sum_{i=1}^k w_i = 1$. Thus a k -GMM is parametrized by $3k$ parameters; namely, the mixing weights, means, and precisions of each component.⁴ We let $\Theta_k = \mathcal{S}_k \times \mathbb{R}^k \times \mathbb{R}_+^k$ be the set of parameters, where \mathcal{S}_k is the simplex in k dimensions. For each $\theta \in \Theta_k$, we identify it canonically with $\theta = (w, \mu, \tau)$ where w, μ , and τ are each vectors of length k , and we let

$$\mathcal{M}_\theta(x) = \sum_{i=1}^k w_i \cdot \mathcal{N}_{\mu_i, \tau_i}(x)$$

⁴Note that there are only $3k - 1$ degrees of freedom since the mixing weights must sum to 1.

be the pdf of the k -GMM with parameters θ .

2.2 Important tools

We now turn our attention to results from prior work.

2.2.1 Density estimation with piecewise polynomials

Our algorithm uses the following result about density estimation of k -GMMs as a subroutine.

Fact 2 ([ADLS15]). *Let $k \geq 1$, $\epsilon > 0$ and $\delta > 0$. There is an algorithm ESTIMATE-DENSITY(k, ϵ, δ) that satisfies the following properties: the algorithm*

- *takes $\tilde{O}((k + \log(1/\delta))/\epsilon^2)$ samples from the unknown distribution with pdf f ,*
- *runs in time $\tilde{O}((k + \log 1/\delta)/\epsilon^2)$, and*
- *returns p_{dens} , an $O(k)$ -piecewise polynomial of degree $O(\log(1/\epsilon))$ such that*

$$\|f - p_{\text{dens}}\|_1 \leq 4 \cdot \text{OPT}_k + \epsilon$$

with probability at least $1 - \delta$, where

$$\text{OPT}_k = \min_{\theta \in \Theta_k} \|f - \mathcal{M}_\theta\|_1.$$

2.2.2 Systems of polynomial inequalities

In order to fit a k -GMM to the density estimate, we solve a carefully constructed system of polynomial inequalities. Formally, a system of polynomial inequalities is an expression of the form

$$S = (Q_1 x^{(1)} \in \mathbb{R}^{n_1}) \dots (Q_v x^{(v)} \in \mathbb{R}^{n_v}) P(y, x^{(1)}, \dots, x^{(v)})$$

where

- the $y = (y_1 \dots, y_\ell)$ are free variables,
- for all $i \in [v]$, the quantifier Q_i is either \exists or \forall ,
- $P(y, x^{(1)}, \dots, x^{(v)})$ is a quantifier-free Boolean formula with m predicates of the form

$$g_i(y, x^{(1)}, \dots, x^{(v)}) \Delta_i 0$$

where each g_i is a real polynomial of degree d , and where the relations Δ_i are of the form $\Delta_i \in \{<, \geq, =, \neq, \leq, <\}$. We call such predicates *polynomial predicates*.

We say that $y \in \mathbb{R}^\ell$ is a λ -*approximate solution* for this system of polynomial inequalities if there exists a $y' \in \mathbb{R}^\ell$ such that y' satisfies the system and $\|y - y'\|_2 \leq \lambda$. We use the following result by Renegar as a black-box:

Fact 3 ([Ren92a, Ren92b]). *Let $0 < \lambda < \eta$ and let S be a system of polynomial inequalities as defined above. Then there is an algorithm SOLVE-POLY-SYSTEM(S, λ, η) that finds a λ -approximate solution if there exists a solution y with $\|y\|_2 \leq \eta$. If no such solution exists, the algorithm returns “NO-SOLUTION”. In any case, the algorithm runs in time*

$$(md)^{2^{O(v)}\ell} \Pi_k n_k \log \log \left(3 + \frac{\eta}{\lambda} \right).$$

2.2.3 Shape-restricted polynomials

Instead of fitting Gaussian pdfs to our density estimate directly, we work with piecewise polynomials as a proxy. Hence we need a good approximation of the Gaussian pdf with a piecewise polynomial. In order to achieve this, we use three pieces: two flat pieces that are constant 0 for the tails of the Gaussian, and a center piece that is given by the Taylor approximation.

Let $T_d(x)$ be the degree- d Taylor series approximation to \mathcal{N} around zero. It is straightforward to show:

Lemma 4. *Let $\epsilon, K > 0$ and let $T_d(x)$ denote the degree- d Taylor expansion of the Gaussian pdf \mathcal{N} around 0. For $d = 2K \log(1/\epsilon)$, we have*

$$\int_{2\sqrt{\log 1/\epsilon}}^{2\sqrt{\log 1/\epsilon}} |\mathcal{N}(x) - T_d(x)| dx \leq O\left(\epsilon^K \sqrt{\log(1/\epsilon)}\right).$$

Definition 5 (Shape-restricted polynomials). *Let K be such that*

$$\int_{-2\sqrt{\log 1/\epsilon}}^{2\sqrt{\log 1/\epsilon}} |\mathcal{N}(x) - T_{2K \log(1/\epsilon)}(x)| dx < \frac{\epsilon}{4}.$$

From Lemma 4 we know that such a K always exists. For any $\epsilon > 0$, let $\tilde{\mathcal{P}}_\epsilon(x)$ denote the piecewise polynomial function defined as follows:

$$\tilde{\mathcal{P}}_\epsilon(x) = \begin{cases} T_{2K \log(1/\epsilon)}(x) & \text{if } x \in [-2\sqrt{\log(1/\epsilon)}, 2\sqrt{\log(1/\epsilon)}] \\ 0 & \text{otherwise} \end{cases}.$$

For any set of parameters $\theta \in \Theta_k$, let

$$P_{\epsilon, \theta}(x) = \sum_{i=1}^k w_i \cdot \tau_i \cdot \tilde{\mathcal{P}}_\epsilon(\tau_i(x - \mu_i)).$$

It is important to note that $P_{\epsilon, \theta}(x)$ is a polynomial *both* as a function of θ and as a function of x . This allows us to fit such shape-restricted polynomials with a system of polynomial inequalities. Moreover, our shape-restricted polynomials are good approximations to GMMs. By construction, we get the following result:

Lemma 6. *Let $\theta \in \Theta_k$. Then $\|\mathcal{M}_\theta - P_{\epsilon, \theta}\|_1 \leq \epsilon$.*

Proof. We have

$$\begin{aligned} \|\mathcal{M}_\theta - P_{\epsilon, \theta}\|_1 &= \int |\mathcal{M}_\theta(x) - P_{\epsilon, \theta}(x)| dx \\ &\stackrel{(a)}{\leq} \sum_{i=1}^k w_i \int |\tau_i \cdot \mathcal{N}(\tau_i(x - \mu_i)) - \tau_i \cdot \tilde{\mathcal{P}}_\epsilon(\tau_i(x - \mu_i))| dx \\ &\stackrel{(b)}{\leq} \sum_{i=1}^k w_i \cdot \|\mathcal{N} - \tilde{\mathcal{P}}_\epsilon\|_1 \\ &\stackrel{(c)}{\leq} \sum_{i=1}^k w_i \cdot \epsilon \\ &\leq \epsilon. \end{aligned}$$

Here, (a) follows from the triangle inequality, (b) from a change of variables, and (c) from the definition of $\tilde{\mathcal{P}}_\epsilon$. \square

2.2.4 \mathcal{A}_K -norm and intersections of k -GMMs

In our system of polynomial inequalities, we must encode the constraint that the shape-restricted polynomials are a good fit to the density estimate. For this, the following notion of distance between two densities will become useful.

Definition 7 (\mathcal{A}_K -norm). *Let \mathfrak{I}_K denote the family of all sets of K disjoint intervals $\mathcal{I} = \{I_1, \dots, I_K\}$. For any measurable function $f : \mathbb{R} \rightarrow \mathbb{R}$, we define the \mathcal{A}_K -norm of f to be*

$$\|f\|_{\mathcal{A}_K} \stackrel{\text{def}}{=} \sup_{\mathcal{I} \in \mathfrak{I}_K} \sum_{I \in \mathcal{I}} \left| \int_I f(x) dx \right|.$$

For functions with few zero-crossings, the \mathcal{A}_K -norm is close to the L_1 -norm. More formally, we have the following properties, which are easy to check:

Lemma 8. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a real function. Then for any $K \geq 1$, we have*

$$\|f\|_{\mathcal{A}_K} \leq \|f\|_1.$$

Moreover, if f is continuous and there are at most $K - 1$ distinct values x for which $f(x) = 0$, then

$$\|f\|_{\mathcal{A}_K} = \|f\|_1.$$

The second property makes the \mathcal{A}_K -norm useful for us because linear combinations of Gaussians have few zeros.

Fact 9 ([KMV10] Proposition 7). *Let f be a linear combination of k Gaussian pdfs with variances $\sigma_1, \dots, \sigma_k$ so that $\sigma_i \neq \sigma_j$ for all $i \neq j$. Then there are at most $2(k - 1)$ distinct values x such that $f(x) = 0$.*

These facts give the following corollary.

Corollary 10. *Let $\theta_1, \theta_2 \in \Theta_k$ and let $K \geq 4k$. Then*

$$\|\mathcal{M}_{\theta_1} - \mathcal{M}_{\theta_2}\|_{\mathcal{A}_K} = \|\mathcal{M}_{\theta_1} - \mathcal{M}_{\theta_2}\|_1.$$

Proof. For any $\gamma > 0$, let $\theta_1^\gamma, \theta_2^\gamma$ be so that $\|\theta_i^\gamma - \theta_i\|_\infty \leq \gamma$ for $i \in \{1, 2\}$, and so that the variances of all the components in $\theta_1^\gamma, \theta_2^\gamma$ are all distinct. Lemma 8 and Fact 9 together imply that $\|\mathcal{M}_{\theta_1^\gamma} - \mathcal{M}_{\theta_2^\gamma}\|_1 = \|\mathcal{M}_{\theta_1^\gamma} - \mathcal{M}_{\theta_2^\gamma}\|_{\mathcal{A}_K}$. Letting $\gamma \rightarrow 0$ the LHS tends to $\|\mathcal{M}_{\theta_1} - \mathcal{M}_{\theta_2}\|_{\mathcal{A}_K}$, and the RHS tends to $\|\mathcal{M}_{\theta_1} - \mathcal{M}_{\theta_2}\|_1$. So we get that $\|\mathcal{M}_{\theta_1} - \mathcal{M}_{\theta_2}\|_{\mathcal{A}_K} = \|\mathcal{M}_{\theta_1} - \mathcal{M}_{\theta_2}\|_1$, as claimed. \square

3 Proper learning in the well-behaved case

In this section, we focus on properly learning a mixture of k Gaussians under the assumption that we have a “well-behaved” density estimate. We study this case first in order to illustrate our use of shape-restricted polynomials and the \mathcal{A}_K -norm. Intuitively, our notion of “well-behavedness” requires that there is a good GMM fit to the density estimate such that the mixture components and the overall mixture distribution live at roughly the same scale. Algorithmically, this allows us to solve our system of polynomial inequalities with sufficient accuracy. In Section 4, we remove this assumption and show that our algorithm works for *all* univariate mixtures of Gaussians and requires no special assumptions on the density estimation algorithm.

3.1 Overview of the Algorithm

The first step of our algorithm is to learn a good piecewise-polynomial approximation p_{dens} for the unknown density f . We achieve this by invoking recent work on density estimation [ADLS15]. Once we have obtained a good density estimate, it suffices to solve the following optimization problem:

$$\min_{\theta \in \Theta_k} \|p_{\text{dens}} - \mathcal{M}_\theta\|_1.$$

Instead of directly fitting a mixture of Gaussians, we use a mixture of shape-restricted piecewise polynomials as a proxy and solve

$$\min_{\theta \in \Theta_k} \|p_{\text{dens}} - P_{\epsilon, \theta}\|_1 .$$

Now all parts of the optimization problem are piecewise polynomials. However, we will see that we cannot directly work with the L_1 -norm without increasing the size of the corresponding system of polynomial inequalities substantially. Hence we work with the \mathcal{A}_K -norm instead and solve

$$\min_{\theta \in \Theta_k} \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K} .$$

We approach this problem by converting it to a system of polynomial inequalities with

1. $O(k)$ free variables: one per component weight, mean, and precision,
2. Two levels of quantification: one for the intervals of the \mathcal{A}_K -norm, and one for the breakpoints of the shape-restricted polynomial. Each level quantifies over $O(k)$ variables.
3. A Boolean expression on polynomials with $k^{O(k)}$ many constraints.

Finally, we use Renegar's algorithm to approximately solve our system in time $(k \log 1/\epsilon)^{O(k^4)}$. Because we only have to consider the well-behaved case, we know that finding a polynomially good approximation to the parameters will yield a sufficiently close approximation to the true underlying distribution.

3.2 Density estimation, rescaling, and well-behavedness

Density estimation As the first step of our algorithm, we obtain an agnostic estimate of the unknown probability density f . For this, we run the density estimation subroutine ESTIMATE-DENSITY(k, ϵ, δ) from Fact 2. Let p'_{dens} be the resulting $O(k)$ -piecewise polynomial. In the following, we condition on the event that

$$\|f - p'_{\text{dens}}\|_1 \leq 4 \cdot \text{OPT}_k + \epsilon .$$

which occurs with probability $1 - \delta$.

