

Identifiability of Nonparametric Mixture Models and Bayes Optimal Clustering

Bryon Aragam Chen Dan Pradeep Ravikumar Eric P. Xing

November 28, 2018

Abstract

Motivated by problems in data clustering, we establish general conditions under which families of nonparametric mixture models are identifiable, by introducing a novel framework involving clustering overfitted *parametric* (i.e. misspecified) mixture models. These identifiability conditions generalize existing conditions in the literature, and are flexible enough to include for example mixtures of Gaussian mixtures. In contrast to the recent literature on estimating nonparametric mixtures, we allow for general nonparametric mixture components, and instead impose regularity assumptions on the underlying mixing measure. As our primary application, we apply these results to partition-based clustering, generalizing the notion of a Bayes optimal partition from classical parametric model-based clustering to nonparametric settings. Furthermore, this framework is constructive so that it yields a practical algorithm for learning identified mixtures, which is illustrated through several examples on real data. The key conceptual device in the analysis is the convex, metric geometry of probability measures on metric spaces and its connection to the Wasserstein convergence of mixing measures. The result is a flexible framework for nonparametric clustering with formal consistency guarantees.

1 Introduction

In data clustering, we provide a grouping of a set of data points, or more generally, a partition of the input space from which the data points are drawn [31]. The many approaches to formalize the learning of such a partition from data include mode clustering [21], density clustering [55, 60, 62, 63], spectral clustering [50, 58, 75], K -means [47, 48, 61], stochastic blockmodels [3, 26, 37, 57], and hierarchical clustering [19, 32, 69], among others. In this paper, we are interested in so-called model-based clustering where the data points are drawn i.i.d. from some distribution, the most canonical instance of which is arguably Gaussian model-based clustering, in which points are drawn from a Gaussian mixture model [7, 23]. This mixture model can then be used to specify a natural partition over the input space, specifically into regions where each of the Gaussian mixture components is most likely. When the Gaussian mixture model is appropriate, this provides a simple, well-defined partition, and has been extended to various parametric and semi-parametric models [11, 27, 73]. However, the extension of this methodology to general nonparametric settings has remained elusive. This is largely due to the extreme non-identifiability of nonparametric mixture models, a problem which is well-studied but for which existing results require strong assumptions [14, 40, 42, 68]. It has been a significant open problem to generalize these assumptions to a more flexible class of nonparametric mixture models.

Unfortunately, without the identifiability of the mixture components, we cannot extend the notion of the input space partition used in Gaussian mixture model clustering. Nonetheless, there are many practical clustering algorithms used in practice, such as K -means and spectral techniques, that do estimate a partition even when the data arises from ostensibly unidentifiable nonparametric mixture models, such as mixtures of sub-Gaussian or log-concave distributions [1, 43, 49, 59]. A crucial motivation for this paper is in addressing this gap between theory and practice: This entails demonstrating that nonparametric mixture models might

Contact: naragam@cs.cmu.edu, cdan@andrew.cmu.edu, pradeepr@cs.cmu.edu, epxing@cs.cmu.edu

actually be identifiable given additional *side information*, such as the number of clusters K and the separation between the mixture components, used for instance by algorithms such as K -means.

Let us set the stage for this problem in some generality. Suppose Γ is a probability measure over some metric space X , and that Γ can be written as a finite mixture model

$$\Gamma = \sum_{k=1}^K \lambda_k \gamma_k, \quad \lambda_k > 0 \quad \text{and} \quad \sum_{k=1}^K \lambda_k = 1, \quad (1)$$

where γ_k are also probability measures over X . The γ_k represent distinct subpopulations belonging to the overall heterogeneous population Γ . Given observations from Γ , we are interested in classifying each observation into one of these K subpopulations *without* labels. When the mixture components γ_k and their weights λ_k are identifiable, we can expect to learn the model (1) from this unlabeled data, and then obtain a partition of X into regions where one of the mixture components is most likely. This can also be cast as using Bayes' rule to classify each observation, thus defining a target partition that we call the *Bayes optimal partition* (see Section 7 for formal details). Thus, in studying these partitions, a key question is *when is the mixture model (1) identifiable?* Motivated by the aforementioned applications to clustering, this question is the focus of this paper. Note that under parametric assumptions such as Gaussianity of the γ_k , it is well-known that the representation (1) is unique and hence identifiable [8, 38, 67]. These results mostly follow from an early line of work on the general identification problem [2, 66, 67, 74].

Such parametric assumptions rarely hold in practice, however, and thus it is of interest to study *non-parametric* mixture models of the form (1), i.e. for which each γ_k comes from a flexible, nonparametric family of probability measures. In the literature on nonparametric mixture models, a common assumption is that the component measures γ_k are multivariate with independent marginals [25, 29, 30, 44, 68], which is particularly useful for statistical problems involving repeated measurements [12, 34]. This model also has deep connections to the algebraic properties of latent structure models [4, 13]. Various other structural assumptions have been considered including symmetry [14, 40], tail conditions [42], and translation invariance [28]. The identification problem in discrete mixture models is also a central problem in topic models which are popular in machine learning [5, 6, 65]. Most notably, this existing literature imposes structural assumptions on the components γ_k (e.g. independence, symmetry), which are difficult to satisfy in clustering problems. Are there reasonable constraints that ensure the uniqueness of (1), while avoiding restrictive modeling assumptions on the γ_k ?

In this paper, we establish a series of positive results in this direction, and as a bonus that arises naturally from our theoretical results, we develop a practical algorithm for nonparametric clustering. In contrast to the existing literature, we allow each γ_k to be an arbitrary probability measure over X . We propose a novel framework for reconstructing nonparametric mixing measures by using simple, overfitted mixtures (e.g. Gaussian mixtures) as mixture density estimators, and then using clustering algorithms to partition the resulting estimators. This construction implies a set of regularity conditions on the mixing measure that suffice to ensure that a mixture model is identifiable. As our main application of interest, we apply this to problems in nonparametric clustering.

In the remainder of this section, we outline our major contributions. We then present a high-level geometric overview of our method in Section 2 before proceeding to the main results of the paper. Section 3 covers the necessary background required for our abstract framework. In Section 4, we present a detailed construction that takes a mixture distribution Γ and outputs its mixing measure Λ , culminating in our main theorem on identifiability. In Section 5 we discuss how to use this construction to define a consistent estimator of the parameter Λ , and then in Section 6 we provide explicit examples of mixture models that satisfy our assumptions. In Section 7 we apply these results to the problem of clustering and prove a consistency theorem for this problem. Section 8 introduces a simple algorithm for nonparametric clustering along with some experiments, and Section 9 concludes the paper with some discussion and extensions.

Contributions. At a high-level, our contributions are the following:

- A new identification criterion for nonparametric mixture models based on a property we call *clusterability* (Definition 4.3);
- Extending model-based clustering to more general nonparametric settings;

- A practical algorithm for nonparametric clustering.

Each of these contributions builds on the previous one, and provides an overall narrative that strengthens the well-known connections between identifiability in mixture models, cluster analysis, and nonparametric density estimation. Our main results can be divided into three main theorems:

1. *Nonparametric identifiability* (Section 4). We formulate a general set of assumptions that guarantee a family of nonparametric mixtures will be identifiable (Theorem 4.1), based on two properties we introduce: *regularity* (Definition 4.1) and *clusterability* (Definition 4.3).
2. *Estimation* (Section 5). We show that a simple clustering procedure will correctly identify the mixing measure that generates Γ as long as the γ_k are sufficiently well-separated in Hellinger distance (Proposition 5.1), and this procedure defines an estimator that consistently recovers the nonparametric clusters given i.i.d. observations from Γ (Theorem 5.4). We also discuss conditions for both Hellinger and uniform convergence of the resulting estimators of the mixture components.
3. *Clustering* (Section 7). We make connections with the so-called *Bayes optimal partition* (Definition 7.1), and extend this notion to general nonparametric settings by using our results on nonparametric mixtures (Theorem 7.2).

Furthermore, we construct explicit examples of nonparametric mixture models that satisfy our assumptions in Section 6. In particular Theorem 6.2 establishes the existence of such families and Figure 4 illustrates some examples. As a final contribution, we invoke this analysis to construct an intuitive algorithm for nonparametric clustering, which is investigated in Section 8.

2 Overview

Before outlining the formal details, we present an intuitive geometric picture of our construction in Figure 1. At a high-level, our strategy for identifying the mixture distribution (1) is the following:

- (1) Approximate Γ with an overfitted mixture of $L \gg K$ Gaussians (Figure 1b);
- (2) Cluster these L Gaussian components into K groups such that each group roughly approximates some γ_k (Figure 1c);
- (3) Use this clustering to define a new mixing measure (Figure 1d);
- (4) Show that this new mixing measure converges to the true mixing measure Λ as $L \rightarrow \infty$.

If the mixing measure constructed by the above procedure converges to Λ , then Λ must be identifiable.

While this procedure makes intuitive sense, one of the main thrusts of this paper is outlining a way to make this procedure well-defined in the sense that it will always return the same mixing measure. This is a surprisingly subtle problem and requires careful consideration of the various spaces involved, so the formal details of this analysis are postponed until Section 4. Furthermore, although we have used mixtures of Gaussians to approximate Γ in this example, our main results will apply to any properly chosen family of base measures.

Of course, this construction is not guaranteed to succeed for arbitrary mixing measures Λ , which will be illustrated by the examples in Section 3.2. Thus, a key aspect of our analysis will be to provide assumptions that ensure the success of this construction. Intuitively, it should be clear that as long as the γ_k are well-separated, the corresponding mixture approximation will consist of Gaussian components that are also well-separated. Unfortunately, this is not quite enough to imply identifiability, as illustrated by Example 5. This highlights some of the subtleties inherent in this construction. In the sequel, we formalize these ideas and introduce the concepts of *regularity* (Section 4.3) and *clusterability* (Section 4.4), which axiomatize the conditions needed in order for Λ to be reconstructed—and hence identified—from Γ . Then in Sections 5 and 6, we discuss the existence of nontrivial mixture distributions that satisfy these conditions as well as how to learn such mixtures from data.

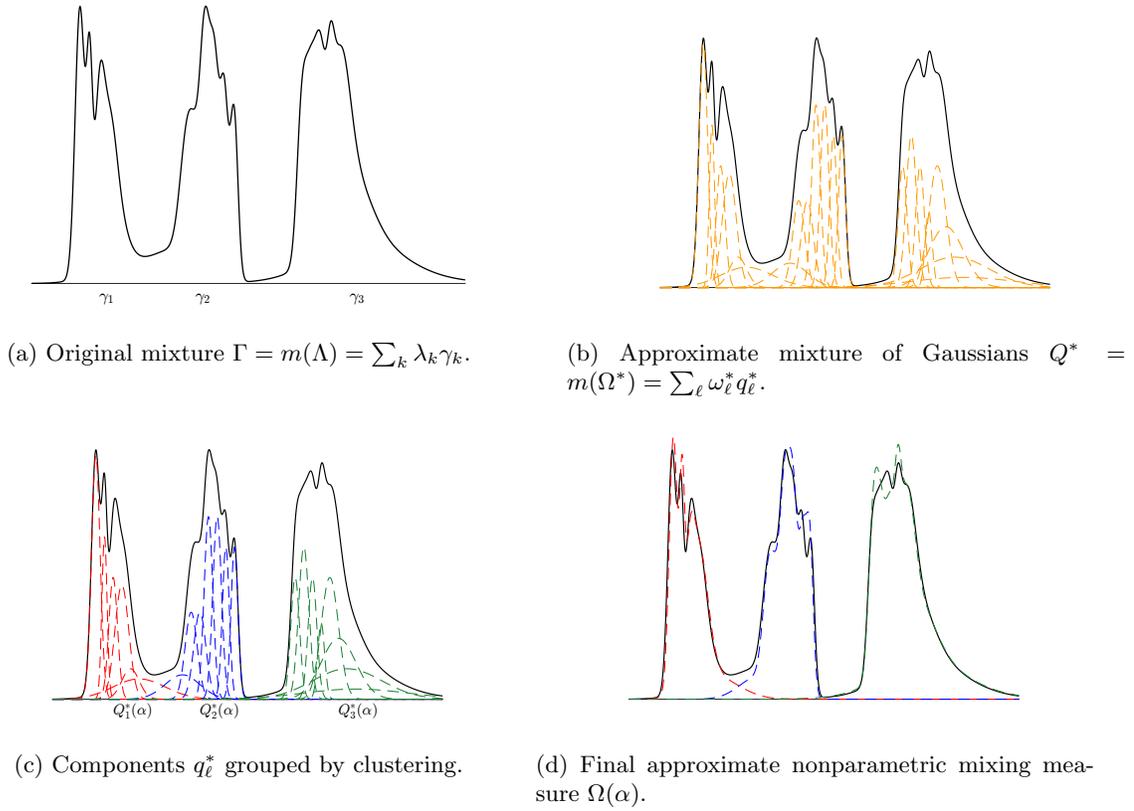


Figure 1: Overview of the method.

3 Preliminaries

Our approach is general, built on the theory of abstract measures on metric spaces [54]. In this section we introduce this abstract setting, outline our notation, and discuss the general problem of identifiability in mixture models. We deliberately include plenty of examples in order to help acquaint the reader with our particular notation and problem setting. For a more thorough introduction to the general topic of mixture models in statistics, see Lindsay [46], Ritter [56], Titterton et al. [70].

