

---

# Robust Model Selection and Nearly-Proper Learning for GMMs

---

**Allen Liu**  
 MIT  
 Cambridge, MA 02139  
 cliu568@mit.edu

**Jerry Li**  
 Microsoft Research  
 Redmond, WA 98052  
 jerrl@microsoft.com

**Ankur Moitra**  
 MIT  
 Cambridge, MA 02139  
 moitra@mit.edu

## Abstract

In learning theory, a standard assumption is that the data is generated from a finite mixture model. But what happens when the number of components is not known in advance? The problem of estimating the number of components, also called *model selection*, is important in its own right but there are essentially no known efficient algorithms with provable guarantees let alone ones that can tolerate adversarial corruptions. In this work, we study the problem of robust model selection for univariate Gaussian mixture models (GMMs). Given  $\text{poly}(k/\epsilon)$  samples from a distribution that is  $\epsilon$ -close in TV distance to a GMM with  $k$  components, we can construct a GMM with  $\tilde{O}(k)$  components that approximates the distribution to within  $\tilde{O}(\epsilon)$  in  $\text{poly}(k/\epsilon)$  time. Thus we are able to approximately determine the minimum number of components needed to fit the distribution within a logarithmic factor. Prior to our work, the only known algorithms for learning arbitrary univariate GMMs either output significantly more than  $k$  components (e.g.  $k/\epsilon^2$  components for kernel density estimates) or run in time exponential in  $k$ . Moreover, by adapting our techniques we obtain similar results for reconstructing Fourier-sparse signals.

## 1 Introduction

Many works in learning theory operate under the assumption that the data is generated from a finite mixture model, and furthermore that the number of components is known in advance. But what happens when the number of components is not known in advance? The problem of estimating the number of components is called *model selection* and has been intensively studied in statistics for over fifty years [Neyman and Scott, 1966]. Indeed, in many scientific applications, it is the central issue. Consider the motivation given by Chen et al. [2004]: In genetics, we might have a continuous-valued trait, like height, that can be measured across a population and we want to understand its genetic basis. But is the underlying genetic mechanism simple or complex? Is it controlled by just a few genes or are there many more genes waiting to be discovered that each have a small effect on it?

From a statistical perspective, what makes model selection challenging is that the standard analysis of the likelihood ratio test breaks down because of lack of regularity and non-identifiability [Hartigan, 1985]. Despite many attempts [Ghosh and Sen, 1984, Lo et al., 2001, Huang et al., 2017] and rejoinders [Jeffries, 2003], even understanding the asymptotic distribution of the likelihood ratio statistics has remained a long-standing challenge in the field [Kasahara and Shimotsu, 2015]. From an algorithmic standpoint, the problem is even more difficult.

In this work, we study the problem of robust model selection for one-dimensional Gaussian mixture models with  $k$  components ( $k$ -GMMs for short). A natural approach for this problem is via *agnostic proper learning*, where the task is to, given samples from an unknown distribution, output the best  $k$ -GMM approximation to this distribution in TV distance. An efficient agnostic proper

learning algorithm, combined with standard tools from hypothesis testing, would immediately yield an algorithm for model selection.

Unfortunately, while there are many efficient algorithms for learning one-dimensional GMMs, they all fall into one of several categories: (1) They assume some strong separation conditions on the components so that the samples can be clustered based on which component they were generated from. (2) They solve the harder problem of learning the parameters of the components, which information-theoretically requires the number of samples to be exponential in  $k$  [Moitra and Valiant, 2010]. (3) They employ brute-force search [Daskalakis and Kamath, 2014, Acharya et al., 2014] or solve a system of polynomial inequalities [Li and Schmidt, 2017], and run in time exponential in  $k$ . (4) They learn an approximation that is either not a GMM, e.g. a piece-wise polynomial approximation [Chan et al., 2013, Acharya et al., 2017] or output a GMM where the number of components is much larger than  $k$  [Wu and Xie, 2018, Devroye and Lugosi, 2012, Bhaskara et al., 2015]. (5) They assume that the components in the GMM have the same or similar variances and means not too far apart so that there is a good approximation to the density with just a logarithmic number of components [Wu and Yang, 2018, Polyanskiy and Wu, 2020]. In all cases, these guarantees are insufficient for efficient model selection, and/or yield a trivial approximation to the number of components in a GMM except in restricted settings. In this work, we ask: Are there efficient algorithms for learning arbitrary one-dimensional GMMs that output an approximation with  $\tilde{O}(k)$  components? Relatedly: Are there efficient algorithms for approximating the number of components in a GMM? We give efficient algorithms whose running time and sample complexity are polynomial in  $k$  for both of these problems, and also the related problem of reconstructing Fourier-sparse signals with an unknown number of frequencies.

## 1.1 Learning and model selection for GMMs

Our main result is a new robust learning algorithm for one-dimensional GMMs. We show:

**Theorem 1.1.** *Let  $k, \epsilon > 0$  be parameters and let  $f$  be a distribution such that  $d_{TV}(\mathcal{M}, f) \leq \epsilon$  for some unknown mixture of Gaussians  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$ . Assume that we are given  $\tilde{O}(k/\epsilon^2)$  samples from  $f$ . Then there is an algorithm that runs in  $\text{poly}(k/\epsilon)$  time and with probability 0.9 (over the random samples), outputs a mixture of  $\tilde{O}(k)$  Gaussians,  $\tilde{\mathcal{M}}$ , such that*

$$d_{TV}(\tilde{\mathcal{M}}, f) \leq \tilde{O}(\epsilon).$$

In contrast to other known learning algorithms (discussed earlier), our learning algorithm works for arbitrary GMMs, runs in polynomial time and uses a polynomial number of samples, and while it does not output a GMM with exactly  $k$  components, it does the next best thing: it outputs a GMM with at most a polylogarithmic factor more components.

As a corollary, we also give an algorithm for robust approximate model selection for GMMs. The connection to model selection is that when our algorithm fails to find a GMM with  $\tilde{O}(k)$  components that fits the data we can be assured that there must more than  $k$  components to begin with. Notice in particular that improper approximations by themselves do not suffice for the model selection problem, as a good improper approximation could exist even if the distribution is far from any GMM with  $\tilde{O}(k)$  components.

**Theorem 1.2.** *Let  $k, \epsilon > 0$  be parameters we are given. Let  $\mathcal{F}_1$  be the family of distributions that are  $\epsilon$ -close to a  $k$ -GMM with  $k$  components (in TV distance). Let  $\mathcal{F}_2$  be the family of distributions that are not  $\tilde{O}(\epsilon)$ -close to any GMM with  $\tilde{O}(k)$  components. There is an algorithm that given  $\text{poly}(k/\epsilon)$  samples from a known distribution  $\mathcal{D}$ , runs in  $\text{poly}(k/\epsilon)$  time, and outputs 1 if  $\mathcal{D} \in \mathcal{F}_1$  and outputs 2 if  $\mathcal{D} \in \mathcal{F}_2$  both with failure probability at most 0.2.*

**Remark.** *Even if the distribution  $\mathcal{D}$  is completely unknown and we are only given samples from it, the above result still holds as long as  $\mathcal{D}$  is somewhat well behaved (note that such an assumption is necessary as hypothesis testing with respect to total variation distance without any assumptions on  $\mathcal{D}$  is impossible). In particular we can use piecewise polynomial approximation [Chan et al., 2013] or kernel density estimates [Terrell and Scott, 1992] to learn a distribution  $\mathcal{D}'$  that is close to  $\mathcal{D}$  that we have an explicit form for and then run the hypothesis test using  $\mathcal{D}'$ .*

## 1.2 Fourier sparse interpolation

Our techniques also immediately apply to the problem of Fourier sparse interpolation, where the goal is to interpolate a signal based on noisy measurements of it at a few points [Chen et al., 2016]. We say that a function  $\mathcal{M}$  is  $(k, C)$  simple if it can be written in the form

$$\mathcal{M}(t) = \sum_{j=1}^k a_j e^{2\pi i \theta_j t},$$

where additionally  $\sum_j |a_j| \leq C$ . In other words, a function is  $(k, C)$  simple if it is  $k$ -sparse in the Fourier domain, and its Fourier coefficients are bounded in  $\ell_1$  by  $C$ .

We consider the following problem. We get query access to a function  $f(t) = \mathcal{M}(t) + \eta(t)$  at any point in the interval  $[-1, 1]$ , where  $\mathcal{M}$  is  $(k, C)$  simple and has all frequencies in the interval  $[-F, F]$ , and  $\eta(t)$  is noise that we will assume is bounded in  $L_2$  norm. The goal is to compute a Fourier-sparse approximation  $\widetilde{\mathcal{M}}(t)$  that is close to  $f(t)$ , in the sense that its error is comparable to that of  $\mathcal{M}(t)$ . Recently Chen et al. [2016] showed how to construct an approximation  $\widetilde{\mathcal{M}}(t)$  that satisfies

$$\|f(t) - \widetilde{\mathcal{M}}(t)\|_2 \leq \|\eta(t)\|_2 + \epsilon \|\mathcal{M}(t)\|_2$$

where the  $L_2$  norm is taken over the interval  $[-1, 1]$ . Their algorithm works for any  $\epsilon > 0$  and uses  $\text{poly}(k, \log 1/\epsilon) \log F$  measurements. Moreover the  $\widetilde{\mathcal{M}}(t)$  that they output is  $\text{poly}(k, \log 1/\epsilon)$ -Fourier sparse. Similarly to the GMM setting, a natural goal is to perform robust interpolation but with tighter bounds on the number of frequencies. We show:

**Theorem 1.3.** *Let  $f, \mathcal{M}$  be as above where  $\mathcal{M}$  is  $(k, 1)$ -simple. Then for any desired accuracy  $\epsilon > 0$  and constant  $c > 0$ , in  $\text{poly}(k, \log 1/\epsilon) \log F$  queries and  $\text{poly}(k/c, \log 1/\epsilon) \log^2 F$  time, we can output a function  $\widetilde{\mathcal{M}}$  such that with probability  $1 - 2^{-\Omega(k)}$ ,*

1.  $\widetilde{\mathcal{M}}$  is  $\widetilde{O}(k)$ -Fourier sparse with  $\|\widetilde{\mathcal{M}}\|_1 \leq \widetilde{O}(k)$
2.  $\int_{-1+c}^{1-c} |\widetilde{\mathcal{M}} - f|^2 \leq \widetilde{O}(\epsilon^2 + \int_{-1}^1 |f - \mathcal{M}|^2)$

**Remark.** Note the constraints  $\|\widehat{\mathcal{M}}\|_1$  and  $\|\widetilde{\mathcal{M}}\|_1$  translate into bounds on the sizes of the coefficients of the exponentials in  $\mathcal{M}$  and  $\widetilde{\mathcal{M}}$  respectively.

The natural open question left by our work is to improve the sparsity bounds, both for interpolation/learning and model selection. In principle it could be possible that there are efficient algorithms for these problems, however it now seems somewhat unlikely. Even without noise, learning a Gaussian mixture model with  $k$  components without a separation condition in time  $\text{poly}(k, 1/\epsilon)$  is open. From our work (see Section 2), we see that even in the well-conditioned case this is equivalent to finding a non-trivially sparse solution to a system of polynomial equations where there seems to be no structure that makes algorithmic search better than brute-force possible. Moreover, this question has already been open for many years, but there hasn't been any progress on proper learning. Thus, we conjecture that both the learning and model selection problems are computationally hard if we are not allowed to relax the number of components.

## 1.3 Related work

There is a vast literature on the three problems we consider. Here we will give a more detailed review of related work.

**Learning Mixtures of Gaussians and Model Selection** Since the pioneering work of Pearson [1894], mixtures of Gaussians have become one of the most ubiquitous and well-studied generative models in both theory and practice. Numerous problems have been studied on the context of learning mixtures of Gaussians, including clustering [Dasgupta, 1999, Vempala and Wang, 2004, Achlioptas and McSherry, 2005, Dasgupta and Schulman, 2007, Arora and Kale, 2007, Kumar and Kannan, 2010, Awasthi and Sheffet, 2012, Mixon et al., 2017, Hopkins and Li, 2018, Kothari et al., 2018, Diakonikolas et al., 2018], learning in the presence of adversarial noise in

high dimensional settings [Diakonikolas et al., 2018, Hopkins and Li, 2018, Kothari et al., 2018, Bakshi et al., 2020, Diakonikolas et al., 2020, Kane, 2021, Liu and Moitra, 2020, 2021], parameter estimation [Kalai et al., 2010, Belkin and Sinha, 2015, Moitra and Valiant, 2010, Hardt and Price, 2015], learning in smoothed settings [Hsu and Kakade, 2013, Anderson et al., 2014, Bhaskara et al., 2014, Ge et al., 2015], and density estimation [Devroye and Lugosi, 2012, Chan et al., 2014, Acharya et al., 2017].

Of particular interest to us is the line of work on proper learning [Feldman et al., 2006, Acharya et al., 2014, Li and Schmidt, 2017, Ashtiani et al., 2018], where the goal is to output a mixture of  $k$ -Gaussians which is close in total variation to the underlying ground truth. Unfortunately, while the sample complexity of these algorithms is usually polynomial, the runtime for all known approaches is exponential in  $k$ . In contrast, our runtimes are polynomial, albeit for a relaxed version of the problem, where the output is allowed to be a mixture of  $k'$  Gaussians, for  $k' > k$ .

For this “semi-proper” regime, efficient algorithms are known, albeit either only for restricted settings, or with significantly worse quantitative results than we achieve. In the “well-conditioned” case, where the means are close together, and the variances of all the components are comparable, the aforementioned work of [Wu and Yang, 2018, Polyanskiy and Wu, 2020] demonstrates that the nonparametric MLE can efficiently obtain an estimate using only logarithmically many pieces. However, the nonparametric MLE is not suited for the general setting, where the means could be far apart, and variances could be very different, and will not converge in general. Moreover, while nonparametric MLE is robust to perturbations in KL, it is not robust to perturbations in total variation distance, as we consider here.

For the general case, by using kernel density estimates, one can achieve  $\epsilon$  approximation using  $k' = O(k/\epsilon^C)$  for some constant  $C$  [Devroye and Lugosi, 2012]. Similarly Bhaskara et al. [2015] achieves  $\epsilon$  error using  $k' = O(k/\epsilon^3)$  pieces. That is, for both of these approaches, they require a number of pieces which scales polynomially with  $1/\epsilon$ . In comparison, our dependence on  $\epsilon$  in terms of the number of pieces is logarithmic.

As discussed previously, there are strong connections between proper learning and model selection [Neyman and Scott, 1966, Hartigan, 1985, Ghosh and Sen, 1984, Lo et al., 2001, Jeffries, 2003, Kasahara and Shimotsu, 2015, Huang et al., 2017]. Related notions have been considered in distribution testing [Parnas et al., 2006, Valiant and Valiant, 2010a,b, 2011, Jiao et al., 2016, 2017, Han et al., 2016] and testing properties of boolean functions [Diakonikolas et al., 2007, Iyer et al., 2021].

**Continuous Time Sparse Fourier Transforms** Sparse Fourier transforms in the continuous setting, also known as sparse Fourier transforms off the grid, has been the subject of intensive study. Indeed, the first algorithm for this problem dates back to Prony [1795]. Modern algorithms include MUSIC [Schmidt, 1982], ESPRIT [Roy et al., 1986], maximum likelihood estimators [Bresler and Macovski, 1986], convex programming based methods [Candès and Fernandez-Granda, 2014] and the matrix pencil method [Moitra, 2015].

Most of these works, especially those that work in a noisy setting, require a frequency gap. Moreover they require more than  $k$  samples (their bound usually depends on the frequency gap), even if the underlying signal is  $k$ -sparse in the Fourier domain. A recent line of work has focused on the problem of improving the sample complexity – in particular getting bounds which only depend on  $k$  with runtimes that are polynomial in  $k$  [Fannjiang and Liao, 2012, Duarte and Baraniuk, 2013, Tang et al., 2013, 2014, Boufounos et al., 2015, Huang and Kakade, 2015, Price and Song, 2015]. The setting where there is no gap and there is noise is particularly challenging. One approach is to relax the definition of a frequency gap, and require it only between “clusters” of frequencies [Batenkov et al., 2020]. Another line of work [Avron et al., 2019, Chen and Price, 2019] shows how to output a hypothesis which is  $k$ -sparse without any gap assumptions and with sample complexity which is polynomial in  $k$ . However these methods run in exponential time. As we previously discussed, the most relevant works to us are Chen et al. [2016] and Chen and Price [2019], which give an algorithm whose running time and sample complexity are polynomial in  $k$  that works without any gap assumptions, but for a relaxation where we are allowed to output a  $\tilde{O}(k^2)$ -Fourier sparse signal.

## 2 Technical overview

We now give an overview of our approach. We will focus on just the GMM case in this overview. Our approach for sparse Fourier interpolation follows a very similar outline. We first present our techniques assuming that we have explicit access to  $f$ . In Section 2.3 we show how to reduce to this case when we are only given samples. In other words, the problem is as follows: we are given a function  $f$ , and we want to find a sparse approximation to  $f$  as a nonnegative sum of Gaussians, i.e. we want to write

$$f \sim a_1 G_1 + \cdots + a_n G_n$$

with  $n$  small, where each  $G_i$  is a Gaussian.

### 2.1 Well-conditioned case

We first solve the “well-conditioned” case. Roughly, we say that a GMM is well-conditioned if the variances of the components are all constant scale and the means are all not too far from zero. Formally, we have the following definition:

**Definition 2.1.** *We say a Gaussian  $G = N(\mu, \sigma^2)$  is  $\delta$ -well-conditioned if  $|\mu| \leq \delta$  and  $|\sigma^2 - 1| \leq \delta$ . Furthermore we say a mixture of Gaussians  $\mathcal{M} = w_1 G_1 + \cdots + w_k G_k$  is  $\delta$ -well-conditioned if all of the components  $G_1, \dots, G_k$  are  $\delta$ -well-conditioned.*

Naturally, our techniques also apply to a shared scaling and/or translation of the components, but we will ignore this for now. Earlier work of [Wu and Yang, 2018, Polyanskiy and Wu, 2020] proved an important structural result that a well-conditioned GMM can be  $\epsilon$ -approximated by a mixture with  $\tilde{O}(1/\epsilon)$  components. However we will want a robust and algorithmic version: In particular, instead of requiring the distribution to be exactly a well-conditioned GMM, we will only require that it be close in total variation distance. Even in this setting, with some level of model misspecification, we want an efficient algorithm for constructing an approximating GMM with few components. To this end, a key result, proved in Section A, is:

**Lemma 2.2.** *Let  $\epsilon > 0$  be a parameter. Assume we are given access to a distribution  $f$  such that  $d_{TV}(f, \mathcal{M}) \leq \epsilon$  where  $\mathcal{M} = w_1 G_1 + \cdots + w_k G_k$  is a 0.5-well-conditioned mixture of Gaussians. Then we can compute, in  $\text{poly}(1/\epsilon)$  time, a mixture  $\tilde{\mathcal{M}}$  of at most  $O(\log 1/\epsilon)$  Gaussians such that  $d_{TV}(\mathcal{M}, \tilde{\mathcal{M}}) \leq \tilde{O}(\epsilon)$ .*

Our approach departs from the moment matching framework of Wu and Yang [2018], Polyanskiy and Wu [2020]. Instead we take the probability density function of any well-conditioned Gaussian  $G_j$ . We can expand it as a Taylor series around 0 of the form

$$G_j(x) = c_{G_j}^{(0)} + \frac{c_{G_j}^{(1)}x}{1!} + \frac{c_{G_j}^{(2)}x^2}{2!} + \dots$$

for some coefficients  $c_{G_j}^{(i)}$ . We can then associate it with the vector  $c_{G_j} = (c_{G_j}^{(0)}, \dots, c_{G_j}^{(\ell-1)})$  of length  $\ell = O(\log 1/\epsilon)$ . Then, for any well-conditioned mixture  $\mathcal{M} = a_1 G_1 + \cdots + a_k G_k$ , we can associate it with the corresponding convex combination of the vectors of its components, i.e., we define  $c_{\mathcal{M}} = a_1 c_{G_1} + \dots + a_k c_{G_k} \in \mathbb{R}^\ell$ .

The point of this is the following implication: if two well-conditioned mixtures get mapped to vectors which are close, then these two mixtures must be close in total variation distance. The intuition is that when we write down the  $L_1$  distance between the two mixtures, because the Taylor coefficients of Gaussians decay exponentially fast, the contribution of terms with degree  $l > O(\log 1/\epsilon)$  to the integral becomes negligible.

Now, we can associate the set of well-conditioned mixtures with a convex body in  $O(\log 1/\epsilon)$ -dimensions, where the vertices of the convex body are given by single Gaussians. Consequently we can use Caratheodory’s theorem to argue that any point within this body can be approximated as an  $O(\log 1/\epsilon)$ -sparse convex combination of the vertices, or equivalently, any well-conditioned mixture can be approximated by a mixture of  $O(\log 1/\epsilon)$  well-conditioned Gaussians.

It remains to demonstrate how to actually find this sparse mixture of Gaussians. Naively, the number of vertices is infinite, as there are infinitely many well-conditioned Gaussians. However, it is not too

hard to show that if we consider a slight coarsening of this body by only taking the vertices to be the vectors associated to the well-conditioned Gaussians which belong in some  $\text{poly}(1/\epsilon)$ -sized net, then the quality of our solution only degrades by constant multiplicative factors. At this point we can appeal to standard results in convex optimization to find the desired sparse approximation. We defer the details of this argument to Section A.

## 2.2 Localization

After solving the well-conditioned case, the next step is to reduce the general case to the well-conditioned case via localization. We begin with an important definition.

**Definition 2.3** (Gaussian Multiplier). *For parameters  $\mu, \sigma$ , we define*

$$M_{\mu, \sigma^2}(x) = e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

i.e. it is a Gaussian scaled so that its maximum value is 1.