Rescaling Since we can solve systems of polynomial inequalities only with bounded precision, we have to post-process the density estimate. For example, it could be the case that some mixture components have extremely large mean parameters μ_i , in which case accurately approximating these parameters could take an arbitrary amount of time. Therefore, we shift and rescale p'_{dens} so that its non-zero part is in $[-1, 1]$ (note that p_{dens} can only have finite support because it consists of a bounded number of pieces).

Let p_{dens} be the scaled and shifted piecewise polynomial. Since the L_1 -norm is invariant under shifting and scaling, it suffices to solve the following problem

$$\min_{\theta \in \Theta_k} \|p_{\text{dens}} - \mathcal{M}_\theta\|_1 .$$

Once we have solved this problem and found a corresponding θ with

$$\|p_{\text{dens}} - \mathcal{M}_{\theta'}\|_1 \leq C$$

for some $C \geq 0$, we can undo the transformation applied to the density estimate and get a $\theta' \in \Theta_k$ such that

$$\|p'_{\text{dens}} - \mathcal{M}_{\theta'}\|_1 \leq C .$$

Well-behavedness While rescaling the density estimate p'_{dens} to the interval $[-1, 1]$ controls the size of the mean parameters μ_i , the precision parameters τ_i can still be arbitrarily large. Note that for a mixture component with very large precision, we also have to approximate the corresponding μ_i very accurately. For clarity of presentation, we ignore this issue in this section and assume that the density estimate is *well-behaved*. This assumption allows us to control the accuracy in Renegar’s algorithm appropriately. We revisit this point in Section 4 and show how to overcome this limitation. Formally, we introduce the following assumption:

Definition 11 (Well-behaved density estimate). *Let p'_{dens} be a density estimate and let p_{dens} be the rescaled version that is supported on the interval $[-1, 1]$ only. Then we say p_{dens} is γ -well-behaved if there is a set of GMM parameters $\theta \in \Theta_k$ such that*

$$\|p_{\text{dens}} - \mathcal{M}_\theta\|_1 = \min_{\theta^* \in \Theta_k} \|p_{\text{dens}} - \mathcal{M}_{\theta^*}\|_1$$

and $\tau_i \leq \gamma$ for all $i \in [k]$.

The well-behaved case is interesting in its own right because components with very high precision parameter, i.e., very spiky Gaussians, can often be learnt by clustering the samples.⁵ Moreover, the well-behaved case illustrates our use of shape-restricted polynomials and the \mathcal{A}_K -distance without additional technical difficulties.

3.3 The \mathcal{A}_K -norm as a proxy for the L_1 -norm

Computing the L_1 -distance between the density estimate p_{dens} and our shape-restricted polynomial approximation $P_{\epsilon, \theta}$ *exactly* requires knowledge of the zeros of the piecewise polynomial $p_{\text{dens}} - P_{\epsilon, \theta}$. In a system of polynomial inequalities, these zeros can be encoded by introducing auxiliary variables. However, note that we cannot simply introduce one variable per zero-crossing without affecting the running time significantly: since the polynomials have degree $O(\log 1/\epsilon)$, this would lead to $O(k \log 1/\epsilon)$ variables, and hence the running time of Renegar’s algorithm would depend exponentially on $O(\log 1/\epsilon)$. Such an exponential dependence on $\log(1/\epsilon)$ means that the running time of solving the system of polynomial inequalities becomes super-polynomial in $\frac{1}{\epsilon}$, while our goal was to avoid any polynomial dependence on $\frac{1}{\epsilon}$ when solving the system of polynomial inequalities.

Instead, we use the \mathcal{A}_K -norm as an *approximation* of the L_1 -norm. Since both $P_{\epsilon, \theta}$ and p_{dens} are close to mixtures of k Gaussians, their difference only has $O(k)$ zero crossings that contribute significantly to the L_1 -norm. More formally, we should have $\|p_{\text{dens}} - P_{\epsilon, \theta}\|_1 \approx \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K}$. And indeed:

Lemma 12. *Let $\epsilon > 0$, $k \geq 2$, $\theta \in \Theta_k$, and $K = 4k$. Then we have*

$$0 \leq \|p_{\text{dens}} - P_{\epsilon, \theta}\|_1 - \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K} \leq 8 \cdot \text{OPT}_k + O(\epsilon).$$

Proof. Recall Lemma 8: for any function f , we have $\|f\|_{\mathcal{A}_K} \leq \|f\|_1$. Thus, we know that $\|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K} \leq \|p_{\text{dens}} - P_{\epsilon, \theta}\|_1$. Hence, it suffices to show that $\|p_{\text{dens}} - P_{\epsilon, \theta}\|_1 \leq 8 \cdot \text{OPT}_k + O(\epsilon) + \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K}$.

We have conditioned on the event that the density estimation algorithm succeeds. So from Fact 2, we know that there is some mixture of k Gaussians $\mathcal{M}_{\theta'}$ so that $\|p_{\text{dens}} - \mathcal{M}_{\theta'}\|_1 \leq 4 \cdot \text{OPT}_k + \epsilon$. By repeated applications of the triangle inequality and Corollary 10, we get

$$\begin{aligned} \|p_{\text{dens}} - P_{\epsilon, \theta}\|_1 &\leq \|p_{\text{dens}} - \mathcal{M}_{\theta'}\|_1 + \|\mathcal{M}_{\theta'} - \mathcal{M}_\theta\|_1 + \|P_{\epsilon, \theta} - \mathcal{M}_\theta\|_1 \\ &\leq 4 \cdot \text{OPT}_k + \epsilon + \|\mathcal{M}_{\theta'} - \mathcal{M}_\theta\|_{\mathcal{A}_K} + \epsilon \\ &\leq 4 \cdot \text{OPT}_k + 2\epsilon + \|\mathcal{M}_{\theta'} - p_{\text{dens}}\|_{\mathcal{A}_K} + \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K} + \|P_{\epsilon, \theta} - \mathcal{M}_\theta\|_{\mathcal{A}_K} \\ &\leq 4 \cdot \text{OPT}_k + 2\epsilon + \|\mathcal{M}_{\theta'} - p_{\text{dens}}\|_1 + \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K} + \|P_{\epsilon, \theta} - \mathcal{M}_\theta\|_1 \\ &\leq 8 \cdot \text{OPT}_k + 4\epsilon + \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K}, \end{aligned}$$

⁵However, very spiky Gaussians can still be very close, which makes this approach challenging in some cases – see Section 4 for details.

as claimed. □

Using this connection between the \mathcal{A}_K -norm and the L_1 -norm, we can focus our attention on the following problem:

$$\min_{\theta \in \Theta_k} \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K} .$$

As mentioned above, this problem is simpler from a computational perspective because we only have to introduce $O(k)$ variables into the system of polynomial inequalities, regardless of the value of ϵ .

When encoding the above minimization problem in a system of polynomial inequalities, we convert it to a sequence of feasibility problems. In particular, we solve $O(\log(1/\epsilon))$ feasibility problems of the form

$$\text{Find } \theta \in \Theta_k \text{ s.t. } \|p_{\text{dens}} - P_{\epsilon, \theta}\|_{\mathcal{A}_K} < \nu . \quad (1)$$

Next, we show how to encode such an \mathcal{A}_K -constraint in a system of polynomial inequalities.

3.4 A general system of polynomial inequalities for encoding closeness in \mathcal{A}_K -norm

In this section, we give a general construction for the \mathcal{A}_K -distance between any fixed piecewise polynomial (in particular, the density estimate) and any piecewise polynomial we optimize over (in particular, our shape-restricted polynomials which we wish to fit to the density estimate). The only restriction we require is that we already have variables for the breakpoints of the polynomial we optimize over. As long as these breakpoints depend only polynomially or rationally on the parameters of the shape-restricted polynomial, this is easy to achieve. Presenting our construction of the \mathcal{A}_K -constraints in this generality makes it easy to adapt our techniques to the general algorithm (without the well-behavedness assumption, see Section 4) and to new classes of distributions (see Section 5).

The setup in this section will be as follows. Let p be a given, fixed piecewise polynomial supported on $[-1, 1]$ with breakpoints c_1, \dots, c_r . Let \mathcal{P} be a set of piecewise polynomials so that for all $\theta \in S \subseteq \mathbb{R}^u$ for some fixed, known S , there is a $P_\theta(x) \in \mathcal{P}$ with breakpoints $d_1(\theta), \dots, d_s(\theta)$ such that

- S is a semi-algebraic set.⁶ Moreover, assume membership in S can be stated as a Boolean formula over R polynomial predicates, each of degree at most D_1 , for some R, D_1 .
- For all $1 \leq i \leq s$, there is a polynomial h_i so that $h_i(d_i(\theta), \theta) = 0$, and moreover, for all θ , we have that $d_i(\theta)$ is the unique real number y satisfying $h_i(y, \theta) = 0$. That is, the breakpoints of P_θ can be encoded as polynomial equality in the θ 's. Let D_2 be the maximum degree of any h_i .
- The function $(x, \theta) \mapsto P_\theta(x)$ is a polynomial in x and θ as long as x is not at a breakpoint of P_θ . Let D_3 be the maximum degree of this polynomial.

Let $D = \max(D_1, D_2, D_3)$.

Our goal then is to encode the following problem as a system of polynomial inequalities:

$$\text{Find } \theta \in S \text{ s.t. } \|p - P_\theta\|_{\mathcal{A}_K} < \nu . \quad (2)$$

In Section 3.5, we show that this is indeed a generalization of the problem in Equation (1), for suitable choices of S and \mathcal{P} .

In the following, let $p_\theta^{\text{diff}} \stackrel{\text{def}}{=} p - P_\theta$. Note that p_θ^{diff} is a piecewise polynomial with breakpoints contained in $\{c_1, \dots, c_r, d_1(\theta), \dots, d_s(\theta)\}$. In order to encode the \mathcal{A}_K -constraint, we use the fact that a system of polynomial inequalities can contain for-all quantifiers. Hence it suffices to encode the \mathcal{A}_K -constraint for a single set of K intervals. We provide a construction for a single \mathcal{A}_K -constraint in Section 3.4.1. In Section 3.4.2, we introduce two further constraints that guarantee validity of the parameters θ and combine these constraints with the \mathcal{A}_K -constraint to produce the full system of polynomial inequalities.

⁶Recall a semi-algebraic set is a set where membership in the set can be described by polynomial inequalities.

3.4.1 Encoding Closeness for a Fixed Set of Intervals

Let $[a_1, b_1], \dots, [a_K, b_K]$ be K disjoint intervals. In this section we show how to encode the following constraint:

$$\sum_{i=1}^K \left| \int_{a_i}^{b_i} p_{\theta}^{\text{diff}}(x) dx \right| \leq \nu.$$

Note that a given interval $[a_i, b_i]$ might contain several pieces of p_{θ}^{diff} . In order to encode the integral over $[a_i, b_i]$ correctly, we must therefore know the current order of the breakpoints (which can depend on θ).

However, once the order of the breakpoints of p_{θ}^{diff} and the a_i and b_i is fixed, the integral over $[a_i, b_i]$ becomes the integral over a fixed set of sub-intervals. Since the integral over a single polynomial piece is still a polynomial, we can then encode this integral over $[a_i, b_i]$ piece-by-piece.

More formally, let Φ be the set of permutations of the variables

$$\{a_1, \dots, a_K, b_1, \dots, b_K, c_1, \dots, c_r, d_1(\theta), \dots, d_s(\theta)\}$$

such that (i) the a_i appear in order, (ii) the b_i appear in order, (iii) a_i appears before b_i , and (iv) the c_i appear in order. Let $t = 2K + r + s$. For any $\phi = (\phi_1, \dots, \phi_t) \in \Phi$, let

$$\text{ordered}_{p, \mathcal{P}}(\phi) \stackrel{\text{def}}{=} \bigwedge_{i=1}^{t-1} (\phi_i \leq \phi_{i+1}).$$

Note that for any fixed ϕ , this is an unquantified Boolean formula with polynomial constraints in the unknown variables. The order constraints encode whether the current set of variables corresponds to ordered variables under the permutation represented by ϕ . An important property of an ordered ϕ is the following: in each interval $[\phi_i, \phi_{i+1}]$, the piecewise polynomial p_{θ}^{diff} has exactly one piece. This allows us to integrate over p_{θ}^{diff} in our system of polynomial inequalities.

Next, we need to encode whether a fixed interval between ϕ_i and ϕ_{i+1} is contained in one of the \mathcal{A}_K -intervals, i.e., whether we have to integrate p_{θ}^{diff} over the interval $[\phi_i, \phi_{i+1}]$ when we compute the \mathcal{A}_K -norm of p_{θ}^{diff} . We use the following expression:

$$\text{is-active}_{p, \mathcal{P}}(\phi, i) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if there is a } j \text{ such that } a_j \text{ appears as or before } \phi_i \text{ in } \phi \\ & \text{and } b_j \text{ appears as or after } \phi_{i+1} \\ 0 & \text{otherwise} \end{cases}.$$

Note that for fixed ϕ and i , this expression is either 0 or 1 (and hence trivially a polynomial).