3.1 Nonparametric mixture models

Let (X, d) be a metric space and $\mathcal{P}(X)$ denote the space of regular Borel probability measures on X with finite r th moments ($r \geq 1$). Define $\mathcal{P}^2(X) = \mathcal{P}(\mathcal{P}(X))$, the space of (infinite) mixing measures over $\mathcal{P}(X)$. In this paper, we study *finite* mixture models, i.e. mixtures with a finite number of atoms. To this end, define for $s \in \{1, 2, \dots\}$

$$\mathcal{P}_s^2(X) := \{\Lambda \in \mathcal{P}^2(X) : |\text{supp}(\Lambda)| \leq s\}, \quad \mathcal{P}_0^2(X) := \bigcup_{s=1}^{\infty} \mathcal{P}_s^2(X),$$

and $\mathfrak{L}_s = \mathfrak{L} \cap \mathcal{P}_s^2(X)$ for any $\mathfrak{L} \subset \mathcal{P}^2(X)$. We consider $\mathcal{P}(X)$ and $\mathcal{P}^2(X)$ as metric spaces by endowing $\mathcal{P}(X)$ with the Hellinger metric ρ and $\mathcal{P}^2(X)$ with the L_r -Wasserstein metric W_r ($r \geq 1$). When $\Lambda \in \mathcal{P}_K^2(X)$ and $\Lambda' \in \mathcal{P}_{K'}^2(X)$, the L_r -Wasserstein distance between Λ and Λ' is given by the optimal value of the transport

problem

$$W_r(\Lambda, \Lambda') = \inf \left\{ \left[\sum_{i,j} \sigma_{ij} \rho^r(\gamma_i, \gamma'_j) \right]^{1/r} : 0 \leq \sigma_{ij} \leq 1, \right. \\ \left. \sum_{i,j} \sigma_{ij} = 1, \sum_i \sigma_{ij} = \lambda'_j, \sum_j \sigma_{ij} = \lambda_i \right\}. \quad (2)$$

where the infimum is taken over all couplings σ , i.e. probability measures on $\mathcal{P}(X) \times \mathcal{P}(X)$ with marginals Λ and Λ' . For more on Wasserstein distances and their importance in mixture models, see Nguyen [51].

Given $\Lambda \in \mathcal{P}_0^2(X)$, define a new probability measure $m(\cdot; \Lambda) \in \mathcal{P}(X)$ by

$$m(A; \Lambda) = \int \gamma(A) d\Lambda(\gamma) = \sum_{k=1}^K \lambda_k \gamma_k, \quad K := |\text{supp}(\Lambda)|, \quad (3)$$

where $\gamma_1, \dots, \gamma_K$ are the mixture components (i.e. a particular enumeration of $\text{supp}(\Lambda)$) and $\lambda_1, \dots, \lambda_K$ are the corresponding weights. Formally, for any Borel set $A \subset X$ we have a function $h_A : \mathcal{P}(X) \rightarrow \mathbb{R}$ defined by $h_A(\gamma) = \gamma(A)$, and $m(A; \Lambda) = \int \gamma(A) d\Lambda(\gamma) = \int h_A d\Lambda$. This uniquely defines a measure called a *mixture distribution* over X . In a slight abuse of notation, we will write $m(\Lambda)$ as shorthand for $m(\cdot; \Lambda)$ when there is no confusion between the arguments. An element γ_k of $\text{supp}(\Lambda)$ is called a *mixture component*. Given a Borel set $\mathfrak{L} \subset \mathcal{P}^2(X)$, define

$$\mathcal{M}(\mathfrak{L}) := \{m(\Lambda) : \Lambda \in \mathfrak{L}\}, \quad (4)$$

i.e. the family of mixture distributions over X induced by \mathfrak{L} , which can be regarded as a formal representation of a statistical mixture model. We will adopt the shorthand $\mathcal{M}(X) = \mathcal{M}(\mathcal{P}^2(X))$ and $\mathcal{M}_s(X) = \mathcal{M}(\mathcal{P}_s^2(X))$.

Remark 3.1. This abstract presentation of mixture models is needed for two reasons: (i) To emphasize that Λ is the statistical parameter of interest, in contrast to the usual parametrization in terms of atoms and weights; and (ii) To emphasize that our approach works for general measures on metric spaces. This will have benefits in the sequel, albeit at the cost of some extra abstraction here at the onset. Note that we will work exclusively with finite mixtures, i.e. $\mathcal{P}_0^2(X)$, a space which should be contrasted with the more complex space of infinite measures $\mathcal{P}^2(X)$.

Remark 3.2. The Hellinger distance ρ can be replaced by any metric on $\mathcal{P}(X)$; our use of the Hellinger distance is purely for conceptual clarity. By contrast, our convergence results are particular to the Wasserstein distance W_r .

Remark 3.3. As a convention, we will use upper case letters for mixture distributions (e.g. Γ, Q) and mixing measures (e.g. Λ, Ω), and lower case letters for mixture components (e.g. γ_k, q_k) and weights (e.g. λ_k, ω_k).

We conclude this subsection with some examples to help fix ideas and notation.

Example 1 (Gaussian mixtures). Let $\mathfrak{G} \subset \mathcal{P}^2(\mathbb{R}^p)$ denote the subset of mixing measures whose support is contained in the family of p -dimensional Gaussian measures. Then $\mathcal{M}(\mathfrak{G})$ is the family of Gaussian mixtures, and $\mathcal{M}_0(\mathfrak{G})$ is the family of finite Gaussian mixtures. It is well-known that $\mathcal{M}_0(\mathfrak{G})$ is identifiable.

Example 2 (Sub-Gaussian mixtures). Let \mathcal{K} be the collection sub-Gaussian measures on \mathbb{R} , i.e.

$$\mathcal{K} = \{\gamma \in \mathcal{P}(\mathbb{R}) : \gamma(\{x : |x| > t\}) \leq e^{1-t^2/c^2} \text{ for some } c > 0 \text{ and all } t > 0\},$$

and $\mathfrak{K} \subset \mathcal{P}^2(\mathbb{R})$ be the subset of mixing measures whose support is a subset of \mathcal{K} . Then $\mathcal{M}(\mathfrak{K})$ is the family of sub-Gaussian mixture models and $\mathcal{M}_0(\mathfrak{K})$ is the family of finite sub-Gaussian mixtures. This is a nonparametric mixture model, since the base measures \mathcal{K} do not belong to a parametric family. Extensions to sub-Gaussian measures on \mathbb{R}^p are natural.

Obviously, these examples can be extended to arbitrary parametric and nonparametric families. Our definition of mixtures over subsets of mixing measures—as opposed to over families of component distributions—makes it easy to encode additional constraints, as in the following example.

Example 3 (Constrained mixtures). Continuing Example 1, suppose we wish to impose additional constraints on the family of mixture distributions. For example, we might be interested in Gaussian mixtures with at most L components, whose means are contained within some compact set $M \subset \mathbb{R}^p$, and whose covariance matrices are contained within another compact set $V \subset \text{PD}(p)$, where $\text{PD}(p)$ is the set of $p \times p$ positive-definite matrices. Define

$$\mathcal{G}_{M,V} := \{\mathcal{N}(a, v) : a \in M, v \in V\},$$

and

$$\mathfrak{G}_{L,M,V} := \{\Lambda \in \mathcal{P}^2(X) : |\text{supp}(\Lambda)| \leq L, \text{supp}(\Lambda) \subset \mathcal{G}_{M,V}\}. \quad (5)$$

Then $\mathcal{M}(\mathfrak{G}_{L,M,V})$ is the desired family of mixture models.

Example 4 (Mixture of regressions). Suppose $\mathbb{P}(Y | Z) = \int \gamma(Z) d\Lambda(\gamma)$ is a mixture model depending on some covariates Z . We assume here that $(Z, Y) \in W \times X$ where (W, d_W) and (X, d_X) are metric spaces. This is a nonparametric extension of the usual mixed linear regression model. To recover the mixed regression model, assume Λ has at most K atoms and $\gamma_k(Z) \sim \mathcal{N}(\langle \theta_k, Z \rangle, \omega_k^2)$, so that

$$\mathbb{P}(Y | Z) = \int \gamma(Z) d\Lambda(\gamma) = \sum_{k=1}^K \lambda_k \mathcal{N}(\langle \theta_k, Z \rangle, \omega_k^2).$$

By further allowing the mixing measure $\Lambda = \Lambda(Z)$ to depend on the covariates, we obtain the nonparametric generalization of a mixture of experts model [15, 39, 41].

3.2 Identifiability in mixture models

Recall that a mixture model $\mathcal{M}(\mathfrak{L})$ is identifiable if the map $m : \mathfrak{L} \rightarrow \mathcal{M}(\mathfrak{L})$ that sends $\Lambda \mapsto m(\Lambda)$ via (3) is injective. For a good overview of this problem from a more classical perspective, see Hunter et al. [40] and Allman et al. [4]. The main purpose of this section is to highlight some of the known subtleties in identifying nonparametric mixture models.

Unsurprisingly, whether or not a specific mixture $m(\Lambda)$ is identified depends on the choice of \mathfrak{L} . If we allow \mathfrak{L} to be all of $\mathcal{P}^2(X)$, then it is easy to see that $\mathcal{M}(\mathfrak{L})$ is not identifiable, and this continues to be true even if the number of components K is known in advance (i.e. $\mathfrak{L} = \mathcal{P}_K^2(X)$). Indeed, for any partition $\{A_k\}_{k=1}^K$ of X and any Borel set $B \subset X$, we can write

$$\Gamma(B) = \sum_{k=1}^K \underbrace{\Gamma(A_k)}_{\tilde{\lambda}_k} \cdot \underbrace{\frac{\Gamma(B \cap A_k)}{\Gamma(A_k)}}_{\tilde{\gamma}_k} = \sum_{k=1}^K \tilde{\lambda}_k \tilde{\gamma}_k(B), \quad (6)$$

and thus there cannot be a unique decomposition of the measure Γ into the sum (1). Although this example allows for arbitrary, pathological decompositions of Γ into conditional measures, the following concrete example shows that solving the nonidentifiability issue is more complicated than simply avoiding certain pathological partitions of the input space.

Example 5 (Sub-Gaussian mixtures are not identifiable). Consider the mixture of three Gaussians $m(\Lambda)$ in Figure 2. We can write $m(\Lambda)$ as a mixture in four ways: In the top panel, $m(\Lambda)$ is represented uniquely as a mixture of three Gaussians. If we allow sub-Gaussian components, however, then the bottom panel shows three equally valid representations of $m(\Lambda)$ as a mixture of *two* sub-Gaussians. Recalling Examples 1 and 2, it follows that $m(\Lambda)$ is identified with respect to \mathfrak{G} , but not with respect to \mathfrak{R} . Indeed, even if we assume the number of components K is known and the component means are well-separated, $m(\Lambda)$ is non-identifiable with respect to \mathfrak{R} : Just take $K = 2$, $|a_1 - a_2| > 0$ and move a_3 arbitrarily far to the right.

Much of the existing literature makes assumptions on the structure of the allowed γ_k , which is evidently equivalent to restricting the supports of the measures in \mathfrak{L} (e.g. Example 1). Our focus, by contrast, will be to allow the components to take on essentially any shape while imposing regularity assumptions on the mixing measures $\Lambda \in \mathfrak{L}$. In this sense, we shift the focus from the properties of the “local” components to the “global” properties of the mixture itself.

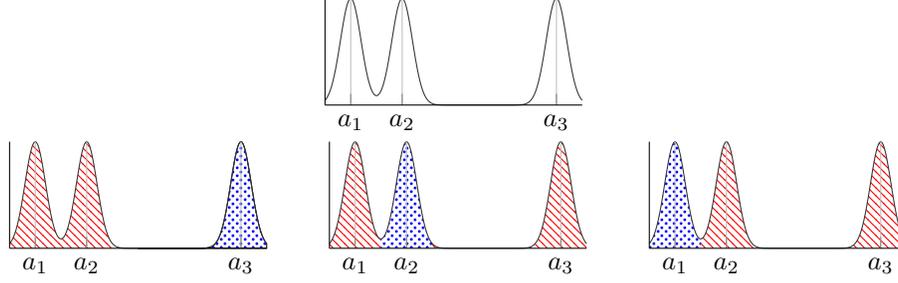


Figure 2: (top) Mixture of three Gaussians. (bottom) Different representations of a mixture of Gaussians as a mixture of two sub-Gaussians. Different fill patterns and colours represent different assignments of mixture components.

4 Nonparametric identifiability

Fix an integer K and let $\mathfrak{L} \subset \mathcal{P}_K^2(X)$ be a family of mixing measures. In particular, we assume that K —the number of nonparametric mixtures—is known; in Section 9 we discuss the case where K is unknown. In this section we study conditions that guarantee the injectivity of the canonical embedding $m : \mathfrak{L} \rightarrow \mathcal{M}(\mathfrak{L})$ using the procedure described in Section 2. Throughout this section, it will be helpful to keep Figure 1 in mind for intuition.

4.1 Projections

We begin by formalizing the first step (1) in our construction from Section 2. In order to ensure that the overfitted mixture approximation is unique, we will be interested in the Hellinger projection of $\Gamma = m(\Lambda)$ onto “well-behaved” families of mixture distributions. Specifically, we will assume in the sequel that $\{\mathfrak{Q}_L\}_{L=1}^\infty$ is an indexed collection of families of mixing measures that satisfies the following:

- (A1) $\mathfrak{Q}_L \subset \mathcal{P}_L^2(X)$ for each L ;
- (A2) $\{\mathfrak{Q}_L\}$ is monotonic, i.e. $\mathfrak{Q}_L \subset \mathfrak{Q}_{L+1}$;
- (A3) $\mathcal{M}(\mathfrak{Q}_L)$ is identifiable for each L .

The purpose of $\{\mathfrak{Q}_L\}$ is to approximate Γ with a sequence of mixing measures of increasing complexity, as quantified by the maximum number of atoms L , which will be taken to be much larger than K in practice. Although our results will apply to generic collections satisfying Conditions (A1)-(A3), we will often be interested in a particular collection induced by a subset $\mathfrak{Q} \subset \mathcal{P}^2(X)$. Specifically, we make the following assumption on \mathfrak{Q} :

- (A) The collection $\{\mathfrak{Q}_L\}_{L=1}^\infty$ defined by $\mathfrak{Q}_L = \mathfrak{Q} \cap \mathcal{P}_L^2(X)$ satisfies Condition (A3) for the family $\mathfrak{Q} \subset \mathcal{P}^2(X)$.

