

Consistent Estimation of Identifiable Nonparametric Mixture Models from Grouped Observations

Alexander Ritchie*
 Department of EECS
 University of Michigan
 Ann Arbor, MI 48109
 aritch@umich.edu

Robert A. Vandermeulen*
 ML group
 Technische Universität Berlin
 10587 Berlin, Germany
 vandermeulen@tu-berlin.de

Clayton Scott
 Departments of EECS, Statistics
 University of Michigan
 Ann Arbor, MI 48109
 clayscot@umich.edu

Abstract

Recent research has established sufficient conditions for finite mixture models to be identifiable from grouped observations. These conditions allow the mixture components to be nonparametric and have substantial (or even total) overlap. This work proposes an algorithm that consistently estimates any identifiable mixture model from grouped observations. Our analysis leverages an oracle inequality for weighted kernel density estimators of the distribution on groups, together with a general result showing that consistent estimation of the distribution on groups implies consistent estimation of mixture components. A practical implementation is provided for paired observations, and the approach is shown to outperform existing methods, especially when mixture components overlap significantly.

1 Introduction

In statistics and machine learning, finite mixture models are often used to describe the distribution of subpopulations within a larger population. A finite mixture model can be written

$$p = \sum_{m=1}^M w_m^* p_m^*, \quad (1)$$

where $w_m^* > 0$ are mixing weights such that $\sum_{m=1}^M w_m^* = 1$, and p_m^* are probability densities. Without additional assumptions, the mixture model p is not identifiable from iid data. Typically, identifiability is ensured by restricting the p_m^* to some family of parametric distributions. Restricting the p_m^* to be Gaussian yields the Gaussian mixture model (GMM) which is identifiable [1, 2].

*Equal contribution.

Most work on estimating mixture models assumes an iid sampling scheme. In this work we examine an alternative sampling scheme where observations occur in iid groups. Each group is generated by sampling a component $m \in [M]$ according to w_m^* , and then drawing N iid observations from p_m^* .

Recent work has shown that *any* finite mixture model is identifiable given grouped observations of sufficient size [3]. In the worst case, any finite mixture model with M components is identifiable given groups of size $N \geq 2M - 1$. It was also shown that, if the underlying components of the mixture model are *jointly irreducible* [4], then the mixture is identifiable given paired observations ($N = 2$). This framework provides a setting where the potential exists to recover nonparametric and highly overlapping mixture components. As of yet, however, no general theory or algorithms are known for this estimation problem.

This work makes the following contributions. We introduce a novel variant of the kernel density estimator that yields statistically consistent estimates of *any identifiable* nonparametric mixture model (NoMM) from grouped observations. To prove this result, we establish an oracle inequality for weighted kernel density estimators. We also establish a general result showing that consistent estimation (with an estimator possessing a natural factored form) of the distribution on groups implies consistent estimates of the underlying components when the NoMM is identifiable. The only additional condition imposed by our theory is that the p_m^* be square integrable. In the case of $N = 2$, we offer an efficient algorithm and demonstrate its effectiveness on several datasets.

We study two applications where paired observations naturally arise. The first is nuclear source detection, where nuclear particles interact with a detector to produce some form of measurement. A critical challenge in this application is to classify incoming particles as belonging to source or background. Because of changing environments, training data are typically not available, and these two classes also have substantial overlap. By positioning two detectors side-by-side, it is possible to simultaneously measure two particles from the same (unknown) class.

We also apply our method to topic modeling of Twitter data. Since tweets usually express a small set of very closely related ideas, words in tweets contain common underlying semantic information. The pairing of words has the potential to encode this semantic information in a way that accounts for context. The proposed method, which operates on continuous word embeddings, allows for flexible modeling of the distributions of topics over words using static word embeddings [5]. Furthermore our method does not require anchor words, allows for substantial overlap of topics without loss of identifiability, and can be trained using documents with as few as two words *without* any document aggregation [6, 7]. While other works have explored topic modeling with word embeddings, which we call continuous topic modeling, most either impose parametric assumptions or are not suited for very short texts. To our knowledge, this is the first work to consider a nonparametric approach to continuous topic modeling of very short texts.

2 Background and Previous Work

Much of the literature concerning NoMMs falls in the category of Bayesian nonparametrics, a thorough summary of which is given in [8]. Typically, mixture models in this setting do not assume that the number of mixture components is known, and instead assume that the mixture components are from a known parametric family of distributions. An in-depth treatment of Bayesian NoMMs (BNoMMs) can be found in [9]. The parametric assumptions on the mixture components have been relaxed in [10], but the identifiability results impose regularity and separation conditions on the

components. We mention BNMMs only for completeness and emphasize that our work considers an alternative setting where the number of mixture components is known, but few to no assumptions are made on the mixture components themselves.

Mixture models are often utilized to solve the clustering problem. Parametric mixture models, such as GMMs, are able to capture overlapping clusters. Most clustering algorithms, however, such as k -means [11, 12], DBSCAN [13], and spectral clustering [14, 15], assume clusters are non-overlapping and hence fail when clusters overlap. The grouped observation setting considered in this work is known in the clustering literature as *clustering with instance-level constraints* [16, 17, 18]. A survey of constrained clustering is given in [19]. Grouped observations correspond to so-called must-link constraints, where two or more observations are known, through expert knowledge or some other means, to belong to the same cluster. Most constrained clustering approaches cannot model overlapping clusters effectively [20].

There is relatively little work on mixture modeling with nonparametric components, and to our knowledge no prior work addresses the incorporation of instance-level constraints in the NoMM setting. Mallapragada et al. [21] use a mixture of kernel density estimators to estimate a NoMM, but do not address identifiability or provide statistical guarantees. Aragam et al. [10] prove identifiability of NoMMs under regularity and separation conditions on the components. They provide a simple algorithm that gives Bayes optimal cluster assignments in the limit, but they does not guarantee on recovery of the underlying components. Schiebinger et al. [22] study kernelized spectral clustering and characterize recoverability of components with small overlap. Zheng and Wu [23] establish consistent estimation of NoMMs under the assumption that mixture components have independent marginals. Bao et al. [24] consider the related problem of “similar-unsupervised” binary classification, which assumes access to unlabeled data in addition to must-link constraints.

In the grouped observation setting, previous works on multi-view models can be adapted to prove identifiability results and give algorithms to recover mixture model components. When the mixture components are linearly independent it has been shown that three observations per group is sufficient to yield identifiability as well as an algorithm to provably recover the components [25, 26]. We note that these approaches require three observations per group, while the proposed method works with as few as two observations per group. This difference amounts to performing kernel density estimation in three times the ambient dimension versus two. With the instability of KDEs in high dimension, the reduction to groups of size two can be very meaningful in practice. Furthermore, in applications like nuclear particle classification, triples may be exceedingly rare or difficult to measure. For discrete data, similar results from nonnegative matrix factorization exist under joint irreducibility with two observations per group [27], and algorithms have been proposed to recover arbitrary mixture models with M components given $2M - 1$ observations per group [28, 3].

3 Notation

For $1 \leq p < \infty$ denote $\|f\|_p := (\int_{\mathbb{R}^d} |f(x)|^p dx)^{1/p}$, and $L^p := \{f : \mathbb{R}^d \rightarrow \mathbb{R} : \|f\|_p < \infty\}$. The transpose of a matrix A will be written A' . Random variables will be referred to by capital letters, and instances of random variables will be referred to by the corresponding lowercase letter. We represent the set of positive integers $\{1, 2, \dots, M\}$ by $[M]$. We let Δ^R be the probability simplex in \mathbb{R}^R . We denote the M -fold Cartesian product of a set with a subscript, e.g., $\Delta_M^R = \underbrace{\Delta^R \times \dots \times \Delta^R}_M$.

4 Problem Statement

We precisely introduce the grouped observation setting, review known identifiability results, and formalize the estimation problem. The standard sampling procedure for a mixture model of the form $p = \sum_{m=1}^M w_m^* p_m^*$ can be viewed as a two step process wherein one samples a mixture component p_m^* with probability w_m^* and then observes one draw from that distribution $X \sim p_m^*$. The *grouped observation setting* considers an alternative sampling scheme where, after selecting a mixture component p_m^* , instead of only drawing a single observation, a group of observations $\mathbf{X} = (X_1, \dots, X_N)$ are drawn iid from p_m^* . As in a standard mixture model, one does not know *a priori* from which mixture component a grouped observation is sampled. Repeating this n times, one's data consists of n groups of N observations per group $\mathbf{X}_1 = (X_{1,1}, \dots, X_{1,N}), \dots, \mathbf{X}_n = (X_{n,1}, \dots, X_{n,N})$. The distribution on groups is $\mathbf{X} \stackrel{iid}{\sim} \sum_{m=1}^M w_m^* p_m^{*\times N}$, where $p_m^{*\times N} : \mathbb{R}_N^d \rightarrow \mathbb{R}$ denotes the product density such that $p_m^{*\times N}(y_1, y_2, \dots, y_N) = p_m^*(y_1)p_m^*(y_2) \dots p_m^*(y_N)$. Note that when $N = 1$ this is simply a standard mixture model.

Vandermeulen and Scott [3] characterized identifiability from grouped observations for mixtures of general probability measures. A mixture model $p = \sum_{m=1}^M w_m^* p_m^*$ is said to be *N-identifiable* if p cannot be expressed $p = \sum_{m=1}^{M'} w'_m p'_m$ for some distinct mixture model such $M' \leq M$ and $\sum_{m=1}^M w_m^* p_m^{*\times N} = \sum_{m=1}^{M'} w'_m p'_m^{*\times N}$. In words, *N-identifiability* of p means there is no other mixture model with M or fewer components that induces the same distribution on groups. They show that a general mixture model is *N-identifiable* from grouped observations provided $N \geq 2M - 1$, and that this cannot be improved without imposing restrictions on the components. The result places no assumptions whatsoever on the components.

In practice, the bound of $2M - 1$ is probably pessimistic, and the most useful cases are likely when N is small, say two or three. The authors of [3] also show that if the p_m^* are *jointly irreducible* (*linearly independent*), then the mixture is *N-identifiable* for $N = 2$ ($N = 3$). A collection of probability densities $\mu_1, \mu_2, \dots, \mu_M$ is said to be *jointly irreducible* (JI) if $\sum_{m=1}^M c_m \mu_m$ is never a valid density whenever some $c_m < 0$. JI is satisfied, for example, if the support of each mixture component has some subset of positive measure that does not intersect the supports of the other mixture components (a continuous analogue of the anchor word assumption). This is not necessary, however; JI is still possible if all densities have the same support. In the remainder of the paper we focus on the setting of $N = 2$, not only because JI provides a flexible nonparametric condition where paired observations suffice, but also because the notation for our estimator becomes cumbersome when $N > 2$. Our theory generalizes easily to $N > 2$, and these details are described in Appendices C, D, and E.

The paired observations $\mathbf{X}_1, \dots, \mathbf{X}_n$ with $\mathbf{X}_i = (X_{i,1}, X_{i,2}) \in \mathbb{R}^d \times \mathbb{R}^d$ are iid and have density

$$q(x, x') := \sum_{m=1}^M w_m^* p_m^*(x)p_m^*(x') \quad x, x' \in \mathbb{R}^d. \quad (2)$$

We assume M is known. Our goal is to consistently estimate w_m^* and p_m^* when p is identifiable.

5 A Weighted Kernel Density Estimator

Our overall strategy is to first devise a consistent estimator of q , the density on pairs, where the estimator has a factorized form reflecting the group sampling scheme. In the next section we prove

that if an estimator for q is consistent, and p is identifiable, then the components comprising our estimator converge to the true components.

Let $k : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function, called a *kernel*, such that $k \geq 0$ and $\int k(x)dx = 1$. An example is the Gaussian kernel $k(x) = (2\pi)^{-1/2} \exp(-\|x\|^2/2)$. For $\sigma > 0$, define $k_\sigma(x, x') := \sigma^{-d}k((x - x')/\sigma)$. We refer to the second argument of k_σ as the *center* of the kernel. A *weighted kernel density estimator* (wKDE) for a density on \mathbb{R}^d , but constructed from the paired observations \mathbf{X}_i , has the form

$$p(x; \theta) = \sum_{r=1}^n \sum_{r'=1}^2 \theta_{r,r'} k_\sigma(x, X_{r,r'}),$$

where $\theta_{r,r'}$ is the element of $\theta = [\theta_{1,1}, \theta_{1,2}, \dots, \theta_{n,1}, \theta_{n,2}]'$ corresponding to the weight of the kernel centered at $X_{r,r'}$. We propose to model the mixture components as wKDEs. Specifically, given n paired observations, we consider estimators of q of the form

$$q_{w,\alpha}(x, x') = \sum_{m=1}^M w_m p(x; \alpha_m) p(x'; \alpha_m), \quad (3)$$

where $w = [w_1, w_2, \dots, w_M]' \in \Delta^M$, $\alpha_m = [\alpha_{m,1,1}, \alpha_{m,1,2}, \dots, \alpha_{m,n,1}, \alpha_{m,n,2}]' \in \Delta^{2n}$ for all $m \in [M]$, with $\alpha_{m,r,r'}$ corresponding to the weight of the kernel centered at $X_{r,r'}$ in the estimate of the m^{th} mixture component, and $\alpha := (\alpha_1, \alpha_2, \dots, \alpha_M) \in \Delta_M^{2n}$.

To select the parameters (w, α) , we propose to minimize the integrated square error (ISE) of $q_{w,\alpha}$ given by $\|q - q_{w,a}\|_2^2 := \int [q(x, x') - q_{w,a}(x, x')]^2 dx dx'$. Expanding the ISE gives

$$\|q - q_{w,a}\|_2^2 = \int q_{w,\alpha}^2(x, x') dx dx' - 2 \int q_{w,\alpha}(x, x') q(x, x') dx dx' + \int q^2(x, x') dx dx' \xrightarrow{\text{const.}}.$$

Since the final term is constant with respect to w and α , we focus on minimizing the first two terms which we call the truncated ISE (TISE) and denote by $J(w, \alpha)$. Substituting the definition of $q_{w,\alpha}$ in the TISE yields

$$J(w, \alpha) := \int q_{w,\alpha}^2(x, x') dx dx' - 2 \sum_{m=1}^M \sum_{r=1}^n \sum_{r'=1}^2 \sum_{s=1}^n \sum_{s'=1}^2 w_m \alpha_{m,r,r'} \alpha_{m,s,s'} h(r, r', s, s'), \quad (4)$$

where $h(r, r', s, s') := \int k_\sigma(x, X_{r,r'}) k_\sigma(x', X_{s,s'}) q(x, x') dx dx'$. Since q is unknown, the ISE and therefore $J(w, a)$ cannot be calculated directly. Noting that $h(r, r', s, s')$ is an expectation, we estimate this term using a hybrid leave-one-out/leave-two-out (LOO/LTO) estimator

$$\hat{h}(r, r', s, s') := \begin{cases} \frac{1}{n-2} \sum_{i \in [n] \setminus \{r,s\}} k_\sigma(X_{i,1}, X_{r,r'}) k_\sigma(X_{i,2}, X_{s,s'}) & r \neq s \\ \frac{1}{n-1} \sum_{i \in [n] \setminus \{r\}} k_\sigma(X_{i,1}, X_{r,r'}) k_\sigma(X_{i,2}, X_{s,s'}) & r = s \end{cases}.$$

In this manner we have the empirical TISE

$$\hat{J}(w, \alpha) := \int q_{w,\alpha}^2(x, x') dx dx' - 2 \sum_{m=1}^M \sum_{r=1}^n \sum_{r'=1}^2 \sum_{s=1}^n \sum_{s'=1}^2 w_m \alpha_{m,r,r'} \alpha_{m,s,s'} \hat{h}(r, r', s, s'). \quad (5)$$

With all the notation in place, our estimate of the nonparametric mixture model is determined by

$$(\hat{w}, \hat{\alpha}) := \arg \min_{w \in \Delta^M, \alpha \in \Delta_M^{2n}} \hat{J}(w, \alpha), \quad (6)$$

where \hat{w}_m are the mixing weights and $p(x; \hat{\alpha}_m)$ are the mixture components for $m \in [M]$. The theoretical results presented in Section 6 concern the behavior of the minimizer of (6). We show not only that the empirical TISE minimizing estimator $\hat{q} := q_{\hat{w}, \hat{\alpha}}$ consistently estimates q , but its components also consistently estimate the underlying mixture model if it is identifiable.

6 Theoretical Results

In this section we state our assumptions and main results. Formal proofs are given in the Appendices C, D, and E. Our overall approach is to first show that the proposed \hat{q} is a consistent estimate of q (Theorems 1 and 2). We then show that if p in (1) is identifiable, then the components $p(x; \hat{\alpha}_m)$ defining \hat{q} are consistent estimates of p_m^* , as are the \hat{w}_m for w_m^* (Theorem 3).

We assume throughout this section that $p_m^* \in L^2$ for all m . We also require that the kernel k satisfy two additional conditions: $k \in L^2$ and $k \leq C_k$ for some constant $C_k < \infty$.

