

IDENTIFIABILITY OF NONPARAMETRIC MIXTURE MODELS AND BAYES OPTIMAL CLUSTERING

BY BRYON ARAGAM¹, CHEN DAN^{2,*}, ERIC P. XING^{2,**} AND PRADEEP RAVIKUMAR^{2,†}

¹*Booth School of Business, University of Chicago, bryon@chicagobooth.edu*

²*Machine Learning Department, School of Computer Science, Carnegie Mellon University, *cdan@andrew.cmu.edu;
**epxing@cs.cmu.edu; †pradeepr@cs.cmu.edu*

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 infinite Gaussian mixtures. In contrast to the recent literature, 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 [34]. The many approaches to formalize the learning of such a partition from data include mode clustering [21], density clustering [57, 62, 64, 65], spectral clustering [53, 60, 77], K -means [50, 51, 63], stochastic blockmodels [3, 27, 40, 59] and hierarchical clustering [19, 35, 72], 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 [8, 25]. 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 semiparametric models [12, 28, 75]. However, the extension of this methodology to general nonparametric settings has remained elusive. This is largely due to the extreme nonidentifiability of nonparametric mixture models, a problem which is well studied but for which existing results require strong assumptions [16, 42, 45, 71]. 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

Received November 2018; revised July 2019.

MSC2020 subject classifications. Primary 62G05, 62H30; secondary 62H12.

Key words and phrases. Mixture models, nonparametric statistics, identifiability, clustering, Bayes optimal partition.

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, 46, 52, 61]. A crucial motivation for this paper is in addressing this gap between theory and practice. This entails demonstrating that nonparametric mixture models might 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

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

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 5 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. Under parametric assumptions such as Gaussianity of the γ_k , it is well known that the representation (1) is unique and, hence, identifiable [9, 41, 70]. These results mostly follow from an early line of work on the general identification problem [2, 69, 70, 76].

Such parametric assumptions rarely hold in practice, however, and thus it is of interest to study *nonparametric* mixture models of the form (1), that is, 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 [24, 32, 33, 47, 71] which is particularly useful for statistical problems involving repeated measurements [14, 37]. 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 [16, 42], tail conditions [45] and translation invariance [29]. The identification problem in discrete mixture models is also a central problem in topic models which are popular in machine learning [5, 6, 67]. 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 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 and present a high-level geometric overview of our method. Section 2 covers the necessary background required

for our abstract framework. In Section 3 we introduce two important concepts, *regularity* and *clusterability*, that are crucial to our identifiability results along with several examples. In Section 4 we show how these concepts are sufficient to ensure identifiability of a nonparametric mixture model and consistency of a minimum distance estimator. In Section 5 we apply these results to the problem of clustering and prove a consistency theorem for this problem. Section 6 introduces a simple algorithm for nonparametric clustering along with some experiments, and Section 7 concludes the paper with some discussion and extensions. All proofs are deferred to the Supplementary Material [7].

Contributions. Our main results can be divided into three main theorems:

1. *Nonparametric identifiability* (Section 4.1). We formulate a general set of assumptions that guarantee a family of nonparametric mixtures will be identifiable (Theorem 4.1), based on two properties introduced in Section 3: *regularity* (Definition 3.1) and *clusterability* (Definition 3.3).

2. *Estimation and consistency* (Section 4.2). 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, and this procedure defines an estimator that consistently recovers the nonparametric clusters given i.i.d. observations from Γ (Theorem 4.3).

3. *Application to nonparametric clustering* (Section 5). We extend the notion of a *Bayes optimal partition* (Definition 5.1) to general nonparametric settings and prove a consistency theorem for recovering such partitions when they are identified (Theorem 5.2).

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. We conclude our study by applying these results to construct an intuitive algorithm for nonparametric clustering which is investigated in Section 6.

Overview. Before outlining the formal details, we present an intuitive geometric picture of our approach in Figure 1. This same example is developed in more detail in the experiments (see Section 6, Figure 4(iv)). 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 1(b));
- (2) Cluster these L Gaussian components into K groups such that each group roughly approximates some γ_k (Figure 1(c));
- (3) Use this clustering to define a new mixing measure (Figure 1(d));
- (4) Show that this new mixing measure converges to the true mixing measure Λ as $L \rightarrow \infty$.

Of course, this construction is not guaranteed to succeed for arbitrary mixing measures Λ which will be illustrated by the examples in Section 2.2. This is a surprisingly subtle problem and requires careful consideration of the various spaces involved. 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. 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.

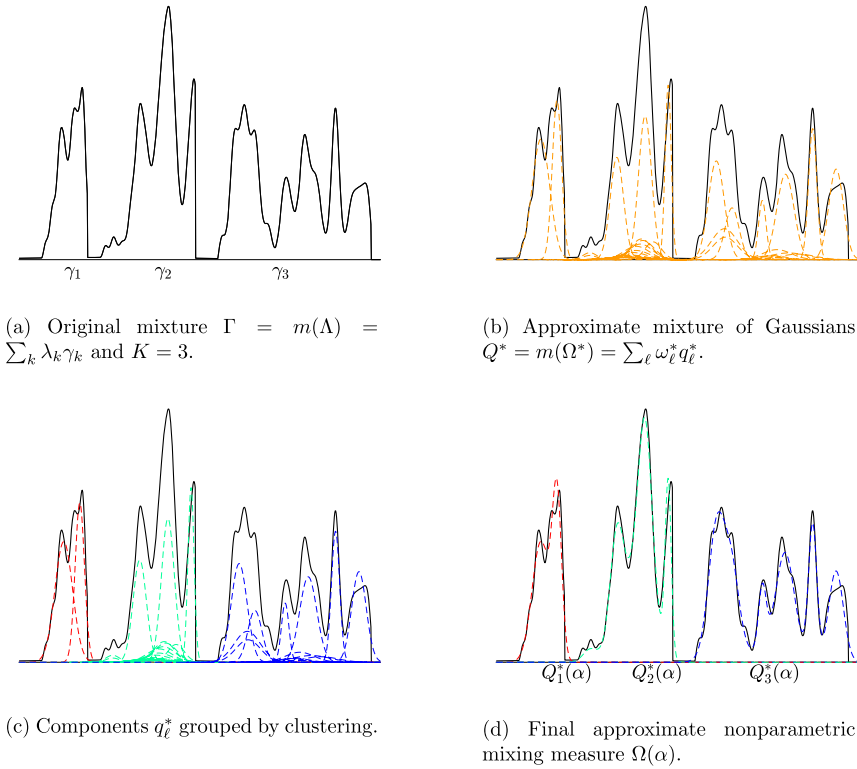


FIG. 1. Overview of the method.

2. Preliminaries. Our approach is general and built on the theory of abstract measures on metric spaces [55]. In this section we introduce this abstract setting, outline our notation and discuss the general problem of identifiability in mixture models. For a more thorough introduction to the general topic of mixture models in statistics, see Lindsay [49], Ritter [58], Titterington et al. [73].

2.1. *Nonparametric mixture models.* Let (X, d) be a metric space and $(\mathcal{P}(X), \rho)$ denote the space of regular Borel probability measures on X with finite r th moments ($r \geq 1$) metrized by a metric ρ . Common choices for ρ include the Hellinger and variational metrics, however, our results will apply to any metric on $\mathcal{P}(X)$. 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, that is, mixtures with a finite number of atoms. To this end, define for $s \in \{1, 2, \dots\}$

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

We treat $\mathcal{P}^2(X)$ as a metric space by endowing it with the L_r -Wasserstein metric W_r ($r \geq 1$). When $\Lambda \in \mathcal{P}_K^2(X)$ and $\Lambda' \in \mathcal{P}_{K'}^2(X)$, this is given by the optimal value of the transport problem

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

where the infimum is taken over all couplings σ , that is, 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 [54].

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

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

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 $\mathcal{L} \subset \mathcal{P}^2(X)$, define in analogy with $\mathcal{P}_s^2(X)$ the subsets of finite mixtures by

$$(4) \quad \mathcal{L}_s = \mathcal{L} \cap \mathcal{P}_s^2(X)$$

and

$$(5) \quad \mathcal{M}(\mathcal{L}) := \{m(\Lambda) : \Lambda \in \mathcal{L}\},$$

that is, the family of mixture distributions over X induced by \mathcal{L} which can be regarded as a formal representation of a statistical mixture model.

REMARK 2.1. 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. For the most part, we will work with finite mixtures, that is, $\mathcal{P}_0^2(X)$, a space which should be contrasted with the more complex space of infinite measures $\mathcal{P}^2(X)$, although some of the examples and proofs will invoke infinite mixtures.

REMARK 2.2. As a convention, we will use upper case letters for mixture distributions (e.g., Γ, \mathcal{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.