Note that if \mathfrak{Q} satisfies Condition (A), then $\{\mathfrak{Q}_L\}$ automatically satisfies Conditions (A1)-(A3). These conditions allow for substantial generality, however, in practice it is often enough to take $\mathfrak{Q} = \mathfrak{G}$, i.e. the set of Gaussian mixing measures (or some subset thereof as in Example 3). Note that in this special case, for any $\Gamma \in \mathcal{P}(X)$ there exist mixing measures $\Omega_L \in \mathfrak{G}_L$ such that $\rho(m(\Omega_L), \Gamma) \rightarrow 0$ as $L \rightarrow \infty$.

In the sequel, assume \mathfrak{Q} is fixed and \mathfrak{Q}_L is defined as in Condition (A). Define the usual ρ -projection by

$$T_L \Gamma = \{Q \in \mathcal{M}(\mathfrak{Q}_L) : \rho(Q, \Gamma) \leq \rho(P, \Gamma) \quad \forall P \in \mathcal{M}(\mathfrak{Q}_L)\}.$$

$T_L \Gamma$ may be empty, set-valued, or unique depending on Γ [16]. We assume \mathfrak{L} is such that $T_L \Gamma$ is unique and well-defined for every $\Gamma \in \mathcal{M}(\mathfrak{L})$, so that the projection map $T_L : \mathcal{M}(\mathfrak{L}) \rightarrow \mathcal{M}(\mathfrak{Q}_L)$ is well-defined. Furthermore, Condition (A3) implies that there exists a well-defined map $M_L : \mathcal{M}(\mathfrak{Q}_L) \rightarrow \mathfrak{Q}_L$ that sends a mixture distribution to its mixing measure. Thus we can unambiguously write $Q^* := T_L \Gamma$ and $\Omega^* = M_L(Q^*)$,

and further define

$$\begin{aligned} T_L(\mathfrak{L}) &= \{T_L[m(\Lambda)] : \Lambda \in \mathfrak{L}\} \subset \mathcal{M}(\mathfrak{Q}_L), \\ M_L(\mathfrak{L}) &= M_L(T_L(\mathfrak{L})) \subset \mathfrak{Q}_L. \end{aligned}$$

An example of the measure Q^* and its mixing measure Ω^* are depicted in Figure 1b.

Remark 4.1. The number of overfitted mixture components L will play an important but largely unheralded role in the sequel. For the most part, we will suppress the dependence of various quantities (e.g. Q^* , Ω^*) on L for notational simplicity. We typically assume that L is sufficiently large in the sense that $L \geq L_0$ for some fixed $L_0 \geq K$, and the phrase “for all sufficiently large L and n ” we mean that there exist $L_0 \geq K$ and $n_0 \geq 0$ such that $L \geq L_0$ and $n \geq n_0$.

We conclude this section with some examples of base families \mathfrak{Q} that satisfy Condition (A).

Example 6 (Gaussian mixtures). An obvious choice for \mathfrak{Q} is \mathfrak{G} , the set of Gaussian mixtures. This has the appealing property of universal approximation: Any $\Gamma \in \mathcal{P}(X)$ can be approximated arbitrarily well by some $Q \in \mathfrak{Q}_L$, as long as L is large enough. In fact, we can limit this family much further while still retaining universal approximation using known results for approximating densities with radial basis functions [17, 52, 53].

Example 7 (Gamma mixtures). Suppose $X = [0, \infty)$ and let \mathfrak{Q} be the family of mixing measures over Gamma distributions. Then any measure $\Gamma \in \mathcal{P}(X)$ can be approximated by a mixture of Gamma distributions [10, 72]. This provides a rich model for censored data on the real line.

Example 8 (Exponential family mixtures). Generalizing Examples 6 and 7, we can take \mathfrak{Q} to be mixtures over an exponential family [8]. In this case, the expressivity of \mathfrak{Q} (and hence \mathfrak{Q}_L) will depend on the choice of exponential family.

4.2 Assignment functions

The projection $Q^* = m(\Omega^*) = \sum_{\ell=1}^L \omega_\ell^* q_\ell^*$ is the best approximation to Γ from $\mathcal{M}(\mathfrak{Q}_L)$, however, it contains many more components L than the true number of *nonparametric* components K . The next step is to find a way to “cluster” the components of Q^* into K subgroups in such a way that each subgroup approximates some γ_k . This is the second step (2) in our construction from Section 2. To formalize this, we introduce the notion of *assignment functions*.

Denote the set of all maps $\alpha : [L] \rightarrow [K]$ by $\mathbb{A}_{L \rightarrow K}$ —a function $\alpha \in \mathbb{A}_{L \rightarrow K}$ represents a particular assignment of L mixture components into K subgroups. Thus, we will call α an *assignment function* in the sequel and a sequence $\{\alpha_L\}$ of assignment functions such that $\alpha_L \in \mathbb{A}_{L \rightarrow K}$ will be called an *assignment sequence*. The set of all assignment sequences is denoted by \mathbb{A}_K^∞ . For any $\Omega \in \mathfrak{Q}_L$, write $Q = m(\Omega) = \sum_{\ell=1}^L \omega_\ell q_\ell$. Given some $\alpha \in \mathbb{A}_{L \rightarrow K}$, define normalizing constants by

$$\varpi_k(\alpha) := \sum_{\ell \in \alpha^{-1}(k)} \omega_\ell, \quad k = 1, \dots, K. \quad (7)$$

Now define

$$\Omega_k(\alpha) := \frac{1}{\varpi_k(\alpha)} \sum_{\ell \in \alpha^{-1}(k)} \omega_\ell \delta_{q_\ell}, \quad Q_k(\alpha) := m(\Omega_k(\alpha)) = \frac{1}{\varpi_k(\alpha)} \sum_{\ell \in \alpha^{-1}(k)} \omega_\ell q_\ell. \quad (8)$$

Here, δ_{q_ℓ} is the point mass concentrated at q_ℓ . Note the normalizing constant $\varpi_k(\alpha)$, which is needed to ensure that $Q_k(\alpha)$ is indeed a probability measure. These quantities define a single, aggregate K -mixture by

$$\Omega(\alpha) := \sum_{k=1}^K \varpi_k(\alpha) \delta_{Q_k(\alpha)}, \quad Q(\alpha) := m(\Omega(\alpha)) = \sum_{k=1}^K \varpi_k(\alpha) Q_k(\alpha). \quad (9)$$

Note that as measures, $Q(\alpha) = m(\Omega(\alpha)) = Q$ for any assignment α and any mixture Q . The difference lies in how we organize the components into K groups: Different choices of α lead to different groupings of the L overfitted components q_ℓ (Figure 1c), and hence different mixing measures $\Omega(\alpha)$ (Figure 1d).

Finally, for any $L \geq 1$, define

$$\mathfrak{Q}_{L \rightarrow K} := \{\Omega(\alpha) : \Omega \in \mathfrak{Q}_L, \alpha \in \mathbb{A}_{L \rightarrow K}\},$$

i.e. $\mathfrak{Q}_{L \rightarrow K}$ is the collection of all mixing measures formed by clustering together the L atoms of some $\Omega \in \mathfrak{Q}_L$ into K groups. Since $Q_k(\alpha) \in \mathcal{M}(\mathfrak{Q}_L)$, $\Omega(\alpha)$ is an atomic mixing measure whose atoms come from $\mathcal{M}_0(\mathfrak{Q})$. Informally, we hope that $Q_k(\alpha)$ is able to approximate γ_k , in a sense that will be made precise in the next section.

4.3 Regular mixtures

Given a nonparametric mixture $m(\Lambda)$, its ρ -proj-ec-tion $Q^* = \sum_{\ell=1}^L \omega_\ell^* q_\ell^*$, and an assignment function α , define $\varpi_k^*(\alpha)$ as in (7) and $Q_k^*(\alpha)$ and $\Omega_k^*(\alpha)$ as in (8). We'd like $Q_k^*(\alpha)$ to approximate γ_k , but this is certainly not guaranteed for any α . The third step (3) in our construction is to find such an assignment. This will be broken into two related assumptions: *Regularity* (present subsection) and *clusterability* (next subsection).

The following notion of regularity encodes the kind of behaviour we seek in an assignment:

Definition 4.1 (Regularity). Suppose $\Lambda \in \mathcal{P}_K^2(X)$ and $\Gamma = m(\Lambda) \in \mathcal{M}_K(X)$. The mixing measure Λ is called \mathfrak{Q}_L -regular if:

- (a) There exists $L_0 \geq 0$ such that the ρ -projection $Q^* = T_L \Gamma$ exists and is unique for each $L \geq L_0$ and $\lim_{L \rightarrow \infty} T_L \Gamma = \Gamma$;
- (b) There exists an assignment sequence $\{\alpha_L\} \in \mathbb{A}_K^\infty$ such that

$$\lim_{L \rightarrow \infty} Q_k^*(\alpha_L) = \gamma_k \quad \text{and} \quad \lim_{L \rightarrow \infty} \varpi_k^*(\alpha_L) = \lambda_k \quad \forall k = 1, \dots, K.$$

When Λ is \mathfrak{Q}_L -regular, we will also call $m(\Lambda)$ \mathfrak{Q}_L -regular.

General conditions under which Definition 4.1(a) holds can be found in [16] and the references therein. See also Section 6.

Definition 4.2 (Regular assignment sequences). Given a regular mixing measure Λ , denote set of all assignment sequences $\{\alpha_L\}$ such that Definition 4.1(b) holds by $\mathbb{A}_K^\infty(\Lambda)$. An arbitrary assignment sequence $\{\alpha_L\} \in \mathbb{A}_K^\infty(\Lambda)$ will be called a *regular assignment sequence*, or Λ -regular when we wish to emphasize the underlying mixing measure.

Let us pause to review what we have developed so far. If a mixing measure Λ is \mathfrak{Q}_L -regular, then the ρ -projections of $m(\Lambda)$ can always be grouped in such a way that each group approximates the nonparametric component γ_k and its mixing weight λ_k . Note that we have not said anything yet about *how* one might find such an assignment, but only that it exists. The problem of identifying α will be discussed in Section 4.4 and Section 5. For now, all we need is that \mathfrak{Q}_L is regular in the sense that $m(\Lambda)$ can be *globally* approximated with some mixture in $\mathcal{M}(\mathfrak{Q}_L)$ (Definition 4.1(a)) and that the parameters of Λ can be *locally* approximated via a sequence of assignment maps $\{\alpha_L\}$ (Definition 4.1(b)).

Remark 4.2. There is always a mixture distribution $Q' = \sum_{\ell=1}^L \omega'_\ell q'_\ell$ and an assignment function α such that $\lim_{L \rightarrow \infty} Q'_k(\alpha) = \gamma_k$ and $\lim_{L \rightarrow \infty} \varpi'_k(\alpha) = \lambda_k$. Taking $\mathfrak{Q} = \mathfrak{G}$, it suffices to approximate each γ_k independently via mixtures of Gaussians Q'_k , $k = 1, \dots, K$ and let $Q' = \sum_{k=1}^K \lambda_k Q'_k$. What the definition of regularity (namely, Definition 4.1(b)) requires, however, is that not just that such an approximation exists, but that this local approximation is achieved specifically by the ρ -projection Q^* . Although this is not always guaranteed, regularity simply asks that Q^* —the closest mixture to Γ —is no *worse* than Q' , which suggests that this condition is fairly weak.

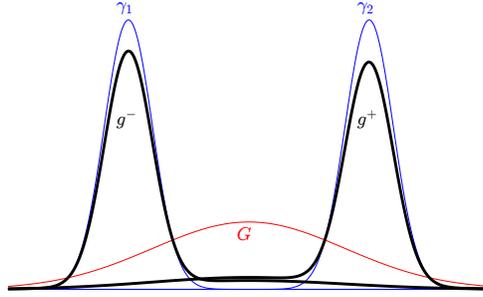


Figure 3: Example of a non-regular mixture from Example 9.

Remark 4.3. If \mathfrak{L} is a \mathfrak{Q}_L -regular family, then Definition 4.1(a) implies

$$\mathcal{M}(\mathfrak{L}) \subset \overline{\bigcup_{L=1}^{\infty} \mathcal{M}(\mathfrak{Q}_L)}.$$

Thus, the expressivity of the collection $\{\mathfrak{Q}_L\}$ constrains how large a regular family can be. Fortunately, for many families such as Gaussian mixtures, it is possible to approximate arbitrary measures; i.e. $\overline{\bigcup_{L=1}^{\infty} \mathcal{M}(\mathfrak{Q}_L)} = \mathcal{P}(X)$. Thus this is not much of a constraint in practice.

In Section 6, we will provide some concrete examples of regular families. For now, we conclude this section with the following (somewhat pathological) example of where regularity *fails*.

Example 9 (Failure of regularity). Let $g_{\pm} \sim \mathcal{N}(\pm a, 1)$ and $G \sim \mathcal{N}(0, \sigma^2)$ where $\sigma^2 > 0$, and define for some $0 < \beta_1 < \beta_2 < 1$,

$$\begin{aligned} \Gamma &= (1 - \beta_1 - \beta_2)g_+ + \beta_2g_- + \beta_1G = \frac{1}{2}\gamma_1 + \frac{1}{2}\gamma_2, \\ \gamma_1 &\propto (1 - \beta_1 - \beta_2)g_+ + \frac{\beta_1}{2}G, \\ \gamma_2 &\propto \beta_2g_- + \frac{\beta_1}{2}G. \end{aligned}$$

See Figure 3. In this example, $K = 2$. If $\mathfrak{Q}_L = \mathfrak{G}_L$, then for any $L > 3$, $Q^* = \Gamma$, and there is no way to cluster the 3 components into 2 mixtures of Gaussians that approximate the γ_k . The problem here is that γ_1 and γ_2 “share” the same Gaussian component G , which evidently cannot be assigned to both γ_1 and γ_2 .

