
Finite mixture models do not reliably learn the number of components

Diana Cai^{*1} Trevor Campbell^{*2} Tamara Broderick³

Abstract

Scientists and engineers are often interested in learning the number of subpopulations (or components) present in a data set. A common suggestion is to use a finite mixture model (FMM) with a prior on the number of components. Past work has shown the resulting FMM component-count posterior is consistent; that is, the posterior concentrates on the true, generating number of components. But consistency requires the assumption that the component likelihoods are perfectly specified, which is unrealistic in practice. In this paper, we add rigor to data-analysis folk wisdom by proving that under even the slightest model misspecification, the FMM component-count posterior *diverges*: the posterior probability of any particular finite number of components converges to 0 in the limit of infinite data. Contrary to intuition, posterior-density consistency is not sufficient to establish this result. We develop novel sufficient conditions that are more realistic and easily checkable than those common in the asymptotics literature. We illustrate practical consequences of our theory on simulated and real data.

1. Introduction

Mixture modeling is a mainstay of statistical machine learning. In applications where the number of mixture components is unknown in advance, a principal inferential goal is to estimate and interpret this number. For example, practitioners might wish to find the number of latent genetic populations (Pritchard et al., 2000; Lorenzen et al., 2006; Huelsenbeck and Andolfatto, 2007; Tonkin-Hill et al., 2019), gene tissue profiles (Yeung et al., 2001; Medvedovic and Sivaganesan, 2002), cell types (Chan et al., 2008; Prab-

hakaran et al., 2016), microscopy groups (Rubin-Delanchy et al., 2015; Griffié et al., 2016), haplotypes (Xing et al., 2006), switching Markov regimes in US dollar exchange rate data (Otranto and Gallo, 2002), gamma-ray burst types (Mukherjee et al., 1998), segmentation regions in an image (e.g., tissue types in an MRI scan (Banfield and Raftery, 1993)), observed humans in radar data (Teklehaymanot et al., 2018), basketball shot selection groups (Hu et al., 2020), or communities in a social network (Geng et al., 2019; Legramanti et al., 2020).

A natural question then is: can we reliably learn the number of latent groups in a data set? To make this question concrete, we focus on a Bayesian approach. Consider the case where the true, generating number of components is known. A natural check on a Bayesian mixture analysis is to establish that the Bayesian posterior on the number of components increasingly concentrates near the truth as the number of data points becomes arbitrarily large. In the remainder, we will focus on this check—though our work has practical implications beyond Bayesian analysis.

A standard Bayesian analysis uses a component-count prior with support on all strictly-positive integers (Miller and Harrison, 2018). Nobile (1994) has shown that the component-count posterior of the resulting *finite mixture model* (FMM) does concentrate at the true number of components. But crucially, this result depends on the assumption that the component likelihoods are perfectly specified. In every application we have listed above, the true generating component likelihoods do not take a convenient parametric form that might be specified in advance. Indeed, some form of misspecification, even if slight, is typical in practice. So we must ask how the component-count posterior behaves when the component likelihoods are misspecified.

Data science folk wisdom suggests that when component likelihoods are misspecified, mixture models will tend to overestimate the number of clusters; see, e.g., Section 7.1 of Frühwirth-Schnatter (2006). This overestimation is apparent in Figure 1, which shows the component-count posterior of a Gaussian mixture model applied to two example gene expression data sets (de Souto et al., 2008; Prabhakaran et al., 2016). In fact, Figure 1 demonstrates an effect far worse than just overestimation: the posterior distribution appears to concentrate for any (large enough) fixed amount of

^{*}Equal contribution ¹Department of Computer Science, Princeton University ²Department of Statistics, University of British Columbia ³CSAIL, Massachusetts Institute of Technology. Correspondence to: Diana Cai <dcai@cs.princeton.edu>, Trevor Campbell <trevor@stat.ubc.ca>, Tamara Broderick <tbroderick@csail.mit.edu>.