Gaussian multipliers will be crucial in the localization step. Now assume that  $f$  can be written as some unknown  $k$ -sparse combination, say

$$f = a_1 G_1 + \cdots + a_k G_k$$

We can then modify  $f$ , e.g. by multiplying by a Gaussian multiplier  $M_{\mu, \sigma^2}$ . Heuristically, this operation changes the coefficients  $a_1, \dots, a_k$  in a predictable way. Namely, the coefficients  $a_j$  of Gaussians  $G_j$  that are far from  $N(\mu, \sigma^2)$  are exponentially attenuated based on the distance to  $N(\mu, \sigma^2)$ . This effectively "localizes" the mixture. More formally,

**Claim 2.4.** *We have the identity*

$$M_{\mu, \sigma^2}(x) N(\mu_1, \sigma_1^2) = \frac{1}{\sqrt{1 + \frac{\sigma_1^2}{\sigma^2}}} e^{-\frac{(\mu_1 - \mu)^2}{2(\sigma_1^2 + \sigma^2)}} N\left(\frac{\mu\sigma_1^2 + \mu_1\sigma^2}{\sigma_1^2 + \sigma^2}, \frac{\sigma_1^2\sigma^2}{\sigma_1^2 + \sigma^2}\right).$$

*Proof.* We prove the above through direct computation.

$$\begin{aligned} M_{\mu, \sigma^2}(x) G_1(x) &= e^{-\frac{(x-\mu)^2}{2\sigma^2} - \frac{(x-\mu_1)^2}{2\sigma_1^2}} \frac{1}{\sigma_1\sqrt{2\pi}} = \frac{1}{\sigma_1\sqrt{2\pi}} \cdot e^{-\frac{1}{2}\left(\left(\frac{1}{\sigma^2} + \frac{1}{\sigma_1^2}\right)x^2 - 2\left(\frac{\mu}{\sigma^2} + \frac{\mu_1}{\sigma_1^2}\right)x + \frac{\mu^2}{\sigma^2} + \frac{\mu_1^2}{\sigma_1^2}\right)} \\ &= \frac{1}{\sigma_1\sqrt{2\pi}} \cdot \exp\left(-\frac{1}{2}\left(\sqrt{\frac{1}{\sigma^2} + \frac{1}{\sigma_1^2}}x - \frac{\frac{\mu}{\sigma^2} + \frac{\mu_1}{\sigma_1^2}}{\sqrt{\frac{1}{\sigma^2} + \frac{1}{\sigma_1^2}}}\right)^2 - \frac{1}{2} \cdot \frac{(\mu_1 - \mu)^2}{\sigma_1^2 + \sigma^2}\right) \\ &= \frac{1}{\sqrt{1 + \frac{\sigma_1^2}{\sigma^2}}} e^{-\frac{(\mu_1 - \mu)^2}{2(\sigma_1^2 + \sigma^2)}} N\left(\frac{\mu\sigma_1^2 + \mu_1\sigma^2}{\sigma_1^2 + \sigma^2}, \frac{\sigma_1^2\sigma^2}{\sigma_1^2 + \sigma^2}\right). \end{aligned}$$

■

The hope is that this will leave us with only components that are not too far from each other – exactly the well-conditioned case which we already know how to solve. If the variances of all of the components are comparable, then this is indeed the case. However, additional complications arise when one of the components  $G_i = N(\mu, \sigma_i^2)$  has variance  $\sigma_i \ll \sigma$  because this component will still have much smaller variance than the others after localizing. Nevertheless, we show that we can carefully localize at different scales, using smaller variance Gaussian multipliers to localize around smaller variance components so that all of the localized mixtures are well-conditioned.

The main remaining question is to select a good family of localizations so that we can then fully reconstruct the original mixture from the localized mixtures. Each localized mixture will cost us  $O(\log 1/\epsilon)$  components, and therefore we must use at most  $\tilde{O}(k)$  different localizations. When all of the variances of the Gaussians are not too dissimilar, we can do so by leveraging the following structural result, which states that one can  $\epsilon$ -approximate the constant function using a sum of evenly spaced Gaussians with variance 1 and spacing  $(\log 1/\epsilon)^{-1/2}$  (or smaller). The intuition behind this observation is that the Fourier transform of a Gaussian is also a Gaussian, which has exponential tail decay.

**Lemma 2.5.** Let  $0 < \epsilon < 0.1$  be a parameter. Let  $c$  be a real number such that  $0 < c \leq (\log 1/\epsilon)^{-1/2}$ . Define

$$f(x) = \sum_{j=-\infty}^{\infty} \frac{c}{\sqrt{2\pi}} M_{cj\sigma,\sigma^2}(x).$$

Then  $1 - \epsilon \leq f(x) \leq 1 + \epsilon$  for all  $x$ .

*Proof.* WLOG  $\sigma = 1$ . Now the function  $f$  is  $c$ -periodic and even, so we may consider its Fourier expansion

$$f(x) = a_0 + 2a_1 \cos\left(\frac{2\pi x}{c}\right) + 2a_2 \cos\left(\frac{4\pi x}{c}\right) + \dots$$

and we will now compute the Fourier coefficients. First note that

$$a_0 = \frac{1}{c} \int_0^c f(x) dx = \frac{1}{\sqrt{2\pi}} \sum_{j=-\infty}^{\infty} \int_{c(j+1)}^{cj} M_{0,1}(x) dx = 1.$$

Next, for any  $j \geq 1$ ,

$$\begin{aligned} a_j &= \frac{1}{c} \int_0^c f(x) \cos\left(\frac{2\pi j x}{c}\right) dx = \frac{1}{\sqrt{2\pi}} \sum_{j=-\infty}^{\infty} \int_{c(j+1)}^{cj} M_{0,1}(x) \cos\left(\frac{2\pi j x}{c}\right) dx \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{1}{2} \left( e^{\frac{-x^2}{2} + \frac{2\pi i j x}{c}} + e^{\frac{-x^2}{2} - \frac{2\pi i j x}{c}} \right) dx = e^{-\frac{2\pi^2 j^2}{c^2}} \end{aligned}$$

where in the above we use the notation  $i = \sqrt{-1}$ . Using the assumption that  $c \leq (\log 1/\epsilon)^{-1/2}$ , it is clear that

$$\sum_{j=1}^{\infty} e^{-\frac{2\pi^2 j^2}{c^2}} \leq \frac{\epsilon}{2}$$

so we deduce that for any  $x$ ,

$$|f(x) - 1| \leq 2(|a_1| + |a_2| + \dots) = 2 \sum_{j=1}^{\infty} e^{-\frac{2\pi^2 j^2}{c^2}} \leq \epsilon.$$

In other words, the function  $f$  is between  $1 - \epsilon$  and  $1 + \epsilon$  everywhere and we are done.  $\blacksquare$

In light of the above lemma, we can use a set of evenly spaced Gaussian multipliers and simply sum the different localized mixtures. Note that it suffices to use  $\tilde{O}(k)$  different localizations because we only need to sum over the Gaussian multipliers that have some nontrivial overlap with one of the  $k$  true components (since for Gaussian multipliers that are far from all of the components, the localized mixture will be approximately 0).

To handle the fully general case, when the variances of the Gaussians are unbounded, we need a generalization of the previous lemma that allows us to  $\epsilon$ -approximate the indicator function of an interval with a sum of  $O(\log^2 1/\epsilon)$  Gaussians. The proof of this generalization is in Section B.

**Definition 2.6** (Significant Interval). For a Gaussian multiplier  $M_{\mu,\sigma^2}$ , we say the  $C$ -significant interval of  $M$  is  $[\mu - C\sigma, \mu + C\sigma]$ . We will use the same terminology for a Gaussian  $N(\mu, \sigma^2)$ .

**Theorem 2.7.** Let  $l$  be a positive real number and  $0 < \epsilon < 0.1$  be a parameter. There is a function  $f$  with the following properties

1.  $f$  can be written a linear combination of Gaussian multipliers

$$f(x) = w_1 M_{\mu_1, \sigma_1^2}(x) + \dots + w_n M_{\mu_n, \sigma_n^2}(x)$$

where  $n = O(\log^2 1/\epsilon)$  and  $0 \leq w_1, \dots, w_n \leq 1$

2. The  $10\sqrt{\log 1/\epsilon}$ -significant intervals of all of the  $M_{\mu_i, \sigma_i^2}$  are contained in the interval  $[-(1 + \epsilon)l, (1 + \epsilon)l]$

3.  $0 \leq f(x) \leq 1 + \epsilon$  for all  $x$
4.  $1 - \epsilon \leq f(x) \leq 1 + \epsilon$  for all  $x$  in the interval  $[-l, l]$
5.  $0 \leq f(x) \leq \epsilon$  for  $x \geq (1 + \epsilon)l$  and  $x \leq -(1 + \epsilon)l$

We combine this structural result with a dynamic program which allows us to efficiently choose the scales at which to localize. Putting all of these pieces together yields our full algorithm, assuming we have access to the pdf of the unknown function. We show how to eliminate the need for pdf access below and present our full algorithm in complete detail in Section C.

### 2.3 Abstracting away the samples

In the previous sections, we have assumed that we have access to the underlying pdf function  $f$ . Typically, however, we only have sample access to the unknown distribution. To rectify this, we will use the improper learner in [Chan et al., 2013] (see Theorem 37) whose output is a piecewise polynomial. We can then only work with this piecewise polynomial, which is an explicit function that we can then perform explicit computations with.

**Definition 2.8.** *A function  $f$  is  $t$ -piecewise degree  $d$  if there is a partition of the real line into intervals  $I_1, \dots, I_t$  and polynomials  $q_1(x), \dots, q_t(x)$  of degree at most  $d$  such that for all  $i \in [t]$ ,  $f(x) = q_i(x)$  on the interval  $I_i$ .*

The work in [Chan et al., 2013] guarantees to learn a piecewise polynomial  $f'$  that is close to  $\mathcal{M}$  in  $L^1$  distance when given  $\tilde{O}(k/\epsilon^2)$  samples (and they also show that this sample complexity is essentially optimal).

**Theorem 2.9** (Chan et al. [2013]). *Let  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  be an unknown mixture of Gaussians and  $f$  a distribution such that  $d_{TV}(f, \mathcal{M}) \leq \epsilon$ . There is an algorithm that, given  $\tilde{O}(k/\epsilon^2)$  samples from  $f$ , runs in  $\text{poly}(k/\epsilon)$  time and returns an  $O(k)$ -piecewise degree  $O(\log 1/\epsilon)$  function  $f'$  such that with 0.9 probability (over the random samples),*

$$\|f' - f\|_1 \leq O(\epsilon).$$

For technical reasons, we will need a few simple post-processing steps after using Theorem 2.9. We can ensure that the output hypothesis  $f'$  is always nonnegative by splitting each polynomial into positive and negative parts and zeroing out the negative parts (since this will not increase the  $L^1$  error). Finally, we can re-normalize so that the output  $f'$  is actually a distribution. This renormalization at most doubles the  $L^1$  error. Thus we have:

**Corollary 2.10.** *Let  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  be an unknown mixture of Gaussians and  $f$  a distribution such that  $d_{TV}(f, \mathcal{M}) \leq \epsilon$ . There is an algorithm that, given  $\tilde{O}(k/\epsilon^2)$  samples from  $\mathcal{D}$ , runs in  $\text{poly}(k/\epsilon)$  time and returns an  $O(k \log 1/\epsilon)$ -piecewise degree  $O(\log 1/\epsilon)$  function  $f'$  such that  $f'$  is a distribution and with 0.9 probability (over the random samples),*

$$d_{TV}(f, f') \leq O(\epsilon).$$

### 2.4 Hypothesis testing for model selection

We now show how our result for model order selection, Theorem 1.2, follows immediately from combining Theorem 1.1 with a standard procedure for testing the TV-distance between two distributions from samples (see Yatracos [1985]).

**Claim 2.11.** *Let  $\mathcal{D}_1, \mathcal{D}_2$  be two distributions for which we have explicitly computable density functions. Let  $\epsilon, \tau > 0$  be parameters. Assume that we are given  $O(1/\epsilon^2 \cdot \log 1/\tau)$  samples from  $\mathcal{D}_1$  and can efficiently sample from  $\mathcal{D}_2$ . Then in  $\text{poly}(1/\epsilon \log 1/\tau)$  time, we can compute  $d$  such that with probability  $1 - \tau$ ,*

$$|d - d_{TV}(\mathcal{D}_1, \mathcal{D}_2)| \leq \epsilon.$$

*Proof of Theorem 1.2.* We can run the algorithm in Theorem 1.1 with parameters  $k, \epsilon$  to obtain an output distribution  $\tilde{\mathcal{M}}$  that is a mixture of  $\tilde{O}(k)$  Gaussians. We can then use Claim 2.11 with parameters  $\epsilon, 0.01$  to measure the TV-distance between  $\tilde{\mathcal{M}}$  and  $\mathcal{D}$  (note that we have explicit access

to the pdf of  $\mathcal{D}$ ) and output 1 or 2 depending on if our estimate of the TV distance is less than  $\tilde{O}(\epsilon)$ . Combining the guarantees of Theorem 1.1 and Claim 2.11 ensures that our output satisfies the desired properties.  $\blacksquare$

## 2.5 Sparse Fourier

We now briefly describe how our techniques can be used for sparse Fourier reconstruction. Recall that the problem is to, given query access to a function  $f$  on  $[-1, 1]$  which is approximately  $k$ -Fourier sparse, approximate it with an  $\tilde{O}(k)$ -Fourier sparse function. As before, we first abstract away the query access, by leveraging the following result from Chen et al. [2016]:

**Theorem 2.12** (Theorem 1.1 in Chen et al. [2016]). *Let  $f$  be a function defined on  $[-1, 1]$  and assume we are given query access to  $f$ . Let  $\mathcal{M}$  be a function that is  $(k, 1)$ -simple and has frequencies in the interval  $[-F, F]$ . Then for any desired accuracy  $\epsilon$ , in  $\text{poly}(k, \log 1/\epsilon) \log F$  samples and  $\text{poly}(k, \log 1/\epsilon) \log^2 F$  time, we can output a function  $f'$  such that with probability  $1 - 2^{-\Omega(k)}$ ,*

1.  $f'$  is  $(\text{poly}(k, \log 1/\epsilon), \exp(\text{poly}(k, \log 1/\epsilon)))$ -simple

2.

$$\int_{-1}^1 |f' - f|^2 \leq O\left(\epsilon^2 + \int_{-1}^1 |f - \mathcal{M}|^2\right).$$

**Remark.** While the bound on the coefficients of  $f'$  is not explicitly stated in Theorem 1.1 in Chen et al. [2016], it immediately follows from the proof.

Our algorithm for postprocessing this into a  $\tilde{O}(k)$ -Fourier sparse signal follows roughly the same steps as in the Gaussian case. First, we show that in a certain ‘well-conditioned’ regime, namely, when the frequencies are not too dissimilar, there is a signal using  $O(\log 1/\epsilon)$  frequencies which approximates the function. To handle the general case, we use localizations based on carefully chosen kernels to reduce every signal to a sum of well-conditioned signals (at least, approximately).

One important distinction between the GMM and sparse Fourier reconstruction setting we highlight is that in the latter, the goal is usually to have runtimes which scale logarithmically with  $1/\epsilon$ , whereas in the GMM setting,  $\text{poly}(1/\epsilon)$  sample complexity and thus runtime is unavoidable. However, our naive method of solving the well-conditioned case required constructing a net of  $\text{poly}(1/\epsilon)$  many Gaussians, and thus required  $\text{poly}(1/\epsilon)$  runtime. To circumvent this difficulty, we demonstrate that in fact this can be improved, and that by being more careful, and choosing the (much smaller) set of vertices based on the Chebyshev points, we can in fact improve this runtime significantly. See Sections D and E for a full treatment of our algorithm.

## 2.6 Paper organization

The remainder of the paper will be devoted to proving Theorem 1.1, our main result for GMMs and Theorem 1.3, our main result for sparse Fourier reconstruction. Due to space constraints, the remaining parts are deferred to the appendix. We first present the proof of our result for GMMs. In Section A, we deal with the well-conditioned case. In Section B, we present some tools for localization which we will then use in Section C to prove our full result for GMMs. We then present the proof of our main result for sparse Fourier reconstruction which follows a very similar outline. We deal with the well-conditioned case in Section D and then the general case in Section E. Appendix F contains several basic tools that will be used throughout the paper.

## Acknowledgments and Disclosure of Funding

AL was supported in part by an NSF Graduate Research Fellowship and a Fannie and John Hertz Foundation Fellowship. AM was supported in part by a Microsoft Trustworthy AI Grant, NSF CAREER Award CCF-1453261, NSF Large CCF1565235, a David and Lucile Packard Fellowship and an ONR Young Investigator Award.

### 3 Ethics and broader impact

Our work is purely theoretical and we do not think there are any ethical issues or potential negative societal impacts.

## References

Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Near-optimal-sample estimators for spherical gaussian mixtures. *arXiv preprint arXiv:1402.4746*, 2014.

Jayadev Acharya, Ilias Diakonikolas, Jerry Li, and Ludwig Schmidt. Sample-optimal density estimation in nearly-linear time. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1278–1289. SIAM, 2017.

Dimitris Achlioptas and Frank McSherry. On spectral learning of mixtures of distributions. In *International Conference on Computational Learning Theory*, pages 458–469. Springer, 2005.

Joseph Anderson, Mikhail Belkin, Navin Goyal, Luis Rademacher, and James Voss. The more, the merrier: the blessing of dimensionality for learning large gaussian mixtures. In *Conference on Learning Theory*, pages 1135–1164. PMLR, 2014.

Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 227–236, 2007.

Hassan Ashtiani, Shai Ben-David, Nicholas JA Harvey, Christopher Liaw, Abbas Mehrabian, and Yaniv Plan. Nearly tight sample complexity bounds for learning mixtures of gaussians via sample compression schemes. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 3416–3425, 2018.

Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir Zandieh. A universal sampling method for reconstructing signals with simple fourier transforms. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1051–1063, 2019.

Pranjal Awasthi and Or Sheffet. Improved spectral-norm bounds for clustering. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 37–49. Springer, 2012.

Ainesh Bakshi, Ilias Diakonikolas, He Jia, Daniel M Kane, Pravesh K Kothari, and Santosh S Vempala. Robustly learning mixtures of  $k$  arbitrary gaussians. *arXiv preprint arXiv:2012.02119*, 2020.

Dmitry Batenkov, Laurent Demanet, Gil Goldman, and Yosef Yomdin. Conditioning of partial nonuniform fourier matrices with clustered nodes. *SIAM Journal on Matrix Analysis and Applications*, 41(1):199–220, 2020.

Mikhail Belkin and Kaushik Sinha. Polynomial learning of distribution families. *SIAM Journal on Computing*, 44(4):889–911, 2015.

Aditya Bhaskara, Moses Charikar, Ankur Moitra, and Aravindan Vijayaraghavan. Smoothed analysis of tensor decompositions. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 594–603, 2014.

Aditya Bhaskara, Ananda Suresh, and Morteza Zadimoghaddam. Sparse solutions to nonnegative linear systems and applications. In *Artificial Intelligence and Statistics*, pages 83–92. PMLR, 2015.

Petros Boufounos, Volkan Cevher, Anna C Gilbert, Yi Li, and Martin J Strauss. What’s the frequency, kenneth?: Sublinear fourier sampling off the grid. *Algorithmica*, 73(2):261–288, 2015.

Yoram Bresler and Albert Macovski. Exact maximum likelihood parameter estimation of superimposed exponential signals in noise. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 34(5):1081–1089, 1986.

Emmanuel J Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on pure and applied Mathematics*, 67(6):906–956, 2014.

Siu-On Chan, Ilias Diakonikolas, Rocco A. Servedio, and Xiaorui Sun. Efficient density estimation via piecewise polynomial approximation, 2013.

Siu-On Chan, Ilias Diakonikolas, Rocco A Servedio, and Xiaorui Sun. Efficient density estimation via piecewise polynomial approximation. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 604–613, 2014.

Hanfeng Chen, Jiahua Chen, and John D Kalbfleisch. Testing for a finite mixture model with two components. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(1):95–115, 2004.

Xue Chen and Eric Price. Active regression via linear-sample sparsification. In *Conference on Learning Theory*, pages 663–695. PMLR, 2019.

Xue Chen, Daniel M Kane, Eric Price, and Zhao Song. Fourier-sparse interpolation without a frequency gap. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 741–750. IEEE, 2016.

Sanjoy Dasgupta. Learning mixtures of gaussians. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 634–644. IEEE, 1999.

Sanjoy Dasgupta and Leonard Schulman. A probabilistic analysis of em for mixtures of separated, spherical gaussians. *Journal of Machine Learning Research*, 8(Feb):203–226, 2007.

Constantinos Daskalakis and Gautam Kamath. Faster and sample near-optimal algorithms for proper learning mixtures of gaussians. In *Conference on Learning Theory*, pages 1183–1213. PMLR, 2014.

Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2012.

Ilias Diakonikolas, Homin K Lee, Kevin Matulef, Krzysztof Onak, Ronitt Rubinfeld, Rocco A Servedio, and Andrew Wan. Testing for concise representations. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07)*, pages 549–558. IEEE, 2007.

Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. List-decodable robust mean estimation and learning mixtures of spherical gaussians. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 1047–1060, 2018.

Ilias Diakonikolas, Samuel B Hopkins, Daniel Kane, and Sushrut Karmalkar. Robustly learning any clusterable mixture of gaussians. *arXiv preprint arXiv:2005.06417*, 2020.

Marco F Duarte and Richard G Baraniuk. Spectral compressive sensing. *Applied and Computational Harmonic Analysis*, 35(1):111–129, 2013.

Albert Fannjiang and Wenjing Liao. Coherence pattern-guided compressive sensing with unresolved grids. *SIAM Journal on Imaging Sciences*, 5(1):179–202, 2012.

Jon Feldman, Rocco A Servedio, and Ryan O’Donnell. Pac learning axis-aligned mixtures of gaussians with no separation assumption. In *International Conference on Computational Learning Theory*, pages 20–34. Springer, 2006.

Rong Ge, Qingqing Huang, and Sham M Kakade. Learning mixtures of gaussians in high dimensions. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 761–770, 2015.

Jayanta K Ghosh and Pranab Kumar Sen. On the asymptotic performance of the log likelihood ratio statistic for the mixture model and related results. Technical report, North Carolina State University. Dept. of Statistics, 1984.

Venkatesan Guruswami and David Zuckerman. Robust fourier and polynomial curve fitting. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 751–759. IEEE, 2016.

Yanjun Han, Jiantao Jiao, and Tsachy Weissman. Minimax rate-optimal estimation of divergences between discrete distributions. In *Proceedings of the 2016 International Symposium on Information Theory and Its Applications*, ISITA ’16, pages 256–260, Washington, DC, USA, 2016. IEEE Computer Society.

Moritz Hardt and Eric Price. Tight bounds for learning a mixture of two gaussians. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 753–760, 2015.

JA Hartigan. A failure of likelihood asymptotics for normal mixtures. In *Proc. Barkeley Conference in Honor of J. Neyman and J. Kiefer*, volume 2, pages 807–810, 1985.

Samuel B Hopkins and Jerry Li. Mixture models, robustness, and sum of squares proofs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1021–1034, 2018.

Daniel Hsu and Sham M Kakade. Learning mixtures of spherical gaussians: moment methods and spectral decompositions. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 11–20, 2013.

Qingqing Huang and Sham M Kakade. Super-resolution off the grid. *arXiv preprint arXiv:1509.07943*, 2015.

Tao Huang, Heng Peng, and Kun Zhang. Model selection for gaussian mixture models. *Statistica Sinica*, pages 147–169, 2017.

Vishnu Iyer, Avishay Tal, and Michael Whitmeyer. Junta distance approximation with sub-exponential queries. In *Electron. Colloquium Comput. Complex.*, volume 28, page 4, 2021.

Neal O Jeffries. A note on ‘testing the number of components in a normal mixture’. *Biometrika*, 90(4):991–994, 2003.

Jiantao Jiao, Yanjun Han, and Tsachy Weissman. Minimax estimation of the  $\ell_1$  distance. In *Proceedings of the 2016 IEEE International Symposium on Information Theory*, ISIT ’16, pages 750–754, Washington, DC, USA, 2016. IEEE Computer Society.

Jiantao Jiao, Kartik Venkat, Yanjun Han, and Tsachy Weissman. Minimax estimation of functionals of discrete distributions. *IEEE Transactions on Information Theory*, 61(5):2835–2885, 2017.

Adam Tauman Kalai, Ankur Moitra, and Gregory Valiant. Efficiently learning mixtures of two gaussians. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 553–562, 2010.

Daniel M Kane. Robust learning of mixtures of gaussians. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1246–1258. SIAM, 2021.

Hiroyuki Kasahara and Katsumi Shimotsu. Testing the number of components in normal mixture regression models. *Journal of the American Statistical Association*, 110(512):1632–1645, 2015.