4.4 Clusterable families

The concept of regularity is a weak condition that summarizes the most basic behaviour that we seek in a mixture distribution $\Gamma = m(\Lambda)$. To exploit this behaviour in order to identify Λ from Γ , we need to impose a slightly stronger assumption.

Definition 4.3 (Clusterable family). A family of mixing measures $\mathfrak{L} \subset \mathcal{P}^2(X)$ is called a \mathfrak{Q}_L -clusterable family, or just a clusterable family, if

- (a) Λ is \mathfrak{Q}_L -regular for all $\Lambda \in \mathfrak{L}$;
- (b) There exists a function $\chi_L : M_L(\mathfrak{L}) \rightarrow \mathbb{A}_{L \rightarrow K}$ such that $\{\chi_L(\Omega^*)\} \in \mathbb{A}_K^{\infty}(\Lambda)$ for every $\Lambda \in \mathfrak{L}$.

The resulting mixture model $\mathcal{M}(\mathfrak{L})$ is called a clusterable mixture model. If Λ belongs to a clusterable family, we shall call both Λ and $\Gamma = m(\Lambda)$ clusterable measures.

More precisely, Definition 4.3(b) means that for every $\Lambda \in \mathfrak{L}$, if we let $\Omega^* = M_L(T_L(m(\Lambda)))$, then $\alpha_L = \chi_L(\Omega^*)$ defines a regular assignment sequence (Definition 4.2). This will be discussed in more detail in the next section; please see equation (10) and Remark 4.4.

In contrast to regularity—which merely asserts the *existence* of a regular assignment sequence for Λ —clusterability takes this requirement one step further by requiring that a regular assignment sequence can in fact be determined from the ρ -projections $T_L(m(\Lambda))$ alone. The terminology “clusterable” is intended to provoke the reader into imagining χ_L as a cluster function that “clusters” the L components and L weights of Q^* together in such a way that $\Omega^*(\alpha)$ approximates Λ .

The problem of constructing a cluster function χ_L is a fascinating one, and will be taken up in Section 5. There, we will show that under a separation condition on the γ_k , regular assignment sequences can be recovered by single-linkage clustering, so this assumption is not vacuous. For the remainder of this section, however, we take this assumption on faith in order to complete our journey to identify Λ from Γ alone.

4.5 Main result

The final step (4) in our construction is to show that the constructed mixing measure $\Omega^*(\alpha_L)$ converges to Λ when $L \rightarrow \infty$. The rationale for introducing the concept of clusterability in the previous section is that this is precisely the condition that ensures this will happen for every $\Lambda \in \mathfrak{L}$. When this is the case, the mixture model $\mathcal{M}(\mathfrak{L})$ is identifiable:

Theorem 4.1. *If \mathfrak{L} is a \mathfrak{Q}_L -clusterable family then the mixture model $\mathcal{M}(\mathfrak{L})$ is identifiable. That is, the canonical embedding $m : \mathfrak{L} \rightarrow \mathcal{M}(\mathfrak{L})$ is a bijection onto $\mathcal{M}(\mathfrak{L})$.*

As illustrated by the cautionary tales from Examples 5 and 9, identification in nonparametric mixtures is a subtle problem, and this theorem thus provides a powerful general condition for identifiability in nonparametric problems. In Section 6 we will construct some explicit examples of mixture models that are clusterable.

The idea behind the proof is to invoke clusterability to obtain a cluster function χ_L which—when combined with the machinery previously introduced—yields a complete roadmap that takes us from a mixture distribution $m(\Lambda)$ to a mixing measure $\Omega(\alpha_L)$ over K atoms. The following diagram summarizes this roadmap:

$$\mathcal{M}(\mathfrak{L}) \xrightarrow{T_L} T_L(\mathfrak{L}) \xrightarrow{M_L} M_L(\mathfrak{L}) \xrightarrow{\chi_L} \mathbb{A}_{L \rightarrow K} \xrightarrow{\alpha_L} \mathfrak{Q}_{L \rightarrow K}. \quad (10)$$

From this roadmap, we can invoke regularity to show that $\Omega(\alpha_L)$ will be close to Λ in Wasserstein distance as L gets large.

Remark 4.4. Let us pause to unpack the sequence of maps given by (10). The functions M_L and α_L are needed for technical reasons to properly identify a mixing measure of interest, and T_L is a well-known projection operator. What’s novel is the cluster function χ_L , which can be interpreted as a cluster function that takes in L “points” and returns an assignment of these L points into K clusters. This cluster assignment is represented by the assignment map $\chi_L(\Omega^*) \in \mathbb{A}_{L \rightarrow K}$.

So far, Theorem 4.1 merely asserts some abstract conditions that guarantee identifiability. These conditions depend crucially on the choice of \mathfrak{Q}_L and having a family \mathfrak{L} that is clusterable. It remains to discuss (a) How to choose \mathfrak{Q}_L , and (b) When \mathfrak{L} is clusterable. The latter issue is the main topic of the next section. For (a), it suffices to use $\mathfrak{Q}_L = \mathfrak{G}_L$ in most situations. This will be corroborated by our experiments in Section 8. If it is known that the data is censored or fat-tailed, then other choices may be more appropriate (e.g. families of Gamma or t -distributions), but this is problem-dependent.

5 Estimation

We now turn our attention to the problem of identifying and learning a mixing measure Λ in practice. This will be broken down into two steps: 1) We first show that under a natural separation condition on the γ_k , regular assignment sequences can be recovered at the population-level (Section 5.1), and 2) We extend these results to the case where we have i.i.d. samples from $m(\Lambda)$ (Section 5.2). Most importantly, the results of this section imply that there exist nontrivial families of clusterable mixtures, and moreover these families can in fact be learned from data. In Section 6 we will provide explicit examples of such families.

5.1 Separation and clusterability

Theorem 4.1 is abstract and relies on the existence of a cluster function that can reconstruct regular assignments from just the overfitted mixing measure Ω^* . In this section, we make these concepts more concrete by constructing an explicit cluster function via a simple distance-based thresholding rule, which is equivalent to performing single-linkage clustering. Thus, this cluster function can be used in practice *without* knowing the optimal threshold in advance.

Given $\Omega \in \mathfrak{Q}_L$ with atoms q_ℓ , define the Hellinger diameter of Ω by

$$\Delta(\Omega) := \sup\{\rho(q, q') : q, q' \in \text{conv}(\text{supp}(\Omega))\}$$

where $\text{conv}(\cdot)$ denotes the convex hull in $\mathcal{P}(X)$. We will be interested in the special case $\Omega = \Omega_k^*(\alpha)$: $\Delta(\Omega_k^*(\alpha))$ quantifies how “compact” the mixture component $Q_k^*(\alpha)$ is. For any $\alpha \in \mathbb{A}_{L \rightarrow K}$, define

$$\eta(\alpha) := \sup_k \Delta(\Omega_k^*(\alpha)) + \sup_k \rho(\gamma_k, Q_k^*(\alpha)). \quad (11)$$

Finally, define the Hellinger distance matrix by

$$D(\Omega) = (\rho(q_i, q_j))_{i,j=1}^L. \quad (12)$$

Our goal is to show that if the atoms of Λ are sufficiently well-separated, then the cluster assignment α can be reconstructed by clustering the distance matrix $D^* = D(\Omega^*) = (\rho(q_i^*, q_j^*))_{i,j=1}^L$ (hence the choice of terminology *clusterable*). More precisely, we make the following definition:

Definition 5.1 (Separation). A mixing measure $\Lambda \in \mathcal{P}_0^2(X)$ is called δ -*separated* if $\inf_{i \neq j} \rho(\gamma_i, \gamma_j) > \delta$ for some $\delta > 0$.

It turns out that separation on the order $\eta(\alpha)$ (cf. (11)) is sufficient to define a cluster function:

Proposition 5.1. *Let $\Lambda \in \mathcal{P}_K^2(X)$. Suppose that the ρ -projection $Q^* = T_L \Gamma$ exists and is unique for some $L \geq K$ and $\alpha \in \mathbb{A}_{L \rightarrow K}$. If Λ is $4\eta(\alpha)$ -separated, then*

$$\alpha(i) = \alpha(j) \iff \rho(q_i^*, q_j^*) \leq \eta(\alpha), \quad (13)$$

$$\alpha(i) \neq \alpha(j) \iff \rho(q_i^*, q_j^*) \geq 2\eta(\alpha). \quad (14)$$

Moreover, α can be recovered by single-linkage clustering on D^* .

Thus, the assignment α can be recovered by single-linkage clustering of D^* *without knowing the optimal threshold* $\eta(\alpha)$.

Now suppose Λ is a regular mixing measure and let $\{\alpha_L\} \in \mathbb{A}_K^\infty(\Lambda)$. Define

$$\eta(\Lambda) := \limsup_{L \rightarrow \infty} \eta(\alpha_L) \quad (15)$$

and note that $\eta(\Lambda)$ is independent of the choice of regular assignment sequence $\{\alpha_L\}$. Moreover, we have $\eta(\Lambda) = \limsup_L \sup_k \Delta(\Omega_k^*(\alpha_L))$ since $\sup_k \rho(\gamma_k, Q_k^*(\alpha_L)) \rightarrow 0$. Thus $\eta(\Lambda)$ is a measure of the “width” of the approximating mixtures $Q_k^*(\alpha_L)$, as measured by their Hellinger diameter. The following corollary, which is an immediate consequence of Theorem 4.1 and Proposition 5.1, establishes that control over $\eta(\Lambda)$ is sufficient for \mathfrak{L} to be clusterable:

Corollary 5.2. *Suppose $\mathfrak{L} \subset \mathcal{P}_K^2(X)$ is a family of regular mixing measures such that for every $\Lambda \in \mathfrak{L}$ there exists $\xi > 0$ such that Λ is $(4 + \xi)\eta(\Lambda)$ -separated. Then \mathfrak{L} is clusterable and hence identifiable.*

Thus, we have a practical separation condition under which a regular mixture model becomes identifiable:

$$\inf_{i \neq j} \rho(\gamma_i, \gamma_j) > (4 + \xi)\eta(\Lambda). \quad (16)$$

The nonparametric components γ_k must be separated by a gap proportional to the Hellinger diameters of the approximating mixtures $Q_k^*(\alpha_L)$ (i.e. in the limit). This highlights the issue in Example 5—although the

means can be arbitrarily separated, as we increase the separation, the diameter of the components continues to increase as well. Thus, the γ_k cannot be chosen in a haphazard way (see also Example 9).

The separation condition (16) is quite weak, but no attempt has been made here to optimize this lower bound. For example, a minor tweak to the proof can reduce the constant of 4 to any constant $b > 2$. Although we expect that a more careful analysis can weaken this condition, our main focus here is to present the main idea behind identifiability and its connection to clusterability and separation, so we save such optimizations for future work. Further, although Proposition 5.1 justifies the use of single-linkage clustering in order to group the components $\{q_\ell^*\}$, one can easily imagine using other clustering schemes. Indeed, since the distance matrix D^* is always well-defined, we could have applied other clustering algorithms such as complete-linkage hierarchical clustering, K -means, or spectral clustering to D^* to define an assignment sequence $\{\alpha_L\}$. Any condition on D^* that ensures a clustering algorithm will correctly reconstruct a regular assignment sequence then yields an identification condition in the spirit of Proposition 5.1. For example, if the means of the overfitted components q_ℓ^* are always well-separated, then simple algorithms such as K -means could suffice to identify a regular assignment sequence. This highlights the advantage of our abstract viewpoint, in which the specific forms of both the assignment sequence $\{\alpha_L\}$ and the cluster functions χ_L are left unspecified.

5.2 Estimation of clusterable mixtures

Our results so far provide a framework for learning nonparametric mixture measures in principle, however, our discussion has so far been restricted to population-level properties of such measures. To complete this circle of ideas, it remains to discuss how to estimate Λ from data.

Suppose $Z^{(1)}, \dots, Z^{(n)} \stackrel{\text{iid}}{\sim} \Gamma$ and for each $L \geq K$, $\widehat{\Omega}_L = \widehat{\Omega}_L(Z^{(1)}, \dots, Z^{(n)})$ is a W_r -consistent estimator of Ω^* . That is, $\{\widehat{\Omega}_L\}$ is a sequence of estimators and for each L , $\lim_{n \rightarrow \infty} W_r(\widehat{\Omega}_{L,n}, \Omega_L^*) = 0$ where we have written $\Omega^* = \Omega_L^*$ to emphasize the dependence on L . For example, $\widehat{\Omega}$ could be the minimum Hellinger distance estimator (MHDE) from Beran [9]. Since L is a known quantity, the corresponding estimation problems are always well-specified, i.e. both $\widehat{\Omega}_L$ and Ω_L^* have the same, known number of components. As usual, for brevity we will omit the dependence of $\Omega^* = \Omega_L^*$ on L and $\widehat{\Omega} = \widehat{\Omega}_{L,n}$ on L and n in the sequel. Write

$$\widehat{Q} := m(\widehat{\Omega}) = \sum_{\ell=1}^L \widehat{\omega}_\ell \widehat{q}_\ell,$$

and note that $W_r(\widehat{\Omega}, \Omega^*) \rightarrow 0$ implies there is a permutation $\sigma : [L] \rightarrow [L]$ such that $\sup_\ell \rho(\widehat{q}_\ell, q_{\sigma(\ell)}^*) \rightarrow 0$. Without loss of generality, assume that the atoms are re-arranged so that $\sup_\ell \rho(\widehat{q}_\ell, q_\ell^*) \rightarrow 0$. Define

$$\varepsilon = \varepsilon_{L,n} := \sup_\ell \rho(\widehat{q}_\ell, q_\ell^*), \quad \delta = \delta_L := \sup_k \rho(Q_k^*(\alpha), \gamma_k). \quad (17)$$

Then ε represents the estimation error, which vanishes as n increases, and δ represents the approximation error, which vanishes as L increases.