With the constructs introduced above, we can now integrate p_{θ}^{diff} over an interval $[\phi_i, \phi_{i+1}]$. It remains to bound the absolute value of the integral for each individual piece. For this, we introduce a set of t new variables ξ_1, \dots, ξ_t which will correspond to the absolute value of the integral in the corresponding piece.

$$\begin{aligned} \mathcal{A}_K\text{-bounded-interval}_{p, \mathcal{P}}(\phi, \theta, \xi, i) &\stackrel{\text{def}}{=} \left(\left(-\xi_i \leq \int_{\phi_i}^{\phi_{i+1}} p_{\theta}^{\text{diff}}(x) dx \right) \wedge \left(\int_{\phi_i}^{\phi_{i+1}} p_{\theta}^{\text{diff}}(x) dx \leq \xi_i \right) \right) \\ &\vee (\text{is-active}_{p, \mathcal{P}}(\phi, i) = 0). \end{aligned}$$

Note that the above is a valid polynomial constraint because p_{θ}^{diff} depends only on θ and x for fixed breakpoint order ϕ and fixed interval $[\phi_i, \phi_{i+1}]$. Moreover, recall that by assumption, $P_{\epsilon, \theta}(x)$ depends polynomially on both θ and x , and therefore the same holds for p_{θ}^{diff} .

We extend the \mathcal{A}_K -check for a single interval to the entire range of p_{θ}^{diff} as follows:

$$\mathcal{A}_K\text{-bounded-fixed-permutation}_{p, \mathcal{P}}(\phi, \theta, \xi) \stackrel{\text{def}}{=} \bigwedge_{i=1}^{t-1} \mathcal{A}_K\text{-bounded-interval}_{p, \mathcal{P}}(\phi, \theta, \xi, i).$$

We now have all the tools to encode the \mathcal{A}_K -constraint for a fixed set of intervals:

$$\mathcal{A}_K\text{-bounded}_{p,\mathcal{P}}(\theta, \nu, a, b, c, d, \xi) \stackrel{\text{def}}{=} \left(\sum_{i=1}^{t-1} \xi_i \leq \nu \right) \wedge \left(\bigwedge_{i=1}^{t-1} (\xi_i \geq 0) \right) \\ \wedge \left(\bigvee_{\phi \in \Phi} \text{ordered}_{p,\mathcal{P}}(\phi) \wedge \mathcal{A}_K\text{-bounded-fixed-permutation}_{p,\mathcal{P}}(\phi, \theta, \xi) \right).$$

By construction, the above constraint now satisfies the following:

Lemma 13. *There exists a vector $\xi \in \mathbb{R}^t$ such that $\mathcal{A}_K\text{-bounded}_{p,\mathcal{P}}(\theta, \nu, a, b, c, d, \xi)$ is true if and only if*

$$\sum_{i=1}^K \left| \int_{a_i}^{b_i} p_{\theta}^{\text{diff}}(x) dx \right| \leq \nu.$$

Moreover, $\mathcal{A}_K\text{-bounded}_{p,\mathcal{P}}$ has less than $6t^{t+1}$ polynomial constraints.

The bound on the number of polynomial constraints follows simply from counting the number of polynomial constraints in the construction described above.

3.4.2 Complete system of polynomial inequalities

In addition to the \mathcal{A}_K -constraint introduced in the previous subsection, our system of polynomial inequalities contains the following constraints:

Valid parameters First, we encode that the mixture parameters we optimize over are valid, i.e., we let

$$\text{valid-parameters}_S(\theta) \stackrel{\text{def}}{=} \theta \in S.$$

Recall this can be expressed as a Boolean formula over R polynomial predicates of degree at most D .

Correct breakpoints We require that the d_i are indeed the breakpoints of the shape-restricted polynomial P_{θ} . By the assumption, this can be encoded by the following constraint:

$$\text{correct-breakpoints}_{\mathcal{P}}(\theta, d) \stackrel{\text{def}}{=} \bigwedge_{i=1}^s (h_i(d_i(\theta), \theta) = 0).$$

The full system of polynomial inequalities We now combine the constraints introduced above and introduce our entire system of polynomial inequalities:

$$S_{K,p,\mathcal{P},S}(\nu) = \forall a_1, \dots, a_K, b_1, \dots, b_K : \\ \exists d_1, \dots, d_s, \xi_1 \dots \xi_t : \\ \text{valid-parameters}_S(\theta) \wedge \text{correct-breakpoints}_{\mathcal{P}}(\theta, d) \wedge \mathcal{A}_K\text{-bounded}_{p,\mathcal{P}}(\theta, \nu, a, b, c, d, \xi).$$

This system of polynomial inequalities has

- two levels of quantification, with $2K$ and $s + t$ variables, respectively,
- u free variables,
- $R + s + 4t^{t+1}$ polynomial constraints,
- and maximum degree D in the polynomial constraints.

Let γ be a bound on the free variables, i.e., $\|\theta\|_2 \leq \gamma$, and let λ be a precision parameter. Then Renegar's algorithm (see Fact 3) finds a λ -approximate solution θ for this system of polynomial inequalities satisfying $\|\theta\|_2 \leq \gamma$, if one exists, in time

$$((R + s + 6t^{t+1})D)^{O(K(s+t)u)} \log \log \left(3 + \frac{\gamma}{\lambda} \right).$$

3.5 Instantiating the system of polynomial inequalities for GMMs

We now show how to use the system of polynomial inequalities developed in the previous subsection for our initial goal: that is, encoding closeness between a well-behaved density estimate and a set of shape-restricted polynomials (see Equation 1). Our fixed piecewise polynomial (p in the subsection above) will be p_{dens} . The set of piecewise polynomials we optimize over (the set \mathcal{P} in the previous subsection) will be the set \mathcal{P}_ϵ of all shape-restricted polynomials $P_{\epsilon,\theta}$. Our S (the domain of θ) will be $\Theta'_k \subseteq \Theta_k$, which we define below. For each $\theta \in S$, we associate it with $P_{\epsilon,\theta}$. Moreover:

- Define

$$\Theta_{k,\gamma} = \left\{ \theta \mid \left(\sum_{i=1}^k w_i = 1 \right) \wedge (\forall i \in [k] : (w_i \geq 0) \wedge (\gamma \geq \tau_i > 0) \wedge (-1 \leq \mu_i \leq 1)) \right\},$$

that is, the set of parameters which have bounded means and variances. S is indeed semi-algebraic, and membership in S can be encoded using $2k + 1$ polynomial predicates, each with degree $D_1 = 1$.

- For any fixed parameter $\theta \in \Theta_k$, the shape-restricted polynomial P_θ has $s = 2k$ breakpoints by definition, and the breakpoints $d_1(\theta), \dots, d_{2k}(\theta)$ of $P_{\epsilon,\theta}$ occur at

$$d_{2i-1}(\theta) = \frac{1}{\tau_i} (\mu_i - 2\tau_i \log(1/\epsilon)) , \quad d_{2i}(\theta) = \frac{1}{\tau_i} (\mu_i + 2\tau_i \log(1/\epsilon)) , \quad \text{for all } 1 \leq i \leq k .$$

Thus, for all parameters θ , the breakpoints $d_1(\theta), \dots, d_{2k}(\theta)$ are the unique numbers so that so that

$$\tau_i \cdot d_{2i-1}(\theta) - (\mu_i - 2\tau_i \log(1/\epsilon)) = 0 , \quad \tau_i \cdot d_{2i}(\theta) - (\mu_i + 2\tau_i \log(1/\epsilon)) = 0 , \quad \text{for all } 1 \leq i \leq k ,$$

and thus each of the $d_1(\theta), \dots, d_{2k}(\theta)$ can be encoded as a polynomial equality of degree $D_2 = 2$.

- Finally, it is straightforward to verify that the map $(x, \theta) \rightarrow P_{\epsilon,\theta}(x)$ is a polynomial of degree $D_3 = O(\log 1/\epsilon)$ in (x, θ) , at any point where x is not at a breakpoint of P_θ .

From the previous subsection, we know that the system of polynomial inequalities $S_{K,p_{\text{dens}},\mathcal{P}_\epsilon,\Theta_{k,\gamma}}(\nu)$ has two levels of quantification, each with $O(k)$ variables, it has $k^{O(k)}$ polynomial constraints, and has maximum degree $O(\log 1/\epsilon)$ in the polynomial constraints. Hence, we have shown:

Corollary 14. *For any fixed ϵ , the system of polynomial inequalities $S_{K,p_{\text{dens}},\mathcal{P}_\epsilon,\Theta_{k,\gamma}}(\nu)$ encodes Equation (1). Moreover, for all $\gamma, \lambda \geq 0$, Renegar's algorithm SOLVE-POLY-SYSTEM($S_{K,p_{\text{dens}},\mathcal{P}_\epsilon,\Theta_{k,\gamma}}(\nu), \lambda, \gamma$) runs in time $(k \log(1/\epsilon))^{O(k^4)} \log \log(3 + \frac{\gamma}{\lambda})$.*

3.6 Overall learning algorithm

We now combine our tools developed so far and give an agnostic learning algorithm for the case of well-behaved density estimates (see Algorithm 1).

3.7 Analysis

Before we prove correctness of LEARN-WELL-BEHAVED-GMM, we introduce two auxiliary lemmas.

An important consequence of the well-behavedness assumption (see Definition 11) are the following robustness properties.

Lemma 15 (Parameter stability). *Fix $2 \geq \epsilon > 0$. Let the parameters $\theta, \theta' \in \Theta_k$ be such that (i) $\tau_i, \tau'_i \leq \gamma$ for all $i \in [k]$ and (ii) $\|\theta - \theta'\|_2 \leq C \frac{1}{\gamma} \left(\frac{\epsilon}{k}\right)^2$, for some universal constant C . Then*

$$\|\mathcal{M}_\theta - \mathcal{M}_{\theta'}\|_1 \leq \epsilon .$$

Algorithm 1 Algorithm for learning a mixture of Gaussians in the well-behaved case.

```

1: function LEARN-WELL-BEHAVED-GMM( $k, \epsilon, \delta, \gamma$ )
2:   ▷ Density estimation. Only this step draws samples.
3:    $p'_{\text{dens}} \leftarrow \text{ESTIMATE-DENSITY}(k, \epsilon, \delta)$ 

4:   ▷ Rescaling
5:   Let  $p_{\text{dens}}$  be a rescaled and shifted version of  $p'_{\text{dens}}$  such that the support of  $p_{\text{dens}}$  is  $[-1, 1]$ .
6:   Let  $\alpha$  and  $\beta$  be such that  $p_{\text{dens}}(x) = p'_{\text{dens}}\left(\frac{2(x-\alpha)}{\beta-\alpha} - 1\right)$ 

7:   ▷ Fitting shape-restricted polynomials
8:    $K \leftarrow 4k$ 
9:    $\nu \leftarrow \epsilon$ 
10:   $\theta \leftarrow \text{SOLVE-POLY-SYSTEM}(S_{K, p_{\text{dens}}, \mathcal{P}_{\epsilon, \Theta_{k, \gamma}}}(\nu), C \frac{1}{\gamma} \left(\frac{\epsilon}{k}\right)^2, 3k\gamma)$ 
11:  while  $\theta$  is “NO-SOLUTION” do
12:     $\nu \leftarrow 2 \cdot \nu$ 
13:     $\theta \leftarrow \text{SOLVE-POLY-SYSTEM}(S_{K, p_{\text{dens}}, \mathcal{P}_{\epsilon, \Theta_{k, \gamma}}}(\nu), C \frac{1}{\gamma} \left(\frac{\epsilon}{k}\right)^2, 3k\gamma)$ 

14:  ▷ Fix the parameters
15:  for  $i = 1, \dots, k$  do
16:    if  $\tau_i \leq 0$ , set  $w_i \leftarrow 0$  and set  $\tau_i$  to be arbitrary but positive.
17:  Let  $W = \sum_{i=1}^k w_i$ 
18:  for  $i = 1, \dots, k$  do
19:     $w_i \leftarrow w_i/W$ 

20:  ▷ Undo the scaling
21:   $w'_i \leftarrow w_i$ 
22:   $\mu'_i \leftarrow \frac{(\mu_i+1)(\beta-\alpha)}{2} + \alpha$ 
23:   $\tau'_i \leftarrow \frac{\tau_i}{\beta-\alpha}$ 
24:  return  $\theta'$ 

```

Before we prove this lemma, we first need a calculation which quantifies the robustness of the standard normal pdf to small perturbations.