EXAMPLE 1 (Parametric mixtures). Let $\mathcal{Q} = \{q_\theta : \theta \in \Theta\}$ be a family of measures parametrized by θ . Then, any mixing measure whose support is contained in \mathcal{Q} defines a parametric mixture distribution. For example, 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 $\mathcal{M}(\mathfrak{G}_0)$ is the family of finite Gaussian mixtures. It is well known that $\mathcal{M}(\mathfrak{G}_0)$ is identifiable [69, 70]. Other examples include certain exponential family mixtures [9] and translation families [69] (i.e., $q_\theta(A) = \mu(A - \theta)$ for some known measure $\mu \in \mathcal{P}(\mathbb{R}^d)$).

EXAMPLE 2 (Sub-Gaussian mixtures). Let \mathcal{K} be the collection sub-Gaussian measures on \mathbb{R} , that is,

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

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 $\mathcal{M}(\mathfrak{K}_0)$ is the family of finite sub-Gaussian mixtures. Since the base measures \mathcal{K} do not belong to a parametric family, this is a nonparametric mixture model. Extensions to sub-Gaussian measures on \mathbb{R}^p are natural.

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 the previous examples, suppose we wish to impose additional constraints on the family of mixture distributions. For example, in Example 1 we might be interested in Gaussian mixtures whose means are contained within some set $A \subset \mathbb{R}^p$, whose covariance matrices are contained within another set $V \subset \text{PD}(p)$ and where $\text{PD}(p)$ is the set of $p \times p$ positive-definite matrices. Define $\mathcal{G}(A, V) := \{\mathcal{N}(a, v) : a \in A, v \in V\}$ and

$$(6) \quad \mathfrak{G}(A, V) := \{\Lambda \in \mathcal{P}^2(X) : \text{supp}(\Lambda) \subset \mathcal{G}(A, V)\}.$$

Then, $\mathcal{M}(\mathfrak{G}(A, V))$ is the desired family of mixture models. A special case of interest is $V = \{v\}$ for some fixed $v \in \text{PD}(p)$, which we denote by $\mathfrak{G}(A, v)$, also known as a convolutional (Gaussian) mixture model. Finite mixtures from these families are denoted by $\mathcal{M}(\mathfrak{G}_0(A, V))$ and $\mathcal{M}(\mathfrak{G}_0(A, v))$.

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, 43, 44].

2.2. Identifiability in mixture models. A mixture model $\mathcal{M}(\mathfrak{L})$ is identifiable if the map $m : \mathfrak{L} \rightarrow \mathcal{M}(\mathfrak{L})$ that sends $\Lambda \mapsto m(\Lambda)$ is injective. For an overview of this problem, see Hunter et al. [42] 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

$$(7) \quad \Gamma(B) = \sum_{k=1}^K \underbrace{\Gamma(A_k)}_{\tilde{\lambda}_k} \cdot \underbrace{\frac{\Gamma(B \cap A_k)}{\Gamma(A_k)}}_{\tilde{\gamma}_k} = \sum_{k=1}^K \tilde{\lambda}_k \tilde{\gamma}_k(B),$$

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.

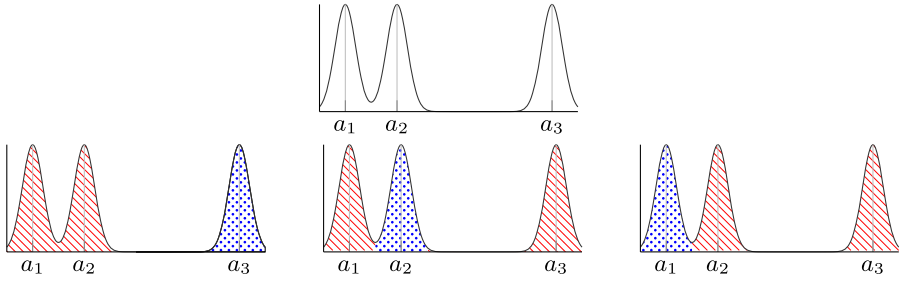


FIG. 2. (Top) Mixture of three Gaussians. (Bottom) Different representations of a mixture of Gaussians as a mixture of two sub-Gaussians. Different fill patterns and colours represent different assignments of mixture components.

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. Indeed, even if we assume the number of components K is known and the component means are well separated, $m(\Lambda)$ is nonidentifiable as a mixture of sub-Gaussians: 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 mixing measures in \mathcal{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 $\Lambda \in \mathcal{L}$. In this sense we shift the focus from the properties of the “local” mixture components to the “global” properties of the mixture itself.

3. Regularity and clusterability. Fix an integer K , and let $\mathcal{L} \subset \mathcal{P}_K^2(X)$ be a family of mixing measures. In particular, we assume that K —the number of nonparametric mixtures—is known; in Section 7 we discuss the case where K is unknown. In this section we study conditions that guarantee the injectivity of the embedding $m : \mathcal{L} \rightarrow \mathcal{M}(\mathcal{L})$, using the procedure described in Section 1. Throughout this section it will be helpful to keep Figure 1 in mind for intuition.

3.1. *Projections.* Let $\{\Omega_L\}_{L=1}^\infty$ be an indexed collection of families of mixing measures that satisfies the following:

- (A1) $\Omega_L \subset \mathcal{P}_L^2(X)$ for each L ;
- (A2) $\{\Omega_L\}$ is monotonic, that is, $\Omega_L \subset \Omega_{L+1}$;
- (A3) $\mathcal{M}(\Omega_L)$ is identifiable for each L .

The purpose of $\{\Omega_L\}$ is to approximate Γ with a sequence of mixture distributions of increasing complexity, as quantified by the maximum number of atoms L , which will be taken to be much larger than K . Although our results apply to generic collections satisfying Conditions (A1)–(A3), in the sequel we will consider the collection induced by a single subset $\Omega \subset \mathcal{P}^2(X)$ and defined by $\Omega_L = \Omega \cap \mathcal{P}_L^2(X)$ (cf. (4)). We make the following assumption on Ω :

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

If \mathfrak{Q} satisfies Condition (A), then $\{\mathfrak{Q}_L\}$ automatically satisfies Conditions (A1)–(A3). Examples of families that satisfy Condition (A) include exponential family mixture models under certain conditions [9], for example, Gaussian or Gamma mixtures [70].

Unless otherwise mentioned, we will assume \mathfrak{Q} satisfies Condition (A), with \mathfrak{Q}_L as defined therein. Define the usual ρ -projection by

$$(8) \quad T_L\Gamma = \{Q \in \mathcal{M}(\mathfrak{Q}_L) : \rho(Q, \Gamma) \leq \rho(P, \Gamma) \forall P \in \mathcal{M}(\mathfrak{Q}_L)\}.$$

As long as \mathfrak{Q} is compact, the projection $T_L\Gamma$ is nonempty. 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. With some abuse of notation, we will write $M_L\Gamma$ for $M_L(T_L\Gamma)$, that is,

$$(9) \quad M_L\Gamma = \{\Omega \in \mathfrak{Q}_L : m(\Omega) \in T_L\Gamma\}.$$

Thus, for any $Q^* \in T_L\Gamma$, we can unambiguously define

$$(10) \quad Q^* = \sum_{\ell=1}^L \omega_\ell^* q_\ell^* = m(\Omega^*) \quad \text{and} \quad \Omega^* = M_L(Q^*).$$

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

REMARK 3.1. We do not assume that $T_L\Gamma$ is unique, that is, there may be more than one projection. This is because $\mathcal{M}(\mathfrak{Q}_L)$ is a nonconvex set. We present our results in this setting, however, it may be simpler on a first reading to consider the special case where the projection is unique, that is, $T_L\Gamma = Q^*$ for each L . In this case, many of the definitions simplify; consider, for example, Definition 3.1 and (18) in the sequel.

REMARK 3.2. 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. In Section 4.2 we discuss how to choose L given the sample size n ; see Corollary 4.4.

3.2. *Assignment functions.* Any projection $Q^* = m(\Omega^*) = \sum_{\ell=1}^L \omega_\ell^* q_\ell^*$, as defined in (10), 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 1. 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

$$(11) \quad \varpi_k(\alpha) := \sum_{\ell \in \alpha^{-1}(k)} \omega_\ell, \quad k = 1, \dots, K.$$

Denote the point mass concentrated at q_ℓ by δ_{q_ℓ} , and define

$$(12) \quad \Omega_k(\alpha) := \frac{1}{\varpi_k(\alpha)} \sum_{\ell \in \alpha^{-1}(k)} \omega_\ell \delta_{q_\ell}, \quad Q_k(\alpha) := m(\Omega_k(\alpha)).$$

These quantities define a single, aggregate K -mixture by

$$(13) \quad \Omega(\alpha) := \sum_{k=1}^K \varpi_k(\alpha) \delta_{Q_k(\alpha)}, \quad Q(\alpha) := m(\Omega(\alpha)) = \sum_{k=1}^K \varpi_k(\alpha) Q_k(\alpha).$$

Since $Q_k(\alpha) \in \mathcal{M}(\Omega_0)$, $\Omega(\alpha)$ is an atomic mixing measure whose atoms come from $\mathcal{M}(\Omega_0)$. 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.