Pravesh K Kothari, Jacob Steinhardt, and David Steurer. Robust moment estimation and improved clustering via sum of squares. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1035–1046, 2018.

Amit Kumar and Ravindran Kannan. Clustering with spectral norm and the k-means algorithm. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 299–308. IEEE, 2010.

Jerry Li and Ludwig Schmidt. Robust and proper learning for mixtures of gaussians via systems of polynomial inequalities. In *Conference on Learning Theory*, pages 1302–1382. PMLR, 2017.

Allen Liu and Ankur Moitra. Settling the robust learnability of mixtures of gaussians. *arXiv preprint arXiv:2011.03622*, 2020.

Allen Liu and Ankur Moitra. Learning gmms with nearly optimal robustness guarantees. *arXiv preprint arXiv:2104.09665*, 2021.

Yungtai Lo, Nancy R Mendell, and Donald B Rubin. Testing the number of components in a normal mixture. *Biometrika*, 88(3):767–778, 2001.

Dustin G Mixon, Soledad Villar, and Rachel Ward. Clustering subgaussian mixtures by semidefinite programming. *Information and Inference: A Journal of the IMA*, 6(4):389–415, 2017.

Ankur Moitra. Super-resolution, extremal functions and the condition number of vandermonde matrices. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 821–830, 2015.

Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of gaussians. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 93–102. IEEE, 2010.

J Neyman and E Scott. On the use of  $c(\alpha)$  tests of composite hypotheses. *Bulletin de L'Institut International de Statistique*, 41:477–497, 1966.

Michal Parnas, Dana Ron, and Ronitt Rubinfeld. Tolerant property testing and distance approximation. *Journal of Computer and System Sciences*, 72(6):1012–1042, 2006.

Karl Pearson. Contributions to the mathematical theory of evolution. *Philosophical Transactions of the Royal Society of London. A*, 185:71–110, 1894.

Yury Polyanskiy and Yihong Wu. Self-regularizing property of nonparametric maximum likelihood estimator in mixture models, 2020. URL <https://arxiv.org/abs/2008.08244>.

Eric Price and Zhao Song. A robust sparse fourier transform in the continuous setting. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 583–600. IEEE, 2015.

R Prony. Essai experimental et analytique sur les lois de la dilabilite des fluides elastiques et sur celles de la force expansive de la vapeur de l'eau et de la vapeur de palcool a differentes temperatures. *Journal de l'Ecole Poly technique*, pages 24–76, 1795.

Theodore J Rivlin. *Chebyshev polynomials*. Courier Dover Publications, 2020.

Robert Roy, Arogyaswami Paulraj, and Thomas Kailath. Esprit—a subspace rotation approach to estimation of parameters of cisoids in noise. *IEEE transactions on acoustics, speech, and signal processing*, 34(5):1340–1342, 1986.

Ralph Otto Schmidt. *A signal subspace approach to multiple emitter location and spectral estimation*. Stanford University, 1982.

Gongguo Tang, Badri Narayan Bhaskar, Parikshit Shah, and Benjamin Recht. Compressed sensing off the grid. *IEEE transactions on information theory*, 59(11):7465–7490, 2013.

Gongguo Tang, Badri Narayan Bhaskar, and Benjamin Recht. Near minimax line spectral estimation. *IEEE Transactions on Information Theory*, 61(1):499–512, 2014.

George R Terrell and David W Scott. Variable kernel density estimation. *The Annals of Statistics*, pages 1236–1265, 1992.

Gregory Valiant and Paul Valiant. A CLT and tight lower bounds for estimating entropy. *Electronic Colloquium on Computational Complexity (ECCC)*, 17(179), 2010a.

Gregory Valiant and Paul Valiant. Estimating the unseen: A sublinear-sample canonical estimator of distributions. *Electronic Colloquium on Computational Complexity (ECCC)*, 17(180), 2010b.

Gregory Valiant and Paul Valiant. Estimating the unseen: An  $n/\log n$ -sample estimator for entropy and support size, shown optimal via new CLTs. In *Proceedings of the 43rd Annual ACM Symposium on the Theory of Computing*, STOC '11, pages 685–694, New York, NY, USA, 2011. ACM.

Santosh Vempala and Grant Wang. A spectral algorithm for learning mixture models. *Journal of Computer and System Sciences*, 68(4):841–860, 2004.

Xuan Wu and Changzhi Xie. Improved algorithms for properly learning mixture of gaussians. In *National Conference of Theoretical Computer Science*, pages 8–26. Springer, 2018.

Yihong Wu and Pengkun Yang. Optimal estimation of gaussian mixtures via denoised method of moments, 2018. URL <https://arxiv.org/abs/1807.07237>.

Yannis G Yatracos. Rates of convergence of minimum distance estimators and kolmogorov’s entropy. *The Annals of Statistics*, pages 768–774, 1985.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]**
  - (b) Did you describe the limitations of your work? **[Yes]**
  - (c) Did you discuss any potential negative societal impacts of your work? **[Yes]**
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]**
  - (b) Did you include complete proofs of all theoretical results? **[Yes]** In the appendix
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[N/A]**
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **[N/A]**
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **[N/A]**
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **[N/A]**
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? **[N/A]**
  - (b) Did you mention the license of the assets? **[N/A]**
  - (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **[N/A]**
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **[N/A]**
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? **[N/A]**
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? **[N/A]**
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **[N/A]**

## Supplementary Materials

### A Well-Conditioned Case: Learning GMMs

We now deal with learning well-conditioned GMMs. We begin by formally specifying the properties that we want the components of the mixture to have. Roughly, we want the components to have comparable variances and the separation between their means cannot be too large compared to the variances. This means that after applying a suitable linear transformation, the components are all not too far from the standard Gaussian  $N(0, 1)$ .

**Definition A.1.** *We say a Gaussian where  $G = N(\mu, \sigma^2)$  is  $\delta$ -well-conditioned if*

- $|\mu| \leq \delta$
- $|\sigma^2 - 1| \leq \delta$

*We say a mixture of Gaussians  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  is  $\delta$ -well-conditioned if all of the components  $G_1, \dots, G_k$  are  $\delta$ -well-conditioned.*

We now state our learning result for well-conditioned mixtures.

**Lemma A.2.** *Let  $\epsilon > 0$  be a parameter. Assume we are given access to a distribution  $f$  such that  $d_{TV}(f, \mathcal{M}) \leq \epsilon$  where  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  is a 0.5-well-conditioned mixture of Gaussians. Then we can compute, in  $\text{poly}(1/\epsilon)$  time, a mixture  $\widetilde{\mathcal{M}}$  of at most  $O(\log 1/\epsilon)$  Gaussians such that  $d_{TV}(\mathcal{M}, \widetilde{\mathcal{M}}) \leq \widetilde{O}(\epsilon)$ .*

**Remark.** *Note that in the well-conditioned case, the number of components in the mixture that we compute does not depend on  $k$ .*

Our algorithm for proving Lemma A.2 can be broken down into two parts. In the first part, we find a mixture of  $\text{poly}(1/\epsilon)$  Gaussians that approximates  $f$ . We then show how to reduce this mixture of  $\text{poly}(1/\epsilon)$  Gaussians to  $O(\log 1/\epsilon)$  Gaussians by using the Taylor series approximation to a Gaussian.

**Lemma A.3.** *Let  $\epsilon > 0$  be a parameter. Assume we are given access to a distribution  $f$  such that  $d_{TV}(f, \mathcal{M}) \leq \epsilon$  where  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  is a 0.5-well-conditioned mixture of Gaussians. Then we can compute, in  $\text{poly}(1/\epsilon)$  time, a mixture of at most  $O(1/\epsilon^2)$  Gaussians that is  $\widetilde{O}(\epsilon)$ -close to  $\mathcal{M}$  in TV distance.*

*Proof.* First, let  $\mathcal{T}$  be the set of all 0.5-well-conditioned Gaussians such that  $\mu$  and  $\sigma^2$  are integer multiples of  $0.1\epsilon$ . Note  $|\mathcal{T}| = O(1/\epsilon^2)$ .

By rounding all of the Gaussians  $G_1, \dots, G_k$  to the nearest element of  $\mathcal{T}$  (this increases our  $L^1$  error by at most  $\epsilon$ ), we may assume that all of the components  $G_1, \dots, G_k$  are actually in  $\mathcal{T}$ . Now note that since  $\|f - w_1 G_1 - \dots - w_k G_k\|_1 \leq 2\epsilon$ , we have for all  $x$ ,

$$|\widehat{f}(x) - w_1 \widehat{G}_1(x) - \dots - w_k \widehat{G}_k(x)| \leq 2\epsilon \quad (1)$$

where  $\widehat{G}_j$  denotes taking the Fourier transform of the pdf of the Gaussian  $G_j$ . Let  $l = \lceil \log 1/\epsilon \rceil$ . We now have,

$$\int_{-l}^l |\widehat{f}(x) - w_1 \widehat{G}_1(x) - \dots - w_k \widehat{G}_k(x)|^2 dx \leq O(l\epsilon^2).$$

Now let all of the Gaussians in  $\mathcal{T}$  be  $G_1, \dots, G_m$  where  $m = |\mathcal{T}|$ . By Lemma F.10 (and splitting into real and imaginary parts), we can compute in  $\text{poly}(1/\epsilon)$  time, nonnegative weights  $\widetilde{w}_1, \dots, \widetilde{w}_m$  with  $\widetilde{w}_1 + \dots + \widetilde{w}_m \leq 1$  such that

$$\int_{-l}^l |\widehat{f}(x) - \widetilde{w}_1 \widehat{G}_1(x) - \dots - \widetilde{w}_m \widehat{G}_m(x)|^2 dx \leq O(l\epsilon^2)$$

which by Cauchy Schwarz implies that

$$\int_{-l}^l |\widehat{f}(x) - \widetilde{w}_1 \widehat{G}_1(x) - \dots - \widetilde{w}_m \widehat{G}_m(x)| dx \leq O(l\epsilon).$$

Now note that since all of the Gaussians  $G_1, \dots, G_m$  are 0.5-well-conditioned, their Fourier transforms  $\widehat{G}_j$  also decay rapidly away from  $[-l, l]$  so combining the above with (1), we deduce that

$$\int_{-\infty}^{\infty} |(\widetilde{w_1} \widehat{G}_1(x) + \dots + \widetilde{w_m} \widehat{G}_m(x)) - w_1 \widehat{G}_1(x) - \dots - w_k \widehat{G}_k(x)| \leq O(l\epsilon).$$

From the Fourier transform of the above we then get for all  $x$

$$|\widetilde{w_1} G_1(x) + \dots + \widetilde{w_m} G_m(x) - w_1 G_1(x) - \dots - w_m G_m(x)| \leq O(l\epsilon)$$

and since all of the Gaussians involved are 0.5-well-conditioned, they all decay rapidly outside the interval  $[-l, l]$  and we conclude

$$\int_{-\infty}^{\infty} |\widetilde{w_1} G_1(x) + \dots + \widetilde{w_m} G_m(x) - w_1 G_1(x) - \dots - w_m G_m(x)| dx \leq O(l^2\epsilon).$$

Finally, note that by the above, we must have  $1 - O(l^2\epsilon) \leq \widetilde{w_1} + \dots + \widetilde{w_m} \leq 1 + O(l^2\epsilon)$  so rescaling to an actual mixture i.e. so that the weights  $\widetilde{w_1} + \dots + \widetilde{w_m} = 1$ , will affect the above error by at most  $O(l^2\epsilon)$ . Thus, we can output this mixture and we are done.  $\blacksquare$

Next, as an immediate consequence of Lemma F.8, a 0.5-well-conditioned Gaussian can be well approximated by its Taylor expansion.

**Corollary A.4.** *Let  $G = N(\mu, \sigma^2)$  be a 0.5-well-conditioned Gaussian. Let  $\epsilon > 0$  be some parameter and let  $l = \lceil \log 1/\epsilon \rceil$ . Then we can compute a polynomial  $P_G(x)$  of degree  $(10l)^2$  such that for all  $x \in [-l, l]$ ,*

$$|G(x) - P_G(x)| \leq O(\epsilon).$$

*Proof.* This follows immediately from using Lemma F.8 and applying the appropriate linear transformation to the polynomial.  $\blacksquare$

We can now complete the proof of Lemma A.2 by using Lemma A.3 and then using Corollary A.4 and Caratheodory to reduce the number of components.

*Proof of Lemma A.2.* By Lemma A.3, we can compute a mixture

$$\widetilde{\mathcal{M}} = \widetilde{w_1} \widetilde{G}_1 + \dots + \widetilde{w_m} \widetilde{G}_m$$

such that  $m = O(1/\epsilon^2)$  and

$$\|\widetilde{\mathcal{M}} - \mathcal{M}\|_1 \leq \widetilde{O}(\epsilon).$$

For each Gaussian  $\widetilde{G}_j$ , let  $P_{\widetilde{G}_j}(x)$  be the polynomial computed in Lemma F.8. Write

$$P_{\widetilde{G}_j}(x) = a_{j,0} + a_{j,1}x + \dots + a_{j,(10l)^2}x^{(10l)^2}.$$

Define the vector

$$v_j = (a_{j,0}, a_{j,1}, \dots, a_{j,(10l)^2}).$$

Now the point  $\widetilde{w_1}v_1 + \dots + \widetilde{w_m}v_m$  is in the convex hull of  $v_1, \dots, v_m$ . By Caratheodory (since the space is  $(10l)^2 + 1$ -dimensional), it must be in the convex hull of some  $(10l)^2 + 1$  of the vertices. Thus, we can compute indices  $i_0, \dots, i_{(10l)^2}$  and nonnegative weights  $w'_0, \dots, w'_{(10l)^2}$  summing to 1 such that

$$\widetilde{w_1}v_1 + \dots + \widetilde{w_m}v_m = w'_0v_{i_0} + \dots + w'_{(10l)^2}v_{i_{(10l)^2}}.$$

The above implies that for all  $x$ ,

$$\widetilde{w_1}P_{\widetilde{G}_1}(x) + \dots + \widetilde{w_m}P_{\widetilde{G}_m}(x) = w'_0P_{\widetilde{G}_{i_0}}(x) + \dots + w'_{(10l)^2}P_{\widetilde{G}_{i_{(10l)^2}}}(x).$$

Now by Corollary A.4 and the fact that all of the Gaussians are 0.5-well-conditioned, meaning that they decay rapidly outside of  $[-l, l]$ , we conclude that if we set

$$\mathcal{M}' = w'_0 \widetilde{G}_{i_0} + \dots + w'_{(10l)^2} \widetilde{G}_{i_{(10l)^2}}$$

then

$$\|\widetilde{\mathcal{M}} - \mathcal{M}'\|_1 \leq \widetilde{O}(\epsilon)$$

and then we have

$$\|\mathcal{M} - \mathcal{M}'\|_1 \leq \|\widetilde{\mathcal{M}} - \mathcal{M}\|_1 + \|\widetilde{\mathcal{M}} - \mathcal{M}'\|_1 \leq \widetilde{O}(\epsilon)$$

as desired.  $\blacksquare$

We can slightly improve Lemma A.2 to work even when we do not have a precise estimate of  $d_{\text{TV}}(f, \mathcal{M})$  since we can just repeatedly decrease our target accuracy until we cannot improve our accuracy further. Recall that we can use Claim 2.11 to test the  $L^1$  distance between two distributions. We now have the following (slight) improvement of Lemma A.2.

**Corollary A.5.** *Let  $\epsilon > 0$  be a parameter. Let  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  be an unknown 0.5-well-conditioned mixture of Gaussians. Assume we are given access to a distribution  $f$ . Then we can compute, in  $\text{poly}(1/\epsilon)$  time, a mixture  $\widetilde{\mathcal{M}}$  of at most  $O(\log 1/\epsilon)$  Gaussians such that with high probability,*

$$d_{\text{TV}}(f, \widetilde{\mathcal{M}}) \leq \epsilon^2 + \text{poly}(\log 1/\epsilon) d_{\text{TV}}(f, \mathcal{M}).$$

*Proof.* We can simply start from  $\epsilon' = 1$  and run the algorithm in Lemma A.2 with parameter  $\epsilon'$  and then estimate  $d_{\text{TV}}(f, \widetilde{\mathcal{M}})$  using Claim 2.11. If  $d_{\text{TV}}(f, \widetilde{\mathcal{M}}) \leq \epsilon' \text{poly}(\log 1/\epsilon)$  then we can decrease  $\epsilon'$  by a factor of 0.9 and repeat. Repeating this process and taking the smallest accuracy  $\epsilon' \geq \epsilon^3$  for which the above check succeeds, we get (from the guarantee of Lemma A.2) that

$$d_{\text{TV}}(f, \widetilde{\mathcal{M}}) \leq \epsilon^2 + \text{poly}(\log 1/\epsilon) d_{\text{TV}}(f, \mathcal{M})$$

and we are done. ■

## B Function Approximations Using Gaussians

In this section, we present several results about approximating functions as a sum of Gaussians. These results will be key building blocks in the localization steps of both of our algorithms. The main result of this section, Theorem B.3, allows us to  $\epsilon$ -approximate the indicator function of an interval as a sum of  $\text{poly}(\log 1/\epsilon)$ -Gaussians.

First, it will be convenient to renormalize Gaussians so that their maximum value is 1. After renormalization, we call them Gaussian multipliers.

**Definition B.1** (Gaussian Multiplier). *For parameters  $\mu, \sigma$ , we define*

$$M_{\mu, \sigma^2}(x) = e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

*i.e. it is a Gaussian scaled so that its maximum value is 1.*

We also introduce the some additional terminology.

**Definition B.2** (Significant Interval). *For a Gaussian multiplier  $M_{\mu, \sigma^2}$ , we say the  $C$ -significant interval of  $M$  is  $[\mu - C\sigma, \mu + C\sigma]$ . We will use the same terminology for a Gaussian  $N(\mu, \sigma^2)$ .*

It will be used repeatedly that for a Gaussian (or Gaussian multiplier),  $1 - \epsilon$ -fraction of its mass is contained in its  $O(\sqrt{\log 1/\epsilon})$ -significant interval. We now state the main result of this section about approximating the indicator function of an interval as a weighted sum of Gaussian multipliers.

**Theorem B.3.** *Let  $l$  be a positive real number and  $0 < \epsilon < 0.1$  be a parameter. There is a function  $f$  with the following properties*

1.  *$f$  can be written a linear combination of Gaussian multipliers*

$$f(x) = w_1 M_{\mu_1, \sigma_1^2}(x) + \dots + w_n M_{\mu_n, \sigma_n^2}(x)$$

*where  $n = O((\log 1/\epsilon)^2)$  and  $0 \leq w_1, \dots, w_n \leq 1$*

2. *The  $10\sqrt{\log 1/\epsilon}$ -significant intervals of all of the  $M_{\mu_i, \sigma_i^2}$  are contained in the interval  $[-(1 + \epsilon)l, (1 + \epsilon)l]$*
3.  *$0 \leq f(x) \leq 1 + \epsilon$  for all  $x$*
4.  *$1 - \epsilon \leq f(x) \leq 1 + \epsilon$  for all  $x$  in the interval  $[-l, l]$*
5.  *$0 \leq f(x) \leq \epsilon$  for  $x \geq (1 + \epsilon)l$  and  $x \leq -(1 + \epsilon)l$*

## B.1 Approximating a Constant Function

Recall Lemma 2.5 (restated below) that allows us to approximate a constant function using an infinite sum of evenly spaced Gaussian multipliers.

**Lemma B.4** (Restated from Lemma 2.5). *Let  $0 < \epsilon < 0.1$  be a parameter. Let  $c$  be a real number such that  $0 < c \leq (\log 1/\epsilon)^{-1/2}$ . Define*

$$f(x) = \sum_{j=-\infty}^{\infty} \frac{c}{\sqrt{2\pi}} M_{cj\sigma, \sigma^2}(x).$$

*Then  $1 - \epsilon \leq f(x) \leq 1 + \epsilon$  for all  $x$ .*

## B.2 Approximating an Interval

The next step in the proof of Theorem B.3 is to show how to approximate an interval using a finite number of Gaussian multipliers i.e. we need to show how to create the sharp transitions at the ends of the interval. In light of Lemma B.4, we can create a function satisfying the last four properties by taking  $\tilde{O}((1/\epsilon)^2)$  evenly spaced Gaussians multipliers with standard deviation  $\epsilon^2 l$ . However, this is too many components and we must reduce the number of components to  $O(\log^2 1/\epsilon)$ . The way we do this is by merging most of these components (all but the ones on the ends) into fewer components with larger standard deviation. We keep iterating this merging process and prove that we can eventually reduce the number of components to  $O(\log^2 1/\epsilon)$ .

First, the following result is an immediate consequence of Lemma B.4. It allows us to approximate a Gaussian with standard deviation  $2\sigma$  as a weighted sum of Gaussians with standard deviation  $\sigma$ .

**Corollary B.5.** *Let  $\epsilon$  be a parameter. Let  $c$  be a real number such that  $0 < c \leq 0.5(\log 1/\epsilon)^{-1/2}$ . Let*

$$g(x) = \sum_{j=-\infty}^{\infty} \frac{\sqrt{2}c}{\sqrt{3\pi}} e^{-\frac{c^2 j^2}{6}} M_{cj\sigma, \sigma^2}(x).$$

*Then for all  $x$ ,*

$$(1 - \epsilon) M_{0, 4\sigma^2}(x) \leq g(x) \leq (1 + \epsilon) M_{0, 4\sigma^2}(x).$$

*Proof.* Lemma B.4 (with  $c \leftarrow \frac{2}{\sqrt{3}}c$ ,  $\sigma \leftarrow \frac{2}{\sqrt{3}}\sigma$ ) implies that the function

$$f(x) = \sum_{j=-\infty}^{\infty} \frac{\sqrt{2}c}{\sqrt{3\pi}} M_{\frac{4}{3}cj\sigma, \frac{4}{3}\sigma^2}(x)$$

is between  $1 - \epsilon$  and  $1 + \epsilon$  everywhere. Now consider

$$\begin{aligned} f(x) \cdot M_{0, 4\sigma^2}(x) &= \sum_{j=-\infty}^{\infty} \frac{\sqrt{2}c}{\sqrt{3\pi}} M_{\frac{4}{3}cj\sigma, \frac{4}{3}\sigma^2}(x) M_{0, 4\sigma^2}(x) = \sum_{j=-\infty}^{\infty} \frac{\sqrt{2}c}{\sqrt{3\pi}} e^{-\frac{x^2 + 3(x - \frac{4}{3}cj\sigma)^2}{8\sigma^2}} \\ &= \sum_{j=-\infty}^{\infty} \frac{\sqrt{2}c}{\sqrt{3\pi}} e^{-\frac{c^2 j^2}{6}} e^{-\frac{(x - cj\sigma)^2}{2\sigma^2}} = \sum_{j=-\infty}^{\infty} \frac{\sqrt{2}c}{\sqrt{3\pi}} e^{-\frac{c^2 j^2}{6}} M_{cj\sigma, \sigma^2}(x). \end{aligned}$$

■

In the next lemma, we show when given a sum of evenly spaced Gaussians with standard deviation  $\sigma$ , we can replace almost all of them (except for ones on the ends) with a sum of fewer evenly spaced Gaussians with standard deviation  $2\sigma$ .