We begin with an oracle inequality, which shows that our estimator selects an approximately optimal member of our model class.

Theorem 1. *Let $\epsilon > 0$ and set $\delta = 8(n^2 - n) \exp\{-\frac{\sigma^{4d}(n-2)\epsilon^2}{8C_k^4}\} + 8n \exp\{-\frac{\sigma^{4d}(n-1)\epsilon^2}{8C_k^4}\}$. With probability at least $1 - \delta$ the following holds: $\|q - q_{\hat{w}, \hat{\alpha}}\|_2^2 \leq \inf_{w \in \Delta^M, \alpha \in \Delta_M^{2n}} \|q - q_{w, \alpha}\|_2^2 + \epsilon$.*

Proof Sketch. The estimators \hat{h} are constructed so that they are sums of independent random variables, allowing us to apply Hoeffding's inequality to show that each \hat{h} concentrates around its h . Then using basic inequalities (triangle inequality, union bound) and the simplex constraints on w and α , we show that $\hat{J}(w, \alpha)$ concentrates around $J(w, \alpha)$ uniformly over the parameter space. \square

The next result uses Theorem 1 to establish that \hat{q} is a consistent estimate of q in the L^1 norm.

Theorem 2. *If $\sigma \rightarrow 0$ and $\frac{n\sigma^{4d}}{\log n} \rightarrow \infty$ as $n \rightarrow \infty$, then $\|q - q_{\hat{w}, \hat{\alpha}}\|_1 \xrightarrow{a.s.} 0$.*

Proof Sketch. We appeal to a result of [29] showing that if $\int \hat{q} = 1$, which it does in our case, then strong consistency (i.e., a.s. convergence) of a density estimator in L^2 implies strong consistency in L^1 . To show strong consistency in L^2 , from Theorem 1 it suffices to exhibit $w \in \Delta^M$ and $\alpha \in \Delta_M^{2n}$ such that $\|q - q_{w, \alpha}\|_1 \xrightarrow{a.s.} 0$. For this we take $w = w^*$ and $\alpha = \alpha^*$ such that each α_m^* is uniform on the data points drawn from p_m^* . This makes $p(\cdot; \alpha_m^*)$ the usual (uniformly weighted) KDE for p_m^* , which is known to be a strongly consistent estimator. Strong consistency of \hat{q} then easily follows. \square

The preceding results hold regardless of whether p in (1) is identifiable. The next result states that if p is identifiable, then the estimates $p(\cdot; \hat{\alpha}_m)$ comprising \hat{q} are consistent estimates of the true components p_m^* , as are the \hat{w}_m of w_m^* . The result is stated for $N \geq 2$.

Theorem 3. *Let $\sum_{m=1}^M w_m p_m$ be an N -identifiable mixture model, and $\sum_{m=1}^M \hat{w}_{m,j} \hat{p}_{m,j}$ be a sequence of mixture models such that $\left\| \sum_{m=1}^M \hat{w}_{m,j} \hat{p}_{m,j}^{\times N} - \sum_{m=1}^M w_m p_m^{\times N} \right\|_1 \rightarrow 0$. Then there is a sequence of permutations σ_j so that $\hat{w}_{\sigma_j(m),j} \rightarrow w_m$ and $\left\| \hat{p}_{\sigma_j(m),j} - p_m \right\|_1 \rightarrow 0$ for all m .*

Proof Sketch. We show that if $\left\| \sum_{m=1}^M \hat{w}_{m,j} \hat{p}_{m,j}^{\times N} - \sum_{m=1}^M w_m p_m^{\times N} \right\|_1 \rightarrow 0$ then the components $\hat{p}_{m,j}$ admit some convergent subsequence, and therefore so do $\hat{p}_{m,j}^{\times N}$. If a subsequence $\hat{p}_{m,j}^{\times N}$ stays away from the components $p_m^{\times N}$ then some subsequence would converge to a component other than some $p_m^{\times N}$. This allows us to construct a mixture model violating N -identifiability, a contradiction. \square

This result has been stated in terms of densities for readability, but Appendix E contains a general measure-theoretic version. We may combine Theorems 2 and 3 to establish the following (returning to the setting of $N = 2$). To our knowledge, this is the first result to establish consistent estimation, under any sampling scheme, of NoMMs with substantial overlap.

Corollary 1. *If $\sigma \rightarrow 0$ and $\frac{n\sigma^{4d}}{\log n} \rightarrow \infty$ as $n \rightarrow \infty$, and p is 2-identifiable (e.g., the p_m^* are jointly irreducible), then $\hat{w}_m \xrightarrow{a.s.} w_m^*$ and $\|p(\cdot; \hat{\alpha}_m) - p_m^*\|_1 \xrightarrow{a.s.} 0$, up to a permutation.*

The significance of the result is that joint irreducibility is both a flexible nonparametric assumption, while ensuring identifiability in the case $N = 2$ for which a practical implementation of \hat{q} is possible. We include an analogous result for *all identifiable NoMMs* in Appendix F.

7 Optimization

In this section we suggest an approach for solving (6). We first consider the problem as presented up to this point, which we call the *full problem*. We then consider an approach for speeding up optimization by heuristically choosing a coresnet as the kernel centers, which we call the *coresnet approach*. In what follows, we assume that $\tilde{k}_\sigma(z_r, z_u) := \int k_\sigma(x, z_r) k_\sigma(x, z_u) dx$ has a closed-form expression or can otherwise be computed efficiently. This assumption is satisfied by many common kernels such as the Gaussian, Cauchy, and Laplacian kernels.

Form of the Optimization Problem. The optimization problem (6) can be written

$$\min_{w \in \Delta^M, \alpha \in \Delta_M^R} \sum_{k=1}^M \sum_{\ell=1}^M w_k w_\ell \left(\alpha'_k G \alpha_\ell \right)^2 - 2 \sum_{m=1}^M w_m (\alpha'_m C \alpha_m), \quad (7)$$

where the matrices $G, C \in \mathbb{R}^{R \times R}$ will be defined shortly. Details are given in Appendix B. In particular, both the full problem and the coresnet approach can be written in the form of (7), differing only in the definitions of R and G, C . We therefore propose to use the same optimization approach for both problems. For the full problem, $R = 2n$ and the matrices G and C have the form

$$G_{a,b} = \tilde{k}_\sigma(X_{\lfloor \frac{a}{2} \rfloor, a \bmod 2}, X_{\lfloor \frac{b}{2} \rfloor, b \bmod 2}) \quad C_{a,b} = \hat{h}(\lfloor \frac{a}{2} \rfloor, a \bmod 2, \lfloor \frac{b}{2} \rfloor, b \bmod 2).$$

Though the problem (7) is nonconvex, we observe that a properly initialized alternating projected stochastic gradient descent (APSGD) procedure produces good solutions in practice.

Pseudocode for the APSGD algorithm for solving (7) is given in Appendix B. We mention that the projections are onto the probability simplex, a decaying step size is used, and stochasticity is introduced via the matrix $C^{(t)}$, which is a mini-batch version of C defined by $C_{a,b}^{(t)} = \frac{1}{|\Omega^{(t)} \setminus \{a,b\}|} \sum_{i \in \Omega^{(t)} \setminus \{a,b\}} k_\sigma(X_{i,1}, X_{\lfloor \frac{a}{2} \rfloor, a \bmod 2}) k_\sigma(X_{i,2}, X_{\lfloor \frac{b}{2} \rfloor, b \bmod 2})$, where $\Omega^{(t)}$ is the index set corresponding to the t^{th} mini-batch.

Coreset Approach. KDEs traditionally center kernels at the location of each observation, i.e., $k_\sigma(\cdot, x_{i,i'})$, where $x_{i,i'}$ is the kernel center. Rather than constraining the wKDE to have kernels centered at the observations, we can formulate the optimization problem with R kernel centers $z_r \in \mathbb{R}^d$ for some suitably chosen z_r , which we take to be our coresset. Further details are given in Appendix B. We note the per-batch computational complexity for our APSGD algorithm is dominated by the gradient calculations and calculating $C^{(t)}$. If we assume $R > M$, the total complexity is $\mathcal{O}(n_e n(M + d)R^2)$ where n_e is the number of training epochs. Thus, choosing $R \ll 2n$ offers a substantial speed-up.

Initialization. We adopt a spectral initialization scheme. We focus on the full problem for concision, but the coresset approach is similar; further details for both are provided in Appendix B. By Lemmas 5.1 and 8.2 of Vandermeulen and Scott [3], one can view the standard KDE on the full sample as a symmetric linear operator $T : L^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d)$. We use the eigenvectors of T , which are wKDEs on \mathbb{R}^d , to form a low-rank approximation of the standard KDE initialize our algorithm. This initialization is a low-rank approximation of the standard KDE.

8 Experiments

In this section we compare our coresset approach against several competing methods on a number of real and highly overlapping synthetic datasets. Datasets are described in Table 1. We call the proposed method Nonparametric Density estimation of Identifiable mixture models from Grouped Observations (NDIGO). All code and synthetic datasets are publicly available.¹ The MAGIC gamma ray detection dataset [30] is publicly available via the UCI machine learning repository. The Russian-troll-tweets Twitter dataset is publicly available through FiveThirtyEight.² For NDIGO and MVLVM, we used a Gaussian kernel in all experiments and Scott’s rule [31] was used for bandwidth selection. For synthetic experiments, R was selected to yield the initialization with the lowest empirical TISE. R was chosen from $\{10, 20, 30, 40, 50\}$ for both moons datasets, and from $\{60, 70, 80, 90, 100\}$ for the Olympic rings and half-disks datasets. We used $R = 200$ for the MAGIC and Twitter datasets.

Several of the methods we compare against do not produce density estimates, so we evaluate the clustering induced by each method. For constrained clustering methods, we compare against constrained spectral clustering (CSC) [32], and constrained GMM (CGMM) [33]. We also compare against the NoMM methods NPMIX of Aragam et al. [10] and MVLVM of Song et al. [34]. MVLVM is our most similar competitor as it considers groups of size three. Each constrained clustering algorithm was given access to all pair information. MVLVM was supplied triplets from the training data. NPMIX does not utilize the pair information in any way. Following the literature, we report the clustering results for the training sample. Out-of-sample results are provided in Appendix A, but we mention NPMIX is the best performer. Parameters for CSC and NPMIX were optimized w.r.t. a separately generated holdout dataset. Average results over ten runs on the synthetic datasets are shown in Figure 1. NDIGO outperforms all methods considered. The synthetic datasets were constructed to have clusters that are non-ellipsoidal in shape with substantial overlap between clusters. The clusterings induced by each method are shown in Figure 1. Performance is measured in terms of the adjusted Rand index (ARI) [35]. We observe that NDIGO gives superior performance

¹ Available Online: <https://github.com/aritchie9590/NDIGO>

² Available Online: <https://github.com/fivethirtyeight/russian-troll-tweets>

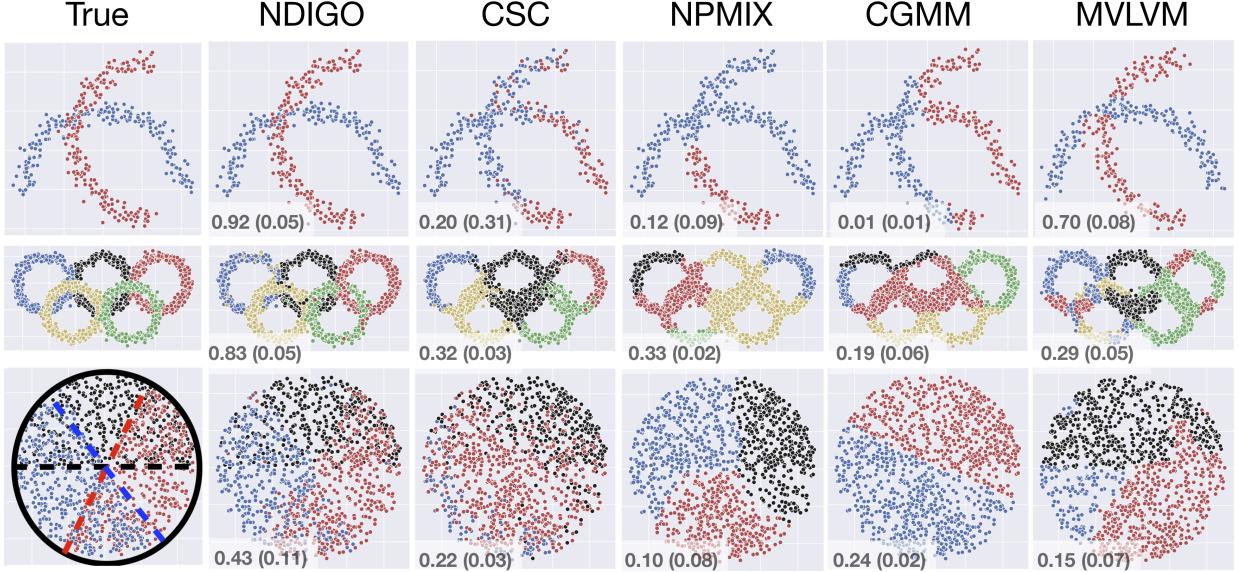


Figure 1: Example cluster assignments of four synthetic datasets by each method. Mean ARI (standard deviation) over 10 runs is shown at the bottom left of each clustering (larger is better). The datasets are overlapping moons (top), Olympic rings (middle), and half-disks (bottom). Half-disks has been annotated to show the true components.

across all experiments, especially when clusters have substantial overlap. Density estimates produced by our method for synthetic datasets are shown in Figure 2.



Figure 2: Component density estimate contours produced NDIGO. From left to right: overlapping moons, Olympic rings, half-disks

Table 1: Description of datasets.
*Quantities after preprocessing.

Dataset (2n)	M/d
Ovlp. Moons (400)	2/2
Olympic Rings (2000)	5/2
Half-disks (1200)	3/2
MAGIC (19,020)	2/10
Twitter (3,382,162*)	-/10*

Results on the MAGIC dataset are shown in Figure 3. The task is to detect gamma radiation events among background radiation. When detecting rare events, the proper performance indicator is given by the receiver operating characteristic (ROC) curve, which plots the true positive rate vs. the false positive rate, parameterized by the threshold of a likelihood ratio test (LRT). Each method was trained using 80% of the available data, and the ROC curve was generated from the remaining 20%. CSC was excluded from this test because it does not produce a density estimate, so a LRT cannot be applied. As an upper bound on possible unsupervised performance, we trained KDEs on each class and plugged the resulting density estimates into an LRT. Previous studies concluded this method, which we call KDE-plugin, is the best approach [30]. We find NDIGO and CGMM perform very similarly in this experiment, outperforming other methods and approaching KDE-plugin.

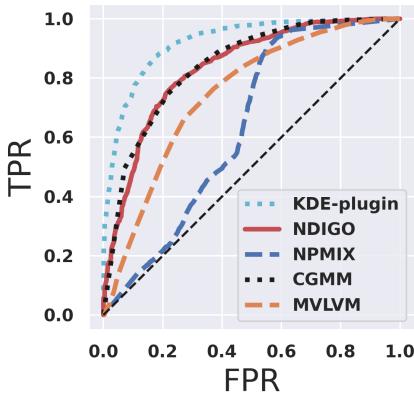


Figure 3: Receiver operating characteristic for MAGIC gamma ray detection dataset.

We applied NDIGO to topic modeling on the Twitter dataset. Results are shown in Tables 2 and 3. Details of data preprocessing are deferred to Appendix A. After preprocessing, the dataset consisted of 1,691,081 pairs of 10-dimensional embedded words where each element in a pair comes from the same tweet. Algorithms for competing methods, as described by their respective authors, could not scale to this experiment. Therefore, we compare to recent methods designed for continuous topic modeling of short texts: LF-DMM [36], and GPU-DMM[37] as implemented by Qiang et al. [38].³ A selection of the top 10 words of topics uncovered by NDIGO is given in Table 2. We find that the discovered topics correspond well to other analyses of the dataset [39]. Using topic coherence (pointwise mutual information) as an evaluation metric [40], we observe that NDIGO is competitive with the competing methods.

9 Conclusion

In this work we introduced a novel variant of the kernel density estimator that yields consistent estimates of any identifiable nonparametric mixture model from grouped observations. We established an oracle inequality for weighted kernel density estimators, and a general consistency result for estimators of the form $q_{w,\alpha}$. Namely, consistent estimation of q implies consistent estimates of the underlying components when the NoMM is identifiable. In the case of $N = 2$, we offer an efficient algorithm and demonstrate its effectiveness on several datasets where traditional approaches fail. Additionally, we show our approach has practical applications in topic modeling with very small documents and nuclear source detection.

References

- [1] J. Anderson, M. Belkin, N. Goyal, L. Rademacher, and J. Voss, “The more, the merrier: the blessing of dimensionality for learning large Gaussian mixtures,” in *Conference on Learning*

³available online: <https://github.com/qiang2100/STTM>

Topic	Selected Top 10 Words
1	dead, man, kill, missing, families, young
2	make, good, better, enough, yet, even, get
3	politics, inside, news, local, police, new, state
4	trial, a*, gentrified, wk, deport, b*
5	businesses, competitive, strength, people, white

Table 3: Mean and standard deviation of topic coherence on Twitter dataset over five experiments.