3.3. 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 (11) and $Q_k^*(\alpha)$ and $\Omega_k^*(\alpha)$ as in (12). We'd like $Q_k^*(\alpha)$ to approximate γ_k , but this is certainly not guaranteed for any α . A key step in our construction is to find such an assignment. Before finding such an assignment, however, we must first ask whether or not such an assignment *exists*. The following notion of regularity encodes this assumption:

DEFINITION 3.1 (Regularity). Suppose $\Lambda \in \mathcal{P}_K^2(X)$ and $\Gamma = m(\Lambda)$. The mixing measure Λ is called Ω -regular if:

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

$$\lim_{L \rightarrow \infty} Q_k^*(\alpha_L) = \gamma_k \quad \text{and} \quad \lim_{L \rightarrow \infty} \varpi_k^*(\alpha_L) = \lambda_k \quad \forall k, \forall Q^* \in T_L \Gamma.$$

When Λ is Ω -regular, we will also call $m(\Lambda)$ Ω -regular.

DEFINITION 3.2 (Regular assignment sequences). Given a regular mixing measure Λ , denote a set of all assignment sequences $\{\alpha_L\}$ such that Definition 3.1(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.

Whether or not a mixing measure is regular depends on both Ω and ρ , although the dependence on ρ will typically be suppressed. When we wish to emphasize this dependence, we will say Λ is Ω -regular under ρ . Clearly, Λ is Ω -regular under the Hellinger metric if and only if it is Ω -regular under the variational metric.

The following examples construct several families of regular mixing measures, as well as an example where regularity fails. Proofs of these claims can be found in Appendix B of the Supplementary Material.

EXAMPLE 6 (Disjoint components). Let $X = \mathbb{R}$. Assume each γ_k has a density f_k with respect to some dominating measure ζ , and there exist disjoint intervals $E_k := [b_k, c_k] \subset \mathbb{R}$ such that $\text{supp}(f_k) \subseteq E_k$. Then, the resulting mixing measure Λ is \mathfrak{G} -regular under both the Hellinger and variational metrics (Lemma B.1 in the Supplementary Material). Furthermore, this example can be generalized to measures on \mathbb{R}^d whose supports are contained in disjoint convex sets.

EXAMPLE 7 (Mixtures of finite mixtures). Fix $\Omega \subset \mathcal{P}^2(X)$ satisfying Condition (A), and assume $\gamma_k = m(P_k)$ for each k , where $P_k \in \Omega_0$, that is, P_k is a finite mixture model, but note that no upper bound is imposed on the number of components in each P_k . Define $B_k := \text{supp}(P_k)$, and assume that B_1, \dots, B_K are disjoint. Then, Λ is Ω -regular under any metric ρ (Lemma B.2 in the Supplementary Material).

EXAMPLE 8 (Mixtures of infinite mixtures). In fact, the previous example can be generalized quite substantially. Let $\Omega \subset \mathcal{P}^2(X)$ be compact and identifiable. Assume that $\gamma_k = m(P_k)$ for each k , where $P_k \in \Omega$. For example, γ_k could be a potentially infinite convolutional mixture (see [54] for details), such as an infinite mixture of Gaussians with $P_k \in \mathfrak{G}(A, v)$ (Example 3). Define $B_k := \text{supp}(P_k)$, and assume that (a) B_1, \dots, B_K are disjoint, compact sets and (b) each B_k is a P -continuity set where $P := \sum_k \lambda_k P_k$. Then, Λ is Ω -regular in both the Hellinger and variational metrics (Lemma B.4 in the Supplementary Material).

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 $0 < \beta_1 < \beta_2 < 1$, $\Gamma = \frac{1}{2}\gamma_1 + \frac{1}{2}\gamma_2$, $\gamma_1 \propto (1 - \beta_1 - \beta_2)g_+ + \frac{\beta_1}{2}G$ and $\gamma_2 \propto \beta_2 g_- + \frac{\beta_1}{2}G$. In this example, $K = 2$. If $\Omega_L = \mathfrak{G}_L$, then for any $L > 3$, $Q^* = \Gamma$, and there is no way to cluster the three components into two 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 .

We conclude by pointing out that, in addition to the concrete examples discussed above, in general the set of regular mixing measures is quite large:

LEMMA 3.1. *Let ρ_{TV} be the variational distance on $\mathcal{P}(\mathbb{R}^p)$, and let W_r be the induced Wasserstein metric on $\mathcal{P}_K^2(\mathbb{R}^p)$. Then, for any $\Lambda \in \mathcal{P}_K^2(\mathbb{R}^p)$ and $\varepsilon > 0$, there exists a \mathfrak{G} -regular mixing measure $\Lambda' \in \mathcal{P}_K^2(\mathbb{R}^p)$ such that $W_r(\Lambda', \Lambda) < \varepsilon$. In particular, the set of \mathfrak{G} -regular mixing measures is dense in $\mathcal{P}_K^2(\mathbb{R}^p)$.*

In fact, the proof is constructive. The family defined in Example 7 with $\Omega = \mathfrak{G}$ is dense in $\mathcal{P}_K^2(\mathbb{R}^p)$.

3.4. *Clusterable families.* If a mixing measure Λ is Ω -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 . We have not said anything yet about *how* one might find such an assignment, only that it exists. The following condition asserts that regular assignments can be determined from the projections Q_L^* :

DEFINITION 3.3 (Clusterable family). A family of mixing measures $\mathfrak{L} \subset \mathcal{P}^2(X)$ is called a Ω -clusterable family, or just a clusterable family, if:

- (a) Λ is Ω -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.

As with regularity, clusterability depends on both Ω and ρ . When we wish to emphasize this dependence, we will say Λ is Ω -clusterable under ρ . 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 . More precisely, Definition 3.3(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 3.2).

3.5. *Separation and clusterability.* In this section we construct an explicit cluster function χ_L via single-linkage clustering.

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

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

where $\text{conv}(\cdot)$ is the convex hull in $\mathcal{P}(X)$. Recalling (13), define for any $\alpha \in \mathbb{A}_{L \rightarrow K}$

$$(14) \quad \eta(\Omega(\alpha)) := \sup_k \Delta(\Omega_k(\alpha)) + \sup_k \rho(\gamma_k, Q_k(\alpha)).$$

We will be interested in the special case $\Omega = \Omega^*$: $\Delta(\Omega_k^*(\alpha))$ quantifies how “compact” the mixture component $Q_k^*(\alpha)$ is and $\eta(\Omega^*(\alpha))$ is a measure of separation between the mixture components γ_k . Finally, define the ρ -distance matrix by

$$(15) \quad D(\Omega) = (\rho(q_i, q_j))_{i,j=1}^L.$$

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 3.4 (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 of the order $\eta(\Omega^*(\alpha))$ (cf. (14)) is sufficient to define a cluster function:

PROPOSITION 3.2. Let $\Lambda \in \mathcal{P}_K^2(X)$. Let $Q^* \in T_L \Gamma$ be a ρ -projection of Γ for some $L \geq K$. Then, for any $\alpha \in \mathbb{A}_{L \rightarrow K}$ such that Λ is $4\eta(\Omega^*(\alpha))$ -separated:

$$(16) \quad \alpha(i) = \alpha(j) \iff \rho(q_i^*, q_j^*) \leq \eta(\Omega^*(\alpha)),$$

$$(17) \quad \alpha(i) \neq \alpha(j) \iff \rho(q_i^*, q_j^*) \geq 2\eta(\Omega^*(\alpha)).$$

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 $\eta(\Omega^*(\alpha))$.

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

$$(18) \quad \eta(\Lambda) := \limsup_{L \rightarrow \infty} \sup_{\Omega^* \in M_L \Gamma} \sup_{\{\alpha_L\} \in \mathbb{A}_K^\infty(\Lambda)} \eta(\Omega^*(\alpha_L)).$$

As a consequence of regularity, the second term in (14) tends to zero as $L \rightarrow \infty$, so that $\eta(\Lambda)$ can be interpreted as a measure of the asymptotic diameter of the approximating mixtures $Q_k^*(\alpha_L)$. For example, when the ρ -projection Q^* is unique (Remark 3.1), the definition in (18) simplifies to $\eta(\Lambda) = \limsup_L \sup_k \Delta(\Omega_k^*(\alpha_L))$. The following corollary, which is an immediate consequence of Proposition 3.2, shows that control over $\eta(\Lambda)$ is sufficient for \mathfrak{L} to be clusterable:

COROLLARY 3.3. 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.