Proposition 5.3. *Let $\Lambda \in \mathcal{P}_K^2(X)$. Suppose that the ρ -projection $Q^* = T_L \Gamma$ exists and is unique for some $L \geq K$. Suppose further that $L, \alpha \in \mathbb{A}_{L \rightarrow K}$, and n satisfy*

$$3\varepsilon - 2\delta < \sup_k \Delta(\Omega_k^*(\alpha)).$$

Define

$$\widehat{\eta} := 2\varepsilon + \sup_k \Delta(\Omega_k^*(\alpha)).$$

If Λ is $4\eta(\alpha)$ -separated, then $\rho(\widehat{q}_i, \widehat{q}_j) \leq \widehat{\eta}$ if and only if $\alpha(i) = \alpha(j)$, and the assignment function α can be recovered by single-linkage clustering on $\widehat{D} = D(\widehat{\Omega})$.

For each L and n , let $\widehat{\alpha} = \widehat{\alpha}_{L,n} \in \mathbb{A}_{L \rightarrow K}$ denote the assignment map induced by Proposition 5.3. With this notation, another way to phrase this result is that for sufficiently large L and n , we have $\widehat{\alpha} = \alpha$. In other words, single-linkage clustering of \widehat{D} yields the same clusters as the assignment α .

Proposition 5.3 is a finite sample result that holds for all sufficiently large L and n . By taking the limit as $L, n \rightarrow \infty$, we can show that Λ is asymptotically learnable.

Theorem 5.4. *Suppose Λ is a regular mixing measure such that Λ is $(4 + \xi)\eta(\Lambda)$ -separated for some $\xi > 0$. Then*

$$\lim_{L \rightarrow \infty} \lim_{n \rightarrow \infty} W_r(\widehat{\Omega}(\widehat{\alpha}_{L,n}), \Lambda) = 0, \quad (18)$$

$$\lim_{L \rightarrow \infty} \lim_{n \rightarrow \infty} \rho(\widehat{Q}_k(\widehat{\alpha}_{L,n}), \gamma_k) = 0. \quad (19)$$

Moreover, for fixed L and α satisfying the conditions of Proposition 5.3, we also have

$$\lim_{n \rightarrow \infty} W_r(\widehat{\Omega}(\widehat{\alpha}_{L,n}), \Omega^*(\alpha_L)) = 0, \quad (20)$$

$$\lim_{n \rightarrow \infty} \rho(\widehat{Q}_k(\widehat{\alpha}_{L,n}), Q_k^*(\alpha_L)) = 0. \quad (21)$$

Thus, we have a Wasserstein consistent estimate of Λ and Hellinger consistent estimates of the component measures γ_k . In applications, it will often be useful to strengthen the latter to *uniform* convergence of the densities (assuming they exist). When the families \mathfrak{Q}_L are equicontinuous, this is guaranteed by Theorem 1 of Sweeting [64]. We store this corollary away here for future use:

Corollary 5.5. *Let \widehat{g}_k be the density of $\widehat{Q}_k(\widehat{\alpha}_{L,n})$ and f_k be the density of γ_k . If the families \mathfrak{Q}_L are equicontinuous for all L , then $\widehat{g}_k \rightarrow f_k$ pointwise and uniformly over compact subsets of X as $L, n \rightarrow \infty$.*

In fact, even weaker assumptions than equicontinuity are possible; see for example Cuevas and Gonzalez-Manteiga [22].

Remark 5.1. It is interesting to inquire how to choose $L = L_n$ as a function of the sample size n in Theorem 5.4. As the proof of Theorem 5.4 shows, this is entirely governed by the rate of convergence of $W_r(\widehat{\Omega}, \Omega^*)$. To the best of our knowledge, such rates have not been established in the literature (e.g. for the MHDE of a misspecified mixture model), and this is an interesting direction for future work.

6 Examples

Corollary 5.2 identifies two key assumptions necessary to identify a mixture model via Theorem 4.1: *Regularity* and *separation*. As Example 9 indicates, these conditions are nontrivial and can fail to hold in practice. Fortunately, it is easy to construct a rich collection of nonparametric mixture models that are both regular and well-separated, which we present here.

Given a Borel set $\mathfrak{Q} \subset \mathcal{P}^2(X)$ and any integer K define

$$\mathfrak{F}(\mathfrak{Q}; K) = \left\{ \sum_{k=1}^K \lambda_k \Lambda_k : \Lambda_k \in \mathfrak{Q}_0, \lambda_k \geq 0, \sum_{k=1}^K \lambda_k = 1, \right. \\ \left. \text{supp}(\Lambda_k) \cap \text{supp}(\Lambda_{k'}) = \emptyset, \sum_{k=1}^K |\text{supp}(\Lambda_k)| < \infty \right\}, \quad (22)$$

where we recall that \mathfrak{Q}_0 is the subset of finite mixing measures in \mathfrak{Q} (Section 3.1). Then $\mathfrak{F}(\mathfrak{Q}; K) \subset \mathcal{P}_K^2(X)$ is the collection of mixing measures whose atoms themselves consist of finite mixture distributions from $\mathcal{M}_0(\mathfrak{Q})$. Note that $\mathfrak{F}(\mathfrak{Q}; K)$ is *not* the same as $\mathcal{P}_K(\mathfrak{Q}_0)$ since (22) also requires that no two Λ_k have overlapping supports (i.e. share a common atom). This assumption precludes the pathology from Example 9 and allows the atoms γ_k themselves to have overlapping supports. Finally, since $\mathcal{M}_0(\mathfrak{Q})$ is a nonparametric family of distributions (since there is no bound on the total number of atoms), $\mathfrak{F}(\mathfrak{Q}; K)$ is a genuine nonparametric mixture model.

Lemma 6.1. *For any $\mathfrak{Q} \subset \mathcal{P}^2(X)$ satisfying Condition (A) and any integer $K \geq 1$, the family $\mathfrak{F}(\mathfrak{Q}; K)$ is \mathfrak{Q}_L -regular. In particular, if $\mathcal{M}(\mathfrak{Q}_L)$ is identifiable for each L , then $\mathfrak{F}(\mathfrak{Q}; K)$ is \mathfrak{Q}_L -regular.*

Example 10 (Mixtures of Gaussian mixtures). By choosing $\Omega = \mathfrak{G}$ in (22), we obtain the family $\mathfrak{F}(\mathfrak{G}; K)$ of *mixtures of Gaussian mixtures*, i.e. a mixture model whose atoms are themselves Gaussian mixtures. Two such examples are depicted in Figure 4. In particular, the atoms γ_k can approximate any distribution on X . It follows that any $\Gamma \in \mathcal{P}(X)$ can be approximated by a mixture distribution from $\mathcal{M}(\mathfrak{F}(\mathfrak{G}; K))$ for some K .

Example 11 (Mixtures of exponential family mixtures). Since finite mixtures of exponential family distributions are identifiable [8], the previous example can be extended to *mixtures of exponential family mixtures*. For example, by taking Ω to be the family of Gamma mixing measures (Example 7), $\mathfrak{F}(\mathfrak{G}; K)$ represents the family of finite mixtures of Gamma mixtures, which can be used to model heavy-tailed and censored data.

Lemma 6.1 indicates that such mixtures are always regular, however, to be identifiable Theorem 4.1 requires that they also be *clusterable*. As noted by the previous examples, the family $\mathfrak{F}(\mathfrak{G}; K)$ can be quite large, so it should not come as a surprise that such mixtures are not in general identifiable. Thus, combining Lemma 6.1 with Corollary 5.2, we have the following identifiability result for subfamilies $\mathfrak{L} \subset \mathfrak{F}(\Omega; K)$:

Theorem 6.2. *Fix $\Omega \subset \mathcal{P}^2(X)$ satisfying Condition (A) and an integer $K \geq 1$. Suppose $\mathfrak{L} \subset \mathfrak{F}(\Omega; K)$ is such that for every $\Lambda \in \mathfrak{L}$ there exists $\xi > 0$ such that Λ is $(4 + \xi)\eta(\Lambda)$ -separated. Then \mathfrak{L} is identifiable.*

Thus, for example, any family of sufficiently well-separated mixtures of Gaussian mixtures is identifiable, and in fact learnable by Theorem 5.4.

Remark 6.1. Since Hellinger separation is a weaker criterion than mean separation, Theorem 6.2 does not require that the mixture distributions in $\mathcal{M}(\mathfrak{L})$ have components with well-separated means. In fact, each γ_k could have identical means (but different variances) and still be well-separated. This is illustrated with a real example in Figure 4b. This suggests that identifiability in mixture models is more general than what is needed in typical clustering applications, where a model such as Figure 4b would usually not be considered to have two distinct clusters. The subtlety here lies in interpreting clustering in $\mathcal{P}(X)$ (i.e. of the q_ℓ^*) vs. clustering in X (i.e. of samples $Z^{(i)} \sim \Gamma$), the latter of which is the interpretation used in data clustering.

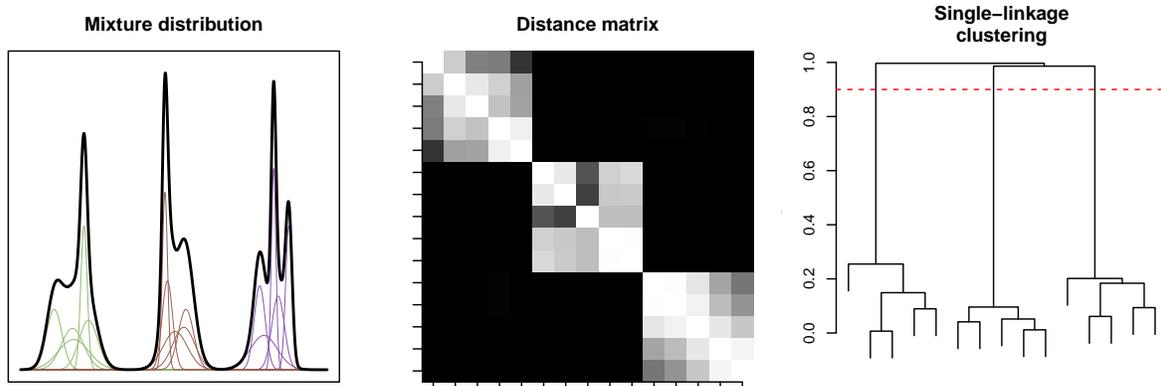
7 Bayes optimal clustering

As an application of the theory developed in Sections 4 and 5, we extend model-based clustering [11, 27] to the nonparametric setting. Given samples from Λ , we seek to partition these samples into K clusters. More generally, Λ defines a partition of the input space X , which can be formalized as a function $c : X \rightarrow [K]$, where K is the number of partitions or “clusters”. First, let us recall the classical Gaussian mixture model (GMM): If $f_1(\cdot; a_1, v_1), \dots, f_K(\cdot; a_K, v_K)$ is a collection of Gaussian density functions, then for any choice of $\lambda_k \geq 0$ such that $\sum_k \lambda_k = 1$ the combination

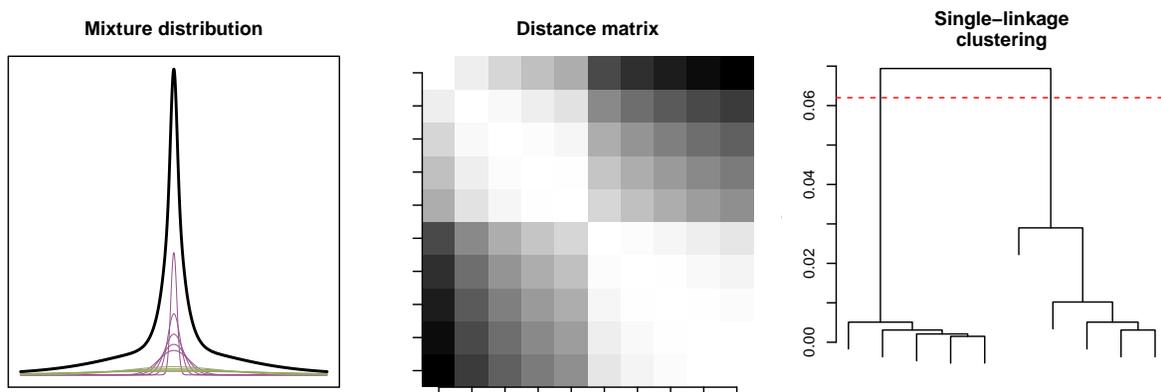
$$F(z) = \sum_{k=1}^K \lambda_k f_k(z; a_k, v_k); \quad z \in \mathbb{R}^d \quad (23)$$

is a GMM. The model (23) is of course equivalent to the integral (3) (see also Example 1), and the Gaussian densities $f_k(z; a_k, v_k)$ can obviously be replaced with any family of parametric densities $f_k(z; \phi_k)$.

Intuitively, the density F has K distinct clusters given by the K Gaussian densities f_k , defining what we call the *Bayes optimal partition* over X into regions where each of the Gaussian components is most likely. It should be obvious that as long as a mixture model $\mathcal{M}(\mathfrak{L})$ is identifiable, the Bayes optimal partition will be well-defined and has a unique interpretation in terms of distinct clusters of the input space X . Thus, the theory developed in the previous sections can be used to extend these ideas to the nonparametric setting. Since the clustering literature is full of examples of datasets that are not well-approximated by parametric mixtures [e.g. 50, 71], there is significant interest in such an extension. In the remainder of this section, we will apply our abstract framework to this problem. First, we discuss identifiability issues with the concept of a Bayes optimal partition (Section 7.1). Then, we provide conditions under which a Bayes optimal partition can be learned from data (Section 7.2).



(a) Example of $\mathfrak{F}(\mathfrak{G}; K)$ with $K = 3$ and well-separated means.