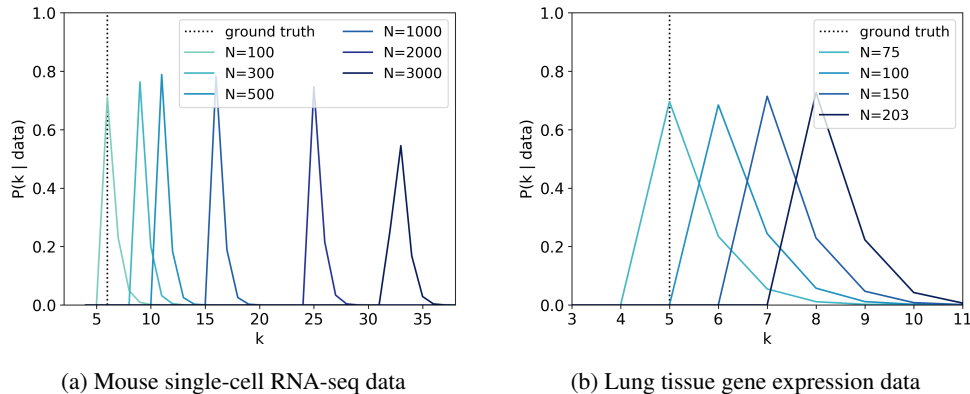


Figure 1. Posterior probability of the number of components k for Gaussian mixture models, fit to (a) mouse cortex single-cell RNA sequencing data and (b) lung tissue gene expression data. Details in Section 7.2.

data, but actually concentrates on increasing values as more data are observed. Therefore, inference is unreliable; the practitioner may draw quite different conclusions depending on how large the data set is.

In the present paper, we add rigor to existing data science folk intuition by proving that this behavior occurs in a wide class of FMMs under arbitrarily small amounts of misspecification. We examine FMMs with essentially any component shape—where we make only mild, realistic, and checkable assumptions on the component likelihoods. Notably, we include univariate and multivariate Gaussian component likelihoods in our theory, but do not restrict only to these shapes. We show that under our assumptions and when the component likelihoods are not perfectly specified, the component-count posterior concentrates strictly away from the true number of components. In fact, we go further to show that the FMM posterior for the number of components *diverges*: for any finite $k \in \mathbb{N}$, the posterior probability that the number of components is k converges to 0 almost surely as the amount of data grows.

We start by introducing FMMs and stating our main result in Section 2. We discuss our assumptions in more detail in Section 3 and prove our result in Section 4. In Section 5 we extend our main theorem to priors that may vary as the data set grows. We discuss related work below and in Section 6. The paper concludes in Section 7 with empirical evidence that the FMM component-count posterior depends strongly on the amount of observed data. Our results demonstrate that, in practice, past estimates of component number may have strongly depended on the size of a particular data set.

Filling a gap in the literature. While recent work has established various asymptotic properties of mixture models, we observe that our results here are not trivial extensions of existing research. First note that, intuitively, as the number of data points grows, the posterior concentrates at the generating density (Schwartz, 1965; Ghosh and Ramamoor-

thi, 2003; Ghosal and van der Vaart, 2017), which can be well-approximated by an infinite mixture due in part to misspecification. However, posterior consistency for the density alone is not enough to guarantee consistency for the model parameters; parameter consistency may not hold under, for instance, a discontinuous mapping from the component parameter to the component density.

Second, note that posterior divergence for the number of components could, in principle, be obtained if parameter consistency for the mixture holds. But existing results on parameter consistency, such as Nguyen (2013), focus on obtaining rates of contraction; thus these results rely on stronger conditions that are typically verified for individual component families by imposing additional constraints, such as (second-order) strong identifiability of the mixture or a compact parameter space (Chen, 1995; Nguyen, 2013; Heinrich and Kahn, 2018). But neither of these constraints are satisfied by common families of interest such as Gaussians with unknown mean and variance. By contrast, our result uses the weakest notion of mixture identifiability (Teicher, 1961) along with a continuity condition on the component family, and we relax the requirement of a compact parameter space. To do so, we develop a novel theoretical condition that requires the component family to have *degenerate limits*. Together, these advances ensure the applicability of our theory to practical likelihood families including, but not limited to, full Gaussians. In fact, the degenerate limits condition and its use in our analysis may be useful for extending other results on posterior asymptotics that currently rely on compact parameter spaces.