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

$$(19) \quad \inf_{i \neq j} \rho(\gamma_i, \gamma_j) > (4 + \xi)\eta(\Lambda).$$

In the limit $L \rightarrow \infty$, the nonparametric components γ_k must be separated by a gap proportional to the ρ -diameters of the approximating mixtures $Q_k^*(\alpha_L)$. 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). Crucially, however, we make no assumptions on the shape of the mixture components.

EXAMPLE 10 (Example of separation). Take $X = \mathbb{R}$, and let $\Omega = \mathfrak{G}(A, v)$ be a family of convolutional mixtures of Gaussians (Example 3). In Example 8 we claimed that as long as the P_k have disjoint supports, this family is $\mathfrak{G}(A, v)$ -regular. To determine when a mixing measure Λ is $\mathfrak{G}(A, v)$ -clusterable, it suffices to check (19). For this, we bound the ρ -diameters $\sup_k \Delta(\Omega_k^*(\alpha_L))$ for large L . If $q_\ell \sim \mathcal{N}(a_\ell, v^2)$ and $q_{\ell'} \sim \mathcal{N}(a_{\ell'}, v^2)$ are both in $\text{supp}(P_k)$, then it is easy to check that as long as

$$|a_\ell - a_{\ell'}| \leq \sqrt{8v^2 \log\left(1 + \frac{\rho_*}{4 - \rho_*}\right)}, \quad \rho_* := \inf_{i \neq j} \rho(\gamma_i, \gamma_j),$$

the separation condition (19) holds.

The separation condition (19) 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 3.2 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 3.2. 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.

4. Identifiability and estimation. We now turn our attention to the problem of identifying and learning a mixing measure Λ from data.

4.1. Identifiability of nonparametric mixtures. According to the next theorem, clusterability is sufficient to identify a nonparametric mixture model.

THEOREM 4.1. *If \mathfrak{L} is a Ω -clusterable family, then the mixture model $\mathcal{M}(\mathfrak{L})$ is identifiable.*

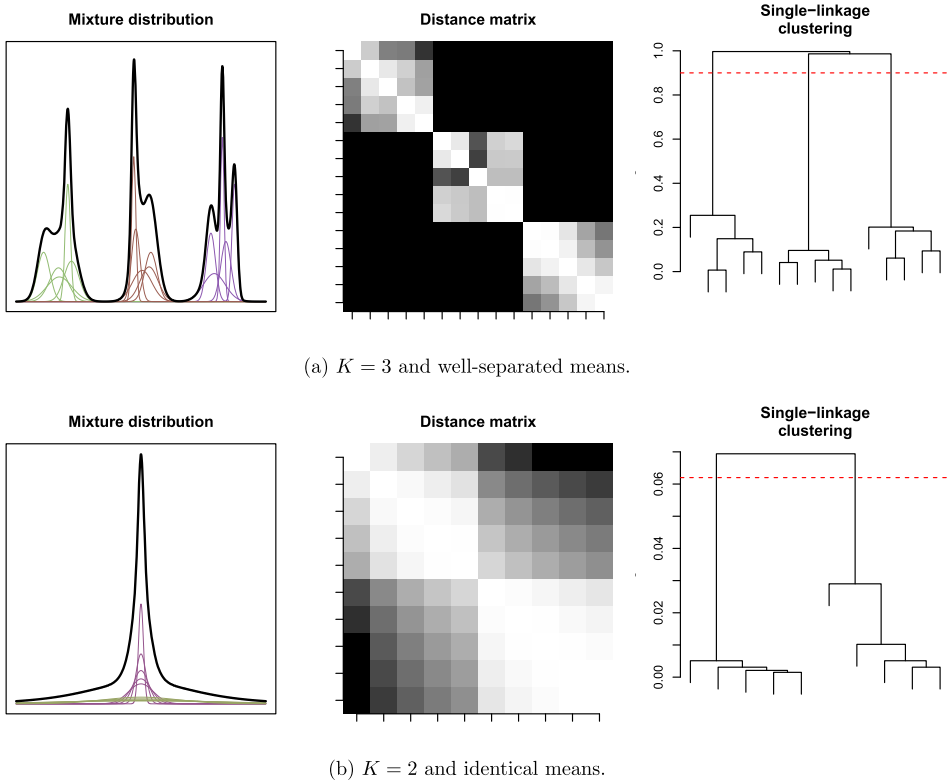


FIG. 3. Illustrating Theorem 4.1 with Example 7. (left) The original mixture distribution (thick black line) is a mixture of finite Gaussian mixtures. Each Gaussian component is coloured according to its membership in different P_k . (middle) The true distance matrix D^* . (right) Results of single-linkage clustering on D^* , cut to find the correct number of clusters.

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.

Two examples of Theorem 4.1 are illustrated in Figure 3. When the means are well separated, as in Figure 3(a), it is easy to see how single-linkage clustering is able to discover a correct assignment. Since ρ -separation is a weaker criterion than mean separation, however, Theorem 4.1 does not require that the mixture distributions in $\mathcal{M}(\mathcal{L})$ 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 in Figure 3(b). This suggests that identifiability in mixture models is more general than what is needed in typical clustering applications, where a model such as Figure 3(b) would 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.

4.2. Estimation of clusterable mixtures. We now discuss how to estimate Λ from data $Z^{(1)}, \dots, Z^{(n)} \stackrel{\text{iid}}{\sim} \Gamma$. Throughout this section we assume that $\Omega^* = \Omega_L^* \in M_L \Gamma$ is arbitrary.

For each $L \geq K$, let $\widehat{\Omega}_{L,n} = \widehat{\Omega}_L(Z^{(1)}, \dots, Z^{(n)})$ be a W_r -consistent estimator of Ω_L^* , where we have written $\widehat{\Omega}_{L,n}$ and Ω_L^* to emphasize the dependence on L and n . That is, $\{\widehat{\Omega}_{L,n}\}$ is a sequence of estimators, and, for each L , $\lim_{n \rightarrow \infty} W_r(\widehat{\Omega}_{L,n}, \Omega_L^*) = 0$. For example, $\widehat{\Omega}_{L,n}$ could be the minimum Hellinger distance estimator (MHDE) from Beran [11] (see Appendix D in the Supplementary Material for details). Since L is a known quantity, the cor-

responding estimation problems are always well specified, that is, both $\widehat{\Omega}_{L,n}$ and Ω_L^* have the same, known number of components. In the sequel we will omit the dependence of $\Omega^* = \Omega_L^*$ on L and $\widehat{\Omega} = \widehat{\Omega}_{L,n}$ on L and n for brevity. Write

$$(20) \quad \widehat{Q} := m(\widehat{\Omega}) = \sum_{\ell=1}^L \widehat{\omega}_\ell \widehat{q}_\ell.$$

Without loss of generality, assume that the atoms are rearranged so that $\sup_\ell \rho(\widehat{q}_\ell, q_\ell^*) \rightarrow 0$ (see Lemma C.3 in the Supplementary Material).

PROPOSITION 4.2. *Let $\Lambda \in \mathcal{P}_K^2(X)$. Let $Q^* \in T_L \Gamma$ be a ρ -projection of Γ for some $L \geq K$. Suppose further that $L, \alpha \in \mathbb{A}_{L \rightarrow K}$, and n satisfy*

$$(21) \quad 3 \sup_\ell \rho(\widehat{q}_\ell, q_\ell^*) - 2 \sup_k \rho(Q_k^*(\alpha), \gamma_k) < \sup_k \Delta(\Omega_k^*(\alpha)).$$

Define

$$\widehat{\eta} := 2 \sup_\ell \rho(\widehat{q}_\ell, q_\ell^*) + \sup_k \Delta(\Omega_k^*(\alpha)).$$

If Λ is $4\eta(\Omega^*(\alpha))$ -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 $\widehat{D} = D(\widehat{\Omega})$.

Proposition 4.2 is a finite sample result that holds as long as L and n satisfy (21) which is guaranteed as long as Λ is Ω -regular (i.e., since in this case the left side tends to zero).

For each L and n , let $\widehat{\alpha} = \widehat{\alpha}_{L,n} \in \mathbb{A}_{L \rightarrow K}$ denote the assignment map defined in Proposition 4.2. With this notation, another way to phrase this result is that under (21), we have $\widehat{\alpha} = \alpha$. In other words, single-linkage clustering of \widehat{D} yields the same clusters as the assignment α . This suggests we use $\widehat{\Omega}(\widehat{\alpha})$ as an estimator of Λ . More precisely:

1. Choose $L \geq K$ sufficiently large;
2. Estimate $\widehat{\Omega} = \widehat{\Omega}(Z^{(1)}, \dots, Z^{(n)})$;
3. Define $\widehat{\alpha} = \widehat{\alpha}(Z^{(1)}, \dots, Z^{(n)})$ by single-linkage clustering on \widehat{D} ;
4. Return $\widehat{\Omega}(\widehat{\alpha})$.

In order for this to be an estimator, we must have a precise rule for selecting $L = L_n$; see Corollary 4.4 below and its discussion for details.