(b) Example of $\mathfrak{F}(\mathfrak{G}; K)$ with $K = 2$ and identical means. See Remark 6.1.

Figure 4: Examples of Theorem 6.2 and Example 10 with mixtures of Gaussian mixtures. (left) Original mixture distribution (thick black line), with Gaussian components coloured according to membership in different Λ_k . (middle) The true distance matrix D^* . (right) Results of single-linkage clustering on D^* , cut to find correct number of clusters.

7.1 Bayes optimal partitions

Throughout the rest of this section, we assume that X is compact and all probability measures are absolutely continuous with respect to some base measure ζ , and hence have density functions. Assume Γ is fixed and write $F = F_\Gamma$ for the density of Γ and f_k for the density of γ_k . Thus whenever $\Gamma \in \mathcal{M}_0(X)$ is a finite mixture we can also write

$$F = \int f_\gamma d\Lambda(\gamma) = \sum_{k=1}^K \lambda_k f_k. \quad (24)$$

For any $\Lambda \in \mathcal{P}_K^2(X)$, define the usual Bayes classifier [e.g. 24]:

$$c_\Lambda(x) := \arg \max_{k \in [K]} \lambda_k f_k(x). \quad (25)$$

Note that c_Λ is only well-defined up to a permutation of the labels (i.e. any labeling of $\text{supp}(\Lambda)$ defines an equivalent classifier). Furthermore, $c_\Lambda(x)$ not properly defined when $\lambda_i f_i(x) = \lambda_j f_j(x)$ for $i \neq j$. To account for this, define an exceptional set

$$E_0 := \bigcup_{i \neq j} \{x \in X : \lambda_i f_i(x) = \lambda_j f_j(x)\}, \quad (26)$$

In principle, E_0 should be small—in fact it will typically have measure zero—hence we will be content to partition $X_0 = X - E_0$. Recall that a *partition* of a space X is a family of subsets $A_k \subset X$ such that $A_k \cap A_{k'} = \emptyset$ for all $k \neq k'$ and $\cup_k A_k = X$. We denote the space of all partitions of X by $\Pi(X)$.

The following definition is standard from the literature [e.g. 18, 27]:

Definition 7.1 (Bayes optimal partition). Define an equivalence relation on X_0 by declaring

$$x \sim y \iff c_\Lambda(x) = c_\Lambda(y). \quad (27)$$

This relation induces a partition on X_0 which we denote by π_Λ or $\pi(\Lambda)$. This partition is known as the *Bayes optimal partition*.

Remark 7.1. Although the function c_Λ is only unique up to a permutation, the partition defined by (27) is always well-defined and independent of the permutation used to label the γ_k .

Given samples from the mixture distribution $\Gamma = m(\Lambda)$, we wish to learn the Bayes optimal partition π_Λ . Unfortunately, there is—yet again—an identifiability issue: If there is more than one mixture measure Λ that represents Γ , the Bayes optimal partition is not well-defined.

Example 12 (Non-identifiability of Bayes optimal partition). In Example 5 and Figure 2, we have four valid representations of Γ as a mixture of sub-Gaussians. In all four cases, each representation leads to a different Bayes optimal partition, even though they each represent the same mixture distribution.

Clearly, if Λ is identifiable, then the Bayes optimal partition is automatically well-defined. Thus the theory from Section 4 immediately implies the following:

Proposition 7.1. *If $\mathcal{M}(\mathcal{L})$ is a clusterable mixture model, then there is a well-defined Bayes optimal partition π_Γ for any $\Gamma \in \mathcal{M}(\mathcal{L})$.*

In particular, whenever $\mathcal{M}(\mathcal{L})$ is clusterable it makes sense to write c_Γ and π_Γ instead of c_Λ and π_Λ , respectively. This provides a useful framework for discussing and analyzing partition-based clustering in nonparametric settings. As discussed previously, a K -clustering of X is equivalent to a function that assigns each $x \in X$ an integer from 1 to K , where K is the number of clusters. Clearly, up to the exceptional set E_0 , (25) is one such function. Thus, the Bayes optimal partition π_Γ can be interpreted as a valid K -clustering.

7.2 Learning partitions from data

Write $\Gamma = m(\Lambda)$ and assume that Λ is identifiable from Γ . Suppose we are given i.i.d. samples $Z^{(1)}, \dots, Z^{(n)} \stackrel{\text{i.i.d.}}{\sim} \Gamma$ and that we seek the Bayes optimal partition $\pi_\Gamma = \pi_\Lambda$. Our strategy will be the following:

1. Use a consistent estimator $\widehat{\Omega}$ to learn Ω^* for some $L \gg K$;
2. Theorem 5.4 guarantees that we can learn a cluster assignment $\widehat{\alpha}_{L,n}$ such that $\widehat{\Omega}(\widehat{\alpha}_{L,n})$ consistently estimates Λ ;
3. Use $\pi(\widehat{\Omega}(\widehat{\alpha}_{L,n}))$ to approximate $\pi_\Lambda = \pi_\Gamma$.

The hope, of course, is that $\pi(\widehat{\Omega}(\widehat{\alpha}_{L,n})) \rightarrow \pi_\Gamma$. There are, however, complications: What do we mean by convergence of partitions? Does $\pi(\widehat{\Omega}(\widehat{\alpha}_{L,n}))$ even converge?

Instead of working directly with the partitions $\pi(\widehat{\Omega}(\widehat{\alpha}_{L,n}))$, we will work with the Bayes classifier (25). Write \widehat{g}_ℓ and \widehat{G} for the densities of \widehat{q}_ℓ and \widehat{Q} , respectively, and

$$\widehat{G}_k(x) := \frac{1}{\widehat{\omega}_k} \sum_{\ell \in \widehat{\alpha}_{L,n}^{-1}(k)} \widehat{\omega}_\ell \widehat{g}_\ell(x), \quad \widehat{\omega}_k := \sum_{\ell \in \widehat{\alpha}_{L,n}^{-1}(k)} \widehat{\omega}_\ell. \quad (28)$$

Then \widehat{G}_k is the density of $\widehat{Q}_k(\widehat{\alpha}_{L,n})$, where here and above we have suppressed the dependence on $\widehat{\alpha}_{L,n}$. Now define the estimated classifier (cf. (25))

$$\widehat{c}_{L,n}(x) := c_{\widehat{\Omega}(\widehat{\alpha}_{L,n})}(x) = \arg \max_{k \in [K]} \widehat{\omega}_k \widehat{G}_k(x). \quad (29)$$

By considering classification functions as opposed to the partitions themselves, we may consider ordinary convergence of the function $\widehat{c}_{L,n}$ to c_Γ , which gives us a convenient notion of consistency for this problem. Furthermore, we can compare partitions by comparing the Bayes optimal equivalence classes $A_k := c^{-1}(k) = \{x \in X : c(x) = k\}$ to the estimated equivalence classes $\widehat{A}_{L,n,k} := \widehat{c}_{L,n}^{-1}(k)$ by controlling $A_k \Delta \widehat{A}_{L,n,k}$, where $F \Delta G = (F - G) \cup (G - F)$ is the usual symmetric difference of two sets. Specifically, we'd like to show that the difference $A_k \Delta \widehat{A}_{L,n,k}$ is small. To this end, define a fattening of E_0 by

$$E_0(t) := \bigcup_{i \neq j} \{x \in X : |\lambda_i f_i(x) - \lambda_j f_j(x)| \leq t\}, \quad t > 0. \quad (30)$$

Then of course $E_0 = E_0(0)$. When the boundaries between classes are sharp, this set will be small, however, if two classes have substantial overlap, then $E_0(t)$ can be large even if t is small. In the latter case, the equivalence classes A_k (and hence the clusters) are less meaningful. The purpose of $E_0(t)$ is to account for sampling error in the estimated partition.

Theorem 7.2. *Assume that $\widehat{G}_k \rightarrow \gamma_k$ uniformly on X as $L, n \rightarrow \infty$ and v is any measure on X . Then there exists a sequence $t_{L,n} \rightarrow 0$ such that $\widehat{c}_{L,n}(x) = c_\Lambda(x)$ for all $x \in X - E_0(t_{L,n})$ and*

$$v \left(\bigcup_{k=1}^K A_k \Delta \widehat{A}_{L,n,k} \right) \leq v(E_0(t_{L,n})) \rightarrow v(E_0). \quad (31)$$

The uniform convergence assumption in Theorem 7.2 may seem strong, however, recall Corollary 5.5, which guarantees uniform convergence whenever Ω_L is equicontinuous. For example, recalling Examples 1 and 3, it is straightforward to show the following:

Corollary 7.3. *Suppose $X \subset \mathbb{R}^d$, Ω is a compact subset of \mathfrak{G} and v is any measure on X . If Λ is a Ω_L -clusterable measure, then there exists a sequence $t_{L,n} \rightarrow 0$ such that $\widehat{c}_{L,n}(x) = c_\Lambda(x)$ for all $x \in X - E_0(t_{L,n})$ and*

$$v \left(\bigcup_{k=1}^K A_k \Delta \widehat{A}_{L,n,k} \right) \leq v(E_0(t_{L,n})) \rightarrow v(E_0). \quad (32)$$

A concrete example of a compact subset of \mathfrak{G} was given in Example 3.

We can interpret Theorem 7.2 as follows: As long as we take L and n large enough and the boundaries between each pair of classes is sharp (in the sense that $v(E_0(t_{L,n}))$ is small), the difference between the true Bayes optimal partition and the estimated partition becomes negligible. In fact, it follows trivially from Theorem 7.2 that $\widehat{c}_{L,n} \rightarrow c_\Lambda$ uniformly on $X - E_0(t)$ for any fixed $t > 0$. Thus, Theorem 7.2 gives rigorous justification to the approximation heuristic outlined above, and establishes precise conditions under which *nonparametric* clusterings can be learned from data.

Remark 7.2. The sequence $t_{L,n}$ is essentially the rate of convergence of $\widehat{G}_k \rightarrow \gamma_k$. It is an interesting question to quantify this convergence rate more precisely, which we have left to future work.

8 Experiments

The theory developed so far suggests an intuitive meta-algorithm for nonparametric clustering. This algorithm can be implemented in just a few lines of code, making it a convenient alternative to more complicated algorithms in the literature. The purpose of this section is merely to illustrate how our theory can be translated into a simple and effective meta-algorithm for nonparametric clustering, which should be understood as a complement to and not a replacement for existing methods that work well in practice.

As in Section 7, we assume we have i.i.d. samples $Z^{(1)}, \dots, Z^{(n)} \stackrel{\text{iid}}{\sim} \Gamma = m(\Lambda)$. Given these samples, we propose the following meta-algorithm:

1. Estimate an overfitted GMM \widehat{Q} with $L \gg K$ components;
2. Define an estimated assignment function $\widehat{\alpha}$ by using single-linkage clustering to group the components of \widehat{Q} together;
3. Use this clustering to define K mixture components $\widehat{Q}_k(\widehat{\alpha})$;
4. Define a partition on X by using Bayes' rule, e.g. (28-29).

Note that Figure 4 already illustrates two examples where this procedure succeeds in the limit as $n \rightarrow \infty$. To further assess the effectiveness of this meta-algorithm in practice, we evaluated its performance on simulated data. In our implementation we used the EM algorithm with regularization and weight clipping to learn the GMM \widehat{Q} in step 1, although clearly any algorithm for learning a GMM can be used in this step.

We call the resulting algorithm NPMIX (for *NonParametric MIX*ture modeling). To illustrate the basic idea, we first implemented four simple one-dimensional models:

- (i) GAUSSGAMMA ($K = 4$): A mixture of two Gaussian distributions, one gamma distribution, and a Gaussian mixture.
- (ii) GUMBEL ($K = 3$): A GMM with three components that has been contaminated with non-Gaussian, Gumbel noise.
- (iii) POLY ($K = 2$): A mixture of two polynomials with non-overlapping supports.
- (iv) SOBOLEV ($K = 3$): A mixture of three random nonparametric densities, generated from random expansions of an orthogonal basis for the Sobolev space $H^1(\mathbb{R})$.

The results are shown in Figure 5. These examples illustrate the basic idea behind the algorithm: Given samples, overfitted mixture components (depicted by dotted lines in Figure 5) are used to approximate the global nonparametric mixture distribution (solid black line). Each of these components is then clustered, with the resulting partition of $X = \mathbb{R}$ depicted alongside the true Bayes optimal partition. In each case, by choosing to cut the cluster tree to produce K components, the induced partitions provided appear to provide sensible and meaningful approximations to the true partitions.

To further validate the proposed algorithm, we implemented the following two-dimensional mixture models and compared the cluster accuracy to existing clustering algorithms on simulated data:

- (v) MOONS ($K = 2$): A version of the classical MOONS dataset in two-dimensions. This model exhibits a classical failure case of spectral clustering, which is known to have difficulties when clusters are unbalanced (i.e. $\lambda_1 \neq \lambda_2$). For this reason, we ran experiments with both balanced and unbalanced clusters.
- (vi) TARGET ($K = 6$): A GMM derived from the TARGET dataset (Figure 8). The GMM has 143 components that are clustered into 6 groups based on the original TARGET dataset from [71].

Visualizations of the results for our method are shown in Figures 6, 7, and 8. One of the advantages of our method is the construction of an explicit partition of the entire input space (in this case, $X = \mathbb{R}^2$), which is depicted in all three figures. Mixture models are known to occasionally lead to unintuitive cluster assignments in the tails, which we observed with the unbalanced MOONS model. This is likely an artifact of the sensitivity of the EM algorithm, and can likely be corrected by using a more robust mixture model estimator in the first step.