Finally, Miller and Harrison (2013; 2014) have shown that typical uses of Dirichlet process mixture models (DPMMs) inconsistently estimate the true, generating number of components. But Miller and Harrison (2013; 2014) focus on the DPMM prior instead of the FMM and on perfectly specified likelihoods. The DPMM is misspecified in a different sense than the one we focus on in the present paper: namely,

the DPMM uses infinitely many components though we assume finitely many generated the data. For this reason, practitioners typically invoke the DPMM posterior on the number of *clusters* (Pella and Masuda, 2006; Huelsenbeck and Andolfatto, 2007), i.e., components represented in the observed data, rather than the component-count posterior directly. Indeed, Miller and Harrison (2018) recommend using the FMM we study here to resolve the difficulties of the DPMM. Finally, observe that the work of Miller and Harrison (2018) demonstrates that nonparametrically estimating component shape with a DPMM would not provide a simple resolution of the FMM divergence issue.

2. Main result

We begin with a brief description of the finite mixture model used in this work. In this section, we provide just enough detail to state Theorem 2.1 and leave the precise probabilistic details for Section 3. Let g be a mixing measure $g := \sum_{j=1}^k p_j \delta_{\theta_j}$ on a parameter space Θ with $p_j \in [0, 1]$ and $\sum_{j=1}^k p_j = 1$, and let $\Psi = \{\psi_\theta : \theta \in \Theta\}$ be a family of component distributions dominated by a σ -finite measure μ . We can express a finite mixture f of the components as

$$f = \int_{\Theta} \psi_\theta dg(\theta) = \sum_{j=1}^k p_j \psi_{\theta_j}.$$

Consider a Bayesian model with prior distribution Π on the set of all mixing measures \mathbb{G} on Θ with finitely many atoms, i.e., $g \sim \Pi$, and likelihood corresponding to conditionally i.i.d. data from $f = \int \psi_\theta dg(\theta)$. The model assumes the likelihood is f , but the model is *misspecified*; i.e., the observations $X_{1:N} := (X_1, \dots, X_N)$ are generated conditionally i.i.d. from a finite mixture f_0 of distributions *not* in Ψ .

Our main result is that under this misspecification of the likelihood, the posterior on the number of components $\Pi(k | X_{1:N})$ diverges; i.e., for any finite $k \in \mathbb{N}$, $\Pi(k | X_{1:N}) \rightarrow 0$ as $N \rightarrow \infty$. We make only two requirements of the mixture model to guarantee this result: (1) the true data-generating distribution f_0 must be arbitrarily well-approximated by finite mixtures of Ψ , and (2) the family Ψ must satisfy mild regularity conditions that hold for popular mixture models (e.g., the family Ψ of Gaussians parametrized by mean and variance). We provide precise definitions of the assumptions needed for Theorem 2.1 to hold in Section 3, and a proof in Section 4.

Theorem 2.1 (Main result). *Suppose observations $X_{1:N}$ are generated i.i.d. from a distribution f_0 that is not a finite mixture of Ψ . Assume that:*

Assumption 3.1: f_0 is in the KL-support of the prior Π ,

Assumption 3.6: Ψ is continuous, is mixture-identifiable, and has degenerate limits.