The following theorem provides conditions under which $\widehat{\Omega}(\widehat{\alpha})$ consistently estimates Λ :

THEOREM 4.3. *Suppose Λ is a regular mixing measure such that Λ is $(4 + \xi)\eta(\Lambda)$ -separated for some $\xi > 0$. Then,*

$$(22) \quad \lim_{L \rightarrow \infty} \lim_{n \rightarrow \infty} W_r(\widehat{\Omega}(\widehat{\alpha}), \Lambda) = 0.$$

In particular, (22) implies that

$$\lim_{L \rightarrow \infty} \lim_{n \rightarrow \infty} \rho(\widehat{Q}_k(\widehat{\alpha}), \gamma_k) = 0 \quad \text{and} \quad \lim_{L \rightarrow \infty} \lim_{n \rightarrow \infty} |\varpi_k^*(\widehat{\alpha}) - \lambda_k| = 0.$$

Thus, we have a Wasserstein consistent estimate of Λ and ρ -consistent estimates of the component measures γ_k . As stated, Theorem 4.3 has an important drawback: Without a rule for choosing $L = L_n$ as a function of the sample size n , $\widehat{\Omega}(\widehat{\alpha})$ is not a proper estimator. This is the cost of abstraction that allows us to state such a theorem for general metric spaces and probability measures. Fortunately, in special cases we can make the dependence on n explicit. Recall the convolutional mixture model described in Example 3. We have already shown that this family is both regular (Example 8) and clusterable (Example 10). Combining these results with a rule for choosing $L = L_n$, the following corollary provides a practical setting in which all of the assumptions laid out in Theorem 4.3 are satisfied:

COROLLARY 4.4. Let $\Lambda = \sum_{k=1}^K \lambda_k P_k$ be as in Example 8 with $P_k \in \mathfrak{G}(A, v)$ for each k . Define $B_k := \text{supp}(P_k)$, and assume that (a) B_1, \dots, B_K are disjoint, compact sets and (b) Each B_k is a P -continuity set where $P := \sum_k \lambda_k P_k$. Define $\rho_* := \inf_{i \neq j} \rho(m(P_i), m(P_j))$, and assume further that

$$\sup_{q_\ell, q_{\ell'} \in \text{supp}(P_k)} |\mathbb{E}q_\ell - \mathbb{E}q_{\ell'}| \leq \sqrt{8v^2 \log\left(1 + \frac{\rho_*}{4 - \rho_*}\right)} \quad \text{for all } k.$$

Then, taking $L_n \asymp n^{2/3} / \log^{1/3} n$, we have

$$(23) \quad \lim_{n \rightarrow \infty} W_r(\widehat{\Omega}(\widehat{\alpha}_{L_n, n}), \Lambda) = 0.$$

The proof follows immediately from Theorem 2 in Nguyen [54] and the results (e.g., Theorem 5) in Genovese and Wasserman [30].

Finally, in applications it will often be useful to strengthen ρ -convergence to uniform convergence of the densities (assuming they exist). When the families Ω_L are equicontinuous, this is guaranteed by Theorem 1 of Sweeting [66]. We store this corollary away here for future use:

COROLLARY 4.5. Let $\widehat{G}_k(\widehat{\alpha})$ be the density of $\widehat{Q}_k(\widehat{\alpha})$ and f_k be the density of γ_k . If the families Ω_L are equicontinuous for all L and $\widehat{Q}_k(\widehat{\alpha})$ converges weakly to γ_k , then $\lim_{L \rightarrow \infty} \lim_{n \rightarrow \infty} \widehat{G}_k(\widehat{\alpha}) = f_k$, where the limits are understood both pointwise and uniformly over compact subsets of X .

The assumption that $\widehat{Q}_k(\widehat{\alpha})$ converges weakly to γ_k restricts the choice of ρ , although it allows most reasonable metrics including Hellinger, variational and Wasserstein, for example. Moreover, even weaker assumptions than equicontinuity are possible [23].

5. Bayes optimal clustering. As an application of the theory developed in Sections 3 and 4, we extend model-based clustering [12, 28] 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

$$(24) \quad F(z) = \sum_{k=1}^K \lambda_k f_k(z; a_k, v_k); \quad z \in \mathbb{R}^d$$

is a GMM. The model (24) 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.

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., [53, 74]), there is significant interest in such an extension. In the remainder of this section, we will apply our framework to this problem. First, we discuss identifiability issues with the concept of a Bayes optimal partition (Section 5.1). Then, we provide conditions under which a Bayes optimal partition can be learned from data (Section 5.2).

5.1. *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 Γ is a finite mixture, we can write

$$(25) \quad F = \int f_\gamma d\Lambda(\gamma) = \sum_{k=1}^K \lambda_k f_k.$$

For any $\Lambda \in \mathcal{P}_K^2(X)$, define the usual Bayes classifier (e.g., [26]):

$$(26) \quad c_\Lambda(x) := \arg \max_{k \in [K]} \lambda_k f_k(x).$$

The classifier 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

$$(27) \quad E_0 := \bigcup_{i \neq j} \{x \in X : \lambda_i f_i(x) = \lambda_j f_j(x)\}.$$

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 $\bigcup_k A_k = X$. We denote the space of all partitions of X by $\Pi(X)$.

The following definition is standard (e.g., [18, 28]):

DEFINITION 5.1 (Bayes optimal partition). Define an equivalence relation on X_0 by declaring

$$(28) \quad x \sim y \iff c_\Lambda(x) = c_\Lambda(y).$$

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 5.1. Although the function c_Λ is only unique up to a permutation, the partition defined by (28) 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 11 (Nonidentifiability 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, Theorem 4.1 immediately implies the following:

COROLLARY 5.1. *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 $\mathcal{M}(\mathfrak{L})$ 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 , (26) is one such function. Thus, the Bayes optimal partition π_Γ can be interpreted as a valid K -clustering.

5.2. *Learning partitions from data.* Write $\Gamma = m(\Lambda)$, and assume that Λ is identifiable from Γ . Suppose we are given i.i.d. samples $Z^{(1)}, \dots, Z^{(n)} \stackrel{\text{iid}}{\sim} \Gamma$ 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 4.3 guarantees that we can learn a cluster assignment $\widehat{\alpha}$ such that $\widehat{\Omega}(\widehat{\alpha})$ consistently estimates Λ ;
3. Use $\pi(\widehat{\Omega}(\widehat{\alpha}))$ to approximate $\pi_\Lambda = \pi_\Gamma$.

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

Instead of working directly with the partitions $\pi(\widehat{\Omega}(\widehat{\alpha}))$, we will work with the Bayes classifier (26). Write \widehat{g}_ℓ and \widehat{G} for the densities of \widehat{q}_ℓ and \widehat{Q} , respectively, and

$$(29) \quad \widehat{G}_k(\widehat{\alpha}) := \frac{1}{\widehat{\omega}_k} \sum_{\ell \in \widehat{\alpha}^{-1}(k)} \widehat{\omega}_\ell \widehat{g}_\ell, \quad \widehat{\omega}_k(\widehat{\alpha}) := \sum_{\ell \in \widehat{\alpha}^{-1}(k)} \widehat{\omega}_\ell.$$

Then, $\widehat{G}_k(\widehat{\alpha})$ is the density of $\widehat{Q}_k(\widehat{\alpha})$, where here and above we have suppressed the dependence on $\widehat{\alpha}$. Now, define the estimated classifier (cf. (26))

$$(30) \quad \widehat{c}(x) := c_{\widehat{\Omega}(\widehat{\alpha})}(x) = \arg \max_{k \in [K]} \widehat{\omega}_k[\widehat{G}_k(\widehat{\alpha})](x).$$

By considering classification functions, as opposed to the partitions themselves, we may consider ordinary convergence of the function \widehat{c} 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 $\widehat{A}_{L,n,k} := \widehat{c}^{-1}(k)$ by controlling $A_k \Delta \widehat{A}_{L,n,k}$, where $A \Delta B = (A - B) \cup (B - A)$ 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

$$(31) \quad E_0(t) := \bigcup_{i \neq j} \{x \in X : |\lambda_i f_i(x) - \lambda_j f_j(x)| \leq t\}, \quad t > 0.$$

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 5.2. *Assume that $\lim_{L \rightarrow \infty} \lim_{n \rightarrow \infty} \widehat{G}_k(\widehat{\alpha}) = f_k$ uniformly on X , and ν is any measure on X . Then, there exists a sequence $t_{L,n} \rightarrow 0$ such that $\widehat{c}(x) = c_\Lambda(x)$ for all $x \in X - E_0(t_{L,n})$ and*

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

As in Corollary 4.4, under the same assumptions we may take $L = L_n \asymp n^{2/3} / \log^{1/3} n$ in Theorem 5.2 when $\mathfrak{Q} = \mathfrak{G}_0(A, \nu)$.