We compared NPMIX against four well-known benchmark algorithms: (i) K -means, (ii) Spectral clustering, (iii) Single-linkage hierarchical clustering, and (iv) A Gaussian mixture model (GMM) with K components. We only considered methods that classify every sample in a dataset (this precludes, e.g. density-based clustering). Moreover, of these four algorithms, only K -means and GMM provide a partition of the entire input space X , which allows for new samples to be classified without re-running the algorithm. All of the methods (including NPMIX) require the specification of the number of clusters K , which was set to the correct number according to the model. In each experiment, we sampled random data from each model and then used each clustering algorithm to classify each sample. To assess cluster accuracy, we computed the adjusted

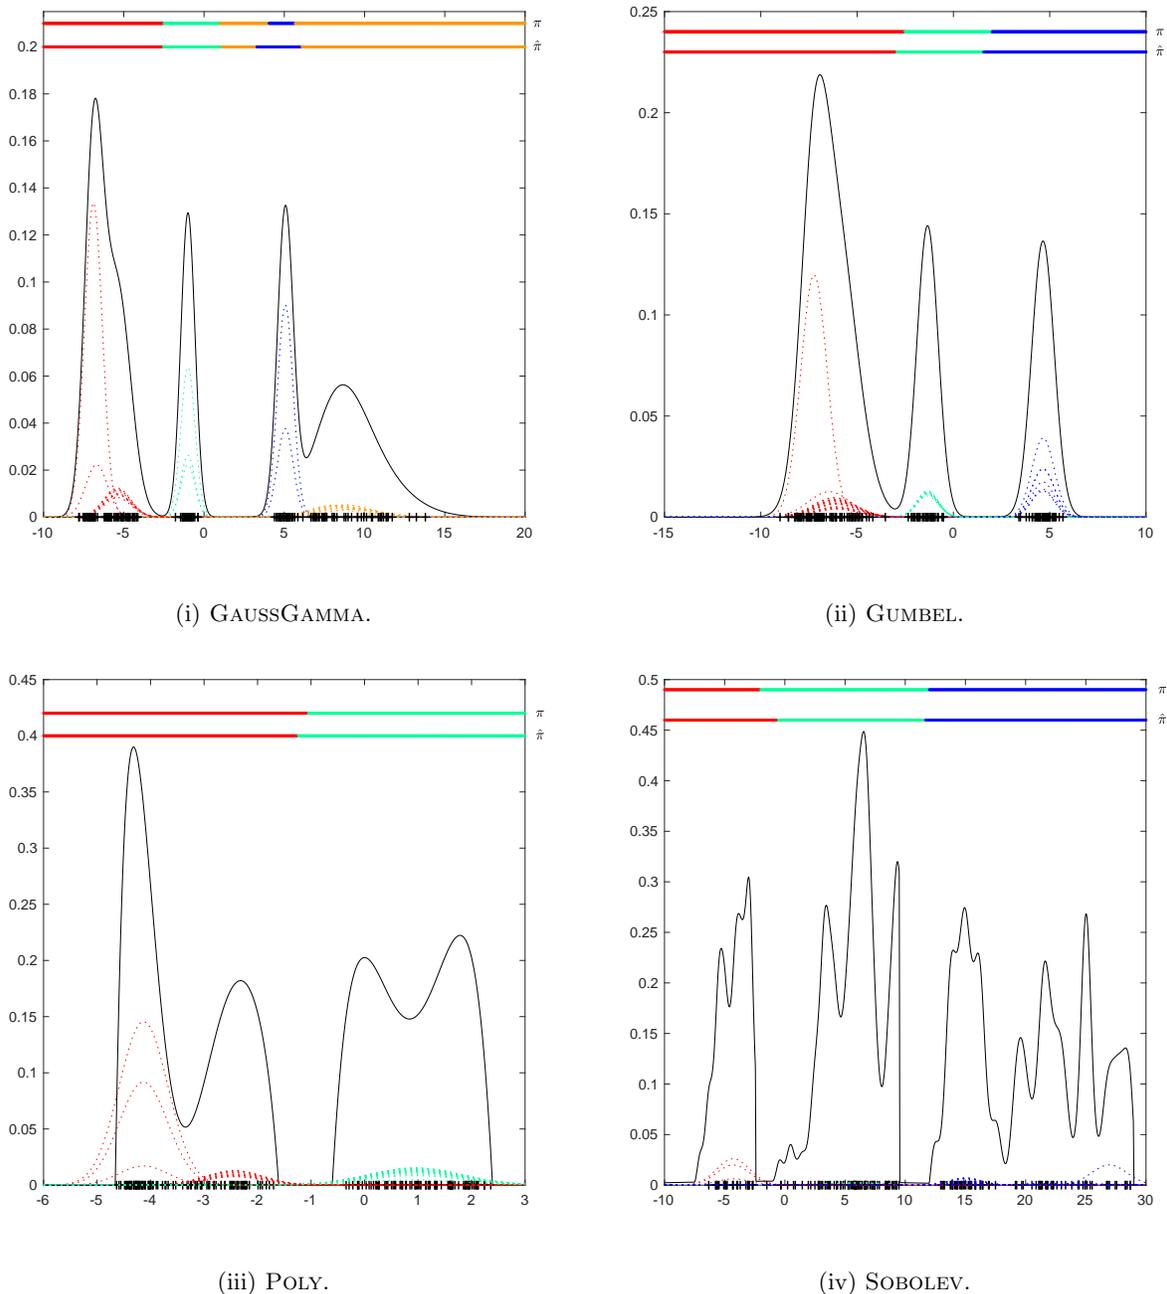


Figure 5: Examples (i)-(iv) of one-dimensional mixture models. The original mixture density is depicted as a solid black line, with the overfitted Gaussian mixture components as dotted lines, coloured according to the cluster they are assigned to. The true Bayes optimal partition π and the estimated partition $\hat{\pi}$ are depicted by the horizontal lines at the top, and the raw data are plotted on the x -axis for reference.

RAND index (ARI) for the clustering returned by each method. ARI is a standard permutation-invariant measure of cluster accuracy in the literature.

The results are shown in Table 1. On the unbalanced MOONS data, NPMIX clearly outperformed each of the four existing methods. On balanced data, K -means, spectral clustering, and GMM improved significantly, with spectral clustering performing quite well on average. All four algorithms were still outperformed by

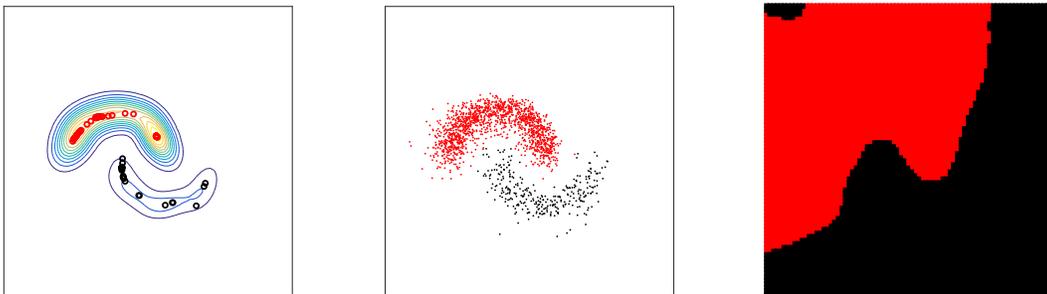


Figure 6: Example of a successful clustering on the unbalanced MOONS mixture model using NPMIX. (Left) Contour plot of overfitted Gaussian mixture approximation, centers marked with \circ 's. (Middle) Original data colour coded by the approximate Bayes optimal partition. (Right) Estimated Bayes optimal partition, visualized as the input space X colour-coded by estimated cluster membership.

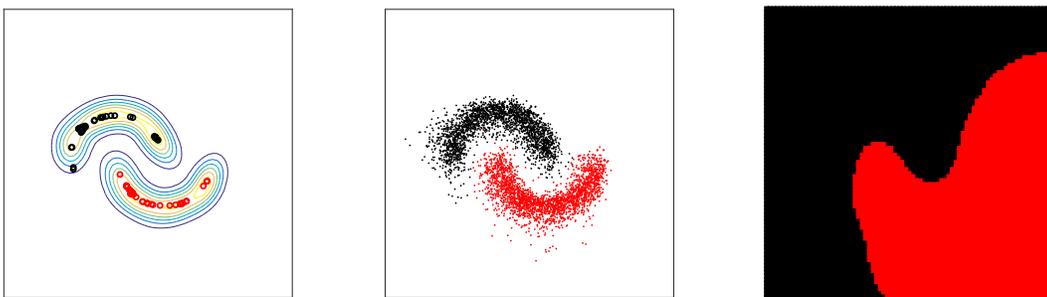


Figure 7: Example of a successful clustering on the balanced MOONS mixture model using NPMIX. (Left) Contour plot of overfitted Gaussian mixture approximation, centers marked with \circ 's. (Middle) Original data colour coded by the approximate Bayes optimal partition. (Right) Estimated Bayes optimal partition, visualized as the input space X colour-coded by estimated cluster membership.

NPMIX. On TARGET, the results were more interesting: Both single-linkage and spectral clustering perform very well on this dataset. NPMIX shows more variance in its performance, as indicated by the high median (0.998) and lower mean (0.696). On 57/100 runs, the ARI for NPMIX was > 0.99 , and on the rest the ARI was < 0.6 . This is likely caused by sensitivity to outliers in the TARGET model, and we expect that this can be corrected by using a more robust algorithm (e.g. instead of the vanilla EM algorithm). As our motivations are mainly theoretical, we leave more detailed fine-tuning of this algorithm and thorough side-by-side comparisons to future work. For example, by using the learned mixture density to remove “background samples” (e.g. as in density-based clustering), this algorithm can be trivially improved.

9 Discussion

We have established a new set of identifiability results for nonparametric mixtures that rely on the notion of *clusterability*. In particular, our results allow for an arbitrary number of components and for each component to take on essentially any shape. The key assumption is separation between the components, which allows simple clustering algorithms such as hierarchical clustering to recover individual mixture components from an overfitted mixture density estimator. Furthermore, we established conditions under which identified mixtures and their partitions can be consistently estimated from data. We also discussed applications to data clustering, including a nonparametric notion of the Bayes optimal partition and an intuitive meta-

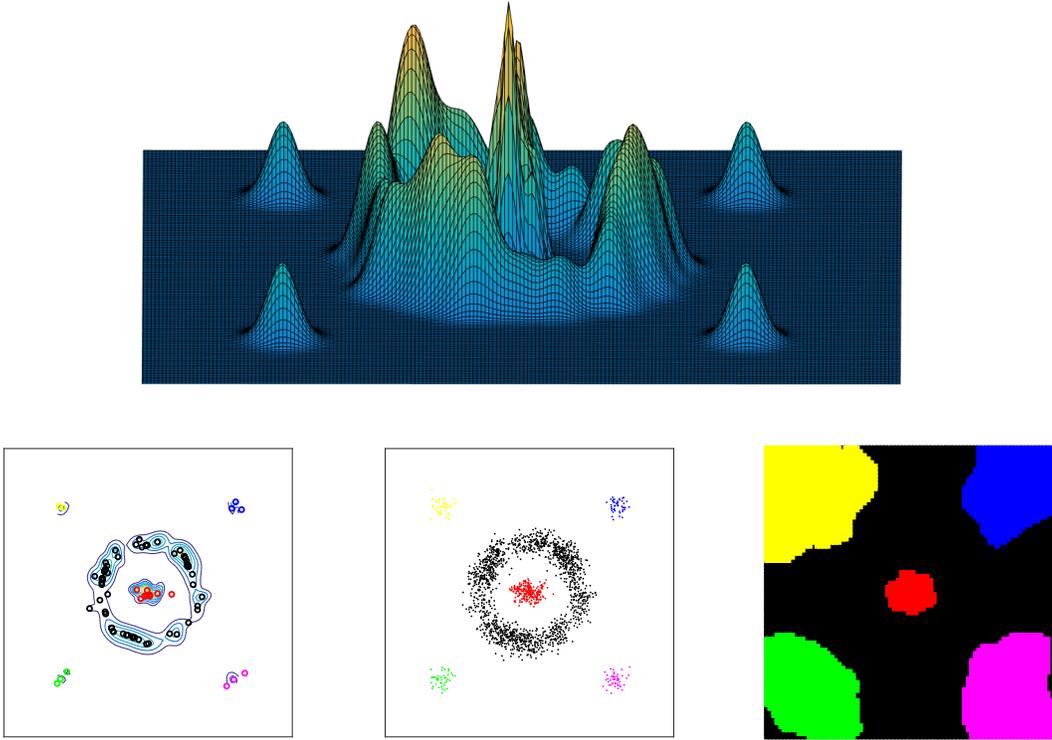


Figure 8: Example of a successful clustering on the TARGET mixture model using NPMIX. (Top) Density plot of the original mixture density. (Left) Contour plot of overfitted Gaussian mixture approximation, centers marked with \circ 's. (Middle) Original data colour coded by the approximate Bayes optimal partition. (Right) Estimated Bayes optimal partition, visualized as the input space X colour-coded by estimated cluster membership.

algorithm for nonparametric clustering.

The assumption that the number of components K is known is of course restrictive in practice, however, this assumption can be substantially relaxed as follows: If K is unknown, simply test whether or not there exists a K such that the separation criterion (16) holds. If such a K exists and is unique, then the resulting K -mixture is identifiable. In practice, however, there may be more than one value of K for which (16) holds. Furthermore, if Λ is identifiable for some K , it may not be the case that Λ is identifiable for $K' < K$ owing to the separation criterion (15) (cf. (11)). Of course, such an exhaustive search may not be practical, in which case it would be interesting to study efficient algorithms for finding such a K .