**Lemma B.6.** *Let  $\epsilon$  be a parameter. Let  $c$  be a real number such that  $0 < c \leq 0.01(\log 1/\epsilon)^{-1/2}$ . Let  $b$  be a positive integer. Consider the function*

$$f(x) = \sum_{j=0}^b \frac{c}{\sqrt{2\pi}} M_{cj\sigma, \sigma^2}(x).$$

Let  $C = \lceil 10^2 c^{-1} \log(1/\epsilon)^{1/2} \rceil$ . There is a function  $g$  of the form

$$g(x) = \sum_{j=0}^{2C} \frac{w_j c}{\sqrt{2\pi}} M_{cj\sigma,\sigma^2}(x) + \sum_{j=b-2C}^b \frac{w_j c}{\sqrt{2\pi}} M_{cj\sigma,\sigma^2}(x) + \sum_{j=\lfloor C/2 \rfloor}^{\lceil (b-C)/2 \rceil} \frac{c}{\sqrt{2\pi}} M_{2cj\sigma,4\sigma^2}(x)$$

where the  $0 \leq w_0, \dots, w_{2C}, w_{b-2C}, \dots, w_b \leq 1$  are weights and

$$\|f - g\|_\infty \leq \epsilon^{10}.$$

*Proof.* Let  $\epsilon' = \epsilon^{100}$ . By Corollary B.5, for any real numbers  $j, x$ ,

$$\left| M_{cj\sigma,4\sigma^2}(x) - \sum_{k=-\infty}^{\infty} \frac{\sqrt{2}c}{\sqrt{3}\pi} e^{-\frac{c^2 k^2}{6}} M_{c(k+j)\sigma,\sigma^2}(x) \right| \leq \epsilon' M_{cj\sigma,4\sigma^2}(x).$$

Now we use the above inequality on each term of the last sum in the expression for  $g(x)$ .

$$\begin{aligned} & \left| \sum_{j=\lfloor C/2 \rfloor}^{\lceil (b-C)/2 \rceil} \frac{c M_{2cj\sigma,4\sigma^2}(x)}{\sqrt{2\pi}} - \sum_{j=\lfloor C/2 \rfloor}^{\lceil (b-C)/2 \rceil} \sum_{k=-\infty}^{\infty} \frac{c^2}{\pi\sqrt{3}} e^{-\frac{c^2 k^2}{6}} M_{c(k+2j)\sigma,\sigma^2}(x) \right| \\ & \leq \epsilon' \sum_{j=\lfloor C/2 \rfloor}^{\lceil (b-C)/2 \rceil} \frac{c M_{2cj\sigma,4\sigma^2}(x)}{\sqrt{2\pi}} \leq 2\epsilon' \end{aligned}$$

where the last step follows from Lemma B.4. Now we rewrite the second sum in the LHS above. Let

$$\begin{aligned} S(x) &= \sum_{j=\lfloor C/2 \rfloor}^{\lceil (b-C)/2 \rceil} \sum_{k=-\infty}^{\infty} \frac{c^2}{\pi\sqrt{3}} e^{-\frac{c^2 k^2}{6}} M_{c(k+2j)\sigma,\sigma^2}(x) \\ &= \sum_{l=-\infty}^{\infty} \frac{c^2}{\pi\sqrt{3}} M_{cl\sigma,\sigma^2}(x) \sum_{j=\lfloor C/2 \rfloor}^{\lceil (b-C)/2 \rceil} e^{-\frac{c^2(l-2j)^2}{6}}. \end{aligned}$$

Define

$$a_l = \sum_{j=\lfloor C/2 \rfloor}^{\lceil (b-C)/2 \rceil} e^{-\frac{c^2(l-2j)^2}{6}}.$$

First, by applying Lemma B.4 with parameters  $c \leftarrow \frac{2}{\sqrt{3}}c, \sigma \leftarrow \sqrt{3}c^{-1}$ , we have that for all real numbers  $l$ ,

$$\left| \frac{\sqrt{3\pi}}{\sqrt{2}c} - \sum_{-\infty}^{\infty} e^{-\frac{c^2(l-2j)^2}{6}} \right| \leq \frac{\epsilon' \sqrt{3\pi}}{\sqrt{2}c}.$$

By the way we chose  $C$ , we deduce that for all integers  $l$  with  $2C \leq l \leq b - 2C$ ,

$$\left| \frac{\sqrt{3\pi}}{\sqrt{2}c} - a_l \right| \leq (2\epsilon') \frac{\sqrt{3\pi}}{\sqrt{2}c} \quad (2)$$

for all integers  $0 \leq l \leq 2C$ , or  $b - 2C \leq l \leq b$ ,

$$a_l \leq \frac{\sqrt{3\pi}}{\sqrt{2}c} (1 + 2\epsilon') \quad (3)$$

and finally for all integers  $l < 0$  or  $l > b$ ,

$$a_l \leq \frac{\sqrt{3\pi}}{\sqrt{2}c} (2\epsilon'). \quad (4)$$

To obtain these inequalities, we simply use the fact that the terms in the sum

$$\sum_{-\infty}^{\infty} e^{-\frac{c^2(l-2j)^2}{6}}$$

decay exponentially when  $j$  is far from  $l/2$  so their total contribution is small.

Now we can set  $w_0, \dots, w_{2C}, w_{b-2C}, \dots, w_b$  in the expression for  $g(x)$  as follows:

$$w_j = \max \left( 0, 1 - \frac{\sqrt{2}c}{\sqrt{3}\pi} a_j \right).$$

It is clear that all of these weights are between 0 and 1. We now have that

$$\|f - g\|_\infty \leq 2\epsilon' + \|f(x) - \left( S(x) + \sum_{j=0}^{2C} \frac{w_j c}{\sqrt{2\pi}} M_{cj\sigma,\sigma^2}(x) + \sum_{j=b-2C}^b \frac{w_j c}{\sqrt{2\pi}} M_{cj\sigma,\sigma^2}(x) \right)\|_\infty$$

The expression inside the norm on the RHS can be rewritten as

$$\begin{aligned} & \sum_{l=2C+1}^{b-2C-1} \left( \frac{c}{\sqrt{2\pi}} - \frac{c^2}{\pi\sqrt{3}} a_l \right) M_{cl\sigma,\sigma^2}(x) + \sum_{l=0}^{2C} \left( \frac{c}{\sqrt{2\pi}} (1 - w_l) - \frac{c^2}{\pi\sqrt{3}} a_l \right) M_{cl\sigma,\sigma^2}(x) \\ & + \sum_{l=b-2C}^b \left( \frac{c}{\sqrt{2\pi}} (1 - w_l) - \frac{c^2}{\pi\sqrt{3}} a_l \right) M_{cl\sigma,\sigma^2}(x) + \sum_{l=-\infty}^{-1} -\frac{c^2}{\pi\sqrt{3}} a_l M_{cl\sigma,\sigma^2}(x) \\ & + \sum_{l=b+1}^{\infty} -\frac{c^2}{\pi\sqrt{3}} a_l M_{cl\sigma,\sigma^2}(x) \end{aligned}$$

and combining (2,3, 4), we deduce that the above has  $L^\infty$  norm at most

$$\|(10\epsilon') \sum_{l=-\infty}^{\infty} \frac{c}{\sqrt{2\pi}} M_{cl\sigma,\sigma^2}(x)\|_\infty \leq 20\epsilon'.$$

where we used Lemma B.4. Thus,  $\|f - g\|_\infty \leq 22\epsilon'$  and we are done.  $\blacksquare$

We can now prove Theorem B.3 by repeatedly applying Lemma B.6.

*Proof of Theorem B.3.* Let  $c = 0.01(\log 1/\epsilon)^{-1/2}$ . Let  $K = \lceil \frac{1+0.5\epsilon}{c\epsilon^2} \rceil$

$$f_0(x) = \sum_{j=-K}^K \frac{c}{\sqrt{2\pi}} M_{cj\epsilon^2 l, \epsilon^4 l^2}(x).$$

Let  $\epsilon' = \epsilon^{10}$ . Using Lemma B.4, (and basic tail decay properties of a Gaussian) we get that

- $0 \leq f_0(x) \leq 1 + \epsilon'$  for all  $x$
- $1 - \epsilon' \leq f_0(x) \leq 1 + \epsilon'$  for all  $x$  in the interval  $[-l, l]$
- $0 \leq f_0(x) \leq \epsilon'$  for  $x \geq (1 + \epsilon)l$  and  $x \leq -(1 + \epsilon)l$

Now we can apply Lemma B.6 to  $f_0(x)$  to obtain

$$\begin{aligned} f_1(x) &= \sum_{j=-K}^{-K+2C} \frac{w_j c}{\sqrt{2\pi}} M_{cj\epsilon^2 l, \epsilon^4 l^2}(x) + \sum_{j=K-2C}^K \frac{w_j c}{\sqrt{2\pi}} M_{cj\epsilon^2 l, \epsilon^4 l^2}(x) \\ &+ \sum_{j=-\lceil (K-C)/2 \rceil}^{\lceil (K-C)/2 \rceil} \frac{c}{\sqrt{2\pi}} M_{2cj\epsilon^2 l, 4\epsilon^4 l^2}(x) \end{aligned}$$

where  $C = \lceil 10^2 c^{-1} \log(1/\epsilon)^{1/2} \rceil$ , the  $w_j$  are weights between 0 and 1, and

$$\|f_1 - f_0\| \leq \epsilon'.$$

Now we can apply Lemma B.6 again on the last sum in the expression for  $f_1$ . We have to do this at most  $10 \log 1/\epsilon$  times before there are at most  $O((\log 1/\epsilon)^2)$  components remaining. It is clear

that in this procedure, the  $10\sqrt{\log 1/\epsilon}$ -significant intervals of all of the Gaussian multipliers always remains in  $[-(1+\epsilon)l, (1+\epsilon)l]$ . Also, the total  $L^\infty$  error incurred over all of the applications of Lemma B.6 is at most  $10\epsilon' \log 1/\epsilon \leq \epsilon^9$ . It is clear that all of the weights are always nonnegative and in the interval  $[0, 1]$ . Thus, the final function  $f$  satisfies

- $0 \leq f(x) \leq 1 + \epsilon$  for all  $x$
- $1 - \epsilon \leq f(x) \leq 1 + \epsilon$  for all  $x$  in the interval  $[-l, l]$
- $0 \leq f(x) \leq \epsilon$  for  $x \geq (1+\epsilon)l$  and  $x \leq -(1+\epsilon)l$

and we are done. ■

In light of Theorem B.3, we may make the following definition.

**Definition B.7.** For parameters  $\epsilon, l$ , let  $\mathcal{I}_{\epsilon, l}$  denote the function computed in Theorem B.3 for parameters  $\epsilon, l/(1+\epsilon)$ . We will also use  $\mathcal{I}_{\epsilon, l}^{(a)}$  to denote the function  $\mathcal{I}_{\epsilon, l}(x-a)$ .

**Remark.** We define  $\mathcal{I}_{\epsilon, l}$  as above because it will be convenient later to be able to say that the significant part of  $\mathcal{I}_{\epsilon, l}$  is contained in the interval  $[-l, l]$ .

## C Nearly-Properly Learning GMMs: Full Version

In this section, we complete the proof of our main result for learning GMMs, Theorem 1.1. We localize the distribution by multiplying by a Gaussian multiplier  $M_{\mu, \sigma^2}$ . Note that the product of two Gaussians is still a Gaussian so multiplying a GMM by a Gaussian multiplier results in a re-weighted mixture of Gaussians. Roughly, we argue that the new weights on components of the mixture that are far away from the multiplier  $M_{\mu, \sigma^2}$  are negligible so the resulting mixture is well-conditioned and we can then use Corollary A.5 to reconstruct the localized distribution. To reconstruct the entire distribution, we show that it suffices to sum together  $\tilde{O}(k)$  different localized reconstructions.

### C.1 Localizing with Gaussian Multipliers

Recall Claim 2.4 (restated below) which gives an explicit formula for what happens when we have a Gaussian  $G_1 = N(\mu_1, \sigma_1^2)$  and we multiply it by a Gaussian multiplier  $M_{\mu, \sigma^2}(x)$ .

**Claim C.1** (Restated from Claim 2.4). *We have the identity*

$$M_{\mu, \sigma^2}(x)N(\mu_1, \sigma_1^2) = \frac{1}{\sqrt{1 + \frac{\sigma_1^2}{\sigma^2}}} e^{-\frac{(\mu_1 - \mu)^2}{2(\sigma_1^2 + \sigma^2)}} N\left(\frac{\mu\sigma_1^2 + \mu_1\sigma^2}{\sigma_1^2 + \sigma^2}, \frac{\sigma_1^2\sigma^2}{\sigma_1^2 + \sigma^2}\right).$$

### C.2 Building Blocks

We first consider reconstructing a GMM  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  after multiplying by a Gaussian multiplier  $M_{\mu, \sigma^2}$ . As a corollary of Claim C.1, we know that when the  $C$ -significant intervals (recall Definition B.2) of a Gaussian  $G_j$  and the multiplier  $M_{\mu, \sigma^2}(x)$  are disjoint for large  $C$ , then the  $L^1$  norm of their product is  $e^{-\Omega(C^2)}$ . In particular this means that after multiplying by  $M_{\mu, \sigma^2}$ , the only components that remain relevant are those that have nontrivial overlap with the multiplier  $M_{\mu, \sigma^2}$ . The only way these components will not form a well-conditioned mixture is if there is some  $G_j$  that is very thin (i.e.  $\sigma_j \ll \sigma$ ) and overlaps with  $M_{\mu, \sigma^2}$ . As long as this doesn't happen, we can apply Corollary A.5. We formalize this below.

**Corollary C.2.** Let  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  be an arbitrary mixture of Gaussians where  $G_i = N(\mu_i, \sigma_i^2)$ . Let  $\epsilon > 0$  be some parameter and let  $l = \lceil \sqrt{\log(1/\epsilon)} \rceil$ . Assume we are given access to a distribution  $f$ . Let  $M_{\mu, \sigma^2}$  be a Gaussian multiplier. Assume that for all  $i \in [k]$ , either  $\sigma_i \geq 4l\sigma$  or the  $10l$ -significant intervals of  $G_i$  and  $M_{\mu, \sigma^2}$  do not intersect. Then in  $\text{poly}(1/\epsilon)$  time and with high probability, we can compute a weighted sum  $\tilde{M}$  of at most  $O(\log(1/\epsilon))$  Gaussians such that

$$\|\tilde{M} - M_{\mu, \sigma^2} f\|_1 \leq \epsilon + \text{poly}(\log(1/\epsilon)) \|M_{\mu, \sigma^2}(\mathcal{M} - f)\|_1.$$

*Proof.* We compute  $M_{\mu, \sigma^2} f$  and let  $C = \|M_{\mu, \sigma^2} f\|_1$ . If  $C \leq \epsilon$  then we may simple output 0. Otherwise, we will apply Corollary A.5 on  $M_{\mu, \sigma^2} f/C$  and multiply the result by  $C$ . We must first verify the conditions of Corollary A.5. Let  $S \subset [k]$  be the indices such that the  $10l$ -significant intervals of  $G_i$  and  $M_{\mu, \sigma^2}$  intersect. First for  $i \notin S$ , by Claim C.1,

$$\|G_i M_{\mu, \sigma^2}\|_1 \leq e^{-\frac{(\mu_i - \mu)^2}{2(\sigma_i^2 + \sigma^2)}} \leq e^{-10l^2} \leq \epsilon^{10}.$$

Let

$$\mathcal{M}' = \sum_{i \in S} w_i G_i.$$

Then we know

$$\left\| \frac{M_{\mu, \sigma^2} f}{C} - \frac{M_{\mu, \sigma^2} \mathcal{M}'}{C} \right\|_1 \leq \epsilon^9 + \frac{\|M_{\mu, \sigma^2}(f - \mathcal{M})\|_1}{C}$$

Next, for  $i \in S$ ,

$$G_i M_{\mu, \sigma^2} = w'_i N\left(\frac{\mu \sigma_i^2 + \mu_i \sigma^2}{\sigma_i^2 + \sigma^2}, \frac{\sigma_i^2 \sigma^2}{\sigma_i^2 + \sigma^2}\right)$$

for some weight  $w'_i$  and since we must have  $\sigma_i \geq 4l\sigma$ , then

$$\begin{aligned} \frac{\sigma^2}{2} &\leq \frac{\sigma_i^2 \sigma^2}{\sigma_i^2 + \sigma^2} \leq \sigma^2 \\ \left| \frac{\mu \sigma_i^2 + \mu_i \sigma^2}{\sigma_i^2 + \sigma^2} - \mu \right| &= \left| \frac{(\mu_i - \mu) \sigma^2}{\sigma_i^2 + \sigma^2} \right| \leq \frac{l(\sigma + \sigma_i) \sigma^2}{\sigma_i^2 + \sigma^2} \leq \frac{\sigma}{2} \end{aligned}$$

Let

$$\mathcal{M}'' = \sum_{i \in S} \frac{w'_i}{\sum_{i \in S} w'_i} N\left(\frac{\mu \sigma_i^2 + \mu_i \sigma^2}{\sigma_i^2 + \sigma^2}, \frac{\sigma_i^2 \sigma^2}{\sigma_i^2 + \sigma^2}\right).$$

Then we deduce, since  $\|M_{\mu, \sigma^2} f/C\|_1 = 1$ , that

$$\begin{aligned} \left\| \frac{M_{\mu, \sigma^2} f}{C} - \mathcal{M}'' \right\|_1 &\leq \left\| \mathcal{M}'' - \frac{M_{\mu, \sigma^2} \mathcal{M}'}{C} \right\|_1 + \left\| \frac{M_{\mu, \sigma^2} f}{C} - \frac{M_{\mu, \sigma^2} \mathcal{M}'}{C} \right\|_1 \\ &\leq 2 \left\| \frac{M_{\mu, \sigma^2} f}{C} - \frac{M_{\mu, \sigma^2} \mathcal{M}'}{C} \right\|_1 \\ &\leq \epsilon^8 + 2 \frac{\|M_{\mu, \sigma^2}(f - \mathcal{M})\|_1}{C} \end{aligned}$$

and further, after applying a suitable linear transformation (taking  $(\mu, \sigma^2) \rightarrow (0, 1)$ ) that the mixture  $\mathcal{M}''$  is 0.5-well-conditioned. Thus, we can apply Corollary A.5 and compute a weighted sum of  $O(\log(1/\epsilon))$  Gaussians,  $\widetilde{\mathcal{M}}$  such that

$$\left\| \frac{M_{\mu, \sigma^2} f}{C} - \widetilde{\mathcal{M}} \right\|_1 \leq \text{poly}(\log(1/\epsilon)) \left( \epsilon^8 + \frac{\|M_{\mu, \sigma^2}(f - \mathcal{M})\|_1}{C} \right).$$

Now we can simply output  $C\widetilde{\mathcal{M}}$  (which is still a weighted sum of  $O(\log(1/\epsilon))$  Gaussians) and we are done.  $\blacksquare$

Recall that Theorem B.3 shows how to express an interval as a sum of Gaussian multipliers. Combining Theorem B.3 with Corollary C.2, we show that we can approximate a GMM over an interval as long as the interval does not overlap with a component that is much thinner than it.

**Lemma C.3.** *Let  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  be an arbitrary mixture of Gaussians where  $G_i = N(\mu_i, \sigma_i^2)$ . Let  $\epsilon > 0$  be some parameter and let  $l = \lceil \sqrt{\log(1/\epsilon)} \rceil$ . Assume we are given access to a distribution  $f$ . Let  $I = [a, b]$  be an interval. Assume that for all  $i \in [k]$ , either  $\sigma_i \geq (b - a)$  or the  $10l$ -significant interval of  $G_i$  does not intersect  $I$ . Then in  $\text{poly}(1/\epsilon)$  time and with high probability, we can compute a weighted sum  $\widetilde{\mathcal{M}}$  of at most  $\text{poly}(\log(1/\epsilon))$  Gaussians such that*

$$\|\widetilde{\mathcal{M}} - f \cdot 1_I\|_1 \leq \text{poly}(\log(1/\epsilon)) (\epsilon + \|1_I(\mathcal{M} - f)\|_1)$$

where  $1_I$  denotes the indicator function of  $I$ .

*Proof.* Consider the function  $\mathcal{I} = \mathcal{I}_{\epsilon, (b-a)/2}^{(a+b)/2}$  (recall Definition B.7). Now note that by Theorem B.3,  $\mathcal{I}$  can be written in the form

$$\mathcal{I} = \widetilde{w}_1 M_{\widetilde{\mu}_1, \widetilde{\sigma}_1^2} + \cdots + \widetilde{w}_n M_{\widetilde{\mu}_n, \widetilde{\sigma}_n^2}$$

where  $n = O(\log^2 1/\epsilon)$ . Furthermore, for all  $i \in [n]$ , we have  $0 \leq \widetilde{w}_i \leq 1$  and  $\widetilde{\sigma}_i \leq (b-a)/(4l)$  and the  $10l$ -significant intervals of  $M_{\widetilde{\mu}_i, \widetilde{\sigma}_i^2}$  are all contained in the interval  $[a, b]$ . Thus we can apply Corollary C.2 on  $M_{\widetilde{\mu}_i, \widetilde{\sigma}_i^2} f$  for all  $i \in [n]$ . Adding the results with the corresponding weights  $\widetilde{w}_1, \dots, \widetilde{w}_n$ , we obtain a function  $\widetilde{\mathcal{M}}$  that is a weighted sum of at most  $\text{poly}(\log(1/\epsilon))$  Gaussians such that

$$\begin{aligned} \|\widetilde{\mathcal{M}} - f\mathcal{I}\|_1 &\leq \text{poly}(\log(1/\epsilon)) \left( \epsilon + \sum_{i=1}^n \widetilde{w}_i \|M_{\widetilde{\mu}_i, \widetilde{\sigma}_i^2}(\mathcal{M} - f)\|_1 \right) \\ &= \text{poly}(\log(1/\epsilon)) (\epsilon + \|\mathcal{I}(\mathcal{M} - f)\|_1). \end{aligned}$$

Thus,

$$\|\widetilde{\mathcal{M}} - \mathcal{M}\mathcal{I}\|_1 \leq \text{poly}(\log(1/\epsilon)) (\epsilon + \|\mathcal{I}(\mathcal{M} - f)\|_1). \quad (5)$$

Next, by the properties in Theorem B.3,

$$\begin{aligned} \|\mathcal{M}(\mathcal{I} - 1_I)\|_1 &\leq \epsilon \int_{-\infty}^{\infty} \mathcal{M} + \int_a^{a+\epsilon(b-a)} \mathcal{M} + \int_{b-\epsilon(b-a)}^b \mathcal{M} \\ &\leq \epsilon + \left( \sum_{j=1}^k w_j \int_a^{a+\epsilon(b-a)} G_j + w_j \int_{b-\epsilon(b-a)}^b G_j \right). \end{aligned}$$

Consider one of the component Gaussians  $G_j$  where  $j \in [k]$ . If the  $10l$ -significant interval of  $G_j$  does not intersect  $[a, b]$  then it is clear that the total mass of  $G_j$  on the interval  $[a, b]$  is at most  $\epsilon$ . Otherwise, we know that the standard deviation of  $G_j$  is at least  $b-a$  which means that its mass on the set  $[a, a+\epsilon(b-a)] \cup [b-\epsilon(b-a), b]$  is at most  $O(\epsilon)$ . Thus we conclude that

$$\|\mathcal{M}(\mathcal{I} - 1_I)\|_1 \leq O(\epsilon). \quad (6)$$

Also note that by the properties in Theorem B.3

$$\|\mathcal{I}(\mathcal{M} - f)\|_1 \leq \epsilon \int_{-\infty}^{\infty} |\mathcal{M} - f| + 2 \int_a^b |\mathcal{M} - f| \leq 2(\epsilon + \|1_I(\mathcal{M} - f)\|_1). \quad (7)$$

Putting together (5, 6, 7), we conclude

$$\|\widetilde{\mathcal{M}} - f \cdot 1_I\|_1 \leq \|1_I(\mathcal{M} - f)\|_1 + \|\widetilde{\mathcal{M}} - \mathcal{M}1_I\|_1 \leq \text{poly}(\log(1/\epsilon)) (\epsilon + \|1_I(\mathcal{M} - f)\|_1)$$

and we are done.  $\blacksquare$

### C.3 Structural Properties

Lemma C.3 allows us to reconstruct the unknown GMM  $\mathcal{M}$  over certain intervals. However, it cannot be applied to an arbitrary interval (because an interval may overlap with a component that is too thin). We will now prove several structural results that will imply that there exist  $\tilde{O}(k)$  intervals for which the conditions of Lemma C.3 are satisfied (i.e. these intervals do not overlap with components that are much thinner than themselves) and such that the union of these intervals contains most of the mass of  $\mathcal{M}$ . Then, to complete the proof of Theorem 1.1, we show how to find such a set of  $\tilde{O}(k)$  intervals using a dynamic program.