Lemma 16. *For all $2 \geq \epsilon > 0$, there is a $\delta_1 = \delta_1(\epsilon) = \frac{\epsilon}{20\sqrt{\log(1/\epsilon)}} \geq O(\epsilon^2)$ so that for all $\delta \leq \delta_1$, we have $\|\mathcal{N}(x) - \mathcal{N}(x + \delta)\|_1 \leq O(\epsilon)$.*

Proof. Note that if $\epsilon > 2$ this claim holds trivially for all choices of δ since the L_1 -distance between two distributions can only ever be 2. Thus assume that $\epsilon \leq 2$. Let I be an interval centered at 0 so that both $\mathcal{N}(x)$ and $\mathcal{N}(x + \delta)$ assign $1 - \frac{\epsilon}{2}$ weight on this interval. By standard properties of Gaussians, we know that $|I| \leq 10\sqrt{\log(1/\epsilon)}$. We thus have

$$\|\mathcal{N}(x) - \mathcal{N}(x + \delta)\|_1 \leq \int_I |\mathcal{N}(x) - \mathcal{N}(x + \delta)| dx + \epsilon.$$

By Taylor’s theorem we have that for all x ,

$$\left| e^{-(x+\delta)^2/2} - e^{-x^2/2} \right| \leq C \cdot \delta$$

for some universal constant $C = \max_{x \in \mathbb{R}} \frac{d}{dx} (e^{-x^2/2}) \leq 1$. Since we choose $\delta_1 \leq \frac{\epsilon}{20\sqrt{\log(1/\epsilon)}}$, we must have that

$$\|\mathcal{N}(x) - \mathcal{N}(x + \delta)\|_1 \leq O(\epsilon),$$

as claimed. □

Proof of Lemma 15. Notice the ℓ_2 guarantee of Renegar's algorithm (see Fact 3) also trivially implies an ℓ_∞ guarantee on the error in the parameters θ ; that is, for all i , we will have that the weights, means, and variances of the two components differ by at most $C\frac{1}{\gamma}\left(\frac{\epsilon}{k}\right)^2$. By repeated applications of the triangle inequality to the quantity in the lemma, it suffices to show the three following claims:

- For any μ, τ ,

$$\|w_1 \mathcal{N}_{\mu, \tau}(x) - w_2 \mathcal{N}_{\mu, \tau}(x)\|_1 \leq \frac{\epsilon}{3k}$$

$$\text{if } |w_1 - w_2| \leq C\frac{1}{\gamma}\left(\frac{\epsilon}{k}\right)^2.$$

- For any $\tau \leq \gamma$,

$$\|\mathcal{N}_{\mu_1, \tau}(x) - \mathcal{N}_{\mu_2, \tau}(x)\|_1 \leq \frac{\epsilon}{3k}$$

$$\text{if } |\mu_1 - \mu_2| \leq C\frac{1}{\gamma}\left(\frac{\epsilon}{k}\right)^2.$$

- For any μ ,

$$\|\mathcal{N}_{\mu, \tau_1}(x) - \mathcal{N}_{\mu, \tau_2}(x)\|_1 \leq \frac{\epsilon}{3k}$$

$$\text{if } |\tau_1 - \tau_2| \leq C\frac{1}{\gamma}\left(\frac{\epsilon}{k}\right)^2.$$

The first inequality is trivial, for C sufficiently small. The second and third inequalities follow from a change of variables and an application of Lemma 16. \square

Recall that our system of polynomial inequalities only considers mean parameters in $[-1, 1]$. The following lemma shows that this restriction still allows us to find a good approximation once the density estimate is rescaled to $[-1, 1]$.

Lemma 17 (Restricted means). *Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a function supported on $[-1, 1]$, i.e., $g(x) = 0$ for $x \notin [-1, 1]$. Moreover, let $\theta^* \in \Theta_k$. Then there is a $\theta' \in \Theta_k$ such that $\mu'_i \in [-1, 1]$ for all $i \in [k]$ and*

$$\|g - \mathcal{M}_{\theta'}\|_1 \leq 5 \cdot \|g - \mathcal{M}_{\theta^*}\|_1.$$

Proof. Let $A = \{i \mid \mu_i^* \in [-1, 1]\}$ and $B = [k] \setminus A$. Let θ' be defined as follows:

- $w'_i = w_i^*$ for all $i \in [k]$.
- $\mu'_i = \mu_i^*$ for $i \in A$ and $\mu'_i = 0$ for $i \in B$.
- $\tau'_i = \tau_i^*$ for all $i \in [k]$.

From the triangle inequality, we have

$$\|g - \mathcal{M}_{\theta'}\|_1 \leq \|g - \mathcal{M}_{\theta^*}\|_1 + \|\mathcal{M}_{\theta^*} - \mathcal{M}_{\theta'}\|_1. \quad (3)$$

Hence it suffices to bound $\|\mathcal{M}_{\theta^*} - \mathcal{M}_{\theta'}\|_1$.

Note that for $i \in B$, the corresponding i -th component has at least half of its probability mass outside $[-1, 1]$. Since g is zero outside $[-1, 1]$, this mass of the i -th component must therefore contribute to the error $\|g - \mathcal{M}_{\theta^*}\|_1$. Let $\mathbb{1}[x \notin [-1, 1]]$ be the indicator function of the set $\mathbb{R} \setminus [-1, 1]$. Then we get

$$\|g - \mathcal{M}_{\theta^*}\|_1 \geq \|\mathcal{M}_{\theta^*} \cdot \mathbb{1}[x \notin [-1, 1]]\|_1 \geq \frac{1}{2} \left\| \sum_{i \in B} w_i^* \cdot \mathcal{N}_{\mu_i^*, \tau_i^*} \right\|_1.$$

For $i \in A$, the mixture components of \mathcal{M}_{θ^*} and $\mathcal{M}_{\theta'}$ match. Hence we have

$$\begin{aligned}
\|\mathcal{M}_{\theta^*} - \mathcal{M}_{\theta'}\|_1 &= \left\| \sum_{i \in B} w_i^* \cdot \mathcal{N}_{\mu_i^*, \tau_i^*} - \sum_{i \in B} w_i' \cdot \mathcal{N}_{\mu_i', \tau_i'} \right\|_1 \\
&\leq \left\| \sum_{i \in B} w_i^* \cdot \mathcal{N}_{\mu_i^*, \tau_i^*} \right\|_1 + \left\| \sum_{i \in B} w_i' \cdot \mathcal{N}_{\mu_i', \tau_i'} \right\|_1 \\
&= 2 \cdot \left\| \sum_{i \in B} w_i^* \cdot \mathcal{N}_{\mu_i^*, \tau_i^*} \right\|_1 \\
&\leq 4 \cdot \|g - \mathcal{M}_{\theta^*}\|_1.
\end{aligned}$$

Combining this inequality with (3) gives the desired result. \square

We now prove our main theorem for the well-behaved case.

Theorem 18. *Let $\delta, \epsilon, \gamma > 0$, $k \geq 1$, and let f be the pdf of the unknown distribution. Moreover, assume that the density estimate p'_{dens} obtained in Line 3 of Algorithm 1 is γ -well-behaved. Then the algorithm LEARN-WELL-BEHAVED-GMM($k, \epsilon, \delta, \gamma$) returns a set of GMM parameters θ' such that*

$$\|\mathcal{M}_{\theta'} - f\|_1 \leq 60 \cdot \text{OPT}_k + \epsilon$$

with probability $1 - \delta$. Moreover, the algorithm runs in time

$$\left(k \cdot \log \frac{1}{\epsilon}\right)^{O(k^4)} \cdot \log \frac{1}{\epsilon} \cdot \log \log \frac{k\gamma}{\epsilon} + \tilde{O}\left(\frac{k}{\epsilon^2}\right).$$

Proof. First, we prove the claimed running time. From Fact 2, we know that the density estimation step has a time complexity of $\tilde{O}\left(\frac{k}{\epsilon^2}\right)$. Next, consider the second stage where we fit shape-restricted polynomials to the density estimate. Note that for $\nu = 3$, the system of polynomial inequalities $S_{p_{\text{dens}}, \mathcal{P}_\epsilon}(\nu)$ is trivially satisfiable because the \mathcal{A}_K -norm is bounded by the L_1 -norm and the L_1 -norm between the two (approximate) densities is at most $2 + O(\epsilon)$. Hence the while-loop in the algorithm takes at most $O(\log \frac{1}{\epsilon})$ iterations. Combining this bound with the size of the system of polynomial inequalities (see Subsection 3.4.2) and the time complexity of Renegar's algorithm (see Fact 3), we get the following running time for solving all systems of polynomial inequalities proposed by our algorithm:

$$\left(k \cdot \log \frac{1}{\epsilon}\right)^{O(k^4)} \cdot \log \log \frac{k\gamma}{\epsilon} \cdot \log \frac{1}{\epsilon}.$$

This proves the stated running time.

Next, we consider the correctness guarantee. We condition on the event that the density estimation stage succeeds, which occurs with probability $1 - \delta$ (Fact 2). Then we have

$$\|f - p'_{\text{dens}}\|_1 \leq 4 \cdot \text{OPT}_k + \epsilon.$$

Moreover, we can assume that the rescaled density estimate p_{dens} is γ -well-behaved. Recalling Definition 11, this means that there is a set of GMM parameters $\theta \in \Theta_k$ such that $\tau_i \leq \gamma$ for all $i \in [k]$ and

$$\begin{aligned}
\|p_{\text{dens}} - \mathcal{M}_\theta\|_1 &= \min_{\theta^* \in \Theta_k} \|p_{\text{dens}} - \mathcal{M}_{\theta^*}\|_1 \\
&= \min_{\theta^* \in \Theta_k} \|p'_{\text{dens}} - \mathcal{M}_{\theta^*}\|_1 \\
&\leq \min_{\theta^* \in \Theta_k} \|p'_{\text{dens}} - f\|_1 + \|f - \mathcal{M}_{\theta^*}\|_1 \\
&\leq 4 \cdot \text{OPT}_k + \epsilon + \min_{\theta^* \in \Theta_k} \|f - \mathcal{M}_{\theta^*}\|_1 \\
&\leq 5 \cdot \text{OPT}_k + \epsilon.
\end{aligned}$$

Applying the triangle inequality again, this implies that

$$\|p_{\text{dens}} - P_{\epsilon, \theta}\|_1 \leq \|p_{\text{dens}} - \mathcal{M}_\theta\|_1 + \|\mathcal{M}_\theta - P_{\epsilon, \theta}\|_1 \leq 5 \cdot \text{OPT}_k + 2\epsilon.$$

This almost implies that $S_{p_{\text{dens}}, \mathcal{P}_\epsilon}(\nu)$ is feasible for $\nu \geq 5 \cdot \text{OPT}_k + 2\epsilon$. However, there are two remaining steps. First, recall that the system of polynomial inequalities restricts the means to lie in $[-1, 1]$. Hence we use Lemma 17, which implies that there is a $\theta^\dagger \in \Theta_k$ such that $\mu_i^\dagger \in [-1, 1]$ and

$$\|p_{\text{dens}} - P_{\epsilon, \theta^\dagger}\|_1 \leq 25 \cdot \text{OPT}_k + 10\epsilon.$$

Moreover, the system of polynomial inequalities works with the \mathcal{A}_K -norm instead of the L_1 -norm. Using Lemma 12, we get that

$$\|p_{\text{dens}} - P_{\epsilon, \theta^\dagger}\|_{\mathcal{A}_K} \leq \|p_{\text{dens}} - P_{\epsilon, \theta^\dagger}\|_1.$$

Therefore, in some iteration when

$$\nu \leq 2 \cdot (25 \cdot \text{OPT}_k + 10\epsilon) = 50 \cdot \text{OPT}_k + 20\epsilon$$

the system of polynomial inequalities $S_{p_{\text{dens}}, \mathcal{P}_\epsilon, \Theta_k, \gamma}(\nu)$ become feasible and Renegar's algorithm guarantees that we find parameters θ' such that $\|\theta' - \theta^\dagger\|_2 \leq \frac{\epsilon}{\gamma}$ for some $\theta^\dagger \in \Theta_k$ and

$$\|p_{\text{dens}} - \mathcal{M}_{\theta'}\|_{\mathcal{A}_K} \leq 50 \cdot \text{OPT}_k + O(\epsilon).$$

Note that we used well-behavedness here to ensure that the precisions in θ^\dagger are bounded by γ . Let θ be the parameters we return. It is not difficult to see that $\|\theta - \theta^\dagger\|_2 \leq \frac{2\epsilon}{\gamma}$. We convert this back to an L_1 guarantee via Lemma 12:

$$\|p_{\text{dens}} - \mathcal{M}_{\theta'}\|_1 \leq 56 \cdot \text{OPT}_k + O(\epsilon).$$

Next, we use parameter stability (Lemma 15) and get

$$\|p_{\text{dens}} - \mathcal{M}_\theta\|_1 \leq 56 \cdot \text{OPT}_k + O(\epsilon).$$

We now relate this back to the unknown density f . Let θ' be the parameters θ scaled back to the original density estimate (see Lines 21 to 23 in Algorithm 1). Then we have

$$\|p'_{\text{dens}} - \mathcal{M}_{\theta'}\|_1 \leq 56 \cdot \text{OPT}_k + O(\epsilon).$$

Using the fact that p'_{dens} is a good density estimate, we get

$$\begin{aligned} \|f - \mathcal{M}_{\theta'}\|_1 &\leq \|f - p'_{\text{dens}}\|_1 + \|p'_{\text{dens}} - \mathcal{M}_{\theta'}\|_1 \\ &\leq 4 \cdot \text{OPT}_k + \epsilon + 56 \cdot \text{OPT}_k + O(\epsilon) \\ &\leq 60 \cdot \text{OPT}_k + O(\epsilon). \end{aligned}$$

As a final step, we choose an internal ϵ' in our algorithm so that the $O(\epsilon')$ in the above guarantee becomes bounded by ϵ . This proves the desired approximation guarantee. \square

4 General algorithm

4.1 Preliminaries

As before, we let p_{dens} be the piecewise polynomial returned by LEARN-PIECEWISE-POLYNOMIAL (see Fact 2). Let I_0, \dots, I_{s+1} be the intervals defined by the breakpoints of p_{dens} . Recall that p_{dens} has degree $O(\log 1/\epsilon)$ and has $s + 2 = O(k)$ pieces. Furthermore, I_0 and I_{s+1} are unbounded in length, and on these intervals p_{dens} is zero. By rescaling and translating, we may assume WLOG that $\cup_{i=1}^s I_i$ is $[-1, 1]$.