NDIGO	LF-DMM	GPU-DMM
0.456 ± 0.069	0.493 ± 0.018	0.435 ± 0.009

Theory, 2014, pp. 1135–1164.

- [2] S. J. Yakowitz and J. D. Spragins, “On the identifiability of finite mixtures,” *The Annals of Mathematical Statistics*, pp. 209–214, 1968.
- [3] R. A. Vandermeulen and C. D. Scott, “An operator theoretic approach to nonparametric mixture models,” *The Annals of Statistics*, vol. 47, no. 5, pp. 2704–2733, Oct. 2019. [Online]. Available: <https://projecteuclid.org/euclid-aos/1564797861>
- [4] G. Blanchard and C. Scott, “Decontamination of mutually contaminated models,” in *Artificial Intelligence and Statistics*, 2014, pp. 1–9.
- [5] J. Pennington, R. Socher, and C. D. Manning, “GloVe: Global vectors for word representation,” in *Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543. [Online]. Available: <http://www.aclweb.org/anthology/D14-1162>
- [6] W. Gao, M. Peng, H. Wang, Y. Zhang, Q. Xie, and G. Tian, “Incorporating word embeddings into topic modeling of short text,” *Knowledge and Information Systems*, vol. 61, no. 2, pp. 1123–1145, 2019.
- [7] J. Qiang, P. Chen, T. Wang, and X. Wu, “Topic modeling over short texts by incorporating word embeddings,” in *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2017, pp. 363–374.
- [8] J. Xuan, J. Lu, and G. Zhang, “A survey on Bayesian nonparametric learning,” *ACM Computing Surveys (CSUR)*, vol. 52, no. 1, pp. 1–36, 2019.
- [9] B. G. Lindsay, “Mixture models: theory, geometry and applications,” in *NSF-CBMS regional conference series in probability and statistics*. JSTOR, 1995, pp. i–163.
- [10] B. Aragam, C. Dan, P. Ravikumar, and E. P. Xing, “Identifiability of nonparametric mixture models and Bayes optimal clustering,” *arXiv preprint arXiv:1802.04397*, 2018.
- [11] E. W. Forgy, “Cluster analysis of multivariate data: efficiency versus interpretability of classifications,” *biometrics*, vol. 21, pp. 768–769, 1965.
- [12] S. Lloyd, “Least squares quantization in PCM,” *IEEE transactions on information theory*, vol. 28, no. 2, pp. 129–137, 1982.
- [13] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, “A density-based algorithm for discovering clusters in large spatial databases with noise.” in *KDD*, vol. 96, no. 34, 1996, pp. 226–231.
- [14] Jianbo Shi and J. Malik, “Normalized cuts and image segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.
- [15] U. Luxburg, “A tutorial on spectral clustering,” *Statistics and Computing*, vol. 17, no. 4, p. 395–416, Dec. 2007. [Online]. Available: <https://doi.org/10.1007/s11222-007-9033-z>
- [16] K. Wagstaff and C. Cardie, “Clustering with instance-level constraints,” *AAAI/IAAI*, vol. 1097, pp. 577–584, 2000.

- [17] K. Wagstaff, C. Cardie, S. Rogers, and S. Schrödl, “Constrained k-means clustering with background knowledge,” in *ICML*, vol. 1, 2001, pp. 577–584.
- [18] K. L. Wagstaff, S. Basu, and I. Davidson, “When is constrained clustering beneficial, and why?” *Ionosphere*, vol. 58, no. 60.1, pp. 62–63, 2006.
- [19] P. Gançarski, B. Crémilleux, G. Forestier, and T. Lampert, “Constrained clustering: Current and new trends,” *A Guided Tour of AI Research*, In press. hal-02548212, 2020.
- [20] J. Scripps and P.-N. Tan, “Constrained overlapping clusters: minimizing the negative effects of bridge-nodes,” *Statistical Analysis and Data Mining: The ASA Data Science Journal*, vol. 3, no. 1, pp. 20–37, 2010.
- [21] P. K. Mallapragada, R. Jin, and A. Jain, “Non-parametric mixture models for clustering,” in *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*. Springer, 2010, pp. 334–343.
- [22] G. Schiebinger, M. J. Wainwright, and B. Yu, “The geometry of kernelized spectral clustering,” *The Annals of Statistics*, vol. 43, no. 2, pp. 819–846, 2015.
- [23] C. Zheng and Y. Wu, “Nonparametric estimation of multivariate mixtures,” *Journal of the American Statistical Association*, pp. 1–16, 2019.
- [24] H. Bao, G. Niu, and M. Sugiyama, “Classification from pairwise similarity and unlabeled data,” in *International Conference on Machine Learning*, 2018, pp. 452–461.
- [25] A. Anandkumar, R. Ge, D. Hsu, S. M. Kakade, and M. Telgarsky, “Tensor decompositions for learning latent variable models,” *Journal of Machine Learning Research*, vol. 15, pp. 2773–2832, 2014. [Online]. Available: <http://jmlr.org/papers/v15/anandkumar14b.html>
- [26] E. S. Allman, C. Matias, and J. A. Rhodes, “Identifiability of parameters in latent structure models with many observed variables,” *Ann. Statist.*, vol. 37, no. 6A, pp. 3099–3132, 12 2009. [Online]. Available: <http://dx.doi.org/10.1214/09-AOS689>
- [27] S. Arora, R. Ge, R. Kannan, and A. Moitra, “Computing a nonnegative matrix factorization – provably,” in *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, ser. STOC ’12. New York, NY, USA: ACM, 2012, pp. 145–162. [Online]. Available: <http://doi.acm.org/10.1145/2213977.2213994>
- [28] Y. Rabani, L. J. Schulman, and C. Swamy, “Learning mixtures of arbitrary distributions over large discrete domains,” in *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ser. ITCS ’14. New York, NY, USA: ACM, 2014, pp. 207–224. [Online]. Available: <http://doi.acm.org/10.1145/2554797.2554818>
- [29] L. Gyorfi and E. Masry, “The L_1 and L_2 strong consistency of recursive kernel density estimation from dependent samples,” *IEEE Transactions on Information Theory*, vol. 36, no. 3, pp. 531–539, 1990.
- [30] R. Bock, A. Chilingarian, M. Gaug, F. Hakl, T. Hengstebeck, M. Jiřina, J. Klaschka, E. Kotrč, P. Savický, and S. Towers, “Methods for multidimensional event classification: a case study

- using images from a cherenkov gamma-ray telescope,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 516, no. 2-3, pp. 511–528, 2004.
- [31] D. W. Scott, *Multivariate density estimation: theory, practice, and visualization*. John Wiley & Sons, 2015, p. 164.
 - [32] X. Wang, B. Qian, and I. Davidson, “On constrained spectral clustering and its applications,” *Data Mining and Knowledge Discovery*, vol. 28, no. 1, pp. 1–30, 2014.
 - [33] S. Basu, I. Davidson, and K. Wagstaff, *Constrained clustering: Advances in algorithms, theory, and applications*. CRC Press, 2008, ch. 3.
 - [34] L. Song, A. Anandkumar, B. Dai, and B. Xie, “Nonparametric estimation of multi-view latent variable models,” in *International Conference on Machine Learning*, 2014, pp. 640–648.
 - [35] L. Hubert and P. Arabie, “Comparing partitions,” *Journal of classification*, vol. 2, no. 1, pp. 193–218, 1985.
 - [36] D. Q. Nguyen, R. Billingsley, L. Du, and M. Johnson, “Improving topic models with latent feature word representations,” *Transactions of the Association for Computational Linguistics*, vol. 3, pp. 299–313, 2015.
 - [37] C. Li, H. Wang, Z. Zhang, A. Sun, and Z. Ma, “Topic modeling for short texts with auxiliary word embeddings,” in *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, 2016, pp. 165–174.
 - [38] J. Qiang, Q. Zhenyu, Y. Li, Y. Yuan, and X. Wu, “Short text topic modeling techniques, applications, and performance: A survey,” *arXiv preprint arXiv:1904.07695*, 2019.
 - [39] D. L. Linvill and P. L. Warren, “Troll factories: Manufacturing specialized disinformation on twitter,” *Political Communication*, pp. 1–21, 2020.
 - [40] D. Newman, J. H. Lau, K. Grieser, and T. Baldwin, “Automatic evaluation of topic coherence,” in *Human language technologies: The 2010 annual conference of the North American chapter of the association for computational linguistics*. Association for Computational Linguistics, 2010, pp. 100–108.
 - [41] J. Kim and C. D. Scott, “Robust kernel density estimation,” *Journal of Machine Learning Research*, vol. 13, no. Sep, pp. 2529–2565, 2012.
 - [42] S. Banach, *Théorie des opérations linéaires*, 1932. [Online]. Available: <http://eudml.org/doc/268537>
 - [43] R. Kadison and J. Ringrose, *Fundamentals of the theory of operator algebras. V1: Elementary theory*, ser. Pure and Applied Mathematics. Elsevier Science, 1983. [Online]. Available: <https://books.google.com/books?id=JbxgKOWu2McC>
 - [44] G. B. Folland, *Real analysis: modern techniques and their applications*, ser. Pure and applied mathematics. Wiley, 1999. [Online]. Available: <http://books.google.com/books?id=uPkYAQAAIAAJ>

- [45] R. Reiss, *Approximate distributions of order statistics: with applications to nonparametric statistics*, ser. Springer series in statistics. Springer, 1989. [Online]. Available: <https://books.google.de/books?id=DxzvAAAAMAAJ>

Appendices

A Additional Experimental Details

In this section we provide details of the Twitter experiment and out-of-sample results for experiments on synthetic datasets.

A.1 Out-of-sample Results

Here we provide out-of-sample results for the synthetic experiments shown in the main paper. These are shown in Table 4. We make the realistic assumption that pair information is not available for out-of-sample data. We generate a test dataset 20% the size of the training set according to the distribution of the training data.

Table 4: Out-of-sample ARI and standard deviation over 10 runs on synthetic datasets.

Dataset	NDIGO	CSC	NPMIX	CGMM	MVLVM
Overlapping moons	0.705 ± 0.091	0.405 ± 0.283	0.131 ± 0.075	0.002 ± 0.011	0.593 ± 0.229
Olympic Rings	0.607 ± 0.058	0.387 ± 0.033	0.367 ± 0.063	0.127 ± 0.012	0.290 ± 0.053
Half-disks	0.221 ± 0.036	0.215 ± 0.038	0.102 ± 0.095	0.185 ± 0.076	0.127 ± 0.062

A.2 Preprocessing of Twitter Dataset

Twitter dataset is publicly available through FiveThirtyEight.⁴ The data consist of tweets, from a variety Russian-troll twitter accounts, tweeted between 2015 and 2018. We considered all tweets from 2016, a total of 878,878. We pre-processed the tweets by removing stop words, punctuation, and hyperlinks, followed by a lemmatization step and the removal of any words that were not contained in the vocabulary of the six billion token GloVe word vectors [5]. For completeness, we mention that lemmatization is a common pre-processing step in natural language processing that removes inflectional differences from words by mapping each inflection to a common base form called the lemma. For example, lemmatization will map each of the words dog, dogs, dog’s, dogs’, and doggy to the word dog. For each tweet, we paired the constituent words uniformly at random without replacement, resulting in 1,691,081 pairs of words where the words of a given pair come from the same tweet. No other information was retained. We emphasize that a given word from a given tweet will *not* be assigned to more than one pair. However, if a given word appears in multiple tweets (which may not all be about the same topic), it will show up in multiple pairs.

For the embedding step we performed PCA on the pre-trained GloVe 50-dimensional embeddings to obtain 10-dimensional vectors which were used to encode the paired words. The kernel centers

⁴ <https://github.com/fivethirtyeight/russian-troll-tweets>

were obtained by running mini-batch k -means with $R = 200$ on a uniform random sample of 10,000 of the 1,691,081 word pairs, from which the matrix G was also calculated. We then trained on the full 1,691,081 word pairs, which are utilized in mini-batches through the matrix $C^{(t)}$.

Algorithms for competing methods, as described by their respective authors, could not scale to this experiment. Therefore, we compare to recent methods designed for continuous topic modeling of short texts: LF-DMM [36], and GPU-DMM[37] as implemented by Qiang et al. [38].⁵ LF-DMM and GPU-DMM were trained on the same preprocessed data as NDIGO, where each of the 1,691,081 word pairs is considered a unique document. Each of these methods were run with default hyperparameters, as described in the documentation for GPU-DMM and LF-DMM.² After training, UCI topic coherence [40] (which measures pointwise mutual information) was used to evaluate performance. UCI topic coherence uses a reference dataset to estimate word co-occurrence probabilities, which is more robust in the short text setting as very common words in a given topic may never be observed to co-occur. A recent Wikipedia article dump was used for the reference dataset, and is provided with our code.

B Optimization Details

In this section we provide details of our algorithm, including the form of the objective function and initialization, for both the full problem and the coresnet approach.

Recall the expression for the ETISE

$$\hat{J}(w, \alpha) \triangleq \int q_{w, \alpha}^2(x, x') dx dx' - 2 \sum_{m=1}^M \sum_{r=1}^n \sum_{r'=1}^2 \sum_{s=1}^n \sum_{s'=1}^2 w_m \alpha_{mrr'} \alpha_{mss'} \hat{h}(r, r', s, s'),$$

where

$$\begin{aligned} \hat{h}(r, r', s, s') &\triangleq \begin{cases} \hat{h}_{\text{LTO}}(r, r', s, s'), & r \neq s \\ \hat{h}_{\text{LOO}}(r, r', s'), & r = s \end{cases} \\ \hat{h}_{\text{LOO}}(r, r', r'') &\triangleq \frac{1}{n-1} \sum_{i \in [n] \setminus \{r\}} k_\sigma(x_{i,1}, x_{r,r'}) k_\sigma(x_{i,2}, x_{r,r''}) \\ \hat{h}_{\text{LTO}}(r, r', s, s') &\triangleq \frac{1}{n-2} \sum_{i \in [n] \setminus \{r,s\}} k_\sigma(x_{i,1}, x_{r,r'}) k_\sigma(x_{i,2}, x_{s,s'}). \end{aligned}$$

For ease of computation, we will rewrite $\hat{J}(w, \alpha)$ in terms of matrix operations. In what follows, we assume that $\tilde{k}_\sigma(z_r, z_u) := \int k_\sigma(x, z_r) k_\sigma(x, z_u) dx$ has a closed-form expression or can otherwise be computed efficiently. Some examples [41] are give in Table 5.

⁵<https://github.com/qiang2100/STTM>

Table 5: Some popular kernel functions and their associated \tilde{k} . Here $\|\cdot\|_2$ is the Euclidean norm.