Then the posterior on the number of components diverges; i.e., for all $k \in \mathbb{N}$,

$$\Pi(k | X_{1:N}) \xrightarrow{N \rightarrow \infty} 0 \quad f_0\text{-a.s.} \quad (1)$$

Note that the conditions of the theorem—although technical—are satisfied by a wide class of models used in practice. Assumption 3.1 requires that the prior Π places enough mass on mixtures near the true generating distribution f_0 . Assumption 3.6 enforces regularity of the component family and is satisfied by many popular models used in practice, such as the multivariate Gaussian family.

Proposition 2.2. *Let $\Psi = \{\mathcal{N}(\nu, \Sigma) : \nu \in \mathbb{R}^d, \Sigma \in \mathbb{S}_{++}^d\}$ be the multivariate Gaussian family, where $\mathbb{S}_{++}^d := \{\Sigma \in \mathbb{R}^{d \times d} : \Sigma = \Sigma^\top, \Sigma \succ 0\}$ is the set of $d \times d$ symmetric, positive definite matrices. Then Ψ satisfies Assumption 3.6.*

Thus, provided that f_0 is in the KL-support of the prior, under a misspecified Gaussian mixture model, our main result implies that the posterior number of components diverges. While Proposition 2.2 is stated for Gaussian component distributions, we generalize it to mixture-identifiable location-scale families Ψ in Proposition C.2.

Additionally, we note that the divergence of the posterior given in Equation (1) is stronger than the behavior described in Miller and Harrison (2013) for DPMMs: namely, Miller and Harrison (2013) show that the posterior probability converges to 0 at the *true* number of components. In contrast, here we show that the posterior probability converges to 0 for any finite number of components. We conjecture that posterior divergence also holds for DPMMs, but the proof is outside of the scope of this paper.

Extension: Priors that vary with N . While the result of Theorem 2.1 assumes that the model uses a fixed prior Π , in practical modeling scenarios one may specify a prior Π_N that depends on the observed data $X_{1:N}$. For instance, these priors can arise in empirical Bayes; see Sections 5 and 7 for examples. In Section 5 we show that if f_0 satisfies a modified KL-support condition with respect to the sequence of priors Π_N , the number of components also diverges in this setting.

Extension: Priors with an upper bound on the number of components. Theorem 2.1 is designed for priors that place full support on any positive integer number of components. One might instead use a prior that has support on at most \tilde{k} components, with \tilde{k} finite. In this case, the posterior number of components will not diverge to infinity but instead typically concentrate on the upper bound, \tilde{k} . A precise statement of this behavior appears in Theorem B.1 (in Appendix B) as an analog of our Theorem 2.1. Theorem B.1 shows that posterior inference does not improve over the baseline estimate of the number of components provided by \tilde{k} . If \tilde{k} is already a good estimate of the number components,

posterior concentration at \tilde{k} does not improve the estimate. In practice \tilde{k} is often chosen as some large upper bound of convenience; then \tilde{k} is not a good estimate of the number of components, and concentration at \tilde{k} is undesirable.

3. Precise setup and assumptions in Theorem 2.1

This section makes the details of the modeling setup and each of the conditions in Theorem 2.1 precise.

3.1. Notation and setup

Let \mathbb{X} and Θ be Polish spaces for the observations and parameters, respectively, and endow both with their Borel σ -algebra. For a topological space (\cdot) , let $\mathcal{C}(\cdot)$ be the bounded continuous functions from (\cdot) into \mathbb{R} , and $\mathcal{P}(\cdot)$ be the set of probability measures on (\cdot) endowed with the weak topology metrized by the Lévy-Prokhorov distance d (Definition A.1). We use $f_i \Rightarrow f$ and $f_i \iff f'_i$ to denote $\lim_{i \rightarrow \infty} d(f_i, f) = 0$ and $\lim_{i \rightarrow \infty} d(f_i, f'_i) = 0$, respectively, for $f_i, f'_i, f \in \mathcal{P}(\cdot)$. We assume that the family of distributions $\Psi = \{\psi_\theta : \theta \in \Theta\}$ is absolutely continuous with respect to a σ -finite base measure μ , i.e., $\psi_\theta \ll \mu$ for all $\theta \in \Theta$, and that for measurable $A \subseteq \mathbb{X}$, $\psi_\theta(A)$ is a measurable function on Θ . Define the measurable mapping $F : \mathcal{P}(\Theta) \rightarrow \mathcal{P}(\mathbb{X})$ from mixing measures to mixtures of Ψ , $F(g) = \int \psi_\theta dg(\theta)$. Let \mathbb{G} be the set of atomic probability measures on Θ with finitely many atoms, and let \mathbb{F} be the set of finite mixtures of Ψ .