It would also be interesting to study convergence rates for the proposed estimators. In particular, there are two important quantities of interest in deriving these rates: The sample size n and the number of overfitted components L . Interestingly, it was only recently that the minimax rate of estimation for *parametric* mixtures was correctly determined [33], which is $n^{1/(4(s-s_0)+2)}$ in the L_1 -Wasserstein metric, where s_0 is the true number of mixture components and s is the number used in estimation. See also [20, 35, 36, 51]. In the general case, this is also related to problems in agnostic learning [45]. In our nonparametric setting, we expect these rates to depend on both L and n . Furthermore, it is necessary to control the distance between the ρ -projection Q^* and Γ , which depends on the choice of L alone. This latter problem will almost certainly require imposing additional regularity conditions on Γ . Understanding these convergence rates is also essential to prescribing the number of components L given the sample size n . In

MOONS (UNBALANCED)	Mean ARI	Median ARI	st. dev.
NPMIX	0.727	0.955	0.284
K -means	0.126	0.124	0.016
Spectral	0.197	0.122	0.232
Single-linkage	0.001	0.001	0.002
GMM	0.079	0.078	$< 10^{-3}$
MOONS (BALANCED)	Mean ARI	Median ARI	st. dev.
NPMIX	0.934	0.972	0.188
K -means	0.502	0.503	0.021
Spectral	0.909	0.910	0.013
Single-linkage	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$
GMM	0.782	0.783	$< 10^{-3}$
TARGET	Mean ARI	Median ARI	st. dev.
NPMIX	0.696	0.998	0.354
K -means	0.081	0.072	0.034
Spectral	0.967	0.975	0.077
Single-linkage	0.824	1.000	0.222
GMM	0.126	0.124	0.002

Table 1: Average and median adjusted RAND index (ARI) for $N = 100$ simulations of three different nonparametric mixture models.

practice, one should take $L < n$, however, a more detailed comparison is left to future work.

Finally, it would be of significant interest to apply existing clustering theory to find new conditions that guarantee clusterability in the same way that Proposition 5.1 shows that separability is sufficient for single-linkage clustering. We have already noted that the separation constant $4\eta(\alpha)$ can be reduced. Furthermore, in simulations we have observed that complete-linkage is often sufficient when working with the proposed NPMIX algorithm. But under what precise conditions on Γ is complete-linkage sufficient? By applying known results from the clustering literature, it may be possible to extend our results to prove deeper identifiability theorems for nonparametric mixtures.

References

- [1] D. Achlioptas and F. McSherry. On spectral learning of mixtures of distributions. In *International Conference on Computational Learning Theory*, pages 458–469. Springer, 2005.
- [2] K. E. Ahmad and E. K. Al-Hussaini. Remarks on the non-identifiability of mixtures of distributions. *Annals of the Institute of Statistical Mathematics*, 34(1):543–544, 1982.
- [3] E. M. Airoldi, T. B. Costa, and S. H. Chan. Stochastic blockmodel approximation of a graphon: Theory and consistent estimation. In *Advances in Neural Information Processing Systems*, pages 692–700, 2013.
- [4] E. S. Allman, C. Matias, and J. A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *Annals of Statistics*, pages 3099–3132, 2009.
- [5] A. Anandkumar, D. Hsu, A. Javanmard, and S. Kakade. Learning linear Bayesian networks with latent variables. In *Proceedings of The 30th International Conference on Machine Learning*, pages 249–257, 2013.
- [6] A. Anandkumar, D. Hsu, M. Janzamin, and S. Kakade. When are overcomplete topic models identifiable? uniqueness of tensor tucker decompositions with structured sparsity. *Journal of Machine Learning Research*, 16:2643–2694, 2015.

- [7] S. Arora and R. Kannan. Learning mixtures of separated nonspherical gaussians. *Annals of Applied Probability*, pages 69–92, 2005.
- [8] O. Barndorff-Nielsen. Identifiability of mixtures of exponential families. *Journal of Mathematical Analysis and Applications*, 12(1):115–121, 1965.
- [9] R. Beran. Minimum hellinger distance estimates for parametric models. *The Annals of Statistics*, pages 445–463, 1977.
- [10] N. Bochkina, J. Rousseau, et al. Adaptive density estimation based on a mixture of gammas. *Electronic Journal of Statistics*, 11(1):916–962, 2017.
- [11] H. H. Bock. Probabilistic models in cluster analysis. *Computational Statistics & Data Analysis*, 23(1): 5–28, 1996.
- [12] S. Bonhomme, K. Jochmans, and J.-M. Robin. Non-parametric estimation of finite mixtures from repeated measurements. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78 (1):211–229, 2016.
- [13] S. Bonhomme, K. Jochmans, J.-M. Robin, et al. Estimating multivariate latent-structure models. *The Annals of Statistics*, 44(2):540–563, 2016.
- [14] L. Bordes, S. Mottelet, P. Vandekerkhove, et al. Semiparametric estimation of a two-component mixture model. *The Annals of Statistics*, 34(3):1204–1232, 2006.
- [15] L. Bordes, I. Kojadinovic, P. Vandekerkhove, et al. Semiparametric estimation of a two-component mixture of linear regressions in which one component is known. *Electronic Journal of Statistics*, 7: 2603–2644, 2013.
- [16] M. Broniatowski and A. Keziou. Minimization of ϕ -divergences on sets of signed measures. *Studia Scientiarum Mathematicarum Hungarica*, 43(4):403–442, 2006.
- [17] C. Calcaterra and A. Boldt. Approximating with gaussians. *arXiv preprint arXiv:0805.3795*, 2008.
- [18] J. E. Chacón et al. A population background for nonparametric density-based clustering. *Statistical Science*, 30(4):518–532, 2015.
- [19] K. Chaudhuri and S. Dasgupta. Rates of convergence for the cluster tree. In *Advances in Neural Information Processing Systems*, pages 343–351, 2010.
- [20] J. Chen. Optimal rate of convergence for finite mixture models. *Annals of Statistics*, pages 221–233, 1995.
- [21] Y.-C. Chen, C. R. Genovese, L. Wasserman, et al. A comprehensive approach to mode clustering. *Electronic Journal of Statistics*, 10(1):210–241, 2016.
- [22] A. Cuevas and W. Gonzalez-Manteiga. Data-driven smoothing based on convexity properties. In *Nonparametric Functional Estimation and Related Topics*, pages 225–240. Springer, 1991.
- [23] S. Dasgupta. Learning mixtures of gaussians. In *Foundations of Computer Science, 1999. 40th Annual Symposium on*, pages 634–644. IEEE, 1999.
- [24] L. Devroye, L. Györfi, and G. Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- [25] X. D’Haultfœuille and P. Février. Identification of mixture models using support variations. *Journal of Econometrics*, 189(1):70–82, 2015.
- [26] J. Eldridge, M. Belkin, and Y. Wang. Graphons, mergeons, and so on! In *Advances in Neural Information Processing Systems*, pages 2307–2315, 2016.
- [27] C. Fraley and A. E. Raftery. Model-based clustering, discriminant analysis, and density estimation. *Journal of the American statistical Association*, 97(458):611–631, 2002.
- [28] E. Gassiat and J. Rousseau. Non parametric finite translation mixtures with dependent regime. *arXiv preprint arXiv:1302.2345*, 2013.

- [29] P. Hall and X.-H. Zhou. Nonparametric estimation of component distributions in a multivariate mixture. *Annals of Statistics*, pages 201–224, 2003.
- [30] P. Hall, A. Neeman, R. Pakyari, and R. Elmore. Nonparametric inference in multivariate mixtures. *Biometrika*, 92(3):667–678, 2005.
- [31] J. A. Hartigan. *Clustering algorithms*, volume 209. Wiley New York, 1975.
- [32] J. A. Hartigan. Consistency of single linkage for high-density clusters. *Journal of the American Statistical Association*, 76(374):388–394, 1981.
- [33] P. Heinrich and J. Kahn. Minimax rates for finite mixture estimation. *arXiv preprint arXiv:1504.03506*, 2015.
- [34] T. Hettmansperger and H. Thomas. Almost nonparametric inference for repeated measures in mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 62(4):811–825, 2000.
- [35] N. Ho and X. Nguyen. On strong identifiability and convergence rates of parameter estimation in finite mixtures. *Electronic Journal of Statistics*, 10(1):271–307, 2016.
- [36] N. Ho and X. Nguyen. Singularity structures and impacts on parameter estimation in finite mixtures of distributions. *arXiv preprint arXiv:1609.02655*, 2016.
- [37] P. W. Holland, K. B. Laskey, and S. Leinhardt. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983.
- [38] H. Holzmann, A. Munk, and T. Gneiting. Identifiability of finite mixtures of elliptical distributions. *Scandinavian journal of statistics*, 33(4):753–763, 2006.
- [39] D. R. Hunter and D. S. Young. Semiparametric mixtures of regressions. *Journal of Nonparametric Statistics*, 24(1):19–38, 2012.
- [40] D. R. Hunter, S. Wang, and T. P. Hettmansperger. Inference for mixtures of symmetric distributions. *The Annals of Statistics*, pages 224–251, 2007.
- [41] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- [42] K. Jochmans, M. Henry, and B. Salanié. Inference on two-component mixtures under tail restrictions. *Econometric Theory*, 33(3):610–635, 2017.
- [43] R. Kannan, H. Salmasian, and S. Vempala. The spectral method for general mixture models. *SIAM Journal on Computing*, 38(3):1141–1156, 2008.
- [44] M. Levine, D. R. Hunter, and D. Chauveau. Maximum smoothed likelihood for multivariate mixtures. *Biometrika*, pages 403–416, 2011.
- [45] J. Li and L. Schmidt. A nearly optimal and agnostic algorithm for properly learning a mixture of k gaussians, for any constant k . *arXiv preprint arXiv:1506.01367*, 2015.
- [46] B. G. Lindsay. Mixture models: theory, geometry and applications. In *NSF-CBMS regional conference series in probability and statistics*, pages i–163. JSTOR, 1995.
- [47] S. Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137, 1982.
- [48] J. MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA, 1967.
- [49] D. G. Mixon, S. Villar, and R. Ward. Clustering subgaussian mixtures by semidefinite programming. *Information and Inference: A Journal of the IMA*, 6(4):389–415, 2017.
- [50] A. Y. Ng, M. I. Jordan, Y. Weiss, et al. On spectral clustering: Analysis and an algorithm. In *NIPS*, volume 14, pages 849–856, 2001.

- [51] X. Nguyen. Convergence of latent mixing measures in finite and infinite mixture models. *The Annals of Statistics*, 41(1):370–400, 2013.
- [52] J. Park and I. W. Sandberg. Universal approximation using radial-basis-function networks. *Neural computation*, 3(2):246–257, 1991.
- [53] J. Park and I. W. Sandberg. Approximation and radial-basis-function networks. *Neural computation*, 5(2):305–316, 1993.
- [54] K. R. Parthasarathy. *Probability measures on metric spaces*, volume 352. American Mathematical Soc., 1967.
- [55] A. Rinaldo and L. Wasserman. Generalized density clustering. *Annals of Statistics*, pages 2678–2722, 2010.
- [56] G. Ritter. *Robust cluster analysis and variable selection*. CRC Press, 2014.
- [57] K. Rohe, S. Chatterjee, and B. Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- [58] G. Schiebinger, M. J. Wainwright, B. Yu, et al. The geometry of kernelized spectral clustering. *The Annals of Statistics*, 43(2):819–846, 2015.
- [59] T. Shi, M. Belkin, and B. Yu. Data spectroscopy: Eigenspaces of convolution operators and clustering. *The Annals of Statistics*, pages 3960–3984, 2009.
- [60] B. Sriperumbudur and I. Steinwart. Consistency and rates for clustering with dbscan. In *Artificial Intelligence and Statistics*, pages 1090–1098, 2012.
- [61] H. Steinhaus. Sur la division des corp materiels en parties. *Bull. Acad. Polon. Sci*, 1(804):801, 1956.
- [62] I. Steinwart. Adaptive density level set clustering. In *COLT*, pages 703–738, 2011.
- [63] I. Steinwart. Fully adaptive density-based clustering. *Annals of Statistics*, 43(5):2132–2167, 2015.
- [64] T. Sweeting. On a converse to scheffé’s theorem. *The Annals of Statistics*, 14(3):1252–1256, 1986.
- [65] J. Tang, Z. Meng, X. Nguyen, Q. Mei, and M. Zhang. Understanding the limiting factors of topic modeling via posterior contraction analysis. In *International Conference on Machine Learning*, pages 190–198, 2014.
- [66] H. Teicher. Identifiability of mixtures. *The annals of Mathematical statistics*, 32(1):244–248, 1961.
- [67] H. Teicher. Identifiability of finite mixtures. *The annals of Mathematical statistics*, pages 1265–1269, 1963.
- [68] H. Teicher. Identifiability of mixtures of product measures. *The Annals of Mathematical Statistics*, 38(4):1300–1302, 1967.
- [69] P. Thomann, I. Steinwart, and N. Schmid. Towards an axiomatic approach to hierarchical clustering of measures. *Journal of Machine Learning Research*, 16:1949–2002, 2015.
- [70] D. M. Titterton, A. F. Smith, and U. E. Makov. *Statistical analysis of finite mixture distributions*. Wiley,, 1985.
- [71] A. Ultsch. Clustering with SOM: U*C. In *Proc. Workshop on Self-Organizing Maps, Paris, France*, pages 75–82, 2005. URL <https://www.uni-marburg.de/fb12/arbeitsgruppen/datenbionik/data>.
- [72] M. Wiper, D. R. Insua, and F. Ruggeri. Mixtures of gamma distributions with applications. *Journal of computational and graphical statistics*, 10(3):440–454, 2001.
- [73] J. H. Wolfe. Pattern clustering by multivariate mixture analysis. *Multivariate Behavioral Research*, 5(3):329–350, 1970.
- [74] S. J. Yakowitz and J. D. Spragins. On the identifiability of finite mixtures. *The Annals of Mathematical Statistics*, pages 209–214, 1968.
- [75] D. Yan, L. Huang, and M. I. Jordan. Fast approximate spectral clustering. In *Proceedings of the*

15th ACM SIGKDD international conference on Knowledge discovery and data mining, pages 907–916.
ACM, 2009.