The uniform convergence assumption in Theorem 5.2 may seem strong, however, recall Corollary 4.5 which guarantees uniform convergence whenever \mathfrak{Q}_L is equicontinuous. For example, recalling Examples 1 and 3, it is straightforward to show the following:

COROLLARY 5.3. *Suppose $X \subset \mathbb{R}^d$, \mathfrak{Q} is a compact subset of \mathfrak{G} , and ν is any measure on X . If Λ is \mathfrak{Q} -clusterable measure under the Hellinger or variational metric, then there exists a sequence $t_{L,n} \rightarrow 0$ such that $\hat{c}(x) = c_\Lambda(x)$ for all $x \in X - E_0(t_{L,n})$ and*

$$(33) \quad \nu \left(\bigcup_{k=1}^K A_k \Delta \hat{A}_{L,n,k} \right) \leq \nu(E_0(t_{L,n})) \rightarrow \nu(E_0).$$

We can interpret Theorem 5.2 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 $\nu(E_0(t_{L,n}))$ is small), the difference between the true Bayes optimal partition and the estimated partition becomes negligible. In fact, it follows trivially from Theorem 5.2 that $\hat{c} \rightarrow c_\Lambda$ uniformly on $X - E_0(t)$ for any fixed $t > 0$. Thus, Theorem 5.2 gives rigorous justification to the approximation heuristic outlined above and establishes precise conditions under which *nonparametric* clusterings can be learned from data.

REMARK 5.2. The sequence $t_{L,n}$ is essentially the rate of convergence of $\hat{G}_k \rightarrow \gamma_k$. It is an interesting question to quantify this convergence rate more precisely, which we have left to future work.

6. 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 5, we assume we have i.i.d. samples $Z^{(1)}, \dots, Z^{(n)} \stackrel{iid}{\sim} \Gamma = m(\Lambda)$. Given these samples, we propose the following meta-algorithm:

1. Estimate an overfitted GMM \hat{Q} with $L \gg K$ components;
2. Define an estimated assignment function $\hat{\alpha}$ by using single-linkage clustering to group the components of \hat{Q} together;
3. Use this clustering to define K mixture components $\hat{Q}_k(\hat{\alpha})$;
4. Define a partition on X by using Bayes' rule, for example, (29)–(30).

Figure 3 has already illustrated 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 \hat{Q} in step 1, although clearly any algorithm for learning a GMM can be used in this step. The details of these experiments can be found in Appendix E of the Supplementary Material.

We call the resulting algorithm NPMIX (for *Nonparametric MIX*tured 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.

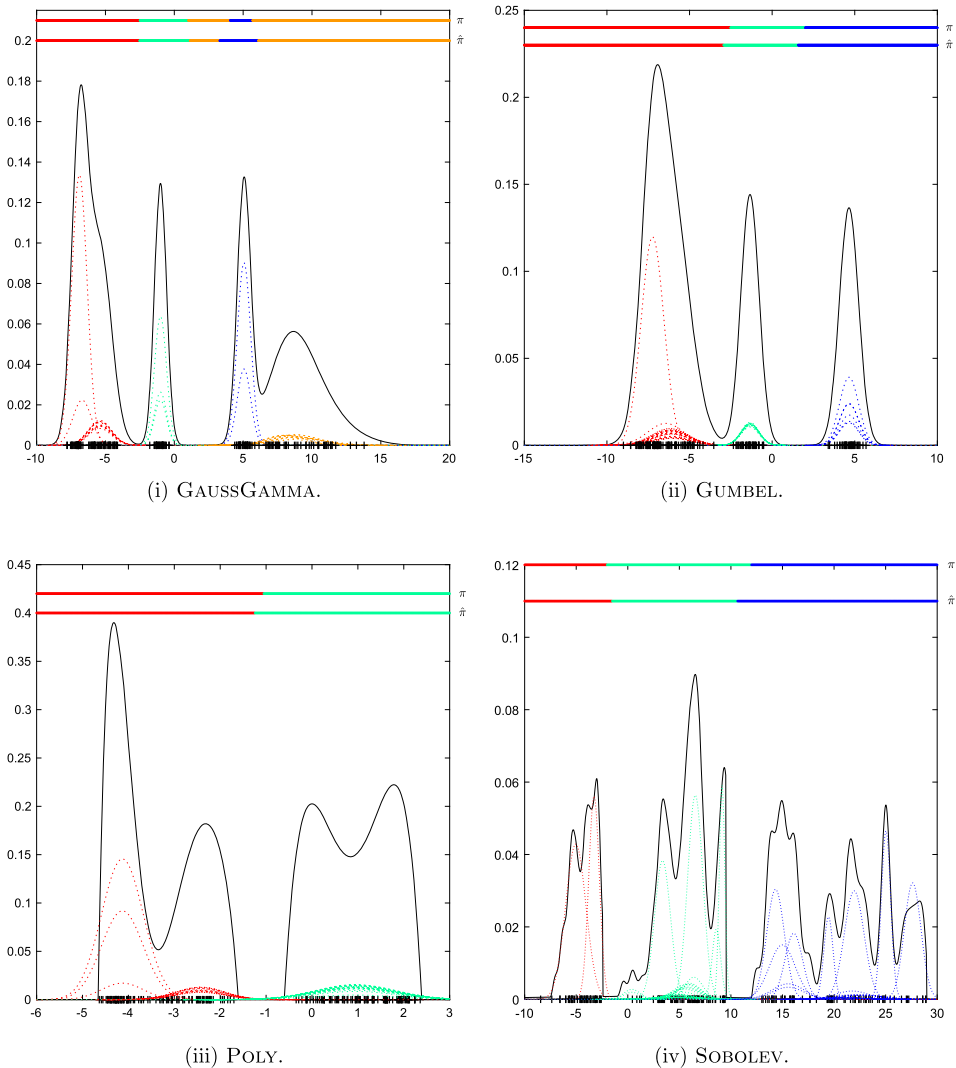


FIG. 4. Examples (i)–(iv) of one-dimensional mixture models. The original mixture density is depicted as a solid black line, with the overfitted Gaussian mixture components as dotted lines and coloured according to the cluster to which they are assigned. The true Bayes optimal partition π and the estimated partition $\hat{\pi}$ are depicted by the horizontal lines at the top, and the raw data are plotted on the x -axis for reference.

(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 nonoverlapping supports.

(iv) SOBOLEV ($K = 3$): A mixture of three random nonparametric densities, generated from random expansions of an orthogonal basis for the Sobolev space $H^1(\mathbb{R})$. This is the same example used in Figure 1.

The results are shown in Figure 4. These examples illustrate the basic idea behind the algorithm. Given samples, overfitted mixture components (depicted by dotted lines in Figure 4) 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, cutting the cluster tree to produce K components provides sensible and meaningful approximations to the true partitions.




FIG. 5. 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.

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 7). The GMM has 143 components that are clustered into six groups based on the original TARGET dataset from [74].

Visualizations of the results for our method are shown in Figures 5, 6 and 7. 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 rerunning 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




FIG. 6. 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.




FIG. 7. 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.

model and then used each clustering algorithm to classify each sample. To assess cluster accuracy, we computed the adjusted 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

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

	Mean ARI	Median ARI	st. dev.
MOONS (UNBALANCED)			
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)			
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			
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

and GMM improved significantly, with spectral clustering performing quite well on average. All four algorithms were still outperformed by 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.

7. 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-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 (19) 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 (19) 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 (18) (cf. (14)). 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 .

As pointed out by a reviewer, there is a connection between the NPMIX algorithm introduced in Section 6 and kernel density estimation (KDE). Indeed, by choosing $L = n$, the overfitted mixture model learned in step 1 is similar to a kernel density estimate with a Gaussian kernel, although not exactly the same since KDE fixes the weights, centers and bandwidth of each kernel unless more sophisticated adaptive bandwidth selection strategies are used. By contrast, a GMM allows these parameters to be learned from the data. Thus, in the limiting case $L = n$, NPMIX is similar to single-linkage clustering applied to a new metric defined via the Wasserstein distance between the n kernels, where this new metric depends crucially on the choice of bandwidth. An important difference in practice is that by taking $L < n$, the NPMIX algorithm denoises the data in the first step, making it less sensitive to outliers. For example, Priebe [56] points out that approximately $L = 30$ Gaussian components suffice to approximate a log-normal density with $n = 10,000$ samples; see also Corollary 4.4. Exploring this connection more deeply is an interesting direction for future work.

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 [36], 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, 38, 39, 54]. In the general case, this is also related to problems in agnostic learning [48]. 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 , for example, as in [30, 31].

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 3.2 shows that separability is sufficient for single-linkage clustering. We have already noted that the separation constant $4\eta(\Lambda)$ 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.

Acknowledgments. The authors would like to thank Larry Wasserman and Sivaraman Balakrishnan for helpful discussions, as well as the anonymous reviewers and Associate Editor for their numerous suggestions and improvements to the paper.