In the Bayesian finite mixture model from Section 2, a mixing measure $g \sim \Pi$ is generated from a prior measure Π on \mathbb{G} , and $f = F(g)$ is a likelihood distribution.

The posterior distribution on the mixing measure is, for all measurable $A \subseteq \mathbb{G}$,

$$\Pi(A | X_{1:N}) = \frac{\int_A \prod_{n=1}^N \frac{df}{d\mu}(X_n) d\Pi(g)}{\int_{\mathbb{G}} \prod_{n=1}^N \frac{df}{d\mu}(X_n) d\Pi(g)}, \quad (2)$$

where $\frac{df}{d\mu}$ is the density of $f = F(g)$ with respect to μ . This posterior on the mixing measure $g \in \mathbb{G}$ induces a posterior on the number of components $k \in \mathbb{N}$ by counting the number of atoms in g , and it also induces a posterior on mixtures $f \in \mathbb{F}$ via the pushforward through the mapping F . We overload the notation $\Pi(\cdot | X_{1:N})$ to refer to all of these posterior distributions and $\Pi(\cdot)$ to refer to prior distributions; the meaning should be clear from context.

3.2. Model assumptions

The first assumption of Theorem 2.1 is that while the true data-generating distribution f_0 is not contained in the model class $f_0 \notin \mathbb{F}$, it lies on the boundary of the model class. In particular, we assume f_0 is in the *KL-support* of the prior

Π . Denote the Kullback-Leibler (KL) divergence between probability measures f_0 and f as

$$\text{KL}(f_0, f) := \begin{cases} \int \log\left(\frac{df_0}{df}\right) df_0 & f_0 \ll f \\ \infty & \text{otherwise} \end{cases}.$$

Assumption 3.1. For all $\epsilon > 0$, the prior distribution Π satisfies

$$\Pi(f \in \mathbb{F} : \text{KL}(f_0, f) < \epsilon) > 0.$$

We use Assumption 3.1 in the proof of Theorem 2.1 primarily to ensure that the Bayesian posterior is consistent for f_0 . Note that Assumption 3.1 is fairly weak in practice. Intuitively, it just requires that the family Ψ is rich enough so that mixtures of Ψ can approximate f_0 arbitrarily well, and that the prior Π places sufficient mass on those mixtures close to f_0 . For Bayesian mixture modeling, Ghosal et al. (1999, Theorem 3), Tokdar (2006, Theorem 3.2), Wu and Ghosal (2008, Theorem 2.3), and Petralia et al. (2012, Theorem 1) provide conditions needed to satisfy Assumption 3.1.

The second assumption of Theorem 2.1 is that the family of component distributions Ψ is well-behaved. This assumption has three stipulations. First, the mapping $\theta \mapsto \psi_\theta$ must be continuous; this condition essentially asserts that similar parameter values θ must result in similar component distributions ψ_θ .

Definition 3.2. The family Ψ is continuous if the map $\theta \mapsto \psi_\theta$ is continuous.

Second, the family Ψ must be *mixture-identifiable*, which guarantees that each mixture $f \in \mathbb{F}$ is associated with a unique mixing measure $G \in \mathbb{G}$.

Definition 3.3 (Teicher (1961; 1963)). The family Ψ is mixture-identifiable if the mapping $F(g) = \int \psi_\theta dg(\theta)$ restricted to finite mixtures $F : \mathbb{G} \rightarrow \mathbb{F}$ is a bijection.