Recall that \mathcal{I} is defined by the set of intervals $\{I_1, \dots, I_s\}$. We know that $s = O(k)$. Intuitively, these intervals capture the different scales at which we need to operate. We formalize this intuition below.

Definition 19. For any Gaussian $\mathcal{N}_{\mu,\tau}$, let $L(\mathcal{N}_{\mu,\tau})$ be the interval centered at μ on which $\mathcal{N}_{\mu,\tau}$ places exactly W of its weight, where $0 < W < 1$ is a universal constant we will determine later. By properties of Gaussians, there is some absolute constant $\omega > 0$ such that $\mathcal{N}_{\mu,\tau}(x) \geq \omega\tau$ for all $x \in L(\mathcal{N}_{\mu,\tau})$.

Definition 20. Say a Gaussian $\mathcal{N}_{\mu,\tau}$ is admissible if (i) $\mathcal{N}_{\mu,\tau}$ places at least $1/2$ of its mass in $[-1, 1]$, and (ii) there is a $J \in \mathcal{I}$ so that $|J \cap L(\mathcal{N}_{\mu,\tau})| \geq 1/(8s\tau)$ and so that

$$\tau \leq \frac{1}{|J|} \cdot \phi,$$

where

$$\phi = \phi(\epsilon, k) \stackrel{\text{def}}{=} \frac{32k}{\omega\epsilon} m(m+1)^2 \cdot (\sqrt{2} + 1)^m,$$

where m is the degree of p_{dens} . We call the interval $J \in \mathcal{I}$ satisfying this property on which $\mathcal{N}_{\mu,\tau}$ places most of its mass its associated interval.

Fix $\theta \in \Theta_k$. We say the ℓ -th component is admissible if the underlying Gaussian is admissible and moreover $w_\ell \geq \epsilon/k$.

Notice that since $m = O(\log(1/\epsilon))$, we have that $\phi(\epsilon, k) = \text{poly}(1/\epsilon, k)$.

Lemma 21 (No Interaction Lemma). Fix $\theta \in \Theta_k$. Let $S_{\text{good}}(\theta) \subseteq [k]$ be the set of $\ell \in [k]$ whose corresponding mixture component is admissible, and let $S_{\text{bad}}(\theta)$ be the rest. Then, we have

$$\|\mathcal{M}_\theta - p_{\text{dens}}\|_1 \geq \left\| \sum_{\ell \in S_{\text{good}}(\theta)} w_\ell \cdot \mathcal{N}_{\mu_\ell, \tau_\ell} - p_{\text{dens}} \right\|_1 + \frac{1}{2} \sum_{\ell \in S_{\text{bad}}(\theta)} w_\ell - 2\epsilon.$$

We briefly remark that the constant $\frac{1}{2}$ we obtain here is somewhat arbitrary; by choosing different universal constants above, one can obtain any fraction arbitrarily close to one, at a minimal loss.

Proof. Fix $\ell \in S_{\text{bad}}(\theta)$, and denote the corresponding component \mathcal{N}_ℓ . Recall that it has mean μ_ℓ and precision τ_ℓ . Let $L_\ell = L(\mathcal{N}_\ell)$.

Let $\mathcal{M}_\theta^{-\ell}(x) = \sum_{i \neq \ell} w_i \mathcal{N}_{\mu_i, \tau_i}(x)$ be the density of the mixture without the ℓ -th component. We will show that

$$\|\mathcal{M}_\theta - p_{\text{dens}}\|_1 \geq \|\mathcal{M}_\theta^{-\ell} - p_{\text{dens}}\|_1 + \frac{1}{2}w_\ell - \frac{2\epsilon}{k}.$$

It suffices to prove this inequality because then we may repeat the argument with a different $\ell' \in S_{\text{bad}}(\theta)$ until we have subtracted out all such ℓ , and this yields the claim in the lemma.

If $w_\ell \leq \epsilon/k$ then this statement is obvious. If \mathcal{N}_ℓ places less than half its weight on $[-1, 1]$, then this is also obvious. Thus we will assume that $w_\ell > \epsilon/k$ and \mathcal{N}_ℓ places at least half its weight on $[-1, 1]$.

Let \mathcal{I}_ℓ be the set of intervals in \mathcal{I} which intersect L_ℓ . We partition the intervals in \mathcal{I}_ℓ into two groups:

1. Let \mathcal{L}_1 be the set of intervals $J \in \mathcal{I}_\ell$ so that $|J \cap L_\ell| \leq 1/(8s\tau_\ell)$.
2. Let \mathcal{L}_2 be the set of intervals $J \in \mathcal{I}_\ell$ not in \mathcal{L}_1 so that

$$\tau_\ell > \frac{1}{|J|} \cdot \phi.$$

By the definition of admissibility, this is indeed a partition of \mathcal{I}_ℓ .

We have

$$\begin{aligned}
\|\mathcal{M}_\theta - p_{\text{dens}}\|_1 &= \|\mathcal{M}_\theta^{-\ell} + \mathcal{N}_\ell - p_{\text{dens}}\|_1 \\
&= \int_{L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| \, dx + \int_{L_\ell^\varepsilon} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| \, dx \\
&\geq \int_{L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| \, dx + \int_{L_\ell^\varepsilon} |\mathcal{M}_\theta^{-\ell}(x) - p_{\text{dens}}(x)| \, dx - w_\ell \int_{L_\ell^\varepsilon} \mathcal{N}_\ell(x) \, dx \\
&\geq \int_{L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| \, dx + \int_{L_\ell^\varepsilon} |\mathcal{M}_\theta^{-\ell}(x) - p_{\text{dens}}(x)| \, dx - (1 - W)w_\ell .
\end{aligned}$$

We split the first term on the RHS into two parts, given by our partition:

$$\begin{aligned}
\int_{L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| \, dx &= \sum_{J \in \mathcal{L}_1} \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| \, dx \\
&\quad + \sum_{J \in \mathcal{L}_2} \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| \, dx .
\end{aligned}$$

We lower bound the contribution of each term separately.

(1) We first bound the first term. Since for each $J \in \mathcal{L}_1$ we have $|J \cap L_\ell| \leq 1/(8s\tau_\ell)$, we know that

$$\int_{J \cap L_\ell} \mathcal{N}_\ell(x) \, dx \leq \frac{1}{8s} \tag{4}$$

and so

$$\begin{aligned}
\sum_{J \in \mathcal{L}_1} \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| \, dx &\geq \sum_{J \in \mathcal{L}_1} \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) - p_{\text{dens}}(x)| \, dx - |\mathcal{L}_1| \cdot w_\ell \cdot \frac{1}{8s} \\
&\geq \sum_{J \in \mathcal{L}_1} \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) - p_{\text{dens}}(x)| \, dx - \frac{1}{8} w_\ell
\end{aligned}$$

since \mathcal{I} and thus \mathcal{L}_1 contains at most s intervals.

(2) We now consider the second term. Fix a $J \in \mathcal{L}_2$, and let p_J be the polynomial which is equal to p_{dens} on J . Since $\int p_{\text{dens}} \leq 1 + \epsilon \leq 2$ (as otherwise its L_1 -distance to the unknown density would be more than ϵ) and p_{dens} is nonnegative, we also know that $\int_J p_J \leq 2$. We require the following fact (see [ADLS15]):

Fact 22. *Let $p(x) = \sum_{j=0}^m c_j x^j$ be a degree- m polynomial so that $p \geq 0$ on $[-1, 1]$ and $\int_{-1}^1 p \leq \beta$. Then $\max_i |c_i| \leq \beta \cdot (m+1)^2 \cdot (\sqrt{2} + 1)^m$.*

Consider the shifted polynomial $q_J(u) = p_J(u \cdot (b_J - a_J)/2 + (b_J + a_J)/2)$ where $J = [a_J, b_J]$. By applying Fact 22 to q_J and noting that $\int_{-1}^1 q_J = (2/|J|) \cdot \int_J p_J$, we conclude that the coefficients of q_J are bounded by

$$\frac{4}{|J|} \cdot (m+1)^2 \cdot (\sqrt{2} + 1)^m$$

and thus

$$|q_J(u)| \leq \frac{4}{|J|} \cdot m(m+1)^2 \cdot (\sqrt{2} + 1)^m$$

for all $u \in [-1, 1]$, and so therefore the same bound applies for $p_J(x)$ for all $x \in J$.

But notice that since we assume that $J \in \mathcal{L}_2$, it follows that for all $x \in J \cap L_\ell$, we have that

$$\mathcal{N}_\ell(x) \geq 8 \frac{k}{\epsilon} p_J(x),$$

and so in particular $w_\ell \mathcal{N}_\ell(x) \geq 8p_J(x)$ for all $x \in J \cap L_\ell$. Hence we have

$$\begin{aligned} \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| dx &= \int_{J \cap L_\ell} \mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_J(x) dx \\ &\geq \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) - p_J(x)| dx + \int_{J \cap L_\ell} \frac{7}{8} w_\ell \mathcal{N}_\ell(x) - p_J(x) dx \\ &\geq \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) - p_J(x)| dx + \frac{3w_\ell}{4} \int_{J \cap L_\ell} \mathcal{N}_\ell(x) dx. \end{aligned}$$

where the second line follows since $\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_J(x) \geq |\mathcal{M}_\theta^{-\ell}(x) - p_J(x)| + \frac{7}{8} w_\ell \mathcal{N}_\ell(x) - p_J(x)$ for all $x \in J \cap L_\ell$.

Thus

$$\begin{aligned} \sum_{J \in \mathcal{L}_2} \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) + w_\ell \mathcal{N}_\ell(x) - p_{\text{dens}}(x)| dx &\geq \\ &\sum_{J \in \mathcal{L}_2} \left(\int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) - p_{\text{dens}}(x)| dx + \frac{3w_\ell}{4} \int_{J \cap L_\ell} \mathcal{N}_\ell(x) dx \right). \end{aligned} \quad (5)$$

Moreover, by Equation (4), we know that

$$\begin{aligned} \sum_{J \in \mathcal{L}_2} \int_{J \cap L_\ell} \mathcal{N}_\ell(x) dx &= \int_{L_\ell} \mathcal{N}_\ell(x) dx - \sum_{J \in \mathcal{L}_1} \int_{J \cap L_\ell} \mathcal{N}_\ell(x) dx \\ &\geq W - \frac{1}{8}, \end{aligned}$$

since \mathcal{L}_1 contains at most s intervals. Thus, the RHS of Equation (5) must be lower bounded by

$$\sum_{J \in \mathcal{L}_2} \int_{J \cap L_\ell} |\mathcal{M}_\theta^{-\ell}(x) - p_{\text{dens}}(x)| dx + \frac{3}{4} \left(W - \frac{1}{8} \right) w_\ell.$$

Putting it all together. Hence, we have

$$\begin{aligned} \int_{L_\ell} |M_\theta(x) - p_{\text{dens}}(x)| dx &= \sum_{J \in \mathcal{L}_1} \int_{J \cap L_\ell} |M_\theta(x) - p_{\text{dens}}(x)| dx + \sum_{J \in \mathcal{L}_2} \int_{J \cap L_\ell} |M_\theta(x) - p_{\text{dens}}(x)| dx \\ &\geq \sum_{J \in \mathcal{L}_1} \int_{J \cap L_\ell} |M_\theta^{-\ell}(x) - p_{\text{dens}}(x)| dx + \sum_{J \in \mathcal{L}_2} \int_{J \cap L_\ell} |M_\theta^{-\ell}(x) - p_{\text{dens}}(x)| dx \\ &\quad + \left[\frac{3}{4} \left(W - \frac{1}{8} \right) - \frac{1}{8} \right] w_\ell \\ &\geq \int_{L_\ell} |M_\theta^{-\ell}(x) - p_{\text{dens}}(x)| dx + \left[\frac{3}{4} \left(W - \frac{1}{8} \right) - \frac{1}{8} \right] w_\ell. \end{aligned}$$

We therefore have

$$\begin{aligned}
\|M_\theta - p_{\text{dens}}\|_1 &= \int_{L_\ell} |M_\theta(x) - p_{\text{dens}}(x)| \, dx + \int_{L_\ell^c} |M_\theta(x) - p_{\text{dens}}(x)| \, dx \\
&\geq \int_{L_\ell} |M_\theta(x) - p_{\text{dens}}(x)| \, dx + \int_{L_\ell^c} |M_\theta^{-\ell}(x) - p_{\text{dens}}(x)| \, dx - \int_{L_\ell^c} w_\ell \mathcal{N}_\ell(x) \, dx \\
&\geq \int_{L_\ell} |M_\theta(x) - p_{\text{dens}}(x)| \, dx + \int_{L_\ell^c} |M_\theta^{-\ell}(x) - p_{\text{dens}}(x)| \, dx - (1 - W)w_\ell \\
&\geq \int_{L_\ell} |M_\theta^{-\ell}(x) - p_{\text{dens}}(x)| \, dx + \left(\frac{7}{4}W - \frac{39}{32}\right)w_\ell + \int_{L_\ell^c} |M_\theta^{-\ell}(x) - p_{\text{dens}}(x)| \, dx \\
&= \|M_\theta^{-\ell} - p_{\text{dens}}\|_1 + \frac{1}{2}w_\ell,
\end{aligned}$$

when we set $W = 55/56$. □

4.2 A parametrization scheme for a single Gaussian

Intuitively, Lemma 21 says that for any $\theta \in \Theta_k$, there are some components which have bounded variance and which can be close to p_{dens} (the components in S_1), and the remaining components, which may have unbounded variance but which will be far away from p_{dens} . Since we are searching for a k -GMM which is close to p_{dens} , in some sense we should not have to concern ourselves with the latter components since they cannot meaningfully interact with p_{dens} . Thus we only need find a suitably robust parametrization for admissible Gaussians.