Chen Dan and Pradeep Ravikumar acknowledge the support of NSF via IIS-1909816 and OAC-1934584.

The material of Bryon Aragam and Eric Xing is based upon work supported by NIH R01GM114311 and P30DA035778. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Institutes of Health.

SUPPLEMENTARY MATERIAL

Supplement to “Identifiability of nonparametric mixture models and Bayes optimal clustering” (DOI: [10.1214/19-AOS1887SUPP](https://doi.org/10.1214/19-AOS1887SUPP); .pdf). This supplement contains proofs of all the main results along with various technical results and experiment details.

REFERENCES

- [1] ACHLIOPTAS, D. and MCSHERRY, F. (2005). On spectral learning of mixtures of distributions. In *International Conference on Computational Learning Theory* 458–469. Springer, Berlin.
- [2] AHMAD, K. E. and AL-HUSSAINI, E. K. (1982). Remarks on the nonidentifiability of mixtures of distributions. *Ann. Inst. Statist. Math.* **34** 543–544. MR0695073 <https://doi.org/10.1007/BF02481052>
- [3] AIROLDI, E. M., COSTA, T. B. and CHAN, S. H. (2013). Stochastic blockmodel approximation of a graphon: Theory and consistent estimation. In *Advances in Neural Information Processing Systems* 692–700.
- [4] ALLMAN, E. S., MATIAS, C. and RHODES, J. A. (2009). Identifiability of parameters in latent structure models with many observed variables. *Ann. Statist.* **37** 3099–3132. MR2549554 <https://doi.org/10.1214/09-AOS689>
- [5] ANANDKUMAR, A., HSU, D., JANZAMIN, M. and KAKADE, S. (2015). When are overcomplete topic models identifiable? Uniqueness of tensor Tucker decompositions with structured sparsity. *J. Mach. Learn. Res.* **16** 2643–2694. MR3450520
- [6] ANANDKUMAR, A., HSU, D., JAVANMARD, A. and KAKADE, S. (2013). Learning linear Bayesian networks with latent variables. In *Proceedings of the 30th International Conference on Machine Learning* 249–257.
- [7] ARAGAM, B., DAN, C., XING, E. P. and RAVIKUMAR, P. (2020). Supplement to “Identifiability of nonparametric mixture models and Bayes optimal clustering.” <https://doi.org/10.1214/19-AOS1887SUPP>.
- [8] ARORA, S. and KANNAN, R. (2005). Learning mixtures of separated nonspherical Gaussians. *Ann. Appl. Probab.* **15** 69–92. MR2115036 <https://doi.org/10.1214/105051604000000512>
- [9] BARNDORFF-NIELSEN, O. (1965). Identifiability of mixtures of exponential families. *J. Math. Anal. Appl.* **12** 115–121. MR0183058 [https://doi.org/10.1016/0022-247X\(65\)90059-4](https://doi.org/10.1016/0022-247X(65)90059-4)
- [10] BASU, A., SHIOYA, H. and PARK, C. (2011). *Statistical Inference: The Minimum Distance Approach. Monographs on Statistics and Applied Probability* **120**. CRC Press, Boca Raton, FL. MR2830561

- [11] BERAN, R. (1977). Minimum Hellinger distance estimates for parametric models. *Ann. Statist.* **5** 445–463. [MR0448700](#)
- [12] BOCK, H. H. (1996). Probabilistic models in cluster analysis. *Comput. Statist. Data Anal.* **23** 5–28.
- [13] BONHOMME, S., JOCHMANS, K. and ROBIN, J.-M. (2016). Estimating multivariate latent-structure models. *Ann. Statist.* **44** 540–563. [MR3476609](#) <https://doi.org/10.1214/15-AOS1376>
- [14] BONHOMME, S., JOCHMANS, K. and ROBIN, J.-M. (2016). Non-parametric estimation of finite mixtures from repeated measurements. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **78** 211–229. [MR3453653](#) <https://doi.org/10.1111/rssb.12110>
- [15] BORDES, L., KOJADINOVIC, I. and VANDEKERKHOVE, P. (2013). Semiparametric estimation of a two-component mixture of linear regressions in which one component is known. *Electron. J. Stat.* **7** 2603–2644. [MR3121625](#) <https://doi.org/10.1214/13-EJS858>
- [16] BORDES, L., MOTTELET, S. and VANDEKERKHOVE, P. (2006). Semiparametric estimation of a two-component mixture model. *Ann. Statist.* **34** 1204–1232. [MR2278356](#) <https://doi.org/10.1214/009053606000000353>
- [17] BRUNI, C. and KOCH, G. (1985). Identifiability of continuous mixtures of unknown Gaussian distributions. *Ann. Probab.* **13** 1341–1357. [MR0806230](#)
- [18] CHACÓN, J. E. (2015). A population background for nonparametric density-based clustering. *Statist. Sci.* **30** 518–532. [MR3432839](#) <https://doi.org/10.1214/15-STS526>
- [19] CHAUDHURI, K. and DASGUPTA, S. (2010). Rates of convergence for the cluster tree. In *Advances in Neural Information Processing Systems* 343–351.
- [20] CHEN, J. H. (1995). Optimal rate of convergence for finite mixture models. *Ann. Statist.* **23** 221–233. [MR1331665](#) <https://doi.org/10.1214/aos/1176324464>
- [21] CHEN, Y.-C., GENOVESE, C. R. and WASSERMAN, L. (2016). A comprehensive approach to mode clustering. *Electron. J. Stat.* **10** 210–241. [MR3466181](#) <https://doi.org/10.1214/15-EJS1102>
- [22] CSISZÁR, I. (1967). Information-type measures of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.* **2** 299–318. [MR0219345](#)
- [23] CUEVAS, A. and GONZÁLEZ MANTEIGA, W. (1991). Data-driven smoothing based on convexity properties. In *Nonparametric Functional Estimation and Related Topics (Spetses, 1990)*. *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.* **335** 225–240. Kluwer Academic, Dordrecht. [MR1154331](#)
- [24] D’HAULTFÈUILLE, X. and FÉVRIER, P. (2015). Identification of mixture models using support variations. *J. Econometrics* **189** 70–82. [MR3397353](#) <https://doi.org/10.1016/j.jeconom.2015.06.022>
- [25] DASGUPTA, S. (1999). Learning mixtures of Gaussians. In *40th Annual Symposium on Foundations of Computer Science (New York, 1999)* 634–644. IEEE Computer Soc., Los Alamitos, CA. [MR1917603](#) <https://doi.org/10.1109/SFFCS.1999.814639>
- [26] DEVROYE, L., GYÖRFI, L. and LUGOSI, G. (1996). *A Probabilistic Theory of Pattern Recognition. Applications of Mathematics (New York)* **31**. Springer, New York. [MR1383093](#) <https://doi.org/10.1007/978-1-4612-0711-5>
- [27] ELDRIDGE, J., BELKIN, M. and WANG, Y. (2016). Graphons, mergeons, and so on! In *Advances in Neural Information Processing Systems* 2307–2315.
- [28] FRALEY, C. and RAFTERY, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *J. Amer. Statist. Assoc.* **97** 611–631. [MR1951635](#) <https://doi.org/10.1198/016214502760047131>
- [29] GASSIAT, E. and ROUSSEAU, J. (2013). Non parametric finite translation mixtures with dependent regime. Preprint. Available at [arXiv:1302.2345](https://arxiv.org/abs/1302.2345).
- [30] GENOVESE, C. R. and WASSERMAN, L. (2000). Rates of convergence for the Gaussian mixture sieve. *Ann. Statist.* **28** 1105–1127. [MR1810921](#) <https://doi.org/10.1214/aos/1015956709>
- [31] GHOSAL, S. and VAN DER VAART, A. W. (2001). Entropies and rates of convergence for maximum likelihood and Bayes estimation for mixtures of normal densities. *Ann. Statist.* **29** 1233–1263. [MR1873329](#) <https://doi.org/10.1214/aos/1013203453>
- [32] HALL, P., NEEMAN, A., PAKYARI, R. and ELMORE, R. (2005). Nonparametric inference in multivariate mixtures. *Biometrika* **92** 667–678. [MR2202653](#) <https://doi.org/10.1093/biomet/92.3.667>
- [33] HALL, P. and ZHOU, X.-H. (2003). Nonparametric estimation of component distributions in a multivariate mixture. *Ann. Statist.* **31** 201–224. [MR1962504](#) <https://doi.org/10.1214/aos/1046294462>
- [34] HARTIGAN, J. A. (1975). *Clustering Algorithms*. Wiley, New York. [MR0405726](#)
- [35] HARTIGAN, J. A. (1981). Consistency of single linkage for high-density clusters. *J. Amer. Statist. Assoc.* **76** 388–394. [MR0624340](#)
- [36] HEINRICH, P. and KAHN, J. (2018). Strong identifiability and optimal minimax rates for finite mixture estimation. *Ann. Statist.* **46** 2844–2870. [MR3851757](#) <https://doi.org/10.1214/17-AOS1641>