In practice, one should always use an identifiable mixture model for clustering; without identifiability, the task of learning the number of components is ill posed. And many models satisfy mixture-identifiability, such as finite mixtures of the multivariate Gaussian family (Yakowitz and Spragins, 1968), the Cauchy family (Yakowitz and Spragins, 1968), the gamma family (Teicher, 1963), the generalized logistic family, the generalized Gumbel family, the Weibull family, and von Mises family (Ho and Nguyen, 2016, Theorem 3.3). A number of authors (e.g. Chen, 1995; Ishwaran et al., 2001; Nguyen, 2013; Ho and Nguyen, 2016; Guha et al., 2019; Heinrich and Kahn, 2018) appeal to stronger notions of identifiability for mixtures than Definition 3.3. But, to show posterior divergence in the present work, we do not require conditions stronger than Definition 3.3.

The third stipulation—that the family Ψ has *degenerate limits*—guarantees that a “poorly behaved” sequence of parameters $(\theta_i)_{i \in \mathbb{N}}$ creates a likewise “poorly behaved” sequence

of distributions $(\psi_{\theta_i})_{i \in \mathbb{N}}$. This condition allows us to rule out such sequences in the proof of Theorem 2.1, and is the essential regularity condition to guarantee that a sequence of finite mixtures of at most k components cannot approximate f_0 arbitrarily closely.

Definition 3.4. A sequence of distributions $(\psi_i)_{i=1}^{\infty}$ is μ -wide if for any closed set C such that $\mu(C) = 0$ and any sequence of distributions $(\phi_i)_{i=1}^{\infty}$ such that $\psi_i \iff \phi_i$,

$$\limsup_{i \rightarrow \infty} \phi_i(C) = 0.$$

Definition 3.5. The family Ψ has degenerate limits if for any tight, μ -wide sequence $(\psi_{\theta_i})_{i \in \mathbb{N}}$, we have that $(\theta_i)_{i \in \mathbb{N}}$ is relatively compact.

The contrapositive of Definition 3.5 provides an intuitive explanation of the condition: as $i \rightarrow \infty$, for any sequence of parameters θ_i that eventually leaves every compact set $K \subseteq \Theta$, either the ψ_{θ_i} become “arbitrarily flat” (not tight) or “arbitrarily peaky” (not μ -wide). For example, consider the family Ψ of Gaussians on \mathbb{R} with Lebesgue measure μ . If the variance of ψ_{θ_i} shrinks as i grows, the sequence of distributions converges weakly to a sequence of point masses (not dominated by the Lebesgue measure). If either the variance or the mean diverges, the distributions flatten out and the sequence is not tight. We use the fact that these are the only two possibilities when a sequence of parameters is poorly behaved (not relatively compact) in the proof of Theorem 2.1.

These three stipulations together yield Assumption 3.6.

Assumption 3.6. The mixture component family Ψ is continuous, is mixture-identifiable, and has degenerate limits.

4. Proof of Theorem 2.1

The proof has two essential steps. The first is to show that the Bayesian posterior is weakly consistent for the mixture f_0 ; i.e., for any weak neighborhood U of f_0 the sequence of posterior distributions satisfies

$$\Pi(U | X_{1:N}) \xrightarrow{N \rightarrow \infty} 1, \quad f_0\text{-a.s.} \quad (3)$$

By Schwartz’s theorem (Theorem A.4), weak consistency for f_0 is guaranteed directly by Assumption 3.1 and the fact that Ψ is dominated by a σ -finite measure μ . The second step is to show that for any finite $k \in \mathbb{N}$, there exists a weak neighborhood U of f_0 containing no mixtures of the family Ψ with at most k components. Together, these steps show that the posterior probability of the set of all k -component mixtures converges to 0 f_0 -a.s. as the amount of observed data grows.