Such a parametrization can be obtained by linearly transforming the domain so that the associated interval gets mapped to $[-1, 1]$. Formally, fix a Gaussian $\mathcal{N}_{\mu, \tau}$ and an interval J . Then it can be written as

$$\mathcal{N}_{\mu, \tau}(x) = \frac{\tilde{\tau}}{|J|/2} \mathcal{N}\left(\tilde{\tau} \cdot \frac{x - \text{mid}(J)}{|J|/2} - \tilde{\mu}\right), \tag{6}$$

for some unique $\tilde{\mu}$ and $\tilde{\tau}$, where for any interval I , we define $\text{mid}(I)$ to denote its midpoint. Call these the *rescaled mean with respect to J* and *rescaled precision with respect to J* of \mathcal{N} , respectively. Concretely, given μ , τ , and an interval J , the rescaled variance and mean with respect to J are defined to be

$$\tilde{\tau} = \frac{|J|}{2} \tau, \quad \tilde{\mu} = \frac{\tilde{\tau}}{|J|/2} (\mu - \text{mid}(J)).$$

For any $\tilde{\mu}, \tilde{\tau}$, we let $\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J}(x)$ denote the function given by the RHS of Equation (6). The following two lemmas says that these rescaled parameters have the desired robustness properties.

Lemma 23. *Let $\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J}$ be an admissible Gaussian with rescaled mean $\tilde{\mu}$ and rescaled precision $\tilde{\tau}$ with respect to its associated interval $J \in \mathcal{I}$. Then $\tilde{\mu} \in [-\frac{2s\phi}{\omega}, \frac{2s\phi}{\omega}]$ and $\sqrt{2\pi} \cdot \omega/(16s) \leq \tilde{\tau} \leq \phi/2$.*

Proof. We first show that $\sqrt{2\pi} \cdot \omega/(16s) \leq \tilde{\tau} \leq \phi/2$. That the rescaled variance is bounded from above follows from a simple change of variables and the definition of admissibility. By the definition of admissibility, we also know that

$$\begin{aligned}
\int_J \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} \, dx &\geq \int_{J \cap L(\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J})} \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} \, dx \\
&\geq \omega\tau \cdot |J \cap L(\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J})| \\
&\geq \frac{\omega}{8s}.
\end{aligned}$$

Furthermore, we trivially have

$$|J| \cdot \frac{\tau}{\sqrt{2\pi}} \geq \int_J \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} dx.$$

Thus, the precision τ must be at least $\sqrt{2\pi}\omega/(8s|J|)$, and so its rescaled precision must be at least $\sqrt{2\pi}\omega/(16s)$, as claimed.

We now show that $\tilde{\mu} \in [-\frac{2s\phi}{\omega}, \frac{2s\phi}{\omega}]$. Because $\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J}$ is an admissible Gaussian with associated interval J , we know that $|J \cap L(\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J})| \geq 1/(8s\tau)$. Moreover, we know that on $J \cap L(\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J})$, we have $\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J}(x) \geq \omega\tau$. Thus in particular

$$\int_J \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} dx \geq \int_{J \cap L(\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J})} \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} dx \geq \frac{\omega}{8s}.$$

Define \tilde{J} to be the interval which is of length $8s|J|/\omega$ around $\text{mid}(J)$. We claim that $\mu \in \tilde{J}$, where μ is the mean of $\mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J}$.

Assume that $\text{mid}(J) \leq \mu$. Let $J_0 = J$ and inductively, for $i < 4s/\omega$, let J_i be the interval with left endpoint at the right endpoint of J_{i-1} and with length $|J|$. That is, the J_i consist of $4s/\omega$ consecutive, non-intersecting copies of J starting at J and going upwards on the number line (for simplicity of exposition we assume that $4s/\omega$ is an integer). Let $J^\dagger = \cup_{i=0}^{(4s/\omega)-1} J_i$. We claim that $\mu \in J^\dagger$. Suppose not. This means that μ is strictly greater than any point in any J_i . In particular, this implies that for all i ,

$$\begin{aligned} \int_{J_i} \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} dx &\geq \int_{J_{i-1}} \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} dx \\ &\geq \int_{J_0} \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} dx \\ &\geq \frac{\omega}{8s}. \end{aligned}$$

But then this would imply that

$$\int_{J^\dagger} \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} dx = \sum_{i=0}^{(4s/\omega)-1} \int_{J_i} \mathcal{N}_{\tilde{\mu}, \tilde{\tau}}^{r, J} dx \geq \frac{1}{2}.$$

Notice that J^\dagger is itself an interval. But any interval containing at least $1/2$ of the weight of any Gaussian must contain its mean, which we assumed did not happen. Thus we conclude that $\mu \in J^\dagger$. Moreover, $J^\dagger \subseteq \tilde{J}$, so $\mu \in \tilde{J}$, as claimed. If $\text{mid}(J) \geq \mu$ then apply the symmetric argument with J_i which are decreasing on the number line instead of increasing.

We have thus shown that $\mu \in \tilde{J}$. It is a straightforward calculation to show that this implies that $\tilde{\mu} \in [-\frac{4s\tau}{\omega}, \frac{4s\tau}{\omega}]$. By the above, we know that $\tau \leq \phi/2$ and thus $\tilde{\mu} \in [-\frac{2s\phi}{\omega}, \frac{2s\phi}{\omega}]$, as claimed. \square

Lemma 24. *For any interval J , and $\tilde{\mu}_1, \tilde{\tau}_1, \tilde{\mu}_2, \tilde{\tau}_2$ so that $|\tilde{\tau}_i| \leq 2\phi$ for $i \in \{1, 2\}$ and $|\tilde{\mu}_1 - \tilde{\mu}_2| + |\tilde{\tau}_1 - \tilde{\tau}_2| \leq O((\epsilon/(\phi k))^2)$, we have*

$$\|\mathcal{N}_{\tilde{\mu}_1, \tilde{\tau}_1}^{r, J}(x) - \mathcal{N}_{\tilde{\mu}_2, \tilde{\tau}_2}^{r, J}(x)\|_1 \leq \epsilon.$$

Proof. This follows by a change of variables and Lemma 16. \square

Moreover, this rescaled parametrization naturally lends itself to approximation by a piecewise polynomial, namely, replace the standard normal Gaussian density function in Equation (6) with \tilde{P}_ϵ . This is the piecewise polynomial that we will use to represent each individual component in the Gaussian mixture.

4.3 A parametrization scheme for k -GMMs

In the rest of this section, our parametrization will often be of the form described above. To distinguish this from the previous notation, for any $\theta \in \Theta_k$, and any set of k intervals J_1, \dots, J_k , we will let $\theta^r \in \Theta_k$ denote the rescaled parameters so that if the i -th component in the mixture represented by θ has parameters w_i, μ_i, τ_i , then the i -th component in the mixture represented by θ^r has parameters $w_i, \tilde{\mu}_i, \tilde{\tau}_i$ so that $\mathcal{N}_{\mu_i, \tau_i} = \mathcal{N}_{\tilde{\mu}_i, \tilde{\tau}_i}^{r, J_i}$. Notice that the transformation between the original and the rescaled parameters is a linear transformation, and thus trivial to compute and to invert.

The final difficulty is that we do not know how many mixture components have associated interval J for $J \in \mathcal{I}$. To deal with this, our algorithm simply iterates over all possible allocations of the mixture components to intervals and returns the best one. There are $O(k)$ possible associated intervals J and k different components, so there are at most $k^{O(k)}$ different possible allocations. In this section, we will see how our parametrization works when we fix an allocation of the mixture components.

More formally, let \mathcal{A} be the set of functions $v : [s] \rightarrow \mathbb{N}$ so that $\sum_{\ell=1}^s v(\ell) = k$. These will represent the number of components ‘‘allocated’’ to exist on the scale of each J_ℓ . For any $v \in \mathcal{A}$, define \mathcal{I}_v to be the set of $I_\ell \in \mathcal{I}$ so that $v(\ell) \neq 0$.

Fix $\theta^r \in \Theta_k$ and $v \in \mathcal{A}$. Decompose θ^r into $(\theta_1^r, \dots, \theta_s^r)$, where θ_ℓ^r contains the rescaled parameters with respect to J_ℓ for the $v(\ell)$ components allocated to interval J_ℓ (note that $v(\ell)$ may be 0 in which case θ_ℓ is the empty set, i.e., corresponds to the parameters for no components). For any $1 \leq \ell \leq s$, let

$$\mathcal{M}_{\ell, \theta_\ell^r}^r(x) = \sum_i w_i \frac{\tilde{\tau}_i}{|I_\ell|/2} \mathcal{N} \left(\tilde{\tau}_i \cdot \frac{x - \text{mid}(I_\ell)}{|I_\ell|/2} - \tilde{\mu}_i \right),$$

where i ranges over the components that θ_j corresponds to, and define $\mathcal{M}_{\theta^r, v}^r(x) = \sum_{\ell=1}^s \mathcal{M}_{\ell, \theta_\ell^r}^r(x)$. Similarly, define

$$P_{\epsilon, \ell, \theta_\ell^r}^r(x) = \sum_i w_i \frac{\tilde{\tau}_i}{|I_\ell|/2} \tilde{P}_\epsilon \left(\tilde{\tau}_i \cdot \frac{x - \text{mid}(I_\ell)}{|J_\ell|/2} - \tilde{\mu}_i \right),$$

and define $P_{\epsilon, \theta^r, v}^r(x) = \sum_{\ell=1}^s P_{\epsilon, \ell, \theta_\ell^r}^r(x)$. Finally, for any v , define $\mathcal{P}_{\epsilon, v}^r$ to be the set of all such $P_{\epsilon, \theta, v}^r$.

We have:

Lemma 25. *For any $\theta^r \in \Theta_k$, we have*

$$\|\mathcal{M}_{\theta^r, v}^r - P_{\epsilon, \theta^r, v}^r\|_1 \leq \epsilon.$$

This follows from roughly the same argument as in the proof of Lemma 6, and so we omit the proof.

We now finally have all the necessary language and tools to prove the following theorem:

Corollary 26. *Fix $2 \geq \epsilon > 0$. There is some allocation $v \in \mathcal{A}$ and a set of parameters $\theta^r \in \Theta_k$ so that $\tilde{\mu}_i \in [-\frac{2s\phi}{\omega}, \frac{2s\phi}{\omega}]$, $1/(8s) \leq \tilde{\tau}_i \leq \phi/2$, and $w_\ell \geq \epsilon/(2k)$ for all i . Moreover,*

$$\|f - \mathcal{M}_{\theta^r, v}^r\|_1 \leq 19 \cdot \text{OPT}_k + O(\epsilon).$$

Proof. Let $\theta^* \in \Theta_k$ be so that $\|f - \mathcal{M}_{\theta^*}\|_1 = \text{OPT}_k$, and let \mathcal{N}_ℓ^* denote its ℓ -th component with parameters w_ℓ^* , μ_ℓ^* , and τ_ℓ^* . Decompose $[k]$ into $S_{\text{good}}(\theta^*)$, $S_{\text{bad}}(\theta^*)$ as in Lemma 21.