First, we define a modified density function for a GMM  $\mathcal{M} = w_1 G_1 + \cdots + w_k G_k$  where we modify each component Gaussian by restricting it to its  $10l$ -significant interval (and making it 0 outside). It is clear that this modified function is close to  $\mathcal{M}$  in  $L^1$ -distance but it will be convenient to use in the analysis later on.

**Definition C.4.** For a mixture of Gaussians  $\mathcal{M} = w_1 G_1 + \cdots + w_k G_k$  where  $G_j = N(\mu_j, \sigma_j^2)$  and a parameter  $l$ , define the function  $\mathcal{M}_{\text{sig}, l}(x)$  to be, at each point  $x \in \mathbb{R}$ , equal to the weighted sum of the components  $G_j$  of  $\mathcal{M}$  such that  $x$  is in the  $10l$ -significant interval of  $G_j$ . Formally,

$$\mathcal{M}_{\text{sig}, l}(x) = \sum_{\substack{j \text{ such that} \\ |x - \mu_j| \leq 10l\sigma_j}} w_j G_j(x).$$

The following claim is immediate from the definition.

**Claim C.5.** *Let  $\epsilon > 0$  be some parameter and let  $l = \lceil \sqrt{\log 1/\epsilon} \rceil$ . Then*

$$\|\mathcal{M} - \mathcal{M}_{\text{sig},l}\|_1 \leq \epsilon$$

*Proof.* The inequality holds because the total mass of a Gaussian outside of its  $10l$ -significant interval is at most  $\epsilon$ .  $\blacksquare$

We now present our first structural result.

**Claim C.6.** *Let  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  be an arbitrary mixture of Gaussians where  $G_i = N(\mu_i, \sigma_i^2)$ . Let  $\epsilon > 0$  be some parameter and let  $l = \lceil \sqrt{\log(1/\epsilon)} \rceil$ . There exist disjoint intervals  $I_1, \dots, I_n$  with lengths, say  $t_1, \dots, t_n$ , where  $n \leq 50kl$  with the following property:*

- For each interval  $I_i$ , for all  $j \in [k]$  either the the  $10l$ -significant interval of  $G_j$  is disjoint from  $I_i$  or  $\sigma_j \geq t_i$
- We have

$$\|\mathcal{M} - (1_{I_1} + \dots + 1_{I_n})\mathcal{M}\|_1 \leq \epsilon$$

*Proof.* Sort the Gaussians by their standard deviations, WLOG  $\sigma_1 \leq \dots \leq \sigma_k$ . Now we will create several intervals  $A_1, A_2, \dots$  and we will also associate each interval with one of the Gaussians  $G_1, \dots, G_k$  which we will call its parent.

First, set  $A_1$  to be the  $10l$ -significant interval of  $G_1$ . Next, we will process the Gaussians  $G_2, \dots, G_k$  in order. For  $G_j$ , assume that the intervals we have created so far are  $A_1, \dots, A_m$  (which will be disjoint by construction). Now consider the  $10l$ -significant interval of  $G_j$ , say  $L_j$ . Note that removing the union of the intervals  $A_1, \dots, A_m$  from  $L_j$  divides  $L_j$  into several (at most  $m+1$ ) disjoint intervals. We label these intervals  $A_{m+1}, A_{m+2}, \dots$  and set all of their parents to be  $G_j$ . We then move onto  $G_{j+1}$  and repeat the above process. The following properties are immediate from the construction:

1. If the parent of  $A_i$  is  $G_j$  then the length of  $A_i$  is at most  $20l\sigma_j$
2. The union of all of the  $A_i$  whose parent is among  $G_1, \dots, G_j$  contains the  $10l$ -significant intervals of all of  $G_1, \dots, G_j$
3. If the parent of  $A_i$  is  $G_j$ , then  $A_i$  is disjoint from the  $10l$ -significant intervals of  $G_1, \dots, G_{j-1}$

Now we claim that at the end of the algorithm, the total number of intervals is at most  $2k$ . To see this, say that after processing  $G_{j-1}$ , the intervals we have created are  $A_1, \dots, A_m$ . Now consider the potential that is the number of intervals  $m$  plus the number of connected components in  $A_1 \cup \dots \cup A_m$ . This potential can increase by at most 2 when processing  $G_j$  so thus the total number of intervals at the end of the execution is at most  $2k$ . We will now assume that the intervals at the end of the execution are  $A_1, \dots, A_{2k}$  (if there are less than  $2k$  intervals then add a bunch of dummy intervals of length 0).

We now describe a post-processing step. For each of  $A_1, \dots, A_{2k}$ , if its parent is  $G_j$  then divide it into intervals of length at most  $\sigma_j$  and assign  $G_j$  as the parent of all of these intervals. By property 1, we can ensure that this creates a total of at most  $50kl$  intervals, say  $I_1, \dots, I_n$  where  $n \leq 50kl$ . We use  $t_1, \dots, t_n$  to denote their lengths. We now prove that this set of intervals satisfies the desired properties. First, note that the following properties are immediate from the construction:

1. If the parent of  $I_i$  is  $G_j$  then  $I_i$  is contained in the  $10l$ -significant interval of  $G_j$  and  $t_i \leq \sigma_j$
2. The union of all of  $I_1, \dots, I_n$  contains the  $10l$ -significant intervals of all of  $G_1, \dots, G_k$
3. If the parent of  $I_i$  is  $G_j$ , then  $I_i$  is disjoint from the  $10l$ -significant intervals of  $G_1, \dots, G_{j-1}$

The first of the desired properties is clear since by construction if the parent of  $I_i$  is  $G_j$ , then  $t_i \leq \sigma_j$  and it must be disjoint from the  $10l$ -significant intervals of all of  $G_1, \dots, G_{j-1}$  (where recall  $G_1, \dots, G_k$  are sorted in increasing order of their standard deviation). Now it remains to verify the second property. Consider the function  $\mathcal{M}_{\text{sig},l}(x)$ . Recall by Claim C.5,

$$\|\mathcal{M} - \mathcal{M}_{\text{sig},l}\|_1 \leq \epsilon \quad (8)$$

Next observe that

$$\mathcal{M}_{\text{sig},l} = (1_{I_1} + \dots + 1_{I_n})\mathcal{M}_{\text{sig},l}.$$

Combining the above, we have

$$\|\mathcal{M} - (1_{I_1} + \dots + 1_{I_n})\mathcal{M}_{\text{sig},l}\|_1 \leq \epsilon.$$

However, note that we have

$$(1_{I_1} + \dots + 1_{I_n})\mathcal{M}_{\text{sig},l} \leq (1_{I_1} + \dots + 1_{I_n})\mathcal{M} \leq \mathcal{M}$$

everywhere along the real line. Thus, we immediately get the desired inequality.  $\blacksquare$

The next structural result shows that the intervals  $I_1, \dots, I_n$  obtained in Claim C.6 are “findable” in the sense that if we draw many samples from  $\mathcal{M}$  (or from a distribution  $f$  that is close to  $\mathcal{M}$ ), then with high probability, there will be samples close to the endpoints of each of  $I_1, \dots, I_n$ . This will mean that algorithmically, it suffices to draw sufficiently many samples and then only consider intervals whose endpoints are given by a pair of samples.

**Claim C.7.** *Let  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  be an arbitrary mixture of Gaussians where  $G_i = N(\mu_i, \sigma_i^2)$ . Let  $\epsilon > 0$  be some parameter and let  $l = \lceil \sqrt{\log(1/\epsilon)} \rceil$ . Let  $f$  be a distribution. Assume we are given samples from  $f$ , say  $x_1, \dots, x_Q$  for some sufficiently large  $Q = \text{poly}(k/\epsilon)$ . Then with high probability, there exists pairs  $\{x_{a_1}, x_{b_1}\}, \dots, \{x_{a_n}, x_{b_n}\}$  such that*

- The intervals  $J_1 = [x_{a_1}, x_{b_1}], \dots, J_n = [x_{a_n}, x_{b_n}]$  are disjoint
- $n \leq 50kl$
- For each interval  $J_i$ , for all  $j \in [k]$  either the the  $10l$ -significant interval of  $G_j$  is disjoint from  $J_i$  or  $\sigma_j \geq |x_{b_i} - x_{a_i}|$
- $\|\mathcal{M} - (1_{J_1} + \dots + 1_{J_n})\mathcal{M}\|_1 \leq 4(\epsilon + \|\mathcal{M} - f\|_1).$

*Proof.* Let  $I_1, \dots, I_n$  be the intervals computed in Claim C.6 applied to the mixture  $\mathcal{M}$  and assume that their lengths are  $t_1, \dots, t_n$ . Let  $C = \lceil (k/\epsilon)^2 \rceil$ . For each interval  $I_i$ , divide it into  $C$  subintervals  $I_i^1, \dots, I_i^C$  of length  $t_i/C$  and assume that these subintervals are sorted in order. We say one of these subintervals is good if

$$\int_{I_i^j} f \geq (\epsilon/k)^{10}.$$

For an index  $i$ , let  $c_i, d_i$  be the smallest and largest index such that  $I_i^{c_i}, I_i^{d_i}$  are good respectively. Then with high probability for all  $i$ , there will be samples, say  $x_{a_i}, x_{b_i}$  in  $I_i^{c_i}$  and  $I_i^{d_i}$ . Now we will form the intervals  $J_i = [x_{a_i}, x_{b_i}]$ . The first two of the desired properties are clear. The third follows from the statement of Claim C.6. It remains to verify the last. Similar to the proof of Claim C.6, we consider the function  $\mathcal{M}_{\text{sig},l}$ . Note that

$$\begin{aligned} \int_{I_i \setminus J_i} \mathcal{M}_{\text{sig},l} &\leq \int_{I_i^1} \mathcal{M}_{\text{sig},l} + \dots + \int_{I_i^{c_i}} \mathcal{M}_{\text{sig},l} + \int_{I_i^{d_i}} \mathcal{M}_{\text{sig},l} + \dots + \int_{I_i^C} \mathcal{M}_{\text{sig},l} \\ &\leq \|1_{I_i}(\mathcal{M}_{\text{sig},l} - f)\|_1 + \int_{I_i^1} f + \dots + \int_{I_i^{c_i-1}} f + \int_{I_i^{c_i+1}} f + \dots + \int_{I_i^C} f \\ &\quad + \int_{I_i^{c_i}} \mathcal{M}_{\text{sig},l} + \int_{I_i^{d_i}} \mathcal{M}_{\text{sig},l} \\ &\leq \|1_{I_i}(\mathcal{M}_{\text{sig},l} - f)\|_1 + (\epsilon/k)^2 + \int_{I_i^{c_i}} \mathcal{M}_{\text{sig},l} + \int_{I_i^{d_i}} \mathcal{M}_{\text{sig},l} \end{aligned}$$

where the last step follows by the minimality and maximality of  $c_i, d_i$ . Now by construction, the only Gaussians among  $G_1, \dots, G_k$  whose  $10l$ -significant intervals intersect  $I_i$  must have standard deviation at least  $t_i$ . Since  $I_i^{c_i}, I_i^{d_i}$  each have length  $(\epsilon/k)^2 t_i$ , we conclude

$$\int_{I_i \setminus J_i} \mathcal{M}_{\text{sig},l} \leq 3(\epsilon/k)^2 + \|1_{I_i}(\mathcal{M}_{\text{sig},l} - f)\|_1.$$

Thus, we have

$$\|(1_{I_1} + \dots + 1_{I_n})\mathcal{M}_{\text{sig},l} - (1_{J_1} + \dots + 1_{J_n})\mathcal{M}_{\text{sig},l}\|_1 \leq \epsilon + \|(\mathcal{M}_{\text{sig},l} - f)\|_1 \leq 2\epsilon + \|(\mathcal{M} - f)\|_1.$$

where we used Claim C.5 in the last step. Now, using the statement of Claim C.6, we deduce

$$\|\mathcal{M} - (1_{J_1} + \dots + 1_{J_n})\mathcal{M}_{\text{sig},l}\|_1 \leq 4\epsilon + \|(\mathcal{M} - f)\|_1.$$

Since

$$(1_{J_1} + \dots + 1_{J_n})\mathcal{M}_{\text{sig},l} \leq (1_{J_1} + \dots + 1_{J_n})\mathcal{M} \leq \mathcal{M}$$

everywhere along the real line, we immediately get the desired inequality.  $\blacksquare$

#### C.4 Finishing the Proof

We can now prove the key lemma and then Theorem 1.1 will follow as an immediate consequence since we can use the improper learner in Corollary 2.10. The lemma states that given explicit access to a distribution  $f$  that is  $\epsilon$ -close to a GMM,  $\mathcal{M}$ , with  $k$  components, we can output a GMM,  $\widetilde{\mathcal{M}}$ , with  $\widetilde{O}(k)$  components that is  $\widetilde{O}(\epsilon)$ -close to  $f$ . At a high-level, the proof involves attempting to reconstruct  $f$  over various intervals using Lemma C.3 and then using a dynamic program to find a union of  $\widetilde{O}(k)$  such intervals that approximates the entire function. We use Claim C.7 to argue that such a solution exists so our dynamic program must find it.

**Lemma C.8.** *Let  $\mathcal{M} = w_1 G_1 + \dots + w_k G_k$  be an arbitrary mixture of Gaussians where  $G_i = N(\mu_i, \sigma_i^2)$ . Let  $\epsilon > 0$  be some parameter. Assume we are given access to a distribution  $f$ . Then in  $\text{poly}(k/\epsilon)$  time and with high probability, we can compute a mixture  $\widetilde{\mathcal{M}}$  of at most  $k \text{poly}(\log(k/\epsilon))$  Gaussians such that*

$$\|\widetilde{\mathcal{M}} - f\|_1 \leq \text{poly}(\log(k/\epsilon))(\epsilon + \|\mathcal{M} - f\|_1).$$

*Proof.* Note that it suffices to compute a weighted sum of Gaussians that satisfies the desired inequality since rescaling such a weighted sum to a mixture will at most increase the  $L^1$  error by a factor of 2. Thus, from now on, we will not worry about ensuring the mixing weights sum to 1.

Let  $\gamma = \epsilon/k$  and  $l = \lceil \sqrt{\log 1/\gamma} \rceil$ . First draw  $Q = \text{poly}(k/\epsilon)$  samples  $x_1, \dots, x_Q$  from  $f$  for sufficiently large  $Q$  that we can apply Claim C.7 with  $\epsilon \leftarrow \gamma$ . While we do not know what the intervals  $J_1, \dots, J_n$  are, we will set up a dynamic program to find a set of at most  $50kl$  intervals that we can reconstruct  $f$  over each one using Lemma C.3 and such that these intervals contain essentially all of the mass of  $f$ .

For each pair  $x_a, x_b$  with  $a, b \in \{1, 2, \dots, Q\}$ , apply Lemma C.3 with parameter  $\epsilon \leftarrow \gamma$  to attempt to approximate  $f$  on the interval  $[x_a, x_b]$ . Let the output obtained be  $\widetilde{\mathcal{M}}_{x_a, x_b}$ . Note that sometimes the algorithm will fail to output a good approximation to  $f$  (because the assumptions of Lemma C.3 fail) but we can ensure that the output is a weighted sum of at most  $\text{poly}(\log 1/\gamma)$  Gaussians. In the proof we will only use the fact that when the assumptions of Lemma C.3 hold, then our approximation to  $f$  restricted to the interval will be accurate.

Now we show how to set up the dynamic program. WLOG the points  $x_1, \dots, x_Q$  are sorted in nondecreasing order. We also use the convention that  $x_0 = -\infty, x_{Q+1} = \infty$ . Now, we maintain the following state for each index  $0 \leq j \leq Q+1$ , and integer  $c \leq 50kl$ : the best approximation to  $f$  from  $(-\infty, x_j]$  using a sum of  $\widetilde{\mathcal{M}}_{x_a, x_b}$  over at most  $c$  intervals. Formally,

**Dynamic Program:** Let  $DP_{j,c}$  be the minimum over all sets  $S$  of  $c$  pairs  $(a^{(1)}, b^{(1)}), \dots, (a^{(c)}, b^{(c)}) \in [j] \times [j]$  such that  $a^{(1)} < b^{(1)} \leq a^{(2)} < \dots < b^{(c)}$  of

$$\begin{aligned} & \|f \cdot 1_{(-\infty, x_j]} - \left( 1_{[x_{a^{(1)}}, x_{b^{(1)}}, \dots, 1_{[x_{a^{(c)}}, x_{b^{(c)}}, \widetilde{\mathcal{M}}_{x_{a^{(1)}}, x_{b^{(1)}}, \dots, \widetilde{\mathcal{M}}_{x_{a^{(c)}}, x_{b^{(c)}}} \right) \|_1 \\ & + \|\widetilde{\mathcal{M}}_{x_{a^{(1)}}, x_{b^{(1)}}} - 1_{[x_{a^{(1)}}, x_{b^{(1)}}, \widetilde{\mathcal{M}}_{x_{a^{(1)}}, x_{b^{(1)}}} \|_1 + \dots + \|\widetilde{\mathcal{M}}_{x_{a^{(c)}}, x_{b^{(c)}}} - 1_{[x_{a^{(c)}}, x_{b^{(c)}}, \widetilde{\mathcal{M}}_{x_{a^{(c)}}, x_{b^{(c)}}} \|_1. \end{aligned}$$

Note that the first term represents the approximation error compared to  $f$ . We must truncate each function  $\widetilde{\mathcal{M}}_{x_{a(1)}, x_{b(1)}}$  to its corresponding interval  $[x_{a(1)}, x_{b(1)}]$  in order for the problem to be solvable via dynamic programming because otherwise previous choices would affect later ones. Thus, we also need to add the additional terms that represent the error from truncation. Note that the  $L^1$  distance can be estimated using Claim 2.11. We can solve the dynamic program in polynomial time because from each state  $DP_{j,c}$ , we simply consider adding all possible intervals among  $x_j, x_{j+1}, \dots, x_Q$  as the next one.

Now we prove that there is a good solution to the dynamic program for which the objective (for  $j = Q + 1$ ) is small. Let  $a_1, b_1, \dots, a_n, b_n$  be the indices obtained in Claim C.7. We claim that setting

$$(a^{(1)}, b^{(1)}) = (a_1, b_1), \dots, (a^{(n)}, b^{(n)}) = (a_n, b_n)$$

results in the objective function being small. Let  $J_i = [a_i, b_i]$  for all  $i$ . Let

$$\widetilde{\mathcal{M}}^{\text{good}} = 1_{[x_{a_1}, x_{b_1}]} \widetilde{\mathcal{M}}_{x_{a_1}, x_{b_1}} + \dots + 1_{[x_{a_n}, x_{b_n}]} \widetilde{\mathcal{M}}_{x_{a_n}, x_{b_n}}.$$

The guarantee from Lemma C.3 implies that for all  $i$ ,

$$\|\widetilde{\mathcal{M}}_{x_{a_i}, x_{b_i}} - 1_{[x_{a_i}, x_{b_i}]} \widetilde{\mathcal{M}}_{x_{a_i}, x_{b_i}}\|_1 \leq \text{poly}(\log 1/\gamma)(\gamma + \|1_{J_i}(\mathcal{M} - f)\|_1) \quad (9)$$

and thus, using the guarantee from Lemma C.3 again, we have

$$\begin{aligned} \|\widetilde{\mathcal{M}}^{\text{good}} - (1_{J_1} + \dots + 1_{J_n})f\| &\leq \|\widetilde{\mathcal{M}}_{x_{a_1}, x_{b_1}} - 1_{J_1} \cdot f\|_1 + \dots + \|\widetilde{\mathcal{M}}_{x_{a_n}, x_{b_n}} - 1_{J_n} \cdot f\|_1 \\ &\quad + \|\widetilde{\mathcal{M}}_{x_{a_1}, x_{b_1}} - 1_{[x_{a_1}, x_{b_1}]} \widetilde{\mathcal{M}}_{x_{a_1}, x_{b_1}}\|_1 + \dots + \|\widetilde{\mathcal{M}}_{x_{a_n}, x_{b_n}} - 1_{[x_{a_n}, x_{b_n}]} \widetilde{\mathcal{M}}_{x_{a_n}, x_{b_n}}\|_1 \\ &\leq n\gamma \text{poly}(\log 1/\gamma) + \text{poly}(\log 1/\gamma) (\|1_{J_1}(\mathcal{M} - f)\|_1 + \dots + \|1_{J_n}(\mathcal{M} - f)\|_1) \\ &\leq \text{poly}(\log 1/\gamma) (\epsilon + \|\mathcal{M} - f\|_1). \end{aligned}$$

However also recall that by Claim C.7,

$$\|\mathcal{M} - (1_{J_1} + \dots + 1_{J_n})\mathcal{M}\|_1 \leq 4(\gamma + \|\mathcal{M} - f\|_1).$$

Combining the above two inequalities, we get

$$\|\widetilde{\mathcal{M}}^{\text{good}} - f\|_1 \leq \text{poly}(\log 1/\gamma) (\epsilon + \|\mathcal{M} - f\|_1).$$

Finally, combining the above with (9) implies that the objective value of the dynamic program is at most  $\text{poly}(\log 1/\gamma) (\epsilon + \|\mathcal{M} - f\|_1)$ . Finally it remains to note that the objective of the dynamic program (for  $j = Q + 1$ ) is an upper bound on

$$\|f - \left( \widetilde{\mathcal{M}}_{x_{a(1)}, x_{b(1)}} + \dots + \widetilde{\mathcal{M}}_{x_{a(c)}, x_{b(c)}} \right)\|_1$$

so thus, we can simply output the solution  $\widetilde{\mathcal{M}} = \widetilde{\mathcal{M}}_{x_{a(1)}, x_{b(1)}} + \dots + \widetilde{\mathcal{M}}_{x_{a(c)}, x_{b(c)}}$  and we are guaranteed to have

$$\|f - \widetilde{\mathcal{M}}\|_1 \leq \text{poly}(\log 1/\gamma) (\epsilon + \|\mathcal{M} - f\|_1).$$

It is clear that  $\widetilde{\mathcal{M}}$  is a weighted sum of at most  $k\text{poly}(\log 1/\gamma)$  Gaussians (since for each interval there are  $\text{poly}(\log 1/\gamma)$  Gaussians and there are at most  $50kl$  total intervals). It is clear that all of the steps run in  $\text{poly}(k/\epsilon)$  time and we are done.  $\blacksquare$

Now we can complete the proof of our main theorem, Theorem 1.1.