- [37] HETTMANSPERGER, T. P. and THOMAS, H. (2000). Almost nonparametric inference for repeated measures in mixture models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **62** 811–825. MR1796294 <https://doi.org/10.1111/1467-9868.00266>
- [38] HO, N. and NGUYEN, X. (2016). On strong identifiability and convergence rates of parameter estimation in finite mixtures. *Electron. J. Stat.* **10** 271–307. MR3466183 <https://doi.org/10.1214/16-EJS1105>
- [39] HO, N. and NGUYEN, X. (2019). Singularity structures and impacts on parameter estimation in finite mixtures of distributions. *SIAM J. Math. Data Sci.* **1** 730–758. MR4016133 <https://doi.org/10.1137/18M122947X>
- [40] HOLLAND, P. W., LASKEY, K. B. and LEINHARDT, S. (1983). Stochastic blockmodels: First steps. *Soc. Netw.* **5** 109–137. MR0718088 [https://doi.org/10.1016/0378-8733\(83\)90021-7](https://doi.org/10.1016/0378-8733(83)90021-7)
- [41] HOLZMANN, H., MUNK, A. and GNEITING, T. (2006). Identifiability of finite mixtures of elliptical distributions. *Scand. J. Stat.* **33** 753–763. MR2300914 <https://doi.org/10.1111/j.1467-9469.2006.00505.x>
- [42] HUNTER, D. R., WANG, S. and HETTMANSPERGER, T. P. (2007). Inference for mixtures of symmetric distributions. *Ann. Statist.* **35** 224–251. MR2332275 <https://doi.org/10.1214/009053606000001118>
- [43] HUNTER, D. R. and YOUNG, D. S. (2012). Semiparametric mixtures of regressions. *J. Nonparametr. Stat.* **24** 19–38. MR2885823 <https://doi.org/10.1080/10485252.2011.608430>
- [44] JACOBS, R. A., JORDAN, M. I., NOWLAN, S. J. and HINTON, G. E. (1991). Adaptive mixtures of local experts. *Neural Comput.* **3** 79–87.
- [45] JOCHMANS, K., HENRY, M. and SALANIÉ, B. (2017). Inference on two-component mixtures under tail restrictions. *Econometric Theory* **33** 610–635. MR3637969 <https://doi.org/10.1017/S0266466616000098>
- [46] KANNAN, R., SALMASIAN, H. and VEMPALA, S. (2008). The spectral method for general mixture models. *SIAM J. Comput.* **38** 1141–1156. MR2421081 <https://doi.org/10.1137/S0097539704445925>
- [47] LEVINE, M., HUNTER, D. R. and CHAUVEAU, D. (2011). Maximum smoothed likelihood for multivariate mixtures. *Biometrika* **98** 403–416. MR2806437 <https://doi.org/10.1093/biomet/asq079>
- [48] LI, J. and SCHMIDT, L. (2015). A nearly optimal and agnostic algorithm for properly learning a mixture of k Gaussians, for any constant k . Preprint. Available at arXiv:1506.01367.
- [49] LINDSAY, B. G. (1995). Mixture models: Theory, geometry and applications. In *NSF-CBMS Regional Conference Series in Probability and Statistics*. JSTOR.
- [50] LLOYD, S. P. (1982). Least squares quantization in PCM. *IEEE Trans. Inform. Theory* **28** 129–137. MR0651807 <https://doi.org/10.1109/TIT.1982.1056489>
- [51] MACQUEEN, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proc. Fifth Berkeley Sympos. Math. Statist. and Probability (Berkeley, Calif., 1965/66)* 281–297. Univ. California Press, Berkeley, CA. MR0214227
- [52] MIXON, D. G., VILLAR, S. and WARD, R. (2017). Clustering subgaussian mixtures by semidefinite programming. *Inf. Inference* **6** 389–415. MR3764529 <https://doi.org/10.1093/imaiai/iax001>
- [53] NG, A. Y., JORDAN, M. I. and WEISS, Y. (2001). On spectral clustering: Analysis and an algorithm. In *NIPS* **14** 849–856.
- [54] NGUYEN, X. (2013). Convergence of latent mixing measures in finite and infinite mixture models. *Ann. Statist.* **41** 370–400. MR3059422 <https://doi.org/10.1214/12-AOS1065>
- [55] PARTHASARATHY, K. R. (1967). *Probability Measures on Metric Spaces. Probability and Mathematical Statistics, No. 3*. Academic Press, New York. MR0226684
- [56] PRIEBE, C. E. (1994). Adaptive mixtures. *J. Amer. Statist. Assoc.* **89** 796–806. MR1294725
- [57] RINALDO, A. and WASSERMAN, L. (2010). Generalized density clustering. *Ann. Statist.* **38** 2678–2722. MR2722453 <https://doi.org/10.1214/10-AOS797>
- [58] RITTER, G. (2015). *Robust Cluster Analysis and Variable Selection. Monographs on Statistics and Applied Probability* **137**. CRC Press, Boca Raton, FL. MR3241693
- [59] ROHE, K., CHATTERJEE, S. and YU, B. (2011). Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.* **39** 1878–1915. MR2893856 <https://doi.org/10.1214/11-AOS887>
- [60] SCHIEBINGER, G., WAINWRIGHT, M. J. and YU, B. (2015). The geometry of kernelized spectral clustering. *Ann. Statist.* **43** 819–846. MR3325711 <https://doi.org/10.1214/14-AOS1283>
- [61] SHI, T., BELKIN, M. and YU, B. (2009). Data spectroscopy: Eigenspaces of convolution operators and clustering. *Ann. Statist.* **37** 3960–3984. MR2572449 <https://doi.org/10.1214/09-AOS700>
- [62] SRIPERUMBUDUR, B. and STEINWART, I. (2012). Consistency and rates for clustering with dbscan. In *Artificial Intelligence and Statistics* 1090–1098.
- [63] STEINHAUS, H. (1956). Sur la division des corps matériels en parties. *Bull. Acad. Polon. Sci. Cl. III.* **4** 801–804. MR0090073
- [64] STEINWART, I. (2011). Adaptive density level set clustering. In *International Conference on Learning Theory* 703–738.

- [65] STEINWART, I. (2015). Fully adaptive density-based clustering. *Ann. Statist.* **43** 2132–2167. MR3396981 <https://doi.org/10.1214/15-AOS1331>
- [66] SWEETING, T. J. (1986). On a converse to Scheffé's theorem. *Ann. Statist.* **14** 1252–1256. MR0856821 <https://doi.org/10.1214/aos/1176350065>
- [67] TANG, J., MENG, Z., NGUYEN, X., MEI, Q. and ZHANG, M. (2014). Understanding the limiting factors of topic modeling via posterior contraction analysis. In *International Conference on Machine Learning* 190–198.
- [68] TEICHER, H. (1960). On the mixture of distributions. *Ann. Math. Stat.* **31** 55–73. MR0121825 <https://doi.org/10.1214/aoms/1177705987>
- [69] TEICHER, H. (1961). Identifiability of mixtures. *Ann. Math. Stat.* **32** 244–248. MR0120677 <https://doi.org/10.1214/aoms/1177705155>
- [70] TEICHER, H. (1963). Identifiability of finite mixtures. *Ann. Math. Stat.* **34** 1265–1269. MR0155376 <https://doi.org/10.1214/aoms/1177703862>
- [71] TEICHER, H. (1967). Identifiability of mixtures of product measures. *Ann. Math. Stat.* **38** 1300–1302. MR0216635 <https://doi.org/10.1214/aoms/1177698805>
- [72] THOMANN, P., STEINWART, I. and SCHMID, N. (2015). Towards an axiomatic approach to hierarchical clustering of measures. *J. Mach. Learn. Res.* **16** 1949–2002. MR3417804
- [73] TITTERINGTON, D. M., SMITH, A. F. M. and MAKOV, U. E. (1985). *Statistical Analysis of Finite Mixture Distributions*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. Wiley, Chichester. MR0838090
- [74] ULTSCH, A. (2005). Clustering with SOM: U*C. In *Proc. Workshop on Self-Organizing Maps* 75–82, Paris, France.
- [75] WOLFE, J. H. (1970). Pattern clustering by multivariate mixture analysis. *Multivar. Behav. Res.* **5** 329–350.
- [76] YAKOWITZ, S. J. and SPRAGINS, J. D. (1968). On the identifiability of finite mixtures. *Ann. Math. Stat.* **39** 209–214. MR0224204 <https://doi.org/10.1214/aoms/1177698520>
- [77] YAN, D., HUANG, L. and JORDAN, M. I. (2009). Fast approximate spectral clustering. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* 907–916. ACM, New York.