By the guarantees of the density estimation algorithm, we know that

$$\left\| \sum_{\ell} w_\ell^* \mathcal{N}_{\mu_\ell^*, \tau_\ell^*} - p_{\text{dens}} \right\|_1 \leq 5\text{OPT}_k + \epsilon.$$

By Lemma 21, this implies that

$$5\text{OPT}_k + \epsilon \geq \left\| \sum_{\ell \in S_{\text{good}}(\theta^*)} w_\ell^* \mathcal{N}_{\mu_\ell^*, \tau_\ell^*} - p_{\text{dens}} \right\|_1 + \frac{1}{2} \sum_{\ell \in S_{\text{bad}}(\theta^*)} w_\ell - 2\epsilon,$$

from which we may conclude the following two inequalities:

$$\left\| \sum_{\ell \in S_{\text{good}}(\theta^*)} w_\ell^* \mathcal{N}_{\mu_\ell^*, \tau_\ell^*} - p_{\text{dens}} \right\|_1 \leq 5 \cdot \text{OPT}_k + 3\epsilon, \quad (7)$$

$$\sum_{\ell \in S_{\text{bad}}(\theta^*)} w_\ell^* \leq 10 \cdot \text{OPT}_k + 6\epsilon. \quad (8)$$

Let θ' be defined so that for all $\ell \in S_{\text{good}}(\theta^*)$, the means and variances of the ℓ -th component in θ' are μ_ℓ^* and τ_ℓ^* , and so that for all $\ell \in S_{\text{bad}}(\theta^*)$, the means and variances of the ℓ -th component in θ' are arbitrary but so that the underlying Gaussian is admissible. Let the weights of the components in θ' be the same as the weights in θ^* .

Then we have

$$\begin{aligned} \|M_{\theta'} - f\|_1 &= \left\| \sum_{\ell \in S_{\text{good}}(\theta^*)} w_\ell^* \mathcal{N}_{\mu_\ell^*, \tau_\ell^*} + \sum_{\ell \in S_{\text{bad}}(\theta^*)} w_\ell^* \mathcal{N}_{\mu'_\ell, \tau'_\ell} - f \right\|_1 \\ &\leq \left\| \sum_{\ell \in S_{\text{good}}(\theta^*)} w_\ell^* \mathcal{N}_{\mu_\ell^*, \tau_\ell^*} - f \right\|_1 + \left\| \sum_{\ell \in S_{\text{bad}}(\theta^*)} w_\ell^* \mathcal{N}_{\mu'_\ell, \tau'_\ell} \right\|_1 \\ &= \left\| \sum_{\ell \in S_{\text{good}}(\theta^*)} w_\ell^* \mathcal{N}_{\mu_\ell^*, \tau_\ell^*} - f \right\|_1 + \sum_{\ell \in S_{\text{bad}}(\theta^*)} w_\ell^* \\ &\leq \left\| \sum_{\ell \in S_{\text{good}}(\theta^*)} w_\ell^* \mathcal{N}_{\mu_\ell^*, \tau_\ell^*} - p_{\text{dens}} \right\|_1 + \|f - p_{\text{dens}}\|_1 + \sum_{\ell \in S_{\text{bad}}(\theta^*)} w_\ell^* \\ &\leq 19 \cdot \text{OPT}_k + O(\epsilon) \end{aligned}$$

where the last line follows from Equation (7), the guarantee of the density estimation algorithm, and Equation (8).

For each $\ell \in [k]$, let $J_\ell \in \mathcal{I}$ denote the interval so that the ℓ -th component of θ' is admissible with respect to J_ℓ . Let θ^r be the rescaling of θ' with respect to J_1, \dots, J_k . Then by Lemma 23, θ^r satisfies that $\tilde{\mu}_i \in [-\frac{2s\phi}{\omega}, \frac{2s\phi}{\omega}]$ and $\sqrt{2\pi} \cdot \omega/(16s) \leq \tilde{\tau}_i \leq \phi/2$ for all i . Let $v \in \mathcal{A}$ be chosen so that $v(i)$ is the number of times that I_i appears in the sequence J_1, \dots, J_k . Then \mathcal{M}_{θ^r} and v satisfies all conditions in the lemma, except possibly that the weights may be too small.

Thus, let θ be the set of parameters whose means and precisions are exactly those of θ^r , but for which the weight of the ℓ -th component is defined to be $w_\ell = \max(\epsilon/(2k), w_\ell^*)$ for all $1 \leq \ell \leq k-1$ and $w_k = 1 - \sum_{\ell=1}^{k-1} w_\ell$. It is easy to see that $\theta \in \Theta_k$; moreover, $\|\mathcal{M}_\theta - \mathcal{M}_{\theta^r}\|_1 \leq \epsilon$. Then it is easy to see that θ and v together satisfy all the conditions of the lemma. \square

4.4 The full algorithm

At this point, we are finally ready to describe our algorithm `LEARN_GMM` which agnostically and properly learns an arbitrary mixture of k Gaussians. Informally, our algorithm proceeds as follows. First, using `ESTIMATE-DENSITY`, we learn a p'_{dens} that with high probability is ϵ -close to the underlying distribution f in L_1 -distance. Then, as before, we may rescale the entire problem so that the density estimate is supported on $[-1, 1]$. Call the rescaled density estimate p_{dens} .

As before, it suffices to find a k -GMM that is close to p_{dens} in \mathcal{A}_K -distance, for $K = 4k - 1$. The following is a direct analog of Lemma 12. We omit its proof because its proof is almost identical to that of Lemma 12.

Lemma 27. *Let $\epsilon > 0, v \in \mathcal{A}, k \geq 2, \theta^r \in \Theta_k$, and $K = 4(k - 1) + 1$. Then we have*

$$0 \leq \|p_{\text{dens}} - P_{\epsilon, \theta^r, v}^r\|_1 - \|p_{\text{dens}} - P_{\epsilon, \theta^r, v}^r\|_{\mathcal{A}_K} \leq 8 \cdot \text{OPT}_k + O(\epsilon).$$

Our algorithm enumerates over all $v \in \mathcal{A}$ and for each v finds a θ^r approximately minimizing

$$\|p_{\text{dens}} - P_{\epsilon, \theta^r, v}^r\|_{\mathcal{A}_K}.$$

Using the same binary search technique as before, we can transform this problem into $\log 1/\epsilon$ feasibility problems of the form

$$\|p_{\text{dens}} - P_{\epsilon, \theta^r, v}^r\|_{\mathcal{A}_K} < \eta. \quad (9)$$

Fix $v \in \mathcal{A}$, and recall $\mathcal{P}_{\epsilon, v}^r$ is the set of all polynomials of the form $P_{\epsilon, \theta^r, v}^r$. Let Θ_k^{valid} denote the set of $\theta^r \in \Theta_k$ so that $\tilde{\mu}_i \in [-\frac{2s\phi}{\omega}, \frac{2s\phi}{\omega}]$, $\sqrt{2\pi}\omega/(8s) \leq \tilde{\tau} \leq \phi/2$, and $w_i \geq \epsilon/(2k)$, for all i . For any $\theta^r \in \Theta_k^{\text{valid}}$, canonically identify it with $P_{\epsilon, \theta^r, v}^r$. By almost exactly the same arguments used in Section 3.5, it follows that the class $\mathcal{P}_{\epsilon, v}^r$, where $\theta \in \Theta_k^{\text{valid}}$, satisfies the conditions in Section 3.4, and that the system of polynomial equations $S_{K, p_{\text{dens}}, \mathcal{P}_{\epsilon, v}^r}(\nu)$ has two levels of quantification (each with $O(k)$ bound variables), has $k^{O(k)}$ polynomial constraints, and has maximum degree $O(\log(1/\epsilon))$. Thus, we have

Corollary 28. *For any fixed ϵ, ν , and for $K = 4k - 1$, we have that $S_{K, p_{\text{dens}}, \mathcal{P}_{\epsilon, \nu}^r, \Theta_k^{\text{valid}}}(\nu)$ encodes Equation (9) ranging over $\theta \in \Theta_k^{\text{valid}}$. Moreover, for all $\gamma, \lambda \geq 0$, $\text{SOLVE-POLY-PROGRAM}(S_{K, p_{\text{dens}}, \mathcal{P}_{\epsilon, \nu}^r, \Theta_k^{\text{valid}}}(\nu), \lambda, \gamma)$ runs in time*

$$(k \log(1/\epsilon))^{O(k^4)} \log \log(3 + \frac{\gamma}{\lambda}).$$

For each v , our algorithm then performs a binary search over η to find the smallest (up to constant factors) η so that Equation (9) is satisfiable for this v , and records both η_v , the smallest η for which Equation (9) is satisfiable for this v , and the output θ_v of the system of polynomial inequalities for this choice of η . We then return θ_{v^*} so that the η_{v^*} is minimal over all $v \in \mathcal{A}$. The pseudocode for `LEARN_GMM` is in Algorithm 2.

The following theorem is our main technical contribution:

Theorem 29. `LEARN_GMM`(k, ϵ, δ) takes $\tilde{O}((k + \log 1/\delta)/\epsilon^2)$ samples from the unknown distribution with density f , runs in time

$$\left(k \log \frac{1}{\epsilon}\right)^{O(k^4)} + \tilde{O}\left(\frac{k}{\epsilon^2}\right),$$

and with probability $1 - \delta$ returns a set of parameters $\theta \in \Theta_k$ so that $\|f - \mathcal{M}_\theta\|_1 \leq 58 \cdot \text{OPT} + \epsilon$.

Proof. The sample complexity follows simply because `ESTIMATE-DENSITY` draws $\tilde{O}((k + \log 1/\delta)/\epsilon^2)$ samples, and these are the only samples we ever use. The running time bound follows because $|\mathcal{A}| = k^{O(k)}$ and from Corollary 28. Thus it suffices to prove correctness.

Let θ be the parameters returned by the algorithm. It was found in some iteration for some $v \in \mathcal{A}$. Let v^*, θ^* be those which are guaranteed by Corollary 26. We have

$$\|p_{\text{dens}} - P_{\epsilon, \theta^*, v^*}^r\|_{\mathcal{A}_K} \leq \|p_{\text{dens}} - f\|_1 + \|f - \mathcal{M}_{\theta^*, v^*}^r\|_1 + \|\mathcal{M}_{\theta^*, v^*}^r - P_{\epsilon, \theta^*, v^*}^r\|_1 \leq 23 \cdot \text{OPT}_k + O(\epsilon).$$

By the above inequalities, the system of polynomial equations is feasible for $\eta \leq 46 \cdot \text{OPT}_k + O(\epsilon)$ in the iteration corresponding to v^* (Corollary 26 guarantees that the parameters θ^* are sufficiently bounded). Hence, for some $\eta_{v^*} \leq \eta$, the algorithm finds some θ' so that there is some θ'' so that $\|\theta' - \theta''\|_2 \leq C_1(\epsilon/(\phi k))^2$, which satisfies $S_{p_{\text{dens}}, \mathcal{P}_{\epsilon, v^*}^r, \Theta_k^{\text{valid}}}(\nu_{v^*})$.

Let θ_1 be the set of parameters computed by the algorithm before rounding the weights back to the simplex (i.e. at Line 11). By our choice of precision in solving the polynomial program, (i.e. by our choice of λ on Line 24 of Algorithm 2), we know that the precisions of the returned mixture are non-negative (so each component is a valid Gaussian). It was found in an iteration corresponding to some $v \in \mathcal{A}$, and there is some $\eta_v \leq \eta_{v^*} \leq 46 \cdot \text{OPT}_k + O(\epsilon)$ and some θ'_1 satisfying the system of polynomial equalities for v and η_v , so

Algorithm 2 Algorithm for proper learning an arbitrary mixture of k Gaussians.

```

1: function LEARNGMM( $k, \epsilon, \delta$ )
2:    $\triangleright$  Density estimation. Only this step draws samples.
3:    $p'_{\text{dens}} \leftarrow \text{ESTIMATE-DENSITY}(k, \epsilon, \delta)$ 

4:    $\triangleright$  Rescaling
5:    $\triangleright p_{\text{dens}}$  is a rescaled and shifted version of  $p'_{\text{dens}}$  such that the support of  $p_{\text{dens}}$  is  $[-1, 1]$ .
6:   Let  $p_{\text{dens}}(x) \stackrel{\text{def}}{=} p'_{\text{dens}} \left( \frac{2(x-\alpha)}{\beta-\alpha} - 1 \right)$ 

7:    $\triangleright$  Fitting shape-restricted polynomials
8:   for  $v \in \mathcal{A}$  do
9:      $\eta_v, \theta_v^r \leftarrow \text{FINDFITGIVENALLOCATION}(p_{\text{dens}}, v)$ 
10:    Let  $\theta$  so that  $\theta^r = \theta_v^r$ , so that  $\eta_v$  is minimal over all  $\eta_v$  (breaking ties arbitrarily).
11:     $\triangleright$  Round weights back to be on the simplex
12:    for  $i = 1, \dots, k-1$  do
13:       $w_i \leftarrow w_i - \epsilon/2k$  (This guarantees that  $\sum_{i=1}^{k-1} w_i \leq 1$ ; see analysis for details)
14:      If  $w_i > 1$ , set  $w_i = 1$ 
15:     $w_k \leftarrow 1 - \sum_{i=1}^{k-1} w_i$ 
16:     $\triangleright$  Undo the scaling
17:     $w'_i \leftarrow w_i$ 
18:     $\mu'_i \leftarrow \frac{(\mu_i+1)(\beta-\alpha)}{2} + \alpha$ 
19:     $\tau'_i \leftarrow \frac{\tau_i}{\beta-\alpha}$ 
20:    return  $\theta'$ 

21: function FINDFITGIVENALLOCATION( $p_{\text{dens}}, v$ )
22:    $\nu \leftarrow \epsilon$ 
23:   Let  $C_1$  be a universal constant sufficiently small.
24:   Let  $\lambda \leftarrow \min(C_1(\epsilon/(\phi k))^2, 1/16s, \epsilon/(4k))$ 
25:    $\triangleright$  This choice of precision provides robustness as needed by Lemma 24, and also ensures that all the
26:   weights and precisions returned must be non-negative.
27:   Let  $\psi \leftarrow 6ks\phi/\omega + 3k\phi/2 + 1$ 
28:    $\theta^r \leftarrow \text{SOLVE-POLY-SYSTEM}(S_{p_{\text{dens}}, \mathcal{P}_{\epsilon, v}^r, \Theta_k^{\text{valid}}}(\nu), \lambda, \psi)$ 
29:   while  $\theta^r$  is “NO-SOLUTION” do
30:      $\nu \leftarrow 2 \cdot \nu$ 
31:      $\theta^r \leftarrow \text{SOLVE-POLY-SYSTEM}(S_{p_{\text{dens}}, \mathcal{P}_{\epsilon, v}^r, \Theta_k^{\text{valid}}}(\nu), \lambda, \psi)$ 
32:   return  $\theta^r, \nu$ 

```

that $\|\theta_1 - \theta'_1\|_2 \leq C_1(\epsilon/(\phi k))^2$. Let θ be the set of rescaled parameters obtained after rounding the weights of θ_1 back to the simplex. It is straightforward to check that $\theta \in \Theta_k$, and moreover, $\|\mathcal{M}_{\theta, v}^r - \mathcal{M}_{\theta'_1, v}^r\|_1 \leq 2\epsilon$, and so $\|P_{\epsilon, \theta, v}^r - P_{\epsilon, \theta'_1, v}^r\|_1 \leq O(\epsilon)$.