Kernel	$k_\sigma(x, x')$	$\tilde{k}_\sigma(x, x')$
Gaussian	$\left(\frac{1}{\sqrt{2\pi}\sigma}\right)^d \exp\left(-\frac{\ x-x'\ _2^2}{2\sigma^2}\right)$	$k_{\sqrt{2}\sigma}(x, x')$
Cauchy	$\left(\frac{1}{\sqrt{\pi}\sigma}\right)^d \left(\frac{\Gamma((1+d)/2)}{\Gamma(1/2)}\right) \left(\frac{\sigma^2 + \ x-x'\ _2^2}{\sigma^2}\right)^{-\frac{1+d}{2}}$	$k_{2\sigma}(x, x')$
Laplacian	$\frac{c_d}{\sigma^d} \exp\left(-\frac{\ x-x'\ _1}{\sigma}\right)$	$\frac{1}{(4\sigma)^d} \prod_{l=1}^d \left(\frac{\sigma + x_l - x'_l }{\sigma}\right) \exp\left(-\frac{\ x-x'\ _1}{\sigma}\right)$

B.1 Full Optimization Problem

We begin by examining the first term of $\hat{J}(w, \alpha)$

$$\begin{aligned}
\int q_{w,\alpha}(x, x')^2 dx dx' &= \int \left(\sum_m w_m \sum_r \sum_{r'} \alpha_{m,r,r'} k_\sigma(x, x_{r,r'}) \sum_s \sum_{s'} \alpha_{m,s,s'} k_\sigma(x, x_{s,s'}) \right) \\
&\quad \times \left(\sum_j w_j \sum_u \sum_{u'} \alpha_{j,u,u'} k_\sigma(x, x_{u,u'}) \sum_v \sum_{v'} \alpha_{j,v,v'} k_\sigma(x, x_{v,v'}) \right) dx dx' \\
&= \sum_{m,j} w_m w_j \sum_{r,r',u,u'} \sum_{s,s',v,v'} \alpha_{m,r,r'} \alpha_{m,s,s'} \alpha_{j,u,u'} \alpha_{j,v,v'} \\
&\quad \times \int k_\sigma(x, x_{r,r'}) k_\sigma(x, x_{u,u'}) dx \int k_\sigma(x', x_{s,s'}) k_\sigma(x', x_{v,v'}) dx' \\
&= \sum_{m,j} w_m w_j \sum_{r,r',u,u'} \alpha_{m,r,r'} \alpha_{j,u,u'} \tilde{k}_\sigma(x_{r,r'}, x_{u,u'}) \sum_{s,s',v,v'} \alpha_{m,s,s'} \alpha_{j,v,v'} \tilde{k}_\sigma(x_{s,s'}, x_{v,v'}) \\
&= \sum_{m,j} w_m w_j \left(\alpha'_m G a_j \right)^2,
\end{aligned}$$

where \times in the first line is scalar multiplication and G is the kernel matrix of the data and is given by

$$G = \begin{bmatrix} \tilde{k}_\sigma(x_{1,1}, x_{1,1}) & \tilde{k}_\sigma(x_{1,1}, x_{1,2}) & \cdots & \cdots & \tilde{k}_\sigma(x_{1,1}, x_{n,1}) & \tilde{k}_\sigma(x_{1,1}, x_{n,2}) \\ \tilde{k}_\sigma(x_{1,2}, x_{1,1}) & \tilde{k}_\sigma(x_{1,2}, x_{1,2}) & \cdots & \cdots & \tilde{k}_\sigma(x_{1,2}, x_{n,1}) & \tilde{k}_\sigma(x_{1,2}, x_{n,2}) \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ \tilde{k}_\sigma(x_{n,1}, x_{1,1}) & \tilde{k}_\sigma(x_{n,1}, x_{1,2}) & \cdots & \cdots & \tilde{k}_\sigma(x_{n,1}, x_{n,1}) & \tilde{k}_\sigma(x_{n,1}, x_{n,2}) \\ \tilde{k}_\sigma(x_{n,2}, x_{1,1}) & \tilde{k}_\sigma(x_{n,2}, x_{1,2}) & \cdots & \cdots & \tilde{k}_\sigma(x_{n,2}, x_{n,1}) & \tilde{k}_\sigma(x_{n,2}, x_{n,2}) \end{bmatrix}.$$

Examining the second term of the ETISE yields

$$\begin{aligned}
\mathbb{E}_q [q_{w,\alpha}] &\approx \sum_{m=1}^M \sum_{r=1}^n \sum_{r'=1}^2 \sum_{s=1}^n \sum_{s'=1}^2 w_m \alpha_{m,r,r'} \alpha_{m,s,s'} \hat{h}(r, r', s, s') \\
&= \begin{cases} \frac{1}{n-2} \sum_m \sum_{r,r'} \sum_{s,s'} w_m \alpha_{m,r,r'} \alpha_{m,s,s'} \sum_{i \in [n] \setminus \{r,s\}} k_\sigma(x_{i,1}, x_{r,r'}) k_\sigma(x_{i,2}, x_{s,s'}), & r \neq s \\ \frac{1}{n-1} \sum_m \sum_{r,r'} \sum_{s,s'} w_m \alpha_{m,r,r'} \alpha_{m,s,s'} \sum_{i \in [n] \setminus \{r\}} k_\sigma(x_{i,1}, x_{r,r'}) k_\sigma(x_{i,2}, x_{r,s'}), & r = s \end{cases} \\
&= \sum_{m=1}^M w_m (\alpha'_m C \alpha_m),
\end{aligned}$$

where C is given by

$$C = \begin{bmatrix} \hat{h}(1, 1, 1, 1) & \hat{h}(1, 1, 1, 2) & \dots & \dots & \hat{h}(1, 1, n, 1) & \hat{h}(1, 1, n, 2) \\ \hat{h}(1, 2, 1, 1) & \hat{h}(1, 2, 1, 2) & \dots & \dots & \hat{h}(1, 2, n, 1) & \hat{h}(1, 2, n, 2) \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ \hat{h}(n, 1, 1, 1) & \hat{h}(n, 1, 1, 2) & \dots & \dots & \hat{h}(n, 1, n, 1) & \hat{h}(n, 1, n, 2) \\ \hat{h}(n, 2, 1, 1) & \hat{h}(n, 2, 1, 2) & \dots & \dots & \hat{h}(n, 2, n, 1) & \hat{h}(n, 2, n, 2) \end{bmatrix}.$$

The diagonal blocks of size two of the matrix C use the leave one out estimator, while the other entries use the leave two out estimator.

B.2 Coreset Approach

In the full problem the matrices $G, C \in \mathbb{R}^{2n \times 2n}$ grow linearly with the data. This can make the proposed optimization problem (6) costly to solve, as the complexity of gradient calculations are quadratic in the dimensions of G, C . Additionally, the complexity of evaluating out-of-sample data is quadratic in n for general KDEs. The motivation of the coresnet approach is to reduce this complexity.

KDEs traditionally center kernels at the location of each observation, i.e., $k_\sigma(\cdot, x_{i,i'})$, where we call $x_{i,i'}$ the kernel center. Rather than constraining the wKDE to have kernels centered at the observations, we can formulate the optimization problem with R kernel centers $z_r \in \mathbb{R}^d$ for some suitably chosen z_r . Additionally, choosing $R \ll n$ will substantially reduce the complexity of gradient calculations and out-of-sample evaluation. The collection of kernel centers z_r will be our coresnet. We don't provide guarantees for the optimality of any particular coresnet. The coresnet could potentially be chosen as the cluster centers output by some clustering algorithm, some suitable subset of the data, or perhaps via some more principled scheme. In all of our experiments, we chose the coresnet to be cluster centers output by mini-batch k -means, where the number of clusters was chosen to be $R > M$.

For the coresnet approach, the ETISE has the same form but the matrices G and C have the form

$$G = \begin{bmatrix} \tilde{k}_\sigma(z_1, z_1) & \cdots & \cdots & \tilde{k}_\sigma(z_1, z_R) \\ \vdots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ \tilde{k}_\sigma(z_R, z_1) & \cdots & \cdots & \tilde{k}_\sigma(z_R, z_R) \end{bmatrix},$$

$$C = \frac{1}{n} \sum_{i=1}^n C_i = \frac{1}{n} \sum_{i=1}^n \begin{bmatrix} k_\sigma(x_{i,1}, z_1)k_\sigma(x_{i,2}, z_1) & \cdots & \cdots & k_\sigma(x_{i,1}, z_1)k_\sigma(x_{i,2}, z_R) \\ \vdots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ k_\sigma(x_{i,1}, z_R)k_\sigma(x_{i,2}, z_1) & \cdots & \cdots & k_\sigma(x_{i,1}, z_R)k_\sigma(x_{i,2}, z_R) \end{bmatrix}$$

This is derived in the same way as the full problem, replacing the kernel centers $x_{r,r'}, x_{s,s'}$ with the coresnet z_r , and using the "leave-one-out" estimator in place of the LOO/LTO estimator \hat{h} .

B.3 Algorithm

Though the problem (7) is nonconvex, we observe that a properly initialized alternating projected stochastic gradient descent (APSGD) procedure produces good solutions in practice. Pseudocode for the APSGD algorithm for solving (7) is given in Algorithm 1. We mention that the projections Π_Δ are onto the probability simplex, a decaying step size $\eta^{(t)}$ is used, and stochasticity is introduced via the matrix $C^{(t)}$, which is a mini-batch version of C defined by $C_{a,b}^{(t)} = \frac{1}{|\Omega^{(t)} \setminus \{a,b\}|} \sum_{i \in |\Omega^{(t)} \setminus \{a,b\}|} k_\sigma(X_{i,1}, X_{\lfloor \frac{a}{2} \rfloor, a \bmod 2})k_\sigma(X_{i,2}, X_{\lfloor \frac{b}{2} \rfloor, b \bmod 2})$, where $\Omega^{(t)}$ is the index set corresponding to the t^{th} mini-batch.

Algorithm 1 Alternating Projected SGD

```

1: init:  $\alpha^{(0)}, w^{(0)}, \eta^{(0)}$ 
2: procedure APSGD( $\alpha^{(0)}, w^{(0)}, \eta^{(0)}$ )
3:   Form  $G$ 
4:   for  $t = 1, 2, \dots$  do
5:     Take a minibatch of paired observations indexed by  $\Omega^{(t)}$ 
6:     Form  $C^{(t)}$  from the minibatch according to the definition of  $C$ 
7:      $w^{(t)} = \Pi_\Delta(w^{(t-1)} - \eta^{(t)} \nabla_w \hat{J}(w^{(t-1)}, \alpha^{(t-1)}))$ 
8:     for  $j = 1, \dots, M$  do
9:        $\alpha_j^{(t)} = \Pi_\Delta(\alpha_j^{(t-1)} - \eta^{(t)} \nabla_{\alpha_j} \hat{J}(w^{(t)}, \alpha_j^{(t-1)}))$ 

```

B.4 Spectral Initialization

We adopt a spectral initialization scheme. First, the initialization is presented for the full problem, then we adapt it to the coresnet approach. The idea here is, given some estimator of q , lets say \tilde{q} , to find a low rank approximation of \tilde{q}

$$\tilde{q}(x, y) \approx \sum_{i=1}^M \lambda_i \psi_i(x) \psi_i(y)$$

and then to use λ_i and ψ_i as starting points for our mixture weights and components. We do this by using the full grouped sample data as an estimate of q which we transform into a linear operator and decompose using a functional eigenvector decomposition.

We begin with a standard KDE applied to our full samples using a product kernel:

$$f_\sigma(y, y') = \frac{1}{2n} \sum_{i=1}^n k_\sigma(y, x_{i,1}) k_\sigma(y', x_{i,2}) + k_\sigma(y, x_{i,2}) k_\sigma(y', x_{i,1}).$$

Note that we include centers at both $(x_{i,1}, x_{i,2})$ and $(x_{i,2}, x_{i,1})$ so our KDE is symmetric in y, y' .

By Lemmas 5.1 and 8.2 of Vandermeulen and Scott [3], f_σ can be viewed as an element of a tensor product space $L^2(\mathbb{R}^d) \otimes L^2(\mathbb{R}^d)$ as follows

$$f_\sigma = \frac{1}{2n} \sum_{i=1}^n k_\sigma(\cdot, x_{i,1}) \otimes k_\sigma(\cdot, x_{i,2}) + k_\sigma(\cdot, x_{i,2}) \otimes k_\sigma(\cdot, x_{i,1}).$$

By the Lemmas referenced above, there is a unitary transformation on the KDE f_σ such that it can be viewed as a linear operator $T : L^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d)$ given by

$$T(g) := \sum_{i=1}^n k_\sigma(\cdot, x_{i,1}) \langle k_\sigma(\cdot, x_{i,2}), g(\cdot) \rangle_{L^2} + k_\sigma(\cdot, x_{i,2}) \langle k_\sigma(\cdot, x_{i,1}), g(\cdot) \rangle_{L^2}, \quad \forall g \in L^2(\mathbb{R}^d)$$

which is symmetric since it includes $k_\sigma(\cdot, x_{i,1}) \langle k_\sigma(\cdot, x_{i,2}), g(\cdot) \rangle_{L^2}$ and $k_\sigma(\cdot, x_{i,2}) \langle k_\sigma(\cdot, x_{i,1}), g(\cdot) \rangle_{L^2}$ terms. We have removed the $1/(2n)$ coefficient since it will not affect the spectral decomposition. For any $g \in L^2$ the quantity $\langle k_\sigma(\cdot, x_{i,i'}), g(\cdot) \rangle_{L^2}$ will be a finite scalar, so $T(g)$ will be a linear combination of the $k_\sigma(\cdot, x_{i,i'})$. Therefore, eigenvectors of the above linear operator will have the form $g(\cdot) = \sum_{j,j'} \beta_{j,j'} k_\sigma(\cdot, x_{j,j'})$ since T applied to any vector must lie in the span of $k_\sigma(\cdot, x_{j,j'})$. Evaluating T on vectors of this form (not necessarily an eigenvector) will yield

$$\begin{aligned} T(g) &:= \sum_{i=1}^n \left\{ k_\sigma(\cdot, x_{i,1}) \langle k_\sigma(\cdot, x_{i,2}), \sum_{j,j'} \beta_{j,j'} k_\sigma(\cdot, x_{j,j'}) \rangle_{L^2} \right. \\ &\quad \left. + k_\sigma(\cdot, x_{i,2}) \langle k_\sigma(\cdot, x_{i,1}), \sum_{j,j'} \beta_{j,j'} k_\sigma(\cdot, x_{j,j'}) \rangle_{L^2} \right\} \\ &= \sum_i \zeta_{i,1} k_\sigma(\cdot, x_{i,1}) + \zeta_{i,2} k_\sigma(\cdot, x_{i,2}) \end{aligned}$$

where $\zeta_{i,1} = \sum_{j,j'} \beta_{j,j'} \tilde{k}_\sigma(x_{i,2}, x_{j,j'})$, $\zeta_{i,2} = \sum_{j,j'} \beta_{j,j'} \tilde{k}_\sigma(x_{i,1}, x_{j,j'})$ and $\tilde{k}_\sigma(y, y')$ is defined in Appendix B.

Define the ordering of the elements of β and ζ by

$$\begin{aligned} \beta &= [\beta_{1,1}, \beta_{1,2}, \beta_{2,1}, \beta_{2,2}, \dots, \beta_{n,1}, \beta_{n,2}]', \\ \zeta &= [\zeta_{1,1}, \zeta_{1,2}, \zeta_{2,1}, \zeta_{2,2}, \dots, \zeta_{n,1}, \zeta_{n,2}]'. \end{aligned}$$

Then we have

$$\zeta = \bar{G}\beta,$$

where

$$\bar{G} = \begin{bmatrix} \tilde{k}_\sigma(x_{1,2}, x_{1,1}) & \tilde{k}_\sigma(x_{1,2}, x_{1,2}) & \tilde{k}_\sigma(x_{1,2}, x_{2,1}) & \cdots & \tilde{k}_\sigma(x_{1,2}, x_{n,1}) & \tilde{k}_\sigma(x_{1,2}, x_{n,2}) \\ \tilde{k}_\sigma(x_{1,1}, x_{1,1}) & \tilde{k}_\sigma(x_{1,1}, x_{1,2}) & \tilde{k}_\sigma(x_{1,1}, x_{2,1}) & \cdots & \tilde{k}_\sigma(x_{1,1}, x_{n,1}) & \tilde{k}_\sigma(x_{1,1}, x_{n,2}) \\ \tilde{k}_\sigma(x_{2,2}, x_{1,1}) & \tilde{k}_\sigma(x_{2,2}, x_{1,2}) & \tilde{k}_\sigma(x_{2,2}, x_{2,1}) & \cdots & \tilde{k}_\sigma(x_{2,2}, x_{n,1}) & \tilde{k}_\sigma(x_{2,2}, x_{n,2}) \\ \tilde{k}_\sigma(x_{2,1}, x_{1,1}) & \tilde{k}_\sigma(x_{2,1}, x_{1,2}) & \tilde{k}_\sigma(x_{2,1}, x_{2,1}) & \cdots & \tilde{k}_\sigma(x_{2,1}, x_{n,1}) & \tilde{k}_\sigma(x_{2,1}, x_{n,2}) \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \tilde{k}_\sigma(x_{n,2}, x_{1,1}) & \tilde{k}_\sigma(x_{n,2}, x_{1,2}) & \tilde{k}_\sigma(x_{n,2}, x_{2,1}) & \cdots & \tilde{k}_\sigma(x_{n,2}, x_{n,1}) & \tilde{k}_\sigma(x_{n,2}, x_{n,2}) \\ \tilde{k}_\sigma(x_{n,1}, x_{1,1}) & \tilde{k}_\sigma(x_{n,1}, x_{1,2}) & \tilde{k}_\sigma(x_{n,1}, x_{2,1}) & \cdots & \tilde{k}_\sigma(x_{n,1}, x_{n,1}) & \tilde{k}_\sigma(x_{n,1}, x_{n,2}) \end{bmatrix}$$

is a row permutation of G defined for the full problem, obtained by exchanging rows corresponding to the first and second elements of each paired observation. The takeaway is that the coefficients of the eigenvectors of T are given by the right eigenvectors of \bar{G} . In particular, we will take the right eigenvectors of \bar{G} corresponding to the M largest real eigenvalues with the intuition that they will capture the dominant modes of T . Note that these eigenvectors will contain real-valued entries by the spectral theorem since T is symmetric. We call these eigenvectors the first, second, and so on. It should be noted that \bar{G} is not a symmetric matrix and so eigenvectors should be found according to, for example, the power iteration or orthogonal iteration. In general the eigenvectors of \bar{G} will have negative entries and not sum to one, so we project these eigenvectors onto the probability simplex to obtain the non-negative weights for initialization where we take α_1 to be the projection of the first eigenvector of \bar{G} , α_2 to be the second, and so on. For the initial w_i , We take w_1 to be the first eigenvalue of \bar{G} , w_2 to be the second and so on, projecting the resulting $w = [w_1, w_2, \dots, w_M]'$ onto the probability simplex.