We provide a proof of the second step. To begin, note that Assumption 3.1 has two additional implications about f_0

beyond Equation (3). First, f_0 must be absolutely continuous with respect to the dominating measure μ ; if it were not, then there exists a measurable set A such that $f_0(A) > 0$ and $\mu(A) = 0$. Since μ dominates Ψ , any $f \in \mathbb{F}$ satisfies $f(A) = 0$. Therefore $\text{KL}(f_0, f) = \infty$, and the prior support condition cannot hold. Second, it implies that f_0 can be arbitrarily well-approximated by finite mixtures under the weak metric, i.e., there exists a sequence of finite mixtures $f_i \in \mathbb{F}$, $i \in \mathbb{N}$ such that $f_i \Rightarrow f_0$ as $i \rightarrow \infty$. This holds because $\sqrt{\frac{1}{2}\text{KL}(f_0, f)} \geq \text{TV}(f_0, f) \geq d(f_0, f)$.

Now suppose the contrary of the claim for the second step, i.e., that there exists a sequence $(f_i)_{i=1}^{\infty}$ of mixtures of at most k components from Ψ such that $f_i \Rightarrow f_0$. By mixture-identifiability, we have a sequence of mixing measures g_i with at most k atoms such that $F(g_i) = f_i$. Suppose first that the atoms of the sequence $(g_i)_{i \in \mathbb{N}}$ either stay in a compact set or have weights converging to 0. More precisely, suppose there exists a compact set $K \subseteq \Theta$ such that

$$g_i(\Theta \setminus K) \rightarrow 0. \quad (4)$$

Decompose each $g_i = g_{i,K} + g_{i,\Theta \setminus K}$ such that $g_{i,K}$ is supported on K and $g_{i,\Theta \setminus K}$ is supported on $\Theta \setminus K$. Define the sequence of probability measures $\hat{g}_{i,K} = \frac{g_{i,K}}{g_{i,K}(\Theta)}$ for sufficiently large i such that the denominator is nonzero. Then Equation (4) implies

$$F(\hat{g}_{i,K}) \Rightarrow f_0.$$

Since Ψ is continuous and mixture-identifiable, the restriction of F to the domain \mathbb{G} is continuous and invertible; and since K is compact, the elements of $(\hat{g}_{i,K})_{i \in \mathbb{N}}$ are contained in a compact set $\mathbb{G}_K \subseteq \mathbb{G}$ by Prokhorov’s theorem (Theorem A.3). Therefore $F(\mathbb{G}_K) = \mathbb{F}_K$ is also compact, and the map F restricted to the domain \mathbb{G}_K is uniformly continuous with a uniformly continuous inverse by Rudin (1976, Theorems 4.14, 4.17, 4.19). Next since $F(\hat{g}_{i,K}) \Rightarrow f_0$, the sequence $F(\hat{g}_{i,K})$ is Cauchy in \mathbb{F}_K ; and since F^{-1} is uniformly continuous on \mathbb{F}_K , the sequence $\hat{g}_{i,K}$ must also be Cauchy in \mathbb{G}_K . Since \mathbb{G}_K is compact, $\hat{g}_{i,K}$ converges in \mathbb{G}_K . Lemma 4.1 below guarantees that the convergent limit g_K is also a mixing measure with at most k atoms; continuity of F implies that $F(g_K) = f_0$, which is a contradiction, since by assumption f_0 is not representable as a finite mixture of Ψ .

Lemma 4.1. Suppose $\phi, (\phi_i)_{i \in \mathbb{N}}$ are Borel probability measures on a Polish space such that $\phi_i \Rightarrow \phi$ and $\text{supp}_i |\text{supp } \phi_i| \leq k \in \mathbb{N}$. Then $|\text{supp } \phi| \leq k$.