*Proof of Theorem 1.1.* We can apply Corollary 2.10 to learn a distribution  $f$  such that  $d_{\text{TV}}(\mathcal{M}, f) \leq O(\epsilon)$ . We can then apply Lemma C.8 using  $f$ . Note that since  $f$  is a piecewise polynomial, we can perform all of the explicit computations with the density function that are used in the proof of Lemma C.8. It is immediate that the output of Lemma C.8 must satisfy

$$d_{\text{TV}}(\widetilde{\mathcal{M}}, \mathcal{M}) \leq \widetilde{O}(\epsilon)$$

so we are done.  $\blacksquare$

## D Well-Conditioned Case: Sparse Fourier Reconstruction

We now move onto our results on sparse Fourier reconstruction. As with GMMs, we will first consider the well-conditioned case. Here, this means a function that has its Fourier support contained in one interval that is not too long i.e., all of its Fourier mass is not too spread out. Note that WLOG, we may assume that this interval is centered at 0 since otherwise we can multiply by a suitable exponential to shift the Fourier support to be around 0. We prove the following statement:

**Lemma D.1.** *Let  $0 < \epsilon < 0.1$  be a parameter and let  $l \geq \lceil \log 1/\epsilon \rceil$  be some parameter. Let  $\mathcal{M}$  be a function such that  $\widehat{\mathcal{M}}$  is supported on  $[-l, l]$  and such that  $\|\widehat{\mathcal{M}}\|_1 \leq 1$ . Also assume that we have access to a function  $f$  such that*

$$\int_{-1}^1 |f(x) - \mathcal{M}(x)|^2 dx \leq \epsilon^2.$$

*There is an algorithm that runs in  $\text{poly}(l)$  time and outputs a function  $\widetilde{\mathcal{M}}$  such that  $\widetilde{\mathcal{M}}$  is  $(O(l), O(l))$ -simple, has Fourier support contained in  $[-l, l]$ , and*

$$\int_{-1}^1 |\widetilde{\mathcal{M}}(x) - f(x)|^2 dx \leq 16\epsilon^2.$$

**Remark.** *Note that in this case, we do not need any constraint on the Fourier sparsity of  $\mathcal{M}$  to guarantee that the output of our algorithm is  $O(\log 1/\epsilon)$ -Fourier sparse. Also, unlike our full result, Theorem 1.3, our output in this case is guaranteed to be a good approximation over the entire interval (instead of a subinterval).*

Our proof will be separated into two parts. The first step will be proving the existence of a function  $\widetilde{\mathcal{M}}$  of the desired form. The second step will be developing an algorithm to actually compute it.

### D.0.1 Existence of a Sparse Approximation

First, we will prove that under the assumptions of Lemma D.1, an approximation  $\widetilde{\mathcal{M}}$  satisfying the desired properties exists. We will also prove that independent of the problem instance, it suffices to only consider a fixed set of  $O(l)$  distinct frequencies given by the Chebyshev points (with suitable rescaling).

The proof relies on first taking the Taylor series of an exponential  $e^{2\pi i \zeta x}$  and arguing that we only need to keep the first  $O(\log 1/\epsilon)$  terms. This essentially lets us represent such an exponential with the coefficients of its Taylor series, which are (up to rescaling)  $(1, \zeta, \zeta^2, \dots)$ . We then use Corollary F.7 to argue that an arbitrary linear combination of such vectors can be replaced with a sparse combination with similarly sized coefficients.

**Lemma D.2.** *Let  $0 < \epsilon < 0.1$  be a parameter and let  $l \geq \lceil \log 1/\epsilon \rceil$  be some parameter. Let  $t_0, \dots, t_{10^2 l}$  be the degree- $10^2 l$  Chebyshev points. Let  $\mathcal{M}$  be a function such that  $\widehat{\mathcal{M}}$  is supported on  $[-l, l]$  and such that  $\|\widehat{\mathcal{M}}\|_1 \leq 1$ . Then there is a function*

$$h(x) = \sum_{j=0}^{10^2 l} c_j e^{2\pi i (lt_j)x}$$

*where  $c_0, \dots, c_{10^2 l}$  are complex numbers such that  $\sum_{j=0}^{10^2 l} |c_j| \leq 200l$  and*

$$\int_{-1}^1 (h(x) - \mathcal{M}(x))^2 dx \leq \epsilon^2.$$

*Proof.* Note that  $\mathcal{M}$  can be written as  $\mathcal{M}(x) = \int_{-l}^l \widehat{\mathcal{M}}(\zeta) e^{2\pi i x \zeta} d\zeta$ . Now consider the Taylor expansion of

$$e^{2\pi i x \zeta} = \sum_{j=0}^{\infty} \frac{(2\pi i x \zeta)^j}{j!}$$

Note that since  $-l \leq \zeta \leq l$ , we have

$$\sum_{j=10^2l+1}^{\infty} \left| \frac{(2\pi i \zeta)^j}{j!} \right| \leq (\epsilon/l)^3.$$

In particular, if we define

$$g_{\zeta}(x) = \sum_{j=0}^{10^2l} \frac{(2\pi i x \zeta)^j}{j!}$$

then over the interval  $x \in [-1, 1]$

$$|e^{2\pi i x \zeta} - g_{\zeta}(x)| \leq (\epsilon/l)^3. \quad (10)$$

Next, for each  $\zeta \in [-l, l]$ , by Corollary F.7, we can write the vector

$$\mathcal{V}_{10^2l}(\zeta) = w_0(\zeta) \mathcal{V}_{10^2l}(t_0 l) + \cdots + w_{10^2l}(\zeta) \mathcal{V}_{10^2l}(t_{10^2l} l)$$

for some real numbers (depending on  $\zeta$ )  $w_0(\zeta), \dots, w_{10^2l}(\zeta)$  with  $\sum |w_j(\zeta)| \leq 200l$ . Thus,

$$g_{\zeta}(x) = w_0(\zeta) g_{t_0 l}(x) + \cdots + w_{10^2l}(\zeta) g_{t_{10^2l} l}(x)$$

for the same weights. Now note that by (10), for all  $x \in [-1, 1]$ ,

$$\left| \mathcal{M}(x) - \sum_{j=0}^{10^2l} g_{t_j l}(x) \left( \int_{-l}^l \widehat{\mathcal{M}}(\zeta) w_j(\zeta) d\zeta \right) \right| = \left| \mathcal{M}(x) - \int_{-l}^l \widehat{\mathcal{M}}(\zeta) g_{\zeta}(x) d\zeta \right| \leq 2l \cdot (\epsilon/l)^3.$$

Note that

$$\sum_{j=0}^{10^2l} \left| \int_{-l}^l \widehat{\mathcal{M}}(\zeta) w_j(\zeta) d\zeta \right| \leq 200l$$

so by (10), for all  $x \in [-1, 1]$ ,

$$\left| \sum_{j=0}^{10^2l} \left( g_{t_j l}(x) - e^{2\pi i x (t_j l)} \right) \left( \int_{-l}^l \widehat{\mathcal{M}}(\zeta) w_j(\zeta) d\zeta \right) \right| \leq 200l (\epsilon/l)^3$$

so therefore for all  $x \in [-1, 1]$ , we have

$$\left| \mathcal{M}(x) - \sum_{j=0}^{10^2l} e^{2\pi i x (t_j l)} \left( \int_{-l}^l \widehat{\mathcal{M}}(\zeta) w_j(\zeta) d\zeta \right) \right| \leq 202l \cdot (\epsilon/l)^3$$

and setting

$$h(x) = \sum_{j=0}^{10^2l} e^{2\pi i x (t_j l)} \left( \int_{-l}^l \widehat{\mathcal{M}}(\zeta) w_j(\zeta) d\zeta \right)$$

immediately leads to the desired conclusion. ■

### D.0.2 Completing the Proof of Lemma D.1

By combining Lemma D.2 and Lemma F.10, we can complete the proof of Lemma D.1.

*Proof of Lemma D.1.* We can separate  $f$  into its real and imaginary parts, say  $f_{\text{re}}, f_{\text{im}}$  and we can separate  $\mathcal{M}$  into its real and imaginary parts  $\mathcal{M}_{\text{re}}, \mathcal{M}_{\text{im}}$ . Now consider the Chebyshev points of degree  $10^2l$ , say  $t_0, \dots, t_{10^2l}$ . We will now apply Lemma F.10 where we consider the set of functions

$$\{f_1, \dots, f_n\} = \{\pm \cos(2\pi t_0 x), \pm \sin(2\pi t_0 x), \dots, \pm \cos(2\pi t_{10^2l} x), \pm \sin(2\pi t_{10^2l} x)\}.$$

The distribution  $\mathcal{D}$  is the uniform distribution on  $[-1, 1]$  and by Lemma D.2, there are coefficients  $a_1, \dots, a_n \geq 0$  with  $a_1 + \dots + a_n \leq O(l)$  (note we can split the complex coefficients  $c_j$  into their real and imaginary parts and then split into positive and negative parts) such that

$$\begin{aligned} \int_{-1}^1 (f_{\text{re}}(x) - (a_1 f_1(x) + \dots + a_n f_n(x)))^2 dx &\leq 2 \int_{-1}^1 (\mathcal{M}_{\text{re}}(x) - (a_1 f_1(x) + \dots + a_n f_n(x)))^2 dx \\ &\quad + 2 \int_{-1}^1 (f_{\text{re}}(x) - \mathcal{M}_{\text{re}}(x))^2 dx \leq 4\epsilon^2 \end{aligned}$$

and similar for the imaginary part of  $f$ . Applying Lemma F.10 to both the real and imaginary part (after rescaling by  $1/(O(l))$ ), adding the results, and rewriting the trigonometric functions using complex exponentials (note the set  $\{t_0, \dots, t_{10^{2l}}\}$  is symmetric around 0 so we can do this) completes the proof.  $\blacksquare$

We can slightly extend Lemma D.1 to work even if we do not know the desired accuracy  $\epsilon$  but only a lower bound on it. It suffices to run the algorithm for Lemma D.1 and repeatedly decrease the target accuracy until our algorithm fails to find the optimal accuracy within a constant factor.

**Corollary D.3.** *Let  $l, \epsilon$  be parameters given to us such that  $l \geq \lceil \log 1/\epsilon \rceil$ . Let  $\mathcal{M}$  be a function such that  $\widehat{\mathcal{M}}$  is supported on  $[-l, l]$  and  $\|\widehat{\mathcal{M}}\|_1 \leq 1$ . Assume that we have access to a function  $g$  defined on  $[-1, 1]$ . There is an algorithm that runs in  $\text{poly}(l)$  time and outputs a function  $\widetilde{\mathcal{M}}$  such that  $\widetilde{\mathcal{M}}$  is  $(O(l), O(l))$ -simple, has Fourier support contained in  $[-l, l]$ , and*

$$\int_{-1}^1 |\widetilde{\mathcal{M}}(x) - f(x)|^2 dx \leq 20 \left( \epsilon^2 + \int_{-1}^1 |f(x) - \mathcal{M}(x)|^2 dx \right).$$

*Proof.* For a target accuracy  $\gamma > \epsilon$  we run the algorithm in Lemma D.1 to get a function  $\widetilde{\mathcal{M}}_\gamma(x)$ . We then check whether

$$\int_{-1}^1 |\widetilde{\mathcal{M}}_\gamma(x) - f(x)|^2 dx \leq 16\gamma^2.$$

Note that the above can be explicitly computed. If the above check passes, we then take  $\gamma \leftarrow 0.99\gamma$ . Taking the smallest  $\gamma$  for which the above succeeds, the guarantee from Lemma D.1 ensures that we have a function  $\widetilde{\mathcal{M}}$  such that

$$\int_{-1}^1 |\widetilde{\mathcal{M}}(x) - f(x)|^2 dx \leq 20 \left( \epsilon^2 + \int_{-1}^1 |f(x) - \mathcal{M}(x)|^2 dx \right).$$

It is clear that we run the routine from Lemma D.1 at most  $O(l)$  times so we are done.  $\blacksquare$

## E Sparse Fourier Reconstruction: Full Version

In this section, we complete the proof of our main result on sparse Fourier reconstruction, Theorem 1.3. The high-level outline of the proof is similar to the proof of Theorem 1.1. The key lemma that goes into the proof is stated below. At a high level, the lemma states that if we know roughly where the Fourier support of the unknown Fourier-sparse signal  $\mathcal{M}$  is located, then we can successfully reconstruct it.

**Lemma E.1.** *Assume we are given  $N, k, \epsilon, c$  with  $0 < \epsilon < 0.1$ . Let  $l = \lceil \log kN/(\epsilon c) \rceil$  be some parameter. Let  $\mathcal{M}$  be a function that is  $(k, 1)$ -simple. Also, assume that we are given a set  $T \subset \mathbb{R}$  of size  $N$  such that all of the support of  $\widehat{\mathcal{M}}$  is within distance 1 of  $N$ . Further, assume we are given access to a function  $f$  such that*

$$\int_{-1}^1 |f(x) - \mathcal{M}(x)|^2 dx \leq \epsilon^2.$$

*There is an algorithm that runs in  $\text{poly}(N, k, l, 1/c)$  time and outputs a function  $\widetilde{\mathcal{M}}$  that is  $(k\text{poly}(l/c), k\text{poly}(l/c))$ -simple and*

$$\int_{-1+c}^{1-c} |\widetilde{\mathcal{M}}(x) - f(x)|^2 dx \leq \epsilon^2 \text{poly}(l).$$

The proof of Lemma E.1 will involve localizing the frequencies and then using Corollary D.3 to reconstruct after localizing. We will do this for  $\text{poly}(N, k, \log 1/\epsilon)$  different localizations (based on the set  $T$  that we are given). We will then select at most  $\tilde{O}(k)$  of these localized reconstructions to add together and output. The intuition behind why we can find such a set of  $\tilde{O}(k)$  localized reconstructions and ignore the rest is that  $\mathcal{M}$  is  $k$ -Fourier sparse so localizations that are far away from the frequencies of  $\mathcal{M}$  can essentially be ignored.

The localization procedure will involve convolving  $f$  by a Gaussian times an exponential (technically we will convolve by a function that approximates a Gaussian times an exponential). Note that this is equivalent to multiplying the Fourier transform by a Gaussian multiplier. This will ensure that frequencies too far away from a certain target frequency will only contribute negligibly and we only need to worry about reconstructing the frequencies that are close to the target frequency.

### E.1 Properties of Localization

In this section, we formalize the localization step and prove several inequalities that will be used in the proof of Lemma E.1. The way we would like to localize the frequencies is by multiplying by a Gaussian multiplier in Fourier space since afterwards, we would be able to essentially neglect any frequencies that are far away from the center of the Gaussian multiplier. This is equivalent to convolving by a Gaussian times an exponential, i.e. a function of the form  $G(x)e^{2\pi i \theta x}$ , in real space. For technical reasons, we will actually define two types of functions, which we call kernels, that approximate functions of the form  $G(x)e^{2\pi i \theta x}$ . The reason that we will need to work with both is that the first type can be computed efficiently while the second is easier to use in the analysis of our algorithm.

We begin with a few definitions.

**Definition E.2.** For a function  $f : \mathbb{R} \rightarrow \mathbb{C}$  and any  $l > 0$ , define  $f^{\text{trunc}(l)}$  to be the function that is equal to  $f$  on  $[-l, l]$  and 0 otherwise.

**Definition E.3.** For parameters  $\mu, l, c$ , we define the function

$$\mathcal{K}_{\mu, l, c} = (1/c)\mathcal{P}_l^{\text{trunc}(l)}(x/c) \cdot e^{2\pi i \mu x}$$

where  $\mathcal{P}_l$  is as defined in Definition F.9.

**Remark.** We call functions of the above form truncated polynomial kernels. The fact that such functions are truncated polynomials will make it easy to explicitly compute convolutions.

**Definition E.4.** For parameters  $\mu, l, a$ , we define the function  $\mathcal{T}_{\mu, l, a}$  as follows. First define

$$\mathcal{S}_{l, a} = M_{0, 1}^{\text{trunc}(l)}(x/a)$$

(recall Definition B.1) and then define

$$\mathcal{T}_{\mu, l, a}(x) = \widehat{\mathcal{S}_{l, a}}(x)e^{2\pi i \mu x}.$$

**Remark.** We call functions of the above form truncated Gaussian kernels. Note that truncated Gaussian kernels are compactly supported in Fourier space, which will be a convenient property in the analysis of our algorithm later on.

Note that both  $\mathcal{K}_{0, l, c}$  and  $\mathcal{T}_{0, l, 2\pi/c}$  are meant to approximate the Gaussian  $N(0, c^2)$ . The fact that  $\mathcal{K}_{0, l, c}$  approximates  $N(0, c^2)$  is clear from the definition (and Lemma F.8). To see why  $\mathcal{T}_{0, l, 2\pi/c}$  approximates the Gaussian, note that if in the definition of  $\mathcal{T}$ , we did not truncate before taking the Fourier transform, then we would get exactly  $N(0, c^2)$ .

We will now prove several inequalities relating to how convolving with the kernels  $\mathcal{K}$  and  $\mathcal{T}$  affect a function. The first set of bounds are an immediate consequence of Lemma F.8.

**Claim E.5.** Let  $l, \epsilon > 0$  be parameters such that  $l \geq \lceil \log 1/\epsilon \rceil$ . Let  $0 < c < 1$  be some constant. Then

$$\begin{aligned} \|\mathcal{K}_{0, l, c}(x) - N(0, c^2)(x)\|_1 &\leq O(\epsilon l) \\ \|\mathcal{K}_{0, l, c}(x) - N(0, c^2)(x)\|_2^2 &\leq O(l\epsilon^2/c). \end{aligned}$$

*Proof.* We know by Lemma F.8 that

$$|(1/c)\mathcal{P}_l(x/c) - N(0, c^2)(x)| \leq \epsilon/c$$

for all  $x \in [-lc, lc]$ . Thus, since  $G$  decays rapidly outside the interval  $[-lc, lc]$  we have

$$\|\mathcal{K}_{0,l,c}(x) - N(0, c^2)(x)\|_1 \leq O(\epsilon l).$$

The second inequality follows by a similar argument.  $\blacksquare$

The next claim formalizes the intuition that  $\mathcal{K}_{\mu,l,c}$  and  $\mathcal{T}_{\mu,l,2\pi/c}$  must be close because they both approximate the same function (the function  $N(0, c^2)e^{2\pi i \mu x}$ ).

**Claim E.6.** *Let  $l$  be some parameter and let  $\epsilon > 0$  be such that  $l \geq \lceil \log 1/\epsilon \rceil$ . Let  $g$  be a function such that  $\|\widehat{g}\|_1 \leq 1$ . Then for any  $\mu, c$ ,*

$$\|\mathcal{K}_{\mu,l,c} * g - \mathcal{T}_{\mu,l,2\pi/c} * g\|_\infty \leq O(\epsilon l),$$

where  $*$  denotes convolution.

*Proof.* First, note that it suffices to prove the above for  $\mu = 0$ . Now let  $G = N(0, c^2)$ . By Claim E.5, we have

$$\|\mathcal{K}_{0,l,c}(x) - G(x)\|_1 \leq O(\epsilon l).$$

Since  $\|\widehat{g}\|_1 \leq 1$ , we know that  $\|g\|_\infty \leq 1$  and thus for all  $x$ ,

$$|\mathcal{K}_{0,l,c} * g(x) - G * g(x)| \leq O(\epsilon l). \quad (11)$$

Finally, observe that the Fourier transform of  $G * g$  is equal to  $\widehat{G}\widehat{g}$ . Note that

$$\widehat{G}(x) = M_{0,(2\pi/c)^2}(x) = M_{0,1}(cx/(2\pi)).$$

By construction,

$$\widehat{\mathcal{T}_{0,l,2\pi/c} * g} = M_{0,1}^{\text{trunc}(l)}(cx/(2\pi))\widehat{g}$$

which is equal to  $\widehat{G}\widehat{g}$  restricted to the interval  $[-2\pi l/c, 2\pi l/c]$ . Using the fact that  $\widehat{G} = M_{0,(2\pi/c)^2}$  decays rapidly outside  $[-2\pi l/c, 2\pi l/c]$ , we have that

$$\|\widehat{\mathcal{T}_{0,l,2\pi/c} * g} - \widehat{G}\widehat{g}\|_1 \leq \epsilon.$$

Thus,  $|\mathcal{T}_{0,l,2\pi/c} * g(x) - G * g(x)| \leq \epsilon$  for all  $x$  and combining with (11), we are done.  $\blacksquare$

## E.2 Decoupling

In our full algorithm, we will reconstruct frequency-localized versions of a function independently for different frequencies  $\theta$  that we localize around. We will then combine our localized reconstructions by adding them. In this section, we prove several inequalities that will allow us to analyze what happens to our estimation error when we add different localized reconstructions together. Recall that convolving by a truncated polynomial kernel  $\mathcal{K}_{\mu,l,c}$  or truncated Gaussian kernel  $\mathcal{T}_{\mu,l,a}$  is approximately equivalent to multiplying the Fourier transform by a Gaussian multiplier centered around  $\mu$ . Lemma B.4 implies that adding up evenly spaced Gaussian multipliers approximates the constant function. Thus, we expect that convolving by an expression of the form  $\sum_\mu \mathcal{K}_{\mu,l,c}$  or  $\sum_\mu \mathcal{T}_{\mu,l,a}$  where the sum is over evenly spaced  $\mu$  should roughly recover the original function. The first two claims here formalize this intuition.

In the first claim, we analyze what happens when we add several localizations obtained by convolving with various truncated polynomial kernels.

**Claim E.7.** *Let  $l$  be some parameter and let  $0 < \epsilon < 0.1$  be such that  $l \geq \lceil \log 1/\epsilon \rceil$ . Let  $g$  be a function. Let  $c$  be some constant. Let  $S$  be a set of integer multiples of  $2\pi/(cl)$ . Then*

$$\int_{-1}^1 \left| \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{K}_{\mu,l,c} * g \right|^2 \leq (1 + O(\epsilon|S|))^2 \int_{-(1+lc)}^{1+lc} |g|^2,$$

*Proof.* Note that  $\mathcal{K}_{\mu,l,c}$  is supported on the interval  $[-lc, lc]$  so we can restrict  $g$  to be supported on  $[-(1+lc), 1+lc]$  and 0 outside the interval. Define the function

$$V(x) = \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{K}_{\mu,l,c}.$$

Then

$$\widehat{V}(x) = \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} \widehat{\mathcal{K}_{0,l,c}}(x - \mu).$$

Next, let  $G$  denote the Gaussian  $G = N(0, c^2)$ . Recall by Claim E.5,  $\|\mathcal{K}_{0,l,c}(x) - G(x)\|_1 \leq O(\epsilon l)$  so  $\|\widehat{\mathcal{K}_{0,l,c}} - \widehat{G}\|_\infty \leq O(\epsilon l)$ . Let

$$U(x) = \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} \widehat{G}(x - \mu).$$

Note that since  $\widehat{G} = M_{0,(2\pi/c)^2}$ , for any  $x$ ,

$$|U(x)| = \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} M_{0,(2\pi/c)^2}(x - \mu) \leq \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in (2\pi)/(cl)\mathbb{Z}} M_{0,(2\pi/c)^2}(x - \mu) \leq 1 + \epsilon$$

where the last inequality follows from Lemma B.4. Also  $\|\widehat{K}_{0,l,c} - \widehat{G}\|_\infty \leq O(\epsilon l)$  so for all  $x$ ,

$$|\widehat{V}(x)| \leq 1 + O(\epsilon|S|).$$

We conclude

$$\int_{-1}^1 |V * g(x)|^2 dx \leq \|V * g\|_2^2 = \|\widehat{V}\widehat{g}\|_2^2 \leq (1 + O(\epsilon|S|))^2 \|\widehat{g}\|_2^2 = (1 + O(\epsilon|S|))^2 \|g\|_2^2$$

To complete the proof recall that we restricted  $g$  to be supported on  $[-(1+lc), 1+lc]$  and we are done.  $\blacksquare$

The next claim is similar to the previous one except we analyze what happens when we add several localizations obtained by convolving with various truncated Gaussian kernels.