Coreset Approach Initialization Initialization for the coresnet approach is similar. Since we assume no relationship between z_i , we don't have the same notion of using paired kernel centers even though our data is still paired. However, we can still write the KDE using a product kernel over the coresnet as

$$f_\sigma(y, y') = \frac{1}{R} \sum_{i=1}^R k_\sigma(y, z_i) k_\sigma(y', z_i).$$

This KDE is symmetric since the same kernel center is used in each term of the product kernel. Again appealing to Lemmas 5.1 and 8.2 of Vandermeulen and Scott [3], f_σ can be viewed as an element of a tensor product space $L^2(\mathbb{R}^d) \otimes L^2(\mathbb{R}^d)$ as follows

$$f_\sigma = \frac{1}{R} \sum_{i=1}^R k_\sigma(\cdot, z_i) \otimes k_\sigma(\cdot, z_i).$$

By the Lemmas referenced above, there is a unitary transformation on the KDE f_σ such that it can be viewed as a linear operator $T : L^2(\mathbb{R}^d) \rightarrow L^2(\mathbb{R}^d)$ given by

$$T(g) = \sum_{i=1}^R k_\sigma(\cdot, z_i) \langle k_\sigma(\cdot, z_i), g(\cdot) \rangle_{L^2}, \quad \forall g \in L^2.$$

By the same argument as the full problem, the eigenvectors of the above operator have the form $g(\cdot) = \sum_j \beta_j k_\sigma(\cdot, z_j)$. Applying T to a vector of this form (not necessarily an eigenvector), we have

$$\begin{aligned} T\left(\sum_j \beta_j k_\sigma(\cdot, z_j)\right) &= \sum_{i=1}^R k_\sigma(\cdot, z_i) \langle k_\sigma(\cdot, z_i), \sum_j \beta_j k_\sigma(\cdot, z_j) \rangle_{L^2} \\ &= \sum_{i=1}^R \zeta_i k_\sigma(\cdot, z_i), \end{aligned}$$

where $\zeta_i = \sum_j \beta_j k_\sigma(\cdot, z_i)$. In this setting we have the standard ordering, $\beta = [\beta_1, \beta_2, \dots, \beta_R]'$ and $\zeta = [\zeta_1, \zeta_2, \dots, \zeta_R]'$. In matrix form, the relationship between β and ζ is given by

$$\zeta = G\beta,$$

where G is as previously defined for the coresnet approach. No row permutation is needed in this setting as both centers of our product kernel are the same. From this point, the initialization scheme is essentially the same as for the full problem, but using the eigenvectors and eigenvalues of G . One key difference is that G is a symmetric matrix, so the eigenvalues and eigenvectors of G can be found using any standard solver.

C Proof and General Form of Theorem 1

Theorem 1 was presented in the main paper for groups of size $N = 2$. Here, we provide the proof for groups of size two, as well as the proof for the more general case of arbitrary group size $N > 2$. One tool we will use is Hoeffding's inequality for independent bounded random variables, which we state here for completeness.

Theorem. *Hoeffding's Inequality: Let V_1, V_2, \dots, V_n be independent bounded random variables such that $a_i \leq V_i \leq b_i$ with probability one. If $S_n = \sum_{i=1}^n V_i$, then for all $t > 0$*

$$P\left\{\left|S_n - \mathbb{E}\{S_n\}\right| \geq t\right\} \leq 2 \exp\left\{-\frac{2t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right\}.$$

C.1 Proof of Theorem 1: Groups of Size Two

We restate Theorem 1 for convenience.

Theorem 1. *Let $\epsilon > 0$ and set $\delta = 8(n^2 - n) \exp\{-\frac{\sigma^{4d}(n-2)\epsilon^2}{8C_k^4}\} + 8n \exp\{-\frac{\sigma^{4d}(n-1)\epsilon^2}{8C_k^4}\}$. With probability at least $1 - \delta$ the following holds:*

$$\|q - q_{\hat{w}, \hat{\alpha}}\|_2^2 \leq \inf_{w \in \Delta^M, \alpha \in \Delta_M^{2n}} \|q - q_{w, \alpha}\|_2^2 + \epsilon.$$

Proof. Our goal is to bound $|J(w, a) - \hat{J}(w, a)|$ uniformly over $w \in \Delta^M$, $\alpha \in \Delta_M^{2n}$. Recall the

following definitions

$$\begin{aligned}
h(r, r', s, s') &:= \int k_\sigma(x, x_{r,r'}) k_\sigma(x', x_{s,s'}) q(x, x') dx dx' \\
\hat{h}(r, r', s, s') &:= \begin{cases} \hat{h}_{\text{LTO}}(r, r', s, s'), & r \neq s \\ \hat{h}_{\text{LOO}}(r, r', s'), & r = s \end{cases} \\
\hat{h}_{\text{LOO}}(r, r', r'') &:= \frac{1}{n-1} \sum_{i \in [n] \setminus \{r\}} k_\sigma(x_{i,1}, x_{r,r'}) k_\sigma(x_{i,2}, x_{r,r''}) \\
\hat{h}_{\text{LTO}}(r, r', s, s') &:= \frac{1}{n-2} \sum_{i \in [n] \setminus \{r,s\}} k_\sigma(x_{i,1}, x_{r,r'}) k_\sigma(x_{i,2}, x_{s,s'}).
\end{aligned}$$

The use of the leave one out (LOO) and leave two out (LTO) estimators above is to ensure independence so that we will be able to apply Hoeffding's inequality. We have

$$\begin{aligned}
P_q \{ \sup_{\substack{w \in \Delta_M^M \\ \alpha \in \Delta_M^{2n}}} |J(w, \alpha) - \hat{J}(w, \alpha)| > \frac{\epsilon}{2} \} \\
&= P_q \left\{ \sup_{\substack{w \in \Delta_M^M \\ \alpha \in \Delta_M^{2n}}} \left| \sum_{m=1}^M w_m \sum_{r=1}^n \sum_{s=1}^n \sum_{r'=1}^2 \sum_{s'=1}^2 \alpha_{m,r,r'} \alpha_{m,s,s'} h(r, r', s, s') \right. \right. \\
&\quad \left. \left. - \sum_{m=1}^M w_m \sum_{r=1}^n \sum_{s=1}^n \sum_{r'=1}^2 \sum_{s'=1}^2 \alpha_{m,r,r'} \alpha_{m,s,s'} \hat{h}(r, r', s, s') \right| > \frac{\epsilon}{4} \right\} \\
&\leq P_q \left\{ \sup_{\substack{w \in \Delta_M^M \\ \alpha \in \Delta_M^{2n}}} \sum_m \sum_{r,s} \sum_{r',s'} w_m \alpha_{m,r,r'} \alpha_{m,s,s'} \left| h(r, r', s, s') - \hat{h}(r, r', s, s') \right| > \frac{\epsilon}{4} \right\} \\
&\leq P_q \left\{ \max_{r,s,r',s'} \left| h(r, r', s, s') - \hat{h}(r, r', s, s') \right| > \frac{\epsilon}{4} \right\} \\
&\leq \sum_{r,s} \sum_{r',s'} P_q \left\{ \left| h(r, r', s, s') - \hat{h}(r, r', s, s') \right| > \frac{\epsilon}{4} \right\}.
\end{aligned}$$

The second step above is due to the triangle inequality, and the penultimate step is due to simplex constraints on w, α . Let $k_i(r, r', s, s') := k_\sigma(x_{i,1}, x_{r,r'}) k_\sigma(x_{i,2}, x_{s,s'})$. Noting that $h(r, r', s, s') =$

$$\mathbb{E}_{(x_{i,1}, x_{i,2}) \sim q} \{k_i(r, r', s, s')\},$$

$$\begin{aligned}
& P_q \left\{ |h(r, r', s, s') - \hat{h}(r, r', s, s')| > \frac{\epsilon}{4} \right\} \\
&= \begin{cases} P_q \left\{ \left| \mathbb{E}_{(x_{i,1}, x_{i,2}) \sim q} \{k_i(r, r', s, s')\} - \frac{1}{n-2} \sum_{i \in [n] \setminus \{r, s\}} k_i(r, r', s, s') \right| > \frac{\epsilon}{4} \right\}, & r \neq s \\ P_q \left\{ \left| \mathbb{E}_{(x_{i,1}, x_{i,2}) \sim q} \{k_i(r, r', s, s')\} - \frac{1}{n-1} \sum_{i \in [n] \setminus \{r\}} k_i(r, r', s, s') \right| > \frac{\epsilon}{4} \right\}, & r = s \end{cases} \\
&= \begin{cases} P_q \left\{ \left| \frac{1}{n-2} \sum_{i \in [n] \setminus \{r, s\}} \mathbb{E}_{(x_{i,1}, x_{i,2}) \sim q} \{k_i(r, r', s, s')\} - k_i(r, r', s, s') \right| > \frac{\epsilon}{4} \right\}, & r \neq s \\ P_q \left\{ \left| \frac{1}{n-1} \sum_{i \in [n] \setminus \{r\}} \mathbb{E}_{(x_{i,1}, x_{i,2}) \sim q} \{k_i(r, r', s, s')\} - k_i(r, r', s, s') \right| > \frac{\epsilon}{4} \right\}, & r = s \end{cases} \\
&= \begin{cases} P_q \left\{ \left| \sum_{i \in [n] \setminus \{r, s\}} \mathbb{E}_{(x_{i,1}, x_{i,2}) \sim q} \{k_i(r, r', s, s')\} - k_i(r, r', s, s') \right| > \frac{(n-2)\epsilon}{4} \right\}, & r \neq s \\ P_q \left\{ \left| \sum_{i \in [n] \setminus \{r\}} \mathbb{E}_{(x_{i,1}, x_{i,2}) \sim q} \{k_i(r, r', s, s')\} - k_i(r, r', s, s') \right| > \frac{(n-1)\epsilon}{4} \right\}, & r = s \end{cases}
\end{aligned}$$

The terms $k_i(r, r', s, s')$ are independent random variables due to use of the LOO/LTO estimator. By assumption, $0 \leq k_i(r, r', s, s') \leq C_k^2 \sigma^{-2d}$ so the k_i are bounded for fixed $\sigma > 0$. We apply Hoeffding's inequality

$$\begin{aligned}
P_q \left\{ |h(r, r', s, s') - \hat{h}(r, r', s, s')| > \frac{\epsilon}{4} \right\} &\leq \begin{cases} 2 \exp\left\{-\frac{2(n-2)^2\epsilon^2}{16(n-2)C_k^4\sigma^{-4d}}\right\}, & r \neq s \\ 2 \exp\left\{-\frac{2(n-1)^2\epsilon^2}{16(n-1)C_k^4\sigma^{-4d}}\right\}, & r = s \end{cases} \\
&\leq \begin{cases} 2 \exp\left\{-\frac{\sigma^{4d}(n-2)\epsilon^2}{8C_k^4}\right\}, & r \neq s \\ 2 \exp\left\{-\frac{\sigma^{4d}(n-1)\epsilon^2}{8C_k^4}\right\}, & r = s \end{cases}. \tag{8}
\end{aligned}$$

Substituting backward we obtain the desired upper bound

$$\begin{aligned}
P_q \left\{ \sup_{\substack{w \in \Delta^M \\ \alpha \in \Delta_M^{2n}}} |J(w, \alpha) - \hat{J}(w, \alpha)| > \frac{\epsilon}{2} \right\} &\leq \sum_{r,s} \sum_{r',s'} \begin{cases} 2 \exp\left\{-\frac{\sigma^{4d}(n-2)\epsilon^2}{8C_k^4}\right\}, & r \neq s \\ 2 \exp\left\{-\frac{\sigma^{4d}(n-1)\epsilon^2}{8C_k^4}\right\}, & r = s \end{cases} \\
&= 8(n^2 - n) \exp\left\{-\frac{\sigma^{4d}(n-2)\epsilon^2}{8C_k^4}\right\} + 8n \exp\left\{-\frac{\sigma^{4d}(n-1)\epsilon^2}{8C_k^4}\right\}.
\end{aligned}$$

Letting $\delta = 8(n^2 - n) \exp\left\{-\frac{\sigma^{4d}(n-2)\epsilon^2}{8C_k^4}\right\} + 8n \exp\left\{-\frac{\sigma^{4d}(n-1)\epsilon^2}{8C_k^4}\right\}$, we have

$$J(w, \alpha) - \frac{\epsilon}{2} \leq \hat{J}(w, \alpha) \leq J(w, \alpha) + \frac{\epsilon}{2} \quad \forall w, \alpha \tag{9}$$

with probability at least $1 - \delta$. Thus, with probability at least $1 - \delta$, for any $w \in \Delta^M, \alpha \in \Delta_M^{2n}$

$$\begin{aligned}
J(\hat{w}, \hat{\alpha}) &\leq \hat{J}(\hat{w}, \hat{\alpha}) + \frac{\epsilon}{2} \\
&\leq \hat{J}(w, \alpha) + \frac{\epsilon}{2} \\
&\leq J(w, \alpha) + \epsilon,
\end{aligned}$$

where $\hat{w}, \hat{\alpha}$ are defined in (6). Then with probability at least $1 - \delta$

$$\hat{J}(\hat{w}, \hat{\alpha}) \leq \inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{2n}}} J(w, \alpha) + \epsilon. \quad (10)$$

Combining (10) with the definition of the ISE shows, with probability at least $1 - \delta$,

$$\|q - q_{\hat{w}, \hat{\alpha}}\|_2^2 \leq \inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{2n}}} \|q - q_{w, \alpha}\|_2^2 + \epsilon.$$

□

C.2 Theorem 1: Arbitrary Group Size

C.2.1 Preliminaries

Before beginning the proof, we start by redefining q , $q_{w, \alpha}$, J , and \hat{J} for arbitrary group size. Once this is done, the proof will follow the same basic steps as the proof for groups of size two.

Suppose we change the problem setup only in the size of the grouped observations. Consider grouped observations of size N . Consider a set of n grouped observations $\mathbf{x}_1, \dots, \mathbf{x}_n$ with $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,N}) \in \mathbb{R}_N^d := \underbrace{\mathbb{R}^d \times \dots \times \mathbb{R}^d}_N$ drawn i.i.d. from

$$q(y_1, y_2, \dots, y_N) = \sum_{m=1}^M w_m^* p_m^*(y_1) p_m^*(y_2) \dots p_m^*(y_N), \quad y_1, y_2, \dots, y_N \in \mathbb{R}^d. \quad (11)$$

Similar to the paired observation setting, a wKDE in this setting will have the form

$$p(y; \theta) = \sum_{r=1}^n \sum_{r'=1}^N \theta_{r,r'} k_\sigma(y, x_{r,r'}).$$

We may write the corresponding estimator of q

$$q_{w, \alpha}(y_1, y_2, \dots, y_N) = \sum_{m=1}^M w_m p(y_1; \alpha_m) p(y_2; \alpha_m) \dots p(y_N; \alpha_m)$$

where $\alpha_m = [\alpha_{m,1,1} \dots \alpha_{m,1,N} \dots \dots \alpha_{m,n,1} \dots \alpha_{m,n,N}]' \in \Delta^{Nn}$ for $m = 1, \dots, M$, with $\alpha_{m,r,r'}$ corresponding to the weight of the kernel centered at $x_{r,r'}$ in the estimate of the m^{th} mixture component.

In what follows we use $\sum_{r,r'} := \sum_{r_1,r'_1} \dots \sum_{r_N,r'_N}$ to ease notation. Similar to the paired sample case, we define

$$\begin{aligned} J(w, \alpha) &:= \int q_{w, \alpha}^2(y_1, \dots, y_N) dy_1 \dots dy_N - 2 \sum_{m, r, r'} \left(\prod_{i \in [N]} w_m \alpha_{m, r_i, r'_i} \right) h(r_1, r'_1, \dots, r_N, r'_N) \\ \hat{J}(w, \alpha) &:= \int q_{w, \alpha}^2(y_1, \dots, y_N) dy_1 \dots dy_N - 2 \sum_{m, r, r'} \left(\prod_{i \in [N]} w_m \alpha_{m, r_i, r'_i} \right) \hat{h}(r_1, r'_1, \dots, r_N, r'_N), \end{aligned}$$

where

$$h(r_1, r'_1, \dots, r_N, r'_N) := \int k_\sigma(y_1, x_{r_1, r'_1}) \dots k_\sigma(y_N, x_{r_N, r'_N}) q(y_1, \dots, y_N) dy_1 \dots dy_N,$$

$$\hat{h} := \hat{h}_{\text{LNO}}(r_1, r'_1, \dots, r_N, r'_N) := \frac{1}{n-N} \sum_{i \in [n] \setminus L(r_1, r_2, \dots, r_N)} k_\sigma(x_{i,1}, x_{r_1, r'_1}) \dots k_\sigma(x_{i,N}, x_{r_N, r'_N}),$$

where $L(r_1, r_2, \dots, r_N)$ is any subset of $[n]$ containing $\{r_1, r_2, \dots, r_N\}$ and having cardinality N . If r_1, r_2, \dots, r_N are not distinct, the additional indices can be chosen arbitrarily. For simplicity, we use a leave- N -out (LNO) estimator \hat{h}_{LNO} rather than a hybrid estimator like we used in the case of paired observations. As in the paired observation setting, we define

$$(\hat{w}, \hat{\alpha}) := \arg \min_{w \in \Delta^M, \alpha \in \Delta_M^{Nn}} \hat{J}(w, \alpha),$$

and similarly define $\hat{q} := q_{\hat{w}, \hat{\alpha}}$. Whenever \hat{w} , $\hat{\alpha}$, q , or $q_{\hat{w}, \hat{\alpha}}$ are referenced in the arbitrary group size setting, we will be referring to these estimators.