We therefore have

$$\begin{aligned}
\|f - M_\theta\|_1 &\leq \|f - p_{\text{dens}}\|_1 + \|p_{\text{dens}} - P_{\epsilon, \theta, v}^r\|_1 + \|P_{\epsilon, \theta, v}^r - \mathcal{M}_{\epsilon, \theta, v}^r\|_1 \\
&\stackrel{(a)}{\leq} 4 \cdot \text{OPT} + \epsilon + 8 \cdot \text{OPT} + O(\epsilon) + \|p_{\text{dens}} - P_{\epsilon, \theta, v}^r\|_{\mathcal{A}_K} + \epsilon \\
&\stackrel{(b)}{\leq} 12 \cdot \text{OPT} + O(\epsilon) + \|p_{\text{dens}} - P_{\epsilon, \theta'_1, v}^r\|_{\mathcal{A}_K} \\
&\stackrel{(c)}{\leq} 58 \cdot \text{OPT} + O(\epsilon),
\end{aligned}$$

where (a) follows from Lemmas 27 and 25, (b) follows from the arguments above, and (c) follows since θ'_1 satisfies the system of polynomial inequalities for $\eta_v \leq 46 \cdot \text{OPT}_k + O(\epsilon)$.

As a final step, we choose an internal ϵ' in our algorithm so that the $O(\epsilon')$ in the above guarantee becomes bounded by ϵ . This proves the desired approximation guarantee and completes the proof. \square

5 Further classes of distributions

In this section, we briefly show how to use our algorithm to properly learn other parametric classes of univariate distributions.

Let \mathcal{C} be a class of parametric distributions on the real line, parametrized by $\theta \in S$ for $S \subseteq \mathbb{R}^u$. For each θ , let $F_\theta \in \mathcal{C}$ denote the pdf of the distribution parametrized by θ in \mathcal{C} . To apply our algorithm in this setting, it suffices to show the following:

1. (*Simplicity of \mathcal{C}*) For any θ_1 and θ_2 , the function $F_{\theta_1} - F_{\theta_2}$ has at most K zero crossings. In fact it also suffices if any two such functions have “essentially” K zero crossings.
2. (*Simplicity of S*) S is a semi-algebraic set.
3. (*Representation as a piecewise polynomial*) For each $\theta \in S$ and any $\epsilon > 0$, there is a piecewise polynomial $P_{\epsilon, \theta}$ so that $\|P_{\epsilon, \theta} - F_\theta\|_1 \leq \epsilon$. Moreover, the map $(x, \theta) \mapsto P_{\epsilon, \theta}(x)$ is jointly polynomial in x and θ at any point so that x is not at a breakpoint of $P_{\epsilon, \theta}$. Finally, the breakpoints of $P_{\epsilon, \theta}$ also depend polynomially on θ .
4. (*Robustness of the Parametrization*) There is some robust parametrization so that we may assume that all “plausible candidate” parameters are $\leq 2^{\text{poly}(1/\epsilon)}$, and moreover, if $\|\theta_1 - \theta_2\| \leq 2^{-\text{poly}(1/\epsilon)}$, then $\|F_{\theta_1} - F_{\theta_2}\| \leq \epsilon$.

Assuming \mathcal{C} satisfies these conditions, our techniques immediately apply. In this paper, we do not attempt to catalog classes of distributions which satisfy these properties. However, we believe such classes are often natural and interesting. We give evidence for this below, where we show that our framework produces proper and agnostic learning algorithms for mixtures of two more types of simple distributions. The resulting algorithms are both sample optimal (up to log factors) and have nearly-linear running time.

5.1 Learning mixtures of simple distribution

As a brief demonstration of the generality of our technique, we show that our techniques give proper and agnostic learning algorithms for mixtures of k exponential distributions and Laplace distributions (in addition to mixtures of k Gaussians) which are *nearly-sample optimal*, and run in time which is *nearly-linear* in the number of samples drawn, for any constant k .

We now sketch a proof of correctness for both classes mentioned above. In general, the robustness condition is arguably the most difficult to verify of the four conditions required. However, it can be verified that for mixtures of simple distributions with reasonable smoothness conditions the appropriate modification of the parametrization we developed in Section 4 will suffice. Thus, for the classes of distributions mentioned, it suffices to demonstrate that they satisfy conditions (1) to (3).

Condition 1: It follows from the work of [Tos06] that the difference of k exponential distributions or k Laplace distributions has at most $2k$ zero crossings.

Condition 2: This holds trivially for the class of mixtures of exponential distributions. We need a bit of care to demonstrate this condition for Laplace distributions since a Laplace distribution with parameters μ, b has the form

$$\frac{1}{2b} e^{-|x-\mu|/b}$$

and thus the Taylor series is not a polynomial in x or the parameters. However, we may sidestep this issue by simply introducing a variable y in the polynomial program which is defined to be $y = |x - \mu|$.

Condition 3: It can easily be shown that a truncated degree $O(\log 1/\epsilon)$ Taylor expansion (as of the form we use for learning k -GMMs) suffices to approximate a single exponential or Laplace distribution, and hence a $O(k)$ -piecewise degree $O(\log 1/\epsilon)$ polynomial suffices to approximate a mixture of k exponential or Laplace distributions up to L_1 -distance ϵ .

Thus for both of these classes, the sample complexity of our algorithm is $\tilde{O}(k/\epsilon^2)$, and its running time is

$$\left(k \log \frac{1}{\epsilon}\right)^{O(k^4)} + \tilde{O}\left(\frac{k}{\epsilon^2}\right),$$

similar to the algorithm for learning k -GMMs. As for k -GMMs, this sample complexity is nearly optimal, and the running time is nearly-linear in the number of samples drawn, if k is constant.

6 Acknowledgements

We thank Jayadev Acharya, Ilias Diakonikolas, Piotr Indyk, Gautam Kamath, Ankur Moitra, Rocco Servedio, and Ananda Theertha Suresh for helpful discussions.

References

- [ABG⁺14] Joseph Anderson, Mikhail Belkin, Navin Goyal, Luis Rademacher, and James R. Voss. The more, the merrier: the blessing of dimensionality for learning large Gaussian mixtures. In *COLT*, 2014.
- [ADLS15] Jayadev Acharya, Ilias Diakonikolas, Jerry Li, and Ludwig Schmidt. Sample-optimal density estimation in nearly-linear time. In submission to FOCS, 2015.
- [AJOS14] Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Near-optimal-sample estimators for spherical Gaussian mixtures. In *NIPS*, 2014.
- [AM05] Dimitris Achlioptas and Frank McSherry. On spectral learning of mixtures of distributions. In *Learning Theory*, volume 3559 of *Lecture Notes in Computer Science*, pages 458–469. 2005.
- [BCM^V14] Aditya Bhaskara, Moses Charikar, Ankur Moitra, and Aravindan Vijayaraghavan. Smoothed analysis of tensor decompositions. In *STOC*, 2014.
- [BDKVDL15] Peter Buhlmann, Petros Drineas, Michael John Kane, and Mark Van Der Laan. *Handbook of Big Data*. Chapman and Hall/CRC Handbooks of Modern Statistical Methods. Chapman and Hall/CRC, 2015.
- [BS10] Mikhail Belkin and Kaushik Sinha. Polynomial learning of distribution families. In *FOCS*, 2010.
- [BSZ15] Aditya Bhaskara, Ananda Theertha Suresh, and Morteza Zadimoghaddam. Sparse solutions to nonnegative linear systems and applications. In *AISTATS*, 2015.
- [BV08] S. Charles Brubaker and Santosh Vempala. Isotropic PCA and affine-invariant clustering. In *FOCS*, 2008.
- [BWY14] Sivaraman Balakrishnan, Martin J. Wainwright, and Bin Yu. Statistical guarantees for the EM algorithm: From population to sample-based analysis. arXiv:1408.2156, 2014.

- [CDSS13] Siu-On Chan, Ilias Diakonikolas, Rocco Servedio, and Xiaorui Sun. Learning mixtures of structured distributions over discrete domains. In *SODA*, 2013.
- [CDSS14] Siu-On Chan, Ilias Diakonikolas, Rocco Servedio, and Xiaorui Sun. Efficient density estimation via piecewise polynomial approximation. In *STOC*, 2014.
- [Das99] Sanjoy Dasgupta. Learning mixtures of Gaussians. In *FOCS*, 1999.
- [DK14] Constantinos Daskalakis and Gautam Kamath. Faster and sample near-optimal algorithms for proper learning mixtures of gaussians. In *COLT*, 2014.
- [DL01] Luc Devroye and Gabor Lugosi. *Combinatorial methods in density estimation*. Springer Series in Statistics, Springer, 2001.
- [FSO06] Jon Feldman, Rocco A. Servedio, and Ryan O’Donnell. PAC learning axis-aligned mixtures of Gaussians with no separation assumption. In *COLT*, 2006.
- [GHK15] Rong Ge, Qingqing Huang, and Sham M. Kakade. Learning mixtures of gaussians in high dimensions. In *STOC*, 2015. arXiv:1503.00424.
- [HK13] Daniel Hsu and Sham M. Kakade. Learning mixtures of spherical Gaussians: Moment methods and spectral decompositions. In *ITCS*, 2013.
- [HP15] Moritz Hardt and Eric Price. Tight bounds for learning a mixture of two Gaussians. In *STOC*, 2015. arXiv:1404.4997.
- [KMV10] Adam Tauman Kalai, Ankur Moitra, and Gregory Valiant. Efficiently learning mixtures of two Gaussians. In *STOC*, 2010.
- [KSS94] Michael J. Kearns, Robert E. Schapire, and Linda M. Sellie. Toward efficient agnostic learning. *Machine Learning*, 17(2-3):115–141, 1994.
- [KSV08] Ravindran Kannan, Hadi Salmasian, and Santosh Vempala. The spectral method for general mixture models. *SIAM Journal on Computing*, 38(3):1141–1156, 2008.
- [Moi14] Ankur Moitra. Algorithmic aspects of machine learning. <http://people.csail.mit.edu/moitra/docs/bookex.pdf>, 2014.
- [MV10] Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of Gaussians. In *FOCS*, 2010.
- [Pea94] Karl Pearson. Contributions to the mathematical theory of evolution. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 185:71–110, 1894.
- [Ren92a] James Renegar. On the computational complexity and geometry of the first-order theory of the reals. Part i: Introduction. Preliminaries. The geometry of semi-algebraic sets. the decision problem for the existential theory of the reals. *Journal of Symbolic Computation*, 13(3):255 – 299, 1992.
- [Ren92b] James Renegar. On the computational complexity of approximating solutions for real algebraic formulae. *SIAM Journal on Computing*, 21(6):1008–1025, 1992.
- [SK01] Arora Sanjeev and Ravi Kannan. Learning mixtures of arbitrary Gaussians. In *STOC*, 2001.
- [Tim63] Aleksandr F. Timan. *Theory of Approximation of Functions of a Real Variable*. Pergamon, New York, 1963.

- [Tos06] Timo Tossavainen. On the zeros of finite sums of exponential functions. *Australian Mathematical Society Gazette*, 33(1):47 – 50, 2006.
- [Vem12] Santosh S. Vempala. Modeling high-dimensional data: Technical perspective. *Commun. ACM*, 55(2):112–112, 2012.
- [VW04] Santosh Vempala and Grant Wang. A spectral algorithm for learning mixture models. *J. Comput. Syst. Sci.*, 68(4):841–860, 2004.