**Claim E.8.** *Let  $a, l, \epsilon$  be parameters such that  $0 < \epsilon < 0.1$  and  $l \geq \lceil \log 1/\epsilon \rceil$ . Let  $g$  be a function whose Fourier support is contained in a set  $S_0 \subset \mathbb{R}$  and such that  $\|\widehat{g}\|_1 \leq 1$ . Let  $S$  be a set of integer multiples of  $a/l$  that contains all multiples within a distance  $l \cdot a$  of  $S_0$ . Then*

$$\|g - \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{T}_{\mu,l,a} * g\|_\infty \leq 2\epsilon.$$

*Proof.* Consider the function

$$A(x) = \sum_{\mu \in S} \frac{1}{l\sqrt{2\pi}} M_{\mu,a^2}(x).$$

For all  $x \in S_0$ , we claim that

$$|A(x) - 1| \leq \epsilon.$$

To see this, note that by Lemma B.4,

$$\left| \sum_{\mu \in (a/l)\mathbb{Z}} \frac{1}{l\sqrt{2\pi}} M_{\mu,a^2}(x) - 1 \right| \leq 0.1\epsilon$$

for all  $x$ . By assumption, the set  $S$  contains all integers that are within  $la$  of the set  $S_0$  so for any  $x \in S_0$ ,

$$\sum_{\mu \in (a/l)\mathbb{Z} \setminus S} \frac{1}{l\sqrt{2\pi}} M_{\mu,a^2}(x) \leq 0.1\epsilon,$$

and we conclude that we must have  $|A(x) - 1| \leq \epsilon$ . Next, we claim that if we define

$$B(x) = \sum_{\mu \in S} \frac{1}{l\sqrt{2\pi}} M_{0,1}^{\text{trunc}(l)}((x - \mu)/a),$$

then we have for all  $x$ ,

$$|B(x) - A(x)| \leq \epsilon.$$

To see this, first note that  $M_{0,1}((x - \mu)/a) = M_{\mu,a^2}(x)$ . Next, using Gaussian tail decay, we have for all  $x$ ,

$$|B(x) - A(x)| \leq \sum_{\substack{\mu \in (a/l)\mathbb{Z} \\ |x - \mu| \geq la}} \frac{1}{l\sqrt{2\pi}} M_{\mu,a^2}(x) \leq \epsilon.$$

Thus, we have

$$|B(x) - 1| \leq 2\epsilon$$

for  $x \in S_0$ . Note that by definition,

$$\widehat{g}(x) - \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} \widehat{\mathcal{T}_{\mu,l,a} * g}(x) = \widehat{g}(x) - \widehat{g}(x) \sum_{\mu \in S} \frac{1}{l\sqrt{2\pi}} M_{0,1}^{\text{trunc}(l)}((x - \mu)/a) = (1 - B(x))\widehat{g}(x)$$

so therefore

$$\|\widehat{g} - \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} \widehat{\mathcal{T}_{\mu,l,a} * g}\|_1 \leq 2\epsilon$$

and we conclude

$$\|g - \frac{1}{l\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{T}_{\mu,l,a} * g\|_\infty \leq 2\epsilon$$

as desired. ■

The last result in this section will allow us to decouple errors from summing over different localizations. Note that naively, if we add together  $n$  estimates with  $L^2$  errors  $\epsilon_1, \dots, \epsilon_n$ , then the resulting  $L^2$  error of the sum could be as large as  $\epsilon_1 + \dots + \epsilon_n$ . If the estimates were "independent" on the other hand, we would expect the  $L^2$  error of the sum to only be  $\sqrt{\epsilon_1^2 + \dots + \epsilon_n^2}$ . We prove that when adding together functions that are frequency-localized at different locations, the error essentially matches the latter bound (up to logarithmic factors). This tighter bound will be necessary in the proof of Lemma E.1.

Note that if we have functions  $g_1, \dots, g_n$  whose Fourier supports are disjoint, then it is immediate that

$$\|g_1 + \dots + g_n\|_2^2 = \|g_1\|_2^2 + \dots + \|g_n\|_2^2.$$

However, in our setting, we need to restrict the functions to the interval  $[-1, 1]$  first, which causes the Fourier supports to no longer be disjoint. Through a few additional arguments we are able to prove an analogous statement for bounding  $\int_{-1}^1 |g_1 + \dots + g_n|^2$ . We do pay some additional losses both in the inequality itself and the bounds of the integral i.e. we need to integrate the individual functions over a slightly larger interval.

**Claim E.9.** *Let  $\alpha, l, \epsilon$  be parameters such that  $\alpha > 1$  and  $l \geq \lceil \log \alpha n / \epsilon \rceil$ . Let  $I_1, \dots, I_n$  be intervals of length at least  $\alpha$  and assume that for any  $x \in \mathbb{R}$ , at most  $l$  of the intervals contain  $x$ . Let  $g_1, \dots, g_n$  be functions such that for all  $j \in [n]$ ,  $\|\widehat{g}_j\|_1 \leq 1$  and  $\widehat{g}_j$  is supported on  $I_j$ . Then*

$$\int_{-1}^1 |g_1 + \dots + g_n|^2 \leq \text{poly}(l) \left( \epsilon^5 + \int_{-(1+\alpha^{-1})}^{(1+\alpha^{-1})} |g_1|^2 + \dots + \int_{-(1+\alpha^{-1})}^{(1+\alpha^{-1})} |g_n|^2 \right).$$

*Proof.* Consider the Gaussian multiplier  $M = M_{\mu, \alpha^{-2}l-100}$  for some  $\mu \in [-1, 1]$ . Now first, we bound

$$\int_{-\infty}^{\infty} M(x)^2 |g_1(x) + \dots + g_n(x)|^2 dx.$$

Define the functions  $h_j = \widehat{Mg_j}$  for all  $j$ . Then by Plancherel,

$$\int_{-\infty}^{\infty} M(x)^2 |g_1(x) + \cdots + g_n(x)|^2 dx = \int_{-\infty}^{\infty} |h_1 + \cdots + h_n|^2 dx \quad (12)$$

On the other hand, note that  $h_j = \widehat{M} * \widehat{g_j}$ . Let  $J_j$  denote the interval containing all points within distance at most  $10\alpha l^{60}$  of the interval  $I_j$ . Let  $h'_j$  be the function  $h_j$  restricted to  $J_j$  (and equal to 0 outside). Recall that the support of  $\widehat{g_j}$  is contained within  $I_j$ . Then we claim that

$$\int_{-\infty}^{\infty} |h_j - h'_j|^2 \leq (\epsilon/(\alpha n))^{100}.$$

This follows because  $|\widehat{M}| = N(0, (2\pi\alpha l^{50})^2)$  and  $\|\widehat{g_j}\|_1 \leq 1$  so for a point  $x$  that is distance  $d$  away from the interval  $I_j$ , we have

$$|h_j(x)| \leq \max_{y \in I_j} |\widehat{M}(x-y)| \leq N(0, (2\pi\alpha l^{50})^2)(d).$$

Also note that  $\|\widehat{g_j}\|_1 \leq 1$ , implies  $\|g_j\|_{\infty} \leq 1$  so

$$\|h_j\|_2^2 = \|Mg_j\|_2^2 \leq \|M\|_2^2 \leq 1.$$

Combining the previous two inequalities over all  $j$ , we have

$$\begin{aligned} \|\|h_1 + \cdots + h_n\|_2 - \|h'_1 + \cdots + h'_n\|_2\| &\leq (\epsilon/(\alpha n))^{49} \\ \|h_1 + \cdots + h_n\|_2 + \|h'_1 + \cdots + h'_n\|_2 &\leq 3n \end{aligned}$$

which implies

$$\int_{-\infty}^{\infty} |h_1 + \cdots + h_n|^2 dx \leq 0.1(\epsilon/(\alpha n))^{10} + \int_{-\infty}^{\infty} |h'_1 + \cdots + h'_n|^2 dx. \quad (13)$$

Now note that since not too many of the intervals  $I_j$  may contain the same point  $x \in \mathbb{R}$ , not too many of the extended intervals  $J_j$  can contain the same point  $x \in \mathbb{R}$ . In particular, at most  $O(l^{70})$  of the extended intervals can contain the same point  $x \in \mathbb{R}$ . In other words, each point  $x \in \mathbb{R}$  is in the support of at most  $O(l^{70})$  of the  $h'_1, \dots, h'_n$ . Thus, by Cauchy Schwarz,

$$\int_{-\infty}^{\infty} |h'_1 + \cdots + h'_n|^2 \leq O(l^{70}) \left( \int_{-\infty}^{\infty} |h'_1|^2 + \cdots + \int_{-\infty}^{\infty} |h'_n|^2 \right). \quad (14)$$

Now we bound

$$\int_{-\infty}^{\infty} |h'_j|^2 \leq \int_{-\infty}^{\infty} |h_j|^2 = \int_{-\infty}^{\infty} M(x)^2 |g_j(x)|^2 dx \leq 0.1(\epsilon/(\alpha n))^{10} + \int_{-(1+\alpha^{-1})}^{1+\alpha^{-1}} M(x)^2 |g_j(x)|^2 dx \quad (15)$$

where the last step holds because  $\|g_j\|_{\infty} \leq 1$  and the multiplier  $M(x)$  is always at most 1 and decays rapidly outside the interval  $[-(1+\alpha^{-1}), 1+\alpha^{-1}]$  since  $\mu \in [-1, 1]$ . Putting everything together (12, 13, 14, 15), we get

$$\int_{-\infty}^{\infty} M(x)^2 |g_1(x) + \cdots + g_n(x)|^2 dx \leq \text{poly}(l) \left( (\epsilon/(\alpha n))^5 + \sum_{j=1}^n \int_{-(1+\alpha^{-1})}^{1+\alpha^{-1}} M(x)^2 |g_j(x)|^2 dx \right).$$

Now summing the above over different multipliers  $M = M_{\mu, \alpha^{-2}l^{-100}}$  i.e. with  $\mu$  uniformly spaced on  $[-1, 1]$  with spacing  $\alpha^{-1}l^{-50}$ , we conclude

$$\begin{aligned} \int_{-1}^1 |g_1(x) + \cdots + g_n(x)|^2 dx &\leq 10 \sum_{\mu} \int_{-1}^1 M_{\mu, \alpha^{-2}l^{-100}}(x)^2 |g_1(x) + \cdots + g_n(x)|^2 dx \\ &\leq \text{poly}(l)\epsilon^5 + \text{poly}(l) \left( \sum_{\mu} \int_{-(1+\alpha^{-1})}^{1+\alpha^{-1}} M_{\mu, \alpha^{-2}l^{-100}}(x)^2 (|g_1|^2 + \cdots + |g_n|^2) \right) \\ &\leq \text{poly}(l)\epsilon^5 + \text{poly}(l) \left( \int_{-(1+\alpha^{-1})}^{1+\alpha^{-1}} |g_1|^2 + \cdots + \int_{-(1+\alpha^{-1})}^{1+\alpha^{-1}} |g_n|^2 \right). \end{aligned}$$

■

### E.2.1 Completing the Proof of Lemma E.1

In this section, we will complete the proof of Lemma E.1. First, we need to introduce some notation. We will carry over all of the notation from the statement of Lemma E.1. We also use the following conventions:

- Let  $S_0 = \{\theta_1, \dots, \theta_k\}$  be the frequencies in the Fourier support of  $\mathcal{M}$
- Let  $\gamma > 0$  be parameter to be chosen later and let  $l' = \lceil \log 1/\gamma \rceil$  (we will ensure  $\gamma$  is sufficiently small i.e.  $\gamma < (\epsilon c/(kN))^K$  for some sufficiently large absolute constant  $K$ )
- Let the function  $r(x)$  be defined as  $r(x) = f(x) - \mathcal{M}(x)$  on the interval  $[-1, 1]$  and  $r(x) = 0$  outside the interval.

Recall that the way we will reconstruct the function is by attempting to localize around each of the points in the given set  $T$  and reconstructing the localized function using Corollary D.3. We then find  $k\text{poly}(l'/c)$  of these localized reconstructions that we can combine to approximate the entire function.

In the first claim, we bound the error of our reconstruction using Corollary D.3 for a given localization. Recall the two types of kernels, the truncated polynomial kernel and the truncated Gaussian kernel, defined in Section E.1. Consider the kernels  $\mathcal{K}_{\mu, l', c/l'}$  and  $\mathcal{T}_{\mu, l', 2\pi l'/c}$  (which, recall, are approximately the same). In the next lemma, we will bound the distance between  $\mathcal{K}_{\mu, l', c/l'} * f$  and  $\mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}$  in terms of  $r(x)$ . The reason we care about these two functions is that the first is something that we can compute since we are given explicit access to  $f$ . On the other hand, the second is Fourier sparse and has bounded Fourier support so it can be plugged into Corollary D.3 (as the unknown function  $\mathcal{M}$ ).

**Claim E.10.** *For any real number  $\mu$ ,*

$$\int_{-1+c}^{1-c} |\mathcal{K}_{\mu, l', c/l'} * f - \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}|^2 \leq \text{poly}(l'/c)\gamma^2 + 4 \int_{-\infty}^{\infty} |M_{0, (2\pi l'/c)^2}(x - \mu)\widehat{r}(x)|^2 dx.$$

*Proof.* Note that since  $\mathcal{K}_{\mu, l', c/l'}$  is supported on  $[-c, c]$ ,

$$\int_{-1+c}^{1-c} |(\mathcal{K}_{\mu, l', c/l'} * f)(x) - (\mathcal{K}_{\mu, l', c/l'} * \mathcal{M})(x)|^2 dx \leq \int_{-1}^1 |(\mathcal{K}_{\mu, l', c/l'} * r)(x)|^2 dx$$

Now the Fourier transform of  $\mathcal{K}_{\mu, l', c/l'} * r$  is  $\widehat{\mathcal{K}_{0, l', c/l'}}(x - \mu)\widehat{r}(x)$  so

$$\int_{-1}^1 |(\mathcal{K}_{\mu, l', c/l'} * r)(x)|^2 dx \leq \int_{-\infty}^{\infty} |\widehat{\mathcal{K}_{0, l', c/l'}}(x - \mu)\widehat{r}(x)|^2 dx$$

We deduce that

$$\begin{aligned} \int_{-1+c}^{1-c} |\mathcal{K}_{\mu, l', c/l'} * f - \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}|^2 &\leq 2 \int_{-1+c}^{1-c} |\mathcal{K}_{\mu, l', c/l'} * \mathcal{M} - \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}|^2 \\ &\quad + 2 \int_{-\infty}^{\infty} |\widehat{\mathcal{K}_{0, l', c/l'}}(x - \mu)\widehat{r}(x)|^2 dx \\ &\leq \text{poly}(l'/c)\gamma^2 + 2 \int_{-\infty}^{\infty} |\widehat{\mathcal{K}_{0, l', c/l'}}(x - \mu)\widehat{r}(x)|^2 dx \end{aligned}$$

where the last inequality follows from Claim E.6.

Note since  $r$  is supported on  $[-1, 1]$  and  $\|r\|_2^2 \leq \epsilon^2 \leq 0.1$ , we must have  $\|r\|_1 \leq 1$  which then implies  $\|\widehat{r}\|_{\infty} \leq 1$ . Together with Claim E.5, if we let  $G = N(0, (c/l')^2)$  then we have

$$\begin{aligned} \int_{-\infty}^{\infty} |\widehat{\mathcal{K}_{0, l', c/l'}}(x - \mu)\widehat{r}(x)|^2 dx &\leq 2 \int_{-\infty}^{\infty} |\widehat{G}(x - \mu)\widehat{r}(x)|^2 dx + 2\|\widehat{r}\|_{\infty}^2 \|\mathcal{K}_{0, l', c/l'} - G\|_2^2 \\ &\leq \text{poly}(l'/c)\gamma^2 + 2 \int_{-\infty}^{\infty} |M_{0, (2\pi l'/c)^2}(x - \mu)\widehat{r}(x)|^2 dx. \end{aligned}$$

and combining with the previous inequality, we get the desired result.  $\blacksquare$

We are now ready to complete the proof of Lemma E.1.

*Proof of Lemma E.1.* First, let  $T'$  be the set of integer multiples of  $2\pi/c$  that are within distance  $(10l')^2/c$  of the set  $T$ . For all  $\mu \in T'$ , do the following. We compute the function  $f^{(\mu)} = \mathcal{K}_{\mu, l', c/l'} * f$ . Next we apply Corollary D.3 (with appropriate rescaling) to compute a function  $h^{(\mu)}$  in  $\text{poly}(l'/c)$  time that has Fourier support in  $[\mu - 2\pi l'^2/c, \mu + 2\pi l'^2/c]$ , is  $(\text{poly}(l'/c), \text{poly}(l'/c))$ -simple and such that

$$\int_{-1+c}^{1-c} |f^{(\mu)} - h^{(\mu)}|^2 \leq 20 \left( \gamma^2 + \int_{-1+c}^{1-c} |f^{(\mu)} - \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}|^2 \right). \quad (16)$$

To see why we can do this, note that

$$\mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M} = M_{0,1}^{\text{trunc}(l')}(c(x - \mu)/(2\pi l')) \widehat{\mathcal{M}}$$

is supported on  $[\mu - 2\pi l'^2/c, \mu + 2\pi l'^2/c]$ . Also it is clear that  $\|\mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}\|_1 \leq \|\widehat{\mathcal{M}}\|_1 \leq 1$ .

Now we choose a set  $U \subset T'$  with  $|U| \leq k(10l')^2$  such that the following quantity is minimized:

$$E_U = \sum_{\mu \in U} \int_{-1+c}^{1-c} |f^{(\mu)} - h^{(\mu)}|^2 + \sum_{\mu \in T' \setminus U} \int_{-1+c}^{1-c} |f^{(\mu)}|^2.$$

Note that this can be done using a simple greedy procedure. First, we obtain a bound on the value  $E_U$  that we compute. Let  $U_0$  be the set of all integer multiples of  $2\pi/c$  that are within distance  $10l'^2/c$  of  $S_0$  (recall that  $S_0$  is the Fourier support of  $\mathcal{M}$  which consist of  $k$  points). By assumption, we know that  $U_0 \subset T'$  and it is clear that  $|U_0| \leq k(10l')^2$ . Note that by definition, for any  $\mu \notin U_0$ , the function  $\mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}$  is identically 0. Now using (16), then Claim E.10,

$$\begin{aligned} E_{U_0} &\leq 20 \sum_{\mu \in U_0} \left( \gamma^2 + \int_{-1+c}^{1-c} |f^{(\mu)} - \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}|^2 \right) + \sum_{\mu \in T' \setminus U_0} \int_{-1+c}^{1-c} |f^{(\mu)}|^2 \\ &\leq 80 \sum_{\mu \in T'} \left( \text{poly}(l'/c) \gamma^2 + \int_{-\infty}^{\infty} |M_{0, (2\pi l'/c)^2}(x - \mu) \widehat{r}(x)|^2 dx \right) \\ &\leq |T'| \text{poly}(l'/c) \gamma^2 + 80 \int_{-\infty}^{\infty} |\widehat{r}(x)|^2 \sum_{\mu \in (2\pi/c)\mathbb{Z}} M_{0, (2\pi l'/c)^2}(x - \mu)^2 dx \\ &\leq \gamma + \text{poly}(l') \|r\|_2^2. \end{aligned}$$

Note that we used the fact that  $\gamma$  is sufficiently small and the tail decay properties of the Gaussian multipliers in the last step. Thus, we can ensure that the error that we compute satisfies  $E_U \leq \gamma + \text{poly}(l') \|r\|_2^2$ . Now we output the function

$$\widetilde{\mathcal{M}} = \sum_{\mu \in U} \frac{1}{l' \sqrt{2\pi}} h^{(\mu)}.$$

It remains to bound the error between  $\widetilde{\mathcal{M}}$  and  $f$ . First we apply Claim E.9 to decouple over all  $\mu \in T'$ . Note that  $h^{(\mu)}$  and  $\mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}$  both have Fourier support contained in the interval  $[\mu - 2\pi l'^2/c, \mu + 2\pi l'^2/c]$ . For distinct  $\mu$  that are integer multiples of  $2\pi/c$ , there are at most  $O(l'^2)$  intervals that contain any point. Also, note that for all  $\mu$ ,  $\|h^{(\mu)}\|_1 \leq \text{poly}(l'/c)$  and  $\|\mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}\|_1 \leq 1$ . Thus, by Claim E.9 (with appropriate rescaling of the functions and the

interval),

$$\begin{aligned}
& \int_{-1+2c}^{1-2c} \left| \widetilde{\mathcal{M}} - \sum_{\mu \in T'} \frac{1}{l' \sqrt{2\pi}} \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M} \right| \\
& \leq \text{poly}(l'/c) \gamma + \text{poly}(l') \left( \sum_{\mu \in U} \int_{-1+c}^{1-c} |h^{(\mu)} - \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}|^2 + \sum_{\mu \in T' \setminus U} \int_{-1+c}^{1-c} |\mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}|^2 \right) \\
& \leq \text{poly}(l'/c) \gamma + \text{poly}(l') \left( \sum_{\mu \in U} \int_{-1+c}^{1-c} |h^{(\mu)} - f^{(\mu)}|^2 + \sum_{\mu \in T' \setminus U} \int_{-1+c}^{1-c} |f^{(\mu)}|^2 \right) \\
& \quad + \text{poly}(l') \sum_{\mu \in T'} \int_{-1+c}^{1-c} |f^{(\mu)} - \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}|^2 \\
& \leq \text{poly}(l'/c) \gamma + \text{poly}(l') \left( E_U + \sum_{\mu \in T'} \left( \text{poly}(l'/c) \gamma^2 + \int_{-\infty}^{\infty} |M_{0, (2\pi l'/c)^2}(x - \mu) \widehat{r}(x)|^2 dx \right) \right) \\
& \leq \text{poly}(l'/c) \gamma + \text{poly}(l') (E_U + \|r\|_2^2) \\
& \leq \text{poly}(l') \epsilon^2
\end{aligned}$$

where we used that  $\gamma = (\epsilon c / (kN))^{O(1)}$  is sufficiently small and Claim E.10 and the last two inequalities follow from the same argument as in the bound for  $E_{U_0}$ . Next, by Claim E.8 (and the definition of  $T'$ ), we have that

$$\|\mathcal{M} - \sum_{\mu \in T'} \frac{1}{l' \sqrt{2\pi}} \mathcal{T}_{\mu, l', 2\pi l'/c} * \mathcal{M}\|_{\infty} \leq O(\epsilon)$$

so overall, we conclude

$$\int_{-1+2c}^{1-2c} |\widetilde{\mathcal{M}} - \mathcal{M}|^2 \leq \text{poly}(l') \epsilon^2$$

from which we immediately deduce

$$\int_{-1+2c}^{1-2c} |\widetilde{\mathcal{M}} - f|^2 \leq \text{poly}(l') \epsilon^2.$$

It is also clear that  $\widetilde{\mathcal{M}}$  is  $(k\text{poly}(l'/c), k\text{poly}(l'/c))$ -simple (since it is a sum of at most  $k(10l')^2$  functions that are  $(\text{poly}(l'/c), \text{poly}(l'/c))$ -simple). Now we are done because  $l' = O(l)$ .  $\blacksquare$

Similar to obtaining Corollary D.3 from Lemma D.1, we can extend Lemma E.1 to work even when we do not know the target accuracy but only a lower bound on it.