C.2.2 Proof of Theorem 1: Arbitrary Group Size

We now state Theorem 1 for arbitrary group size.

Theorem 1a. *Given grouped observations of size N , let $\epsilon > 0$ and $\delta = 2(Nn)^N \exp\left\{-\frac{\sigma^{2Nd}(n-N)\epsilon^2}{8C_k^{2N}}\right\}$. With probability at least $1 - \delta$ the following holds:*

$$\|q - q_{\hat{w}, \hat{\alpha}}\|_2^2 \leq \inf_{w \in \Delta^M, \alpha \in \Delta_M^{Nn}} \|q - q_{w, \alpha}\|_2^2 + \epsilon.$$

Proof. The proof proceeds as in the paired observation setting. In particular,

$$\begin{aligned} P_q \left\{ \sup_{\substack{w \in \Delta^M \\ \alpha \in \Delta_M^{Nn}}} |J(w, \alpha) - \hat{J}(w, \alpha)| > \frac{\epsilon}{2} \right\} \\ \leq P_q \left\{ \sup_{\substack{w \in \Delta^M \\ \alpha \in \Delta_M^{Nn}}} \sum_{m, r, r'} \prod_{i \in [N]} w_m \alpha_{m, r_i, r'_i} \left| h(r_1, r'_1, \dots, r_N, r'_N) - \hat{h}(r_1, r'_1, \dots, r_N, r'_N) \right| > \frac{\epsilon}{4} \right\} \\ \leq P_q \left\{ \max_{r_1, r'_1, \dots, r_N, r'_N} \left| h(r_1, r'_1, \dots, r_N, r'_N) - \hat{h}(r_1, r'_1, \dots, r_N, r'_N) \right| > \frac{\epsilon}{4} \right\} \\ \leq \sum_{r, r'} P_q \left\{ \left| h(r_1, r'_1, \dots, r_N, r'_N) - \hat{h}(r_1, r'_1, \dots, r_N, r'_N) \right| > \frac{\epsilon}{4} \right\} \end{aligned}$$

The first step above is due to the triangle inequality, and the penultimate step is due to simplex constraints on w, α . Let $k_i(r_1, r'_1, \dots, r_N, r'_N) := k_\sigma(x_{i,1}, x_{r_1, r'_1}) k_\sigma(x_{i,2}, x_{r_2, r'_2}) \dots k_\sigma(x_{i,N}, x_{r_N, r'_N})$. Noting that $h(r_1, r'_1, \dots, r_N, r'_N) = \mathbb{E}_q\{k_\sigma(x_{i,1}, x_{r_1, r'_1}) \dots k_\sigma(x_{i,N}, x_{r_N, r'_N})\}$, we have

$$\begin{aligned} P_q \left\{ \left| h(r_1, r'_1, \dots, r_N, r'_N) - \hat{h}(r_1, r'_1, \dots, r_N, r'_N) \right| > \frac{\epsilon}{4} \right\} \\ = P_q \left\{ \left| \sum_{i \in [n] \setminus L(r_1, \dots, r_N)} \mathbb{E}_{(x_{i,1}, \dots, x_{i,N}) \sim q} \{k_i(r_1, r'_1, \dots, r_N, r'_N)\} \right. \right. \\ \left. \left. - k_i(r_1, r'_1, \dots, r_N, r'_N) \right| > \frac{(n-N)\epsilon}{4} \right\} \end{aligned}$$

The terms $k_i(r_1, r'_1, \dots, r_N, r'_N)$ are independent random variables due to use of the LNO estimator. By assumption, $0 \leq k_i(r_1, r'_1, \dots, r_N, r'_N) \leq C_k^N \sigma^{-Nd}$ so the k_i are bounded for fixed $\sigma > 0$. We apply Hoeffding's inequality

$$\begin{aligned} P_q \left\{ \left| h(r_1, r'_1, \dots, r_N, r'_N) - \hat{h}(r_1, r'_1, \dots, r_N, r'_N) \right| > \frac{\epsilon}{4} \right\} &\leq 2 \exp \left\{ - \frac{2(n-N)^2 \epsilon^2}{16(n-N)C_k^{2N} \sigma^{-2Nd}} \right\} \\ &= 2 \exp \left\{ - \frac{\sigma^{2Nd}(n-N)\epsilon^2}{8C_k^{2N}} \right\}. \end{aligned}$$

Substituting backward we obtain the desired upper bound

$$\begin{aligned} P_q \left\{ \sup_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{Nn}}} |J(w, \alpha) - \hat{J}(w, \alpha)| > \frac{\epsilon}{2} \right\} &\leq \sum_{r_1, r'_1} \dots \sum_{r_N, r'_N} 2 \exp \left\{ - \frac{\sigma^{2Nd}(n-N)\epsilon^2}{8C_k^{2N}} \right\} \\ &= 2(Nn)^N \exp \left\{ - \frac{\sigma^{2Nd}(n-N)\epsilon^2}{8C_k^{2N}} \right\} \end{aligned}$$

From here the proof is identical to the paired observation case, but with

$$\delta = 2(Nn)^N \exp \left\{ - \frac{\sigma^{2Nd}(n-N)\epsilon^2}{8C_k^{2N}} \right\}.$$

□

D Proof and General Form of Theorem 2

In this section we give the proof of Theorem 2 for groups of size two, before extending it to groups of arbitrary size. For readability, we first present some intermediate results to be used in the main proofs.

D.1 Intermediate Results

We first prove two supporting results.

Lemma 1. *For any $1 \leq p < \infty$, any $f, g \in L^p$, and any integer $a \geq 2$,*

$$\|f^{\times a} - g^{\times a}\|_p \leq \|f\|_p^{a-1} \|f - g\|_p + \|g\|_p \|f^{\times(a-1)} - g^{\times(a-1)}\|_p,$$

where $f^{\times a}(y_1, y_2, \dots, y_a) := f(y_1)f(y_2)\dots f(y_a)$.

Proof. Let $f, g \in L^p$, $1 \leq p < \infty$. Then

$$\begin{aligned}
\|f^{\times a} - g^{\times a}\|_p &= \|f^{\times a} - f^{\times(a-1)} \times g + f^{\times(a-1)} \times g - g^{\times a}\|_p \\
&\leq \|f^{\times a} - f^{\times(a-1)} \times g\|_p + \|f^{\times(a-1)} \times g - g^{\times a}\|_p \\
&= \left(\int |f(x_1) \cdots f(x_{a-1})(f(x_a) - g(x_a))|^p dx_1 \cdots dx_a \right)^{\frac{1}{p}} \\
&\quad + \left(\int |g(x_a)(f(x_1) \cdots f(x_{a-1}) - g(x_1) \cdots g(x_{a-1}))|^p dx_1 \cdots dx_a \right)^{\frac{1}{p}} \\
&= \|f^{\times(a-1)}\|_p \|f - g\|_p + \|g\|_p \|f^{\times(a-1)} - g^{\times(a-1)}\|_p \\
&= \|f\|_p^{a-1} \|f - g\|_p + \|g\|_p \|f^{\times(a-1)} - g^{\times(a-1)}\|_p.
\end{aligned}$$

□

We have the following corollary.

Corollary 2. For any $1 \leq p < \infty$, any $f, g \in L^p$, and any integer $a \geq 2$,

$$\|f^{\times a} - g^{\times a}\|_p \leq \left(\sum_{b=1}^a \|f\|_p^{a-b} \|g\|_p^{b-1} \right) \|f - g\|_p.$$

Proof. The proof is by induction. Lemma 1 provides the base of the recursion for $a = 2$. Now suppose the statement is true for $a \geq 2$. To prove the statement for $a + 1$, we apply Lemma 1 again, together with the induction hypothesis, to get

$$\begin{aligned}
\|f^{\times(a+1)} - g^{\times(a+1)}\|_p &\leq \|f\|_p^a \|f - g\|_p + \|g\|_p \|f^{\times a} - g^{\times a}\|_p \\
&\leq \|f\|_p^a \|f - g\|_p + \|g\|_p \left(\sum_{b=1}^a \|f\|_p^{a-b} \|g\|_p^{b-1} \right) \|f - g\|_p \\
&= \left(\sum_{b=1}^{a+1} \|f\|_p^{a+1-b} \|g\|_p^{b-1} \right) \|f - g\|_p.
\end{aligned}$$

This completes the proof. □

D.2 Proof of Theorem 2: Groups of Size Two

We restate Theorem 2 for convenience.

Theorem 2. If $\sigma \rightarrow 0$ and $\frac{n\sigma^{4d}}{\log n} \rightarrow \infty$ as $n \rightarrow \infty$, then $\|q - q_{\hat{w}, \hat{\alpha}}\|_1 \xrightarrow{a.s.} 0$.

Proof. Lemma 3.1 of [29] states that if $\int \hat{q} = 1$ and $\|\hat{q} - q\|_2 \xrightarrow{a.s.} 0$, then $\|\hat{q} - q\|_1 \xrightarrow{a.s.} 0$. Since $\int \hat{q} = 1$ in our case, our strategy is to show $\|\hat{q} - q\|_2 \xrightarrow{a.s.} 0$. To do this it suffices to show that

$$\|q - \hat{q}\|_2^2 - \inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{N,n}}} \|q - q_{w, \alpha}\|_2^2 \xrightarrow{a.s.} 0 \quad (12)$$

and

$$\inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{Nn}}} \|q - q_{w,\alpha}\|_2 \xrightarrow{a.s.} 0. \quad (13)$$

To show (12), by the Borel-Cantelli lemma, it suffices to show that for all $\epsilon > 0$,

$$\sum_{n=1}^{\infty} P_q \left(\|q - \hat{q}\|_2^2 - \inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{2n}}} \|q - q_{w,\alpha}\|_2^2 \geq \epsilon \right) < \infty.$$

Thus let $\epsilon > 0$. By Theorem 1a, the probability in question is at most

$$\begin{aligned} \delta &= 2(nN)^N \exp \left\{ -\frac{(n-N)\sigma^{2Nd}\epsilon^2}{8C_k^{2N}} \right\} \\ &= 8 \exp \left\{ -2 \log n \left(\frac{(n-2)\sigma^{4d}\epsilon^2}{16C_k^4 \log n} - 1 \right) \right\}. \end{aligned}$$

By assumption on the growth of n and σ , there exists N_ϵ such that for all $n \geq N_\epsilon$,

$$\frac{(n-2)\sigma^{4d}\epsilon^2}{16C_k^4 \log n} \geq 2.$$

For such n we have

$$\delta \leq 8 \exp \{-2 \log n\} = \frac{8}{n^2}$$

which is summable.

To show (13), let w^* be the true mixing weights from (2). For $i = 1, \dots, n$ let e_i be the $m \in [M]$ such that $X_i = (x_{i,1}, x_{i,2}) \stackrel{i.i.d.}{\sim} p_m^*$. Define

$$\begin{aligned} n_m &= |\{i : e_i = m\}|, & m &= 1, \dots, M \\ \alpha_{m,i,1}^* &= \alpha_{m,i,2}^* = \begin{cases} \frac{1}{2n_m}, & e_i = m \\ 0, & \text{otherwise} \end{cases}, & m &= 1, \dots, M \end{aligned}$$

With this "oracle" assignment of weights, $p(x; \alpha_m^*)$ is just the regular KDE for p_m^* . Therefore, we may apply known results for consistency of standard KDEs. In particular, we will apply Theorem 3.1 of [29] which implies

$$\|p(\cdot; \alpha_m^*) - p_m^*\|_2 \xrightarrow{a.s.} 0 \text{ as } n_m \rightarrow \infty \quad (14)$$

provided $k \in L^2$ and $\sum_n \frac{1}{n^2 \sigma_n^d} < \infty$. Both of these conditions are satisfied by assumption in our setting. Furthermore, as $n \rightarrow \infty$ we have $\frac{n_m}{n} \rightarrow w_m^*$ almost surely, and therefore $n_m \rightarrow \infty$ almost surely.

Finally, we have

$$\begin{aligned}
\inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{2n}}} \|q - q_{w,\alpha}\|_2 &\leq \|q - q_{w^*, \alpha^*}\|_2 \\
&= \left\| \sum_{m=1}^M w_m^* (p_m^* \times p_m^* - p(\cdot; \alpha_m^*) \times p(\cdot; \alpha_m^*)) \right\|_2 \\
&\leq \sum_{m=1}^M w_m^* \|p_m^* \times p_m^* - p(\cdot; \alpha_m^*) \times p(\cdot; \alpha_m^*)\|_2 \\
&\leq \sum_{m=1}^M w_m^* (\|p_m^*\|_2 + \|p(\cdot; \alpha_m^*)\|_2) \|p_m^* - p(\cdot; \alpha_m^*)\|_2 \\
&\leq \sum_{m=1}^M w_m^* 3 \|p_m^*\|_2 \|p_m^* - p(\cdot; \alpha_m^*)\|_2 \\
&\xrightarrow{a.s.} 0 \text{ as } n \rightarrow \infty,
\end{aligned}$$

where the fourth step uses Lemma 1 and the fifth step holds for n sufficiently large (a.s.). This completes the proof. \square

D.3 Proof of Theorem 2: Arbitrary Group Size

We consider the problem for arbitrary group size as described in Section C.2.1 of this document. The proof of Theorem 2 for arbitrary group size is similar to the proof for groups of size two. The main difference will be in use of Theorem 1a rather than Theorem 1 to invoke the Borel-Cantelli lemma.

Theorem 2a. *Given grouped observations of size $N \in \mathbb{Z}^+$, if $\sigma \rightarrow 0$ and $\frac{n\sigma^{2Nd}}{\log n} \rightarrow \infty$ as $n \rightarrow \infty$ then $\|q - \hat{q}\|_1 \xrightarrow{a.s.} 0$ as $n \rightarrow \infty$.*

Proof. We will appeal to Lemma 3.1 of [29] as we did for groups of size two. Namely, if $\int \hat{q} = 1$ and $\|\hat{q} - q\|_2 \xrightarrow{a.s.} 0$, then $\|\hat{q} - q\|_1 \xrightarrow{a.s.} 0$. Our strategy again is to show $\|\hat{q} - q\|_2 \xrightarrow{a.s.} 0$. To do this it suffices to show that

$$\|q - \hat{q}\|_2^2 - \inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{Nn}}} \|q - q_{w,\alpha}\|_2^2 \xrightarrow{a.s.} 0 \quad (15)$$

and

$$\inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{Nn}}} \|q - q_{w,\alpha}\|_2 \xrightarrow{a.s.} 0. \quad (16)$$

To show (15), by the Borel-Cantelli lemma, it suffices to show that for all $\epsilon > 0$,

$$\sum_{n=1}^{\infty} P_q \left(\|q - \hat{q}\|_2^2 - \inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{Nn}}} \|q - q_{w,\alpha}\|_2^2 \geq \epsilon \right) < \infty.$$

Thus let $\epsilon > 0$. By Theorem 1a, the probability in question is at most

$$\begin{aligned}\delta &= 2(Nn)^N \exp\left\{-\frac{\sigma^{2Nd}(n-N)\epsilon^2}{8C_k^{2N}}\right\} \\ &= 2N^N \exp\left\{N \log n + \left(\frac{(n-N)\sigma^{2Nd}\epsilon^2}{8C_k^{2N}}\right)\right\} \\ &= 2N^N \exp\left\{-N \log n \left(\frac{(n-N)\sigma^{2Nd}\epsilon^2}{8C_k^{2N}N \log n} - 1\right)\right\}.\end{aligned}$$

By assumption on the growth of n and σ , there exists N_ϵ such that for all $n \geq N_\epsilon$,

$$\frac{(n-N)\sigma^{2Nd}\epsilon^2}{8C_k^{2N}N \log n} \geq 2.$$

For such n we have

$$\delta \leq 2N^N \exp\{-N \log n\} = \frac{2N^N}{n^N}$$

which is summable for $N > 1$.

To show (16), let w^* be the true mixing weights from (11). For $i = 1, \dots, n$ let e_i be the $m \in [M]$ such that $X_i = (x_{i,1}, x_{i,2}, \dots, x_{i,N}) \stackrel{i.i.d.}{\sim} p_m^*$. Define

$$\begin{aligned}n_m &= |\{i : e_i = m\}|, & m &= 1, \dots, M \\ \alpha_{m,i,j}^* &= \begin{cases} \frac{1}{n_m N}, & e_i = m \\ 0, & \text{otherwise} \end{cases}, & m &= 1, \dots, M, \quad j = 1, \dots, N\end{aligned}$$

We are again using an “oracle” assignment of weights, so $p(x; \alpha_m^*)$ is just the regular KDE for p_m^* . Therefore, we may again apply Theorem 3.1 of [29] which implies

$$\|p(\cdot; \alpha_m^*) - p_m^*\|_2 \xrightarrow{a.s.} 0 \text{ as } n_m \rightarrow \infty \tag{17}$$

provided $k \in L^2$ and $\sum_n \frac{1}{n^2 \sigma_n^d} < \infty$. Both of these conditions are satisfied by assumption in our setting. Furthermore, as $n \rightarrow \infty$ we have $\frac{n_m}{n} \rightarrow w_m^*$ almost surely, and therefore $n_m \rightarrow \infty$ almost surely.

Finally, we have

$$\begin{aligned}
\inf_{\substack{w \in \Delta_M \\ \alpha \in \Delta_M^{Nn}}} \|q - q_{w,\alpha}\|_2 &\leq \|q - q_{w^*,\alpha^*}\|_2 \\
&= \left\| \sum_{m=1}^M w_m^* (p_m^{*\times N} - p(\cdot; \alpha_m^*)^{\times N}) \right\|_2 \\
&\leq \sum_{m=1}^M w_m^* \|p_m^{*\times N} - p(\cdot; \alpha_m^*)^{\times N}\|_2 \\
&\leq \sum_{m=1}^M w_m^* \left(\sum_{b=1}^N \|p_m^*\|_2^{N-b} \|p(\cdot; \alpha_m^*)\|_2^{b-1} \right) \|p_m^* - p(\cdot; \alpha_m^*)\|_2 \\
&\leq \sum_{m=1}^M w_m^* \left(\sum_{b=1}^N 2^{b-1} \|p_m^*\|_2^{N-1} \right) \|p_m^* - p(\cdot; \alpha_m^*)\|_2 \\
&\xrightarrow{a.s.} 0 \text{ as } n \rightarrow \infty,
\end{aligned}$$

where the penultimate step uses (17) and Corollary 2, and the final step holds for n sufficiently large (a.s.). This completes the proof. \square

E Background on the Grouped Sample Setting and Proof of Theorem 3

Here we prove Theorem 3. We will be proving a general and more technical version of this theorem, Theorem 4, from which Theorem 3 is a direct consequence. First we will introduce some background to the problem setting which was introduced in [3]. **This section uses its own notation which does not extend to other appendices or main text.**

E.1 Identifiability in the Grouped Sample Setting

We will be concerned with probability measures on a measurable space (Ω, \mathcal{F}) . Let δ be the Dirac measure. Let \mathcal{D} be the set of probability measures on (Ω, \mathcal{F}) . We call a probability measure on \mathcal{D} of the form

$$\mathcal{P} = \sum_{i=1}^m a_i \delta_{\mu_i}$$

a *mixture of measures* [3]. For all mixtures of measures we will assume that $a_i > 0$ for all i and $\mu_i \neq \mu_j$ when $i \neq j$ so that m is the number of distinct mixture components. The *grouped sample* setting from [3] considers the situation where samples come in groups of size n by first sampling a random measure component from a mixture of measures $\gamma \sim \mathcal{P}$, which is then sampled iid n times. So one has access to samples of the form $\mathbf{X} = (X_1, \dots, X_n)$ with $X_1, \dots, X_n \stackrel{iid}{\sim} \gamma$. In this situation the identifiability of \mathcal{P} depends on whether the distribution of \mathbf{X} is uniquely determined by \mathcal{P} and the number of samples per group n . To this end [3] introduced the V_n operator which maps a

mixture of measures to the distribution of \mathbf{X} :

$$V_n \left(\sum_{i=1}^m a_i \delta_{\mu_i} \right) = \sum_{i=1}^m a_i \mu_i^{\times n},$$

where $\mu^{\times n}$ denotes the product measure n times. We note that $n = 1$ corresponds to a typical mixture model where each mixture component is sampled once after being selected and there is no grouped sample structure. For the grouped sample setting [3] introduces the following notion of identifiability.

Definition 1. A mixture of measures, $\mathcal{P} = \sum_{i=1}^m a_i \delta_{\mu_i}$, is called n -identifiable if there does not exist a different mixture of measures $\mathcal{Q} = \sum_{j=1}^{m'} b_j \delta_{\nu_j}$, with $m' \leq m$, such that $V_n(\mathcal{P}) = V_n(\mathcal{Q})$.

A completely rigorous mathematical treatment of the previous notions is a bit involved and can be found in [3]. In [3] it is shown that if the mixture components are jointly irreducible then a mixture of measures is 2-identifiable, if they are linearly independent then they are 3-identifiable, and that any mixture of measures with m components is $(2m - 1)$ -identifiable.

E.2 Notation

Before we state and prove the main theorem of this section we need to first introduce some notation.

Let S_m be the symmetric group over m symbols. Abusing notation slightly we will let the elements of S_m be a group action on $[m]$ as well as \mathbb{R}^m . On \mathbb{R}^m it is defined as the following

$$\sigma \left([x_1, \dots, x_m]^T \right) = [x_{\sigma(1)}, \dots, x_{\sigma(m)}]^T. \quad (18)$$

We also let S_m be an operator where $S_m \cdot x$ is the orbit of x , i.e.

$$S_m \cdot x \triangleq \{\sigma(x) : \sigma \in S_m\}. \quad (19)$$

Recall that for a pair of Hilbert spaces H, H' the direct sum $H \oplus H'$ is a Hilbert space with elements of the form $x \oplus x'$ and inner product defined as $\langle x \oplus x', y \oplus y' \rangle = \langle x, y \rangle + \langle x', y' \rangle$. For a pair of Banach spaces B, B' we define the direct sum via the norm $\|b \oplus b'\|_{B \oplus B'} \triangleq \|b\|_B + \|b'\|_{B'}$ which is itself a Banach space ([42] p. 183).

For a pair of Hilbert spaces H, H' let $H \otimes H'$ be the tensor product of these two spaces and $h \otimes h'$ be the tensor product of vectors $h \in H$ and $h' \in H'$. For a vector in a Hilbert space h let $h^{\otimes n}$ denote the tensor power, i.e. $\underbrace{h \otimes \dots \otimes h}_{n \text{ times}}$.

In the following the space of finite signed measures is equipped with the total variation topology and unadorned norms refer to the total variation norm on finite signed measures, which forms a Banach space. Norms for various Lebesgue spaces will have the associated subscript. Finally we note that for two Hilbert spaces of square-integrable functions over σ -finite measure spaces $L^2(\Omega, \mathcal{F}, \mu)$ and $L^2(\Omega', \mathcal{F}', \mu')$ we have that $L^2(\Omega, \mathcal{F}, \mu) \otimes L^2(\Omega', \mathcal{F}', \mu') \cong L^2(\Omega \times \Omega', \mathcal{F} \times \mathcal{F}', \mu \times \mu')$ via an isomorphism $f \otimes f' \mapsto f \times f'$ ([43] Example 2.6.11) and we will use a 2 subscript for both norms.

E.3 Full Theorem Statement and Proof

The following is the full general version of Theorem 3 and the main result of this section.

Theorem 4. *Let (Ω, \mathcal{F}) be a measurable space, $\mathcal{P} = \sum_{j=1}^m a_j \delta_{\mu_j}$ a mixture of measures on that space which is n -identifiable, and $\mathcal{P}_i = \sum_{j=1}^{m'_i} b_{i,j} \delta_{\nu_{i,j}}$ a sequence of mixtures of measures with $m'_i \leq m$ for all i , such that $V_n(\mathcal{P}_i) \rightarrow V_n(\mathcal{P})$. Then $m'_i \rightarrow m$ and there exists a sequence of permutations σ_i such that $\sigma_i(b_i) \rightarrow a$ and $\nu_{i,\sigma_i(j)} \rightarrow \mu_j$ for all j .*

Essentially this says that as one finds grouped sample distributions $V_n(\mathcal{Q}_i)$ which approach the true grouped sample distribution $V_n(\mathcal{P})$ the mixture of measures \mathcal{Q}_i will automatically recover the true mixing weights and components from \mathcal{P} so long as \mathcal{P} is n -identifiable. In other words, one simply needs to fit the grouped distribution $V_n(\mathcal{P})$ well to get a good estimate of the mixture components. Theorem 3 from the main text is a direct consequence of Theorem 4.

Corollary 3 (Theorem 3). *Let $\sum_{m=1}^M w_m p_m$ be an N -identifiable mixture model, and $\sum_{m=1}^M \hat{w}_{m,j} \hat{p}_{m,j}$ be a sequence of mixture models such that $\left\| \sum_{m=1}^M \hat{w}_{m,j} \hat{p}_{m,j}^{\times N} - \sum_{m=1}^M w_m p_m^{\times N} \right\|_1 \rightarrow 0$. Then there is a sequence of permutations σ_j so that $\hat{w}_{\sigma_j(m),j} \rightarrow w_m$ and $\left\| \hat{p}_{\sigma_j(m),j} - p_m \right\|_1 \rightarrow 0$ for all m .*

We introduce some preliminary results before proving Theorem 4. The following lemma will be needed for our proof.

Lemma 2. *Let \mathcal{P} and \mathcal{Q} be mixtures of measures, then $\|V_{n'}(\mathcal{P}) - V_{n'}(\mathcal{Q})\| \leq \|V_n(\mathcal{P}) - V_n(\mathcal{Q})\|$ for all $n' \leq n$.*

Proof of Lemma 2. From [44] (Section 3.1 Exercise 7a) we have the following

$$\begin{aligned}
& \|V_n(\mathcal{P}) - V_n(\mathcal{Q})\| \\
&= \sup \left\{ \sum_{i=1}^k |(V_n(\mathcal{P}) - V_n(\mathcal{Q}))(E_i)| : \right. \\
&\quad \left. k \in \mathbb{N}, E_1, \dots, E_k \in \mathcal{F}^{\times n} \text{ are disjoint, and } \bigcup_{i=1}^k E_i = \Omega^{\times n} \right\} \\
&\geq \sup \left\{ \sum_{i=1}^k \left| (V_n(\mathcal{P}) - V_n(\mathcal{Q}))(E_i \times \Omega^{\times n-n'}) \right| : \right. \\
&\quad \left. k \in \mathbb{N}, E_1, \dots, E_k \in \mathcal{F}^{\times n'} \text{ are disjoint, and } \bigcup_{i=1}^k E_i = \Omega^{\times n'} \right\} \\
&= \sup \left\{ \sum_{i=1}^k |(V_{n'}(\mathcal{P}) - V_{n'}(\mathcal{Q}))(E_i)| : \right. \\
&\quad \left. k \in \mathbb{N}, E_1, \dots, E_k \in \mathcal{F}^{\times n'} \text{ are disjoint, and } \bigcup_{i=1}^k E_i = \Omega^{\times n'} \right\} \\
&= \|V_{n'}(\mathcal{P}) - V_{n'}(\mathcal{Q})\|.
\end{aligned}$$

□

The following lemma is the main workhorse in the proof of Theorem 4.

Lemma 3. *Let (Ω, \mathcal{F}) be a measurable space, $\mathcal{P} = \sum_{i=1}^m a_i \delta_{\mu_i}$ a mixture of measures on that space, $n \in \mathbb{N}$, and $\mathcal{P}_i = \sum_{j=1}^{m'} b_j \delta_{\nu_{i,j}}$ a sequence of mixtures of measures (m' is fixed) with such that $V_n(\mathcal{P}_i) \rightarrow V_n(\mathcal{P})$ (b does not depend on i). Then there exists a subsequence i_k and a collection of probability measures $\nu_1, \dots, \nu_{m'}$ such that $\nu_{i_k,j} \rightarrow \nu_j$ for all j and $V_n(\mathcal{P}) = V_n\left(\sum_{j=1}^{m'} b_j \delta_{\nu_j}\right)$.*

Proof of Lemma 3. We will use bold symbols to represent elements that depend on i , e.g. $\boldsymbol{\nu}_j = \nu_{i,j}$. Let $\bar{\mu} = \sum_{k=1}^m a_k \mu_k$. By the Lebesgue-Radon-Nikodym Theorem ([44] Theorem 3.8) there exists series of measures $\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_{m'}$ and $\boldsymbol{\rho}_1, \dots, \boldsymbol{\rho}_{m'}$ such that $\boldsymbol{\nu}_k = \boldsymbol{\lambda}_k + \boldsymbol{\rho}_k$ with $\boldsymbol{\lambda}_k \perp \bar{\mu}$ and $\boldsymbol{\rho}_k \ll \bar{\mu}$ for all $k \in [m']$.

For some fixed ℓ let \mathbf{A}_ℓ be the sequence of measurable sets such that $\boldsymbol{\lambda}_\ell(\cdot \cap \mathbf{A}_\ell) = \boldsymbol{\lambda}_\ell$ and $\bar{\mu}(\mathbf{A}_\ell) = 0$, this is possible since $\boldsymbol{\lambda}_\ell \perp \bar{\mu}$. From Lemma 2 we have that

$$\left\| \sum_{k=1}^m a_k \mu_k - \sum_{j=1}^{m'} b_j \boldsymbol{\nu}_j \right\| \rightarrow 0 \Rightarrow \left| \sum_{k=1}^m a_k \mu_k(\mathbf{A}_\ell) - \sum_{j=1}^{m'} b_j \boldsymbol{\nu}_j(\mathbf{A}_\ell) \right| \rightarrow 0 \quad (20)$$

$$\Rightarrow \left| b_\ell \boldsymbol{\rho}_\ell(\mathbf{A}_\ell) + b_\ell \boldsymbol{\lambda}_\ell(\mathbf{A}_\ell) + \sum_{j \in [m'] \setminus \{\ell\}} b_j \boldsymbol{\nu}_j(\mathbf{A}_\ell) \right| \rightarrow 0 \quad (21)$$

$$\Rightarrow \left| b_\ell \boldsymbol{\lambda}_\ell(\mathbf{A}_\ell) + \sum_{j \in [m'] \setminus \{\ell\}} b_j \boldsymbol{\nu}_j(\mathbf{A}_\ell) \right| \rightarrow 0. \quad (22)$$

Because all of the summands inside the absolute value on the last line are positive we have that $\|\boldsymbol{\lambda}_\ell\| \rightarrow 0$ and thus $\|\boldsymbol{\rho}_\ell\| \rightarrow 1$. Eventually in our sequence we must have that $\|\boldsymbol{\rho}_\ell\| > 0$, so eventually in our subsequence we can define $\boldsymbol{\nu}'_\ell = \boldsymbol{\rho}_\ell / \|\boldsymbol{\rho}_\ell\|$ which is now a sequence of probability measures which are absolutely continuous with respect to $\bar{\mu}$ and $\|\boldsymbol{\nu}'_\ell - \boldsymbol{\nu}_\ell\| \rightarrow 0$.

From this we have that there exists sequences of probability measures $\boldsymbol{\nu}'_1, \dots, \boldsymbol{\nu}'_{m'}$ such that $\|\boldsymbol{\nu}_k - \boldsymbol{\nu}'_k\| \rightarrow 0$ and $\boldsymbol{\nu}'_k \ll \bar{\mu}$ for all $k \in [m']$. Lemma 3.3.7 in [45] states that, for probability measures over the same domain $\xi_1, \dots, \xi_d, \gamma_1, \dots, \gamma_d$ that $\left\| \prod_{j=1}^d \xi_j - \prod_{k=1}^d \gamma_k \right\| \leq \sum_{k=1}^d \|\xi_k - \gamma_k\|$. It follows therefore that $\left\| \boldsymbol{\nu}'_k^{\times n} - \boldsymbol{\nu}_k^{\times n} \right\| \rightarrow 0$ for all k and

$$\left\| \sum_{k=1}^m a_k \mu_k^{\times n} - \sum_{j=1}^{m'} b_j \boldsymbol{\nu}'_j^{\times n} \right\| \rightarrow 0. \quad (23)$$

For some fixed ℓ let \mathbf{q}'_ℓ be the Radon-Nikodym derivative of $\boldsymbol{\nu}'_\ell$ with respect to $\bar{\mu}$. Let $\mathbf{B}_\ell = \mathbf{q}'_\ell^{-1}([2/b_\ell, \infty))$. We have the following

$$\sum_{k=1}^{m'} b_k \boldsymbol{\nu}'_k(\mathbf{B}_\ell) \geq b_\ell \boldsymbol{\nu}'_\ell(\mathbf{B}_\ell) \quad (24)$$

$$\geq b_\ell \int_{\mathbf{B}_1} 2/b_\ell d\bar{\mu} \quad (25)$$

$$\geq 2\bar{\mu}(\mathbf{B}_\ell). \quad (26)$$

From Lemma 2 applied to (23) we have that $\left| \sum_{k=1}^{m'} b_k \nu'_k(\mathbf{B}_\ell) - \bar{\mu}(\mathbf{B}_\ell) \right| \rightarrow 0$, and it follows that $\bar{\mu}(\mathbf{B}_\ell) \rightarrow 0$ because $\left| \sum_{k=1}^{m'} b_k \nu'_k(\mathbf{B}_\ell) - \bar{\mu}(\mathbf{B}_\ell) \right| \geq \bar{\mu}(\mathbf{B}_\ell)$. Now we have that $\sum_{k=1}^{m'} b_k \nu'_k(\mathbf{B}_\ell) \rightarrow 0$ and thus $\nu'_\ell(\mathbf{B}_\ell) \rightarrow 0$.