**Corollary E.11.** *Assume we are given  $N, k, \epsilon, c$  with  $0 < \epsilon < 0.1$ . Let  $l = \lceil \log kN/(\epsilon c) \rceil$  be some parameter. Let  $\mathcal{M}$  be a function that is  $(k, 1)$ -simple. Also, assume that we are given a set  $T \subset \mathbb{R}$  of size  $N$  such that all of the support of  $\widetilde{\mathcal{M}}$  is within distance 1 of  $N$ . Further, assume we are given access to a function  $f$ . There is an algorithm that runs in  $\text{poly}(N, k, l, 1/c)$  time and outputs a function  $\widetilde{\mathcal{M}}$  that is  $(k\text{poly}(l/c), k\text{poly}(l/c))$ -simple and such that*

$$\int_{-1+c}^{1-c} |\widetilde{\mathcal{M}}(x) - f(x)|^2 dx \leq \epsilon^2 + \text{poly}(l) \int_{-1}^1 |f(x) - \mathcal{M}(x)|^2 dx.$$

*Proof.* This will be the exact same argument as the proof of Corollary D.3. For a target accuracy  $\epsilon' > \epsilon$  we run the algorithm in Lemma E.1 to get a function  $\widetilde{\mathcal{M}}_{\epsilon'}(x)$ . We then check whether

$$\int_{-1+c}^{1-c} |\widetilde{\mathcal{M}}_{\epsilon'}(x) - f(x)|^2 dx \leq \text{poly}(l) \epsilon'^2.$$

If the check passes, we take  $\epsilon' \leftarrow 0.99\epsilon'$  and repeat the above until we find the smallest  $\epsilon'$  (up to a constant factor) for which the check passes. The guarantee of Lemma E.1 implies that for this  $\epsilon'$ , we can just output  $\widetilde{\mathcal{M}}_{\epsilon'}(x)$  and it is guaranteed to satisfy the desired inequality.  $\blacksquare$

### E.3 Proof of Main Theorem

Using Theorem 2.12 and Corollary E.11, we can prove our main theorem, Theorem 1.3. The main thing that we need to prove is that the frequencies in the function  $f'$  computed by Theorem 2.12 cover (within distance  $\text{poly}(k, \log 1/\epsilon)$ ) all of the frequencies in  $\mathcal{M}$ . This will then let us use the frequencies in  $f'$  to construct a set  $T$  of size  $\text{poly}(k, \log 1/\epsilon)$  that covers all frequencies in  $\mathcal{M}$  to within distance 1 that we can then plug into Corollary E.11. We will need the following technical lemma from Chen et al. [2016].

**Lemma E.12.** [Lemma 5.1 in Chen et al. [2016]] For any  $k$ -Fourier sparse signal  $g : \mathbb{R} \rightarrow \mathbb{C}$ ,

$$\max_{x \in [-1, 1]} |g(x)|^2 \leq O(k^4 \log^3 k) \int_{-1}^1 |g(x)|^2 dx.$$

The above lemma roughly says that the mass of a  $k$ -Fourier sparse function cannot be too concentrated. We now finish the proof of Theorem 1.3.

*Proof of Theorem 1.3.* We first apply Theorem 2.12 to compute a function  $f'$ , such that

1.  $f'$  is  $(\text{poly}(k, \log 1/\epsilon), \exp(\text{poly}(k, \log 1/\epsilon)))$ - simple
2. 
$$\int_{-1}^1 |f' - f|^2 \leq O\left(\epsilon^2 + \int_{-1}^1 |f - \mathcal{M}|^2\right).$$

Let  $L = (k \log 1/\epsilon)^K$  for some sufficiently large absolute constant  $K$ . Let  $\gamma = e^{-L}$ . Now apply Claim E.8 on the function  $f'$  with parameters  $a \leftarrow L, l \leftarrow L, \epsilon \leftarrow \gamma$ . Let  $S \subset L\mathbb{Z}$  be the set of all integer multiples of  $L$  that are within distance  $L^3$  of the Fourier support of  $f'$ . We have

$$\|f' - \frac{1}{L\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{T}_{\mu, L, L^2} * f'\|_{\infty} \leq 2\|\widehat{f'}\|_1 \gamma \quad (17)$$

Next, we apply Claim E.6 on the function  $(\mathcal{M} - f')$ . We deduce that for any  $\mu$ ,

$$\|\mathcal{K}_{\mu, L, 2\pi/L^2} * (\mathcal{M} - f') - \mathcal{T}_{\mu, L, L^2} * (\mathcal{M} - f')\|_{\infty} \leq O\left(\gamma L \left(\|\widehat{f'}\|_1 + \|\widehat{\mathcal{M}}\|_1\right)\right). \quad (18)$$

Finally, by Claim E.7 applied to the function  $(\mathcal{M} - f')$  (with parameters  $\epsilon \leftarrow \gamma, l \leftarrow L, c \leftarrow (2\pi)/L^2$ ), we have

$$\int_{-1}^1 \left| \frac{1}{L\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{K}_{\mu, L, 2\pi/L^2} * (\mathcal{M} - f') \right|^2 \leq (1 + O(\gamma|S|))^2 \int_{-1-2\pi/L}^{1+2\pi/L} |\mathcal{M} - f'|^2 \leq 2 \int_{-1}^1 |\mathcal{M} - f'|^2.$$

Note that in the last step we use Lemma E.12 and the fact that  $\mathcal{M} - f'$  is  $\text{poly}(k, \log 1/\epsilon)$ -Fourier sparse so choosing  $L = (k \log 1/\epsilon)^{O(1)}$  sufficiently large ensures that

$$\int_{-1-2\pi/L}^{1+2\pi/L} |\mathcal{M} - f'|^2 \leq 1.1 \int_{-1}^1 |\mathcal{M} - f'|^2.$$

Define the functions

$$\begin{aligned} A(x) &= f' - \frac{1}{L\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{T}_{\mu, L, L^2} * \mathcal{M} \\ B(x) &= \frac{1}{L\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{K}_{\mu, L, 2\pi/L^2} * (\mathcal{M} - f') \end{aligned}$$

Note  $\|A\|_{\infty}, \|B\|_{\infty} \leq |S| \left(\|\widehat{f'}\|_1 + \|\widehat{\mathcal{M}}\|_1\right)$ . Combining (17, 18) we have

$$\int_{-1}^1 |A(x)|^2 - \int_{-1}^1 |B(x)|^2 \leq \gamma \text{poly}\left(L, |S|, \|\widehat{f'}\|_1 + \|\widehat{\mathcal{M}}\|_1\right)$$

However, we proved that  $\int_{-1}^1 |B(x)|^2 \leq 2 \int_{-1}^1 |\mathcal{M} - f'|^2$  so choosing  $L = (k \log 1/\epsilon)^{O(1)}$  sufficiently large and since  $\gamma = e^{-L}$ , we conclude

$$\int_{-1}^1 |A(x)|^2 \leq \epsilon^2 + 2 \int_{-1}^1 |\mathcal{M} - f'|^2 \leq O\left(\epsilon^2 + \int_{-1}^1 |\mathcal{M} - f|^2\right) \quad (19)$$

where we are using the guarantee from Theorem 2.12. Now note that the function

$$\mathcal{M}' = \frac{1}{L\sqrt{2\pi}} \sum_{\mu \in S} \mathcal{T}_{\mu, L, L^2} * \mathcal{M}$$

is  $k$ -Fourier sparse and has  $\|\widehat{\mathcal{M}'}\|_1 \leq 2\|\widehat{\mathcal{M}}\|_1$  by Lemma B.4. Furthermore, all of its Fourier support is within distance  $\text{poly}(L)$  of the Fourier support of  $f'$  (by the construction of the set  $S$ ). Thus, we can apply Corollary E.11 on  $f'$  (where we treat the unknown function as  $\mathcal{M}'$ ) with  $N = \text{poly}(L) = \text{poly}(k, \log 1/\epsilon)$  and recover a function  $\widetilde{\mathcal{M}}$  such that

$$\begin{aligned} \int_{-1+c}^{1-c} |\widetilde{\mathcal{M}} - f'|^2 &\leq \epsilon^2 + \text{poly}(\log k/(\epsilon c)) \int_{-1}^1 |f' - \mathcal{M}'|^2 = \epsilon^2 + \text{poly}(\log k/(\epsilon c)) \int_{-1}^1 |A(x)|^2 \\ &= \text{poly}(\log k/(\epsilon c)) \left( \epsilon^2 + \int_{-1}^1 |\mathcal{M} - f|^2 \right) \end{aligned}$$

where the last step is from (19). Since  $\int_{-1}^1 |f' - f|^2 \leq O\left(\epsilon^2 + \int_{-1}^1 |f - \mathcal{M}|^2\right)$ , the above implies

$$\int_{-1+c}^{1-c} |\widetilde{\mathcal{M}} - f|^2 \leq \text{poly}(\log k/(\epsilon c)) \left( \epsilon^2 + \int_{-1}^1 |\mathcal{M} - f|^2 \right)$$

and we are done. ■

#### E.4 Implementation of Computations

In the proof of Theorem 1.3, we use the result from Chen et al. [2016] to obtain an approximation  $f'$  that is written as a sum of  $\text{poly}(k, \log 1/\epsilon)$  exponentials and has coefficients bounded by  $\exp(\text{poly}(k, \log 1/\epsilon))$ . We then perform explicit computations using this function in our algorithm to eventually compute a sparser approximation with smaller coefficients. Here we briefly explain why these explicit computations can all be implemented efficiently. Note that all of the functions that we perform computations on can be written as sums of polynomials multiplied by exponentials i.e.

$$P_1(x)e^{2\pi i \theta_1 x} + \cdots + P_n(x)e^{2\pi i \theta_n x} \quad (20)$$

where there are at most  $\text{poly}(k, \log 1/\epsilon)$  terms in the sum, all of the polynomials have degree at most  $\text{poly}(k, \log 1/\epsilon)$  and all of the coefficients are bounded by  $\exp(\text{poly}(k, \log 1/\epsilon))$ . To see this, note that convolving by a polynomial  $P(x)$  truncated to an interval (recall the truncated polynomial kernel in Definition E.3) preserves a function of the form in (20) (only increasing the degrees of the polynomials by  $\deg(P)$ ). All other computations that we need such as computing the exact value, adding and multiplying and integrating over some interval can clearly be done explicitly in  $\text{poly}(k, \log 1/\epsilon)$  time and to  $\exp(\text{poly}(k, \log 1/\epsilon))^{-1}$  accuracy for functions of the form specified in (20).

## F Basic Tools

In this section, we have a few basic tools that are used repeatedly throughout the paper.

### F.1 Chebyshev Polynomials

Here we will introduce several basic results about the Chebyshev polynomials, which have algorithmic applications in a wide variety of settings Rivlin [2020], Guruswami and Zuckerman [2016].

**Definition F.1** (Chebyshev Polynomials). *The Chebyshev Polynomials are a family of polynomials defined as follows:  $T_0(x) = 1, T_1(x) = x$  and for  $n \geq 2$ ,*

$$T_n(x) = 2xT_{n-1}(x) - T_{n-2}(x).$$

**Fact F.2.** *The Chebyshev polynomials satisfy the following property:*

$$T_n(\cos \theta) = \cos n\theta.$$

As an immediate consequence of the above, we have a few additional properties.

**Fact F.3.** *The Chebyshev polynomials satisfy the following properties:*

1.  $T_n(x)$  has degree  $n$  and leading coefficient  $2^{n-1}$
2. For  $x \in [-1, 1]$ ,  $T_n(x) \in [-1, 1]$
3.  $T_n(x)$  has  $n$  zeros all in the interval  $[-1, 1]$
4. There are  $n + 1$  values of  $x$  for which  $T_n(x) = \pm 1$ , all in the interval  $[-1, 1]$

In light of the previous properties, we make the following definition.

**Definition F.4.** *For an integer  $n$ , we define the Chebyshev points of degree  $n$ , say  $t_0, \dots, t_n$ , as the points in the interval  $[-1, 1]$  where the Chebyshev polynomial satisfies  $T_n(t_j) = \pm 1$ . Note that the Chebyshev points are exactly*

$$\{\cos 0, \cos \frac{\pi}{n}, \dots, \cos \frac{(n-1)\pi}{n}, \cos \pi\}.$$

Next, we have a result saying that if we have a bound on the value of a degree- $n$  polynomial at all of the degree  $n$  Chebyshev points, then we can bound the value over the entire interval  $[-1, 1]$ . Similar results are used in Rivlin [2020], Guruswami and Zuckerman [2016], but there does not appear to be a directly usable reference.

**Claim F.5.** *Let  $P(x)$  be a polynomial of degree at most  $n$  with real coefficients. Let  $t_0, \dots, t_n$  be the Chebyshev points of degree  $n$ . Assume that  $|P(t_j)| \leq 1$  for  $j = 0, 1, \dots, n$ . Then  $|P(x)| \leq 2n$  for all  $x \in [-1, 1]$ .*

*Proof.* By Lagrange interpolation, we may write

$$P(x) = \frac{P(t_0)(x - t_1) \cdots (x - t_n)}{(t_0 - t_1) \cdots (t_0 - t_n)} + \cdots + \frac{P(t_n)(x - t_0) \cdots (x - t_{n-1})}{(t_n - t_0) \cdots (t_n - t_{n-1})}.$$

Thus, it suffices to upper bound the quantity

$$F(x) = \left| \frac{(x - t_1) \cdots (x - t_n)}{(t_0 - t_1) \cdots (t_0 - t_n)} \right| + \cdots + \left| \frac{(x - t_0) \cdots (x - t_{n-1})}{(t_n - t_0) \cdots (t_n - t_{n-1})} \right|$$

on the interval  $[-1, 1]$ . Note that by Lagrange interpolation on  $T_n(x)$ , we have

$$T_n(x) = \frac{T_n(t_0)(x - t_1) \cdots (x - t_n)}{(t_0 - t_1) \cdots (t_0 - t_n)} + \cdots + \frac{T_n(t_n)(x - t_0) \cdots (x - t_{n-1})}{(t_n - t_0) \cdots (t_n - t_{n-1})}.$$

Also note that  $T_n(t_j) = (-1)^{n-j}$  which has the same sign as  $(t_j - t_0) \cdots (t_j - t_{j-1})(t_j - t_{j+1}) \cdots (t_j - t_n)$ . Thus,

$$\left| \frac{1}{(t_0 - t_1) \cdots (t_0 - t_n)} \right| + \cdots + \left| \frac{1}{(t_n - t_0) \cdots (t_n - t_{n-1})} \right| = 2^{n-1},$$

since the leading coefficient of  $T_n(x)$  is  $2^{n-1}$ . Now we will upper bound

$$M = \max(|(x - t_1) \cdots (x - t_n)|, \dots, |(x - t_0) \cdots (x - t_{n-1})|)$$

and once we do this, we will have a bound on  $F(x)$  since  $F(x) \leq 2^{n-1}M$ . Define the polynomial

$$Q(x) = (x - t_0)(x - t_1) \cdots (x - t_n) = \frac{\sqrt{x^2 - 1}}{2^n} \left( (x + \sqrt{x^2 - 1})^n - (x - \sqrt{x^2 - 1})^n \right).$$

To see why the last equality is true, note that the RHS has roots at  $t_0, \dots, t_n$  and is a monic polynomial of degree  $n + 1$  so it must be equal to  $(x - t_0) \cdots (x - t_n)$ . Now,

$$M = \max \left( \left| \frac{Q(x)}{x - t_0} \right|, \dots, \left| \frac{Q(x)}{x - t_n} \right| \right) \leq \max |Q'(x)|$$

where the last step holds by the mean value theorem (because  $Q(t_j) = 0$  for all  $j$ ). Now note that

$$Q(\cos \theta) = -\frac{\sin \theta \sin(n\theta)}{2^{n-1}}$$

so

$$Q'(\cos \theta) = \frac{n \cos n\theta}{2^{n-1}} + \frac{\cos \theta \sin(n\theta)}{\sin(\theta) 2^{n-1}}$$

and from the above it is clear that

$$|Q'(\cos \theta)| \leq \frac{n}{2^{n-1}} + \frac{n}{2^{n-1}} = \frac{n}{2^{n-2}}.$$

Now we are done because

$$\max_{x \in [-1, 1]} |P(x)| \leq F(x) \leq 2^{n-1} M \leq 2n.$$

■

It turns out that we can restate the above result in terms of convex hulls of points on the moment curve. This reformulation is the version that is useful in our algorithms.

**Definition F.6.** For a real number  $x$ , we define the moment vector  $\mathcal{V}_n(x) = (1, x, \dots, x^n)$ .

**Corollary F.7.** Let  $t_0, t_1, \dots, t_n$  be the Chebyshev points of degree  $n$ . Then for any  $x \in [-1, 1]$ , the point  $\mathcal{V}_n(x)$  is contained in the convex hull of the points

$$\{\pm 2n\mathcal{V}_n(t_0), \dots, \pm 2n\mathcal{V}_n(t_n)\}.$$

*Proof.* Assume for the sake of contradiction that the above is not true. Then there must be a separating hyperplane. Assume that this hyperplane is given by  $a \cdot x = b$  where  $a$  is a vector and  $b$  is a real number. Now WLOG  $b \geq 0$  and we must have

$$\begin{aligned} a \cdot \mathcal{V}_n(x) &\geq b \\ |a \cdot \mathcal{V}_n(t_j)| &\leq \frac{b}{2n} \quad \forall j \end{aligned}$$

However, applying Claim F.5 with  $P(x) = \frac{2n(a \cdot \mathcal{V}_n(x))}{b}$  gives a contradiction. Thus, no separating hyperplane can exist and we are done. ■

## F.2 Approximating a Gaussian with a Polynomial

We will also need to approximate Gaussians with polynomials. This is a somewhat standard result which we state below.

**Lemma F.8.** Let  $G = N(0, 1)$  be the standard Gaussian. Let  $l$  be some parameter. Then we can compute a polynomial  $P(x)$  of degree  $(10l)^2$  such that for all  $x \in [-2l, 2l]$ ,

$$|G(x) - P(x)| \leq e^{-l}.$$

*Proof.* Write

$$G(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}.$$

and now we can write the Taylor expansion

$$e^{-\frac{x^2}{2}} = \sum_{m=0}^{\infty} \frac{\left(-\frac{x^2}{2}\right)^m}{m!} = \sum_{m=0}^{\infty} \frac{(-1)^m x^{2m}}{2^m m!}$$

Now define

$$P(x) = \sum_{m=0}^{(10l)^2} \frac{(-1)^m x^{2m}}{2^m m!}.$$

For  $x \in [-2l, 2l]$ , we have

$$\begin{aligned}
|G(x) - P(x)| &\leq \left| \sum_{m=(10l)^2+1}^{\infty} \frac{(-1)^m x^{2m}}{2^m m!} \right| \leq \sum_{m=10^2 l+1}^{\infty} \frac{(2l)^{2m}}{2^m m!} \leq \sum_{m=10^2 l+1}^{\infty} \left( \frac{(2l)^2}{2m/3} \right)^m \\
&\leq \sum_{m=10^2 l+1}^{\infty} \frac{1}{2^m} \leq e^{-l}.
\end{aligned}$$

■

In light of the above, we use the following notation.

**Definition F.9.** We will use  $\mathcal{P}_l(x)$  to denote the polynomial computed in Lemma F.8 for parameter  $l$ . Note that  $\mathcal{P}_l$  is a polynomial of degree  $(10l)^2$  and for  $G = N(0, 1)$ , we have

$$|G(x) - \mathcal{P}_l(x)| \leq e^{-l}$$

for  $x \in [-2l, 2l]$ .

### F.3 Linear Regression

Recall that at the core of the problems we are studying, we are given some function  $f$  and want to approximate it as a weighted sum  $a_1 f_1 + \dots + a_n f_n$  of some functions  $f_1, \dots, f_n \in \mathcal{F}$  for some family of functions  $\mathcal{F}$ . The result below allows us to solve the problem of computing the coefficients if we already know the components  $f_1, \dots, f_n$  that we want to use. The precise technical statement is slightly more complicated in order to incorporate the various types of additional constraints that we may want to impose on the coefficients  $a_1, \dots, a_n$ .

**Lemma F.10.** Let  $\mathcal{D}$  be a distribution on  $\mathbb{R}$  that we are given. Also assume that we are given functions  $f, f_1, \dots, f_n, g, g_1, \dots, g_n : \mathbb{R} \rightarrow \mathbb{R}$ . Assume that there are nonnegative coefficients  $a_1, \dots, a_n$  such that  $a_1 + \dots + a_n \leq 1$  and

$$\int_{-\infty}^{\infty} (f(x) - a_1 f_1(x) - \dots - a_n f_n(x))^2 \mathcal{D}(x) dx + \int_{-\infty}^{\infty} (g(x) - a_1 g_1(x) - \dots - a_n g_n(x))^2 \mathcal{D}(x) dx \leq \epsilon^2$$

for some parameter  $\epsilon > 0$ . Then there is an algorithm that runs in  $\text{poly}(n, \log 1/\epsilon)$  time and outputs nonnegative coefficients  $b_1, \dots, b_n$  such that  $b_1 + \dots + b_n \leq 1$  and

$$\int_{-\infty}^{\infty} (f(x) - b_1 f_1(x) - \dots - b_n f_n(x))^2 \mathcal{D}(x) dx + \int_{-\infty}^{\infty} (g(x) - b_1 g_1(x) - \dots - b_n g_n(x))^2 \mathcal{D}(x) dx \leq 2\epsilon^2.$$

*Proof.* Let  $v = (1, b_1, \dots, b_n)$ . Note that we can write

$$\int_{-\infty}^{\infty} (f(x) - b_1 f_1(x) - \dots - b_n f_n(x))^2 \mathcal{D}(x) dx + \int_{-\infty}^{\infty} (g(x) - b_1 g_1(x) - \dots - b_n g_n(x))^2 \mathcal{D}(x) dx = v^T M v$$

where  $M$  is a matrix whose entries are  $\int_{-\infty}^{\infty} (f(x) f_j(x) + g(x) g_j(x)) \mathcal{D}(x) dx$  in the first row and column and the other entries are  $\int_{-\infty}^{\infty} (f_i(x) f_j(x) + g_i(x) g_j(x)) \mathcal{D}(x) dx$ . Since all of these functions are given to us, we can explicitly compute  $M$ . Also note that clearly  $M$  is positive semidefinite. Thus, we can compute its positive semidefinite square root, say  $N$ . Now

$$v^T M v = \|Nv\|_2^2$$

so it remains to solve  $\min_v \|Nv\|_2^2$  which is a convex optimization problem that we can solve efficiently (the size of the problem is  $\text{poly}(n)$ ). ■