Because $\nu'_\ell \left(\mathbf{q}'_\ell^{-1}([2/b_\ell, \infty)) \right) \rightarrow 0$ and therefore $\nu'_\ell(\mathbf{B}_\ell^C) \rightarrow 1$, for sufficiently large i we can now define a sequence of probability measures ν''_ℓ via $\nu''_\ell(A) = \nu'_\ell(A \cap \mathbf{B}_\ell^C) / \nu'_\ell(\mathbf{B}_\ell^C)$. We have that

$$\|\nu'_\ell - \nu''_\ell\| = \|(\nu'_\ell(\mathbf{B}_\ell \cap \cdot) + \nu'_\ell(\mathbf{B}_\ell^C \cap \cdot)) - (\nu''_\ell(\mathbf{B}_\ell \cap \cdot) + \nu''_\ell(\mathbf{B}_\ell^C \cap \cdot))\| \quad (27)$$

$$\leq \|\nu'_\ell(\mathbf{B}_\ell \cap \cdot) - \nu''_\ell(\mathbf{B}_\ell \cap \cdot)\| + \|\nu'_\ell(\mathbf{B}_\ell^C \cap \cdot) - \nu''_\ell(\mathbf{B}_\ell^C \cap \cdot)\| \quad (28)$$

$$= \nu'_\ell(\mathbf{B}_\ell) + \|\nu'_\ell(\mathbf{B}_\ell^C \cap \cdot) - \nu'_\ell(\mathbf{B}_\ell^C \cap \cdot) / \nu'_\ell(\mathbf{B}_\ell^C)\| \quad (29)$$

$$= \nu'_\ell(\mathbf{B}_\ell) + |1 - 1/\nu'_\ell(\mathbf{B}_\ell^C)| \|\nu'_\ell(\mathbf{B}_\ell \cap \cdot)\| \quad (30)$$

which goes to zero, so $\|\nu''_\ell - \nu'_\ell\| \rightarrow 0$. Note that ν''_ℓ is a sequence of probability measures with Radon-Nikodym derivatives $\mathbf{q}''_\ell \triangleq \mathbf{q}'_\ell \mathbf{1}_{\mathbf{B}_\ell^C} / \nu'_\ell(\mathbf{B}_\ell^C)$ ($\mathbf{1}$ is the indicator function) and thus

$$\sup_x \mathbf{q}''_\ell(x) = \sup_x \mathbf{q}'_\ell(x) \mathbf{1}_{\mathbf{B}_\ell^C}(x) / \nu'_\ell(\mathbf{B}_\ell^C) \leq 2/(b_\ell \nu'_\ell(\mathbf{B}_\ell^C))$$

and since $\nu'_\ell(\mathbf{B}_\ell^C) \rightarrow 1$ eventually $\|\mathbf{q}''_\ell\|_\infty \leq 3/b_\ell$. From this we have that $\mathbf{q}''_\ell \in L^1(\Omega, \mathcal{F}, \bar{\mu}) \cap L^\infty(\Omega, \mathcal{F}, \bar{\mu})$ and $\|\mathbf{q}''_\ell\|_\infty$ is a bounded sequence. From Hölders's Inequality we have that

$$= \|\mathbf{q}''_\ell\|_2^2 = \|\mathbf{q}''_\ell \mathbf{q}''_\ell\|_1^2 \leq \|\mathbf{q}''_\ell\|_1 \|\mathbf{q}''_\ell\|_\infty = \|\mathbf{q}''_\ell\|_\infty$$

so \mathbf{q}''_ℓ is a bounded sequence in $L^2(\Omega, \mathcal{F}, \bar{\mu})$.

We now define $\nu''_1, \dots, \nu''_{m'}, \mathbf{q}''_1, \dots, \mathbf{q}''_{m'}$ similarly. There exists β such that $\|\mathbf{q}''_j\|_\infty \leq \beta$ and $\|\mathbf{q}''_j\|_2 \leq \beta$ along the whole series and for all j . Let p_1, \dots, p_m be the radon Nikodym derivatives for μ_1, \dots, μ_m with respect to $\bar{\mu}$, again these are in $L^1(\Omega, \mathcal{F}, \bar{\mu}) \cap L^2(\Omega, \mathcal{F}, \bar{\mu}) \cap L^\infty(\Omega, \mathcal{F}, \bar{\mu})$. To see this note that $p_i \leq 1/a_i$ otherwise we have that

$$\begin{aligned} \mu_i(p_i^{-1}((1/a_i, \infty))) &= \int_{p_i^{-1}((1/a_i, \infty))} p_i d\bar{\mu} \\ &> \int_{p_i^{-1}((1/a_i, \infty))} 1/a_i d\bar{\mu} \\ &> \sum_j \int_{p_i^{-1}((1/a_i, \infty))} 1/a_i a_j d\mu_j \\ &\geq \mu_i(p_i^{-1}((1/a_i, \infty))) \end{aligned}$$

a contradiction. Now we have

$$\left\| \sum_{k=1}^m a_k p_k^{\times n} - \sum_{j=1}^{m'} b_j \mathbf{q}_j^{\prime \prime \times n} \right\|_1 \rightarrow 0. \quad (31)$$

and Lemma 2 implies

$$\left\| \sum_{k=1}^m a_k p_k^{\times 2} - \sum_{j=1}^{m'} b_j \mathbf{q}_j^{\prime \prime \times 2} \right\|_1 \rightarrow 0. \quad (32)$$

From Hölder's Inequality ($\|f\|_2^2 \leq \|f\|_1 \|f\|_\infty$) we have that

$$\left\| \sum_{k=1}^m a_k p_k^{\times 2} - \sum_{j=1}^{m'} b_j \mathbf{q}_j''^{\times 2} \right\|_2 \rightarrow 0 \quad (33)$$

and

$$\left\| \sum_{k=1}^m a_k p_k^{\otimes 2} - \sum_{j=1}^{m'} b_j \mathbf{q}_j''^{\otimes 2} \right\|_2 \rightarrow 0. \quad (34)$$

Let $S = \text{span}(\{p_1, \dots, p_m\})$ and $\ell \in [m']$ be arbitrary. We have that $\mathbf{q}_\ell'' = \text{proj}_S(\mathbf{q}_\ell'') + \text{proj}_{S^\perp}(\mathbf{q}_\ell'')$, noting that the summands in the decomposition are both L^2 bounded sequences. So now we have that

$$\left\langle \sum_{k=1}^{m'} b_k \mathbf{q}_k''^{\otimes 2} - \sum_{j=1}^m a_j p_j^{\otimes 2}, \text{proj}_{S^\perp}(\mathbf{q}_\ell'')^{\otimes 2} \right\rangle \rightarrow 0 \quad (35)$$

$$\Rightarrow \left\langle \sum_{j=1}^{m'} b_j \mathbf{q}_j''^{\otimes 2}, \text{proj}_{S^\perp}(\mathbf{q}_\ell'')^{\otimes 2} \right\rangle \rightarrow 0 \quad (36)$$

$$\Rightarrow b_\ell \left\langle \text{proj}_{S^\perp}(\mathbf{q}_\ell'')^{\otimes 2}, \text{proj}_{S^\perp}(\mathbf{q}_\ell'')^{\otimes 2} \right\rangle + \sum_{j \in [m'] \setminus \{\ell\}} b_j \langle \mathbf{q}_j'', \text{proj}_{S^\perp}(\mathbf{q}_j'') \rangle^2 \rightarrow 0 \quad (37)$$

$$\Rightarrow b_\ell \left\| \text{proj}_{S^\perp}(\mathbf{q}_\ell'') \right\|_2^4 \rightarrow 0. \quad (38)$$

From this we have that $\left\| \text{proj}_S(\mathbf{q}_k'') - \mathbf{q}_k'' \right\|_2 \rightarrow 0$ for all k . Since $\bigoplus_{j=1}^{m'} \text{proj}_S(\mathbf{q}_j'')$ is a L^2 bounded sequence on a finite dimensional space by the Bolzano-Weierstrass theorem it has a convergent subsequence which converges to $\bigoplus_{j=1}^{m'} q_j''$ so $\mathbf{q}_j'' \rightarrow q_j''$ in L^2 . From Hölder's Inequality we have that, along this subsequence

$$\left\| \mathbf{q}_k'' - q_k'' \right\|_1 \leq \left\| \mathbf{q}_k'' - q_k'' \right\|_2 \|1\|_2 \leq \left\| \mathbf{q}_k'' - q_k'' \right\|_2 \sqrt{\int 1^2 d\bar{\mu}} = \left\| \mathbf{q}_k'' - q_k'' \right\|_2 \rightarrow 0 \quad (39)$$

so q_k'' is a probability density for all k , since they must be nonnegative to converge and integrate to one. Now we have that

$$\sum_{j=1}^m a_j p_j^{\times n} = \sum_{k=1}^{m'} b_k q_k''^{\times n}. \quad (40)$$

And defining ν_k as the probability measure associated with q_k'' we have that there exists a subsequence such that $\|\nu_k - \nu_k\| \rightarrow 0$ for all k and

$$\sum_{j=1}^m a_j \mu_j^{\times n} = \sum_{k=1}^{m'} b_k \nu_k^{\times n}. \quad (41)$$

□

We can now prove Theorem 4.

Proof of Theorem 4. To help lighten notation we will simply bold some elements which depend on the sequence \mathcal{P}_i . Let $\mathcal{P}_i = \sum_{j=1}^{m'_i} \mathbf{b}_j \delta_{\boldsymbol{\nu}_j}$ be a sequence of mixtures of measures ($\mathbf{b}_j, \boldsymbol{\nu}_j$ are functions of i) such that $V_n(\mathcal{P}_i) \rightarrow V_n(\mathcal{P})$.

We define $\tilde{\mathbf{b}}$ a sequence in Δ^m so that $\tilde{\mathbf{b}}_j = \mathbf{b}_j$ for $j \leq m'_i$ and $\tilde{\mathbf{b}}_k = 0$ for $k > m'_i$. Consider the case where there exists no sequence of permutations such that $\sigma(\tilde{\mathbf{b}}) \rightarrow a$. From this it would follow that there exists a subsequence on i and $\varepsilon > 0$ such that $\|\tilde{\mathbf{b}} - \sigma(a)\| > \varepsilon$ for all $\sigma \in S_m$. The space

$$\Delta^m \cap \left(\bigcap_{\sigma \in S_m} \text{ball}(\sigma(a), \varepsilon)^C \right) \quad (42)$$

is compact (the ball is open) so there exists a sub-subsequence of i where $\tilde{\mathbf{b}}$ converges to a point $b \notin S_m \cdot a$. Let $I \subset [m]$ be the indices of b which are nonzero and $m' = \max(I)$. For sufficiently large i along our sub-subsequence we have that $m'_i \geq m'$ and furthermore

$$\left\| \sum_{j=1}^{m'_i} \mathbf{b}_j \boldsymbol{\nu}_j^{\times n} - \sum_{k \in I} b_k \boldsymbol{\nu}_k^{\times n} \right\| \leq \left\| \sum_{j \in I} \mathbf{b}_j \boldsymbol{\nu}_j^{\times n} - \sum_{k \in I} b_k \boldsymbol{\nu}_k^{\times n} \right\| + \left\| \sum_{j \in I^C} \mathbf{b}_j \boldsymbol{\nu}_j^{\times n} \right\| \quad (43)$$

$$\leq \left\| \sum_{j \in I} (\mathbf{b}_j - b_j) \boldsymbol{\nu}_j^{\times n} \right\| + \left| \sum_{j \in I^C} \mathbf{b}_j \right| \quad (44)$$

$$\leq \sum_{j \in I} |\mathbf{b}_j - b_j| + \left| \sum_{j \in I^C} \mathbf{b}_j \right| \rightarrow 0 \quad (45)$$

and therefore

$$\left\| \sum_{k=1}^m a_k \mu_k^{\times n} - \sum_{j \in I} b_j \boldsymbol{\nu}_j^{\times n} \right\| \rightarrow 0. \quad (46)$$

From Lemma 3 we have that there exists a subsequence of this sub-subsequence such that for $k \in I$ there exists probability measures ν_k with $\|\nu_k - \boldsymbol{\nu}_k\| \rightarrow 0$ and

$$\sum_{k=1}^m a_k \mu_k^{\times n} = \sum_{j \in I} b_j \boldsymbol{\nu}_j^{\times n}. \quad (47)$$

If $|I| < m$ or $\nu_j = \boldsymbol{\nu}_k$ for any $k \neq j$ and $j, k \in I$ then we have clearly violated identifiability since we can construct a mixture of measures \mathcal{P}' with fewer components than \mathcal{P} and $V_n(\mathcal{P}') = V_n(\mathcal{P})$. If $|I| = m$ (i.e. $I = [m]$) and ν_j are all distinct we have also arrived at a contradiction since letting $\mathcal{P}' = \sum_{j=1}^m b_j \delta_{\nu_j} \neq \mathcal{P}$ because there exists no σ such that $\sigma(b) = a$ and $V_n(\mathcal{P}') = V_n(\mathcal{P})$, contradicting identifiability.

So we have that for sufficiently large i that $m'_i = m$ and there exists at least one sequence σ such that $\sigma(\mathbf{b}) \rightarrow a$. So let $\left\| \sum_{k=1}^m a_k \mu_k^{\times n} - \sum_{j=1}^m \mathbf{b}_j \boldsymbol{\nu}_j^{\times n} \right\| \rightarrow 0$. From what we have just shown, we

can permute the indices and, without loss of generality, we can assume that $\mathbf{b} \rightarrow a$. So now we have that $\left\| \sum_{i=1}^m a_i \mu_i^{\times n} - \sum_{j=1}^m a_j \nu_j^{\times n} \right\| \rightarrow 0$.

Let $\tilde{S}_m \subset S_m$ be the subgroup of permutations such that $\sigma(a) = a$ for $\sigma \in \tilde{S}_m$ (also known as the *stabilizer* of a). Note that if a_1, \dots, a_m are distinct then \tilde{S}_m only contains the identity. We proceed by contradiction: suppose there exists no sequence of permutations $\sigma \in \mathbb{N}^{\tilde{S}_m}$ such that $\nu_{\sigma(k)} \rightarrow \mu_k$ for all k . From this it follows that there exists a subsequence and a $\varepsilon > 0$, such that $\bigoplus_{k=1}^m \nu_k$ does not lie in $\bigcap_{\sigma \in \tilde{S}_m} (\text{ball}(\bigoplus_{k=1}^m \mu_{\sigma(k)}), \varepsilon)^C$. From Lemma 3 there exists probability measures, ν_1, \dots, ν_m such that for some subsequence $\|\nu_k - \nu_k\| \rightarrow 0$ for all k and

$$\sum_{j=1}^m a_j \mu_j^{\times n} = \sum_{k=1}^m a_k \nu_k^{\times n}.$$

Because $\bigcap_{\sigma \in \tilde{S}_m} (\text{ball}(\bigoplus_{k=1}^m \mu_{\sigma(k)}), \varepsilon)^C$ is closed we have $\bigoplus_{j=1}^m \nu_j \in \bigcap_{\sigma \in \tilde{S}_m} (\text{ball}(\bigoplus_{k=1}^m \mu_{\sigma(k)}), \varepsilon)^C$ and there exists no $\sigma \in \tilde{S}_m$ such that $\nu_{\sigma(k)} = \mu_k$ for all k so. Setting $\mathcal{P}' = \sum_{k=1}^m a_k \delta_{\nu_k}$ we have that $\mathcal{P}' \neq \mathcal{P}$ but $V_n(\mathcal{P}') = V_n(\mathcal{P})$, a contradiction. \square

F General Version of Corollary 1

Here we present the general version of Corollary 1 which guarantees recovery of the true mixture components using our estimator for any mixture model, provided there are a sufficient number of samples per group. For a mixture model $p = \sum_{m=1}^M w_m^* p_m^*$, using the estimator \hat{q} from Section C.2.1 to estimate (11):

$$q(y_1, y_2, \dots, y_N) = \sum_{m=1}^M w_m^* p_m^*(y_1) p_m^*(y_2) \dots p_m^*(y_N), \quad y_1, y_2, \dots, y_N \in \mathbb{R}^d.$$

combining Theorem 2a and Theorem 3 gives the following result.

Corollary 4. *If $\sigma \rightarrow 0$ and $\frac{n\sigma^{2Nd}}{\log n} \rightarrow \infty$ as $n \rightarrow \infty$, and p is N -identifiable (e.g. $N = 2M - 1$), then $\hat{w}_m \xrightarrow{a.s.} w_m^*$ and $\|p(\cdot; \hat{\alpha}_m) - p_m^*\|_1 \xrightarrow{a.s.} 0$, up to a permutation.*