Proof. Suppose $|\text{supp } \phi| > k$. Then we can find $k + 1$ distinct points $x_1, \dots, x_{k+1} \in \text{supp } \phi$. Pick any metric ρ on the Polish space, and denote the minimum pairwise distance between the points 2ϵ . Then for each point $j = 1, \dots, k + 1$ define the bounded, continuous function

$h_j(x) = 0 \vee (1 - \epsilon^{-1} \rho(x, x_j))$. Since $x_j \in \text{supp } \phi$, we have that $\int h_j d\phi > 0$. Weak convergence $\phi_i \Rightarrow \phi$ therefore implies $\min_{j=1, \dots, k+1} \liminf_{i \rightarrow \infty} \int h_j d\phi_i > 0$. But the h_j are nonzero on disjoint sets, and each ϕ_i only has k atoms; the pigeonhole principle yields a contradiction. \square

Now we consider the remaining case: for all compact sets $K \subseteq \Theta$, $g_i(\Theta \setminus K) \not\rightarrow 0$. Therefore there exists a sequence of parameters $(\theta_i)_{i=1}^{\infty}$ that is not relatively compact such that $\limsup_{i \rightarrow \infty} g_i(\{\theta_i\}) > 0$. By Assumption 3.6, the sequence $(\psi_{\theta_i})_{i \in \mathbb{N}}$ is either not tight or not μ -wide. If $(\psi_{\theta_i})_{i \in \mathbb{N}}$ is not tight then $f_i = F(g_i)$ is not tight, and by Prokhorov's theorem f_i cannot converge to a probability measure, which contradicts $f_i \Rightarrow f_0$. If $(\psi_{\theta_i})_{i \in \mathbb{N}}$ is not μ -wide then $f_i = F(g_i)$ is not μ -wide. Denote $(\phi_i)_{i \in \mathbb{N}}$ to be the singular sequence associated with $(f_i)_{i \in \mathbb{N}}$ and C to be the closed set such that $\limsup_{i \rightarrow \infty} \phi_i(C) > 0$, $\mu(C) = 0$, and $\phi_i \iff f_i$ per Definition 3.4. Since $f_0 \ll \mu$, $f_0(C) = 0$. But $f_i \Rightarrow f_0$ implies that $\phi_i \Rightarrow f_0$, so $\limsup_{i \rightarrow \infty} \phi_i(C) = f_0(C) = 0$ by the Portmanteau theorem (Theorem A.2). This is a contradiction.

5. Extension to priors that vary with N

Our main result (i.e., Theorem 2.1) applies to the setting of a fixed prior Π . However, it is often natural to specify a prior distribution that changes with N (e.g., Roeder and Wasserman, 1997; Richardson and Green, 1997; and Miller and Harrison, 2018, Section 7.2.1). Corollary 5.2 below demonstrates that a result nearly identical to Theorem 2.1 holds for priors that are allowed to vary with N , provided that f_0 is in the KL-support of the *sequence* of priors Π_N . The only difference is that our result in this case is slightly weaker: we show that the posterior number of components diverges in probability rather than almost surely.

Assumption 5.1. *For all $\epsilon > 0$, the sequence of prior distributions Π_N satisfies*

$$\liminf_{N \rightarrow \infty} \Pi_N(f : \text{KL}(f_0, f) < \epsilon) > 0.$$

Corollary 5.2. *Suppose in the setting of Theorem 2.1 we replace Assumption 3.1 with Assumption 5.1. Then the posterior on the number of components diverges in f_0 -probability: i.e., for all $k \in \mathbb{N}$,*

$$\Pi(k | X_{1:N}) \xrightarrow{N \rightarrow \infty} 0 \quad \text{in } f_0\text{-probability.}$$

Proof. Since for any $\epsilon > 0$, $\liminf_{N \rightarrow \infty} \Pi_N(f : \text{KL}(f_0, f) < \epsilon) > 0$, Ghosal and van der Vaart (2017, Theorem 6.17, Lemma 6.26, and Example 6.20) imply that the posterior is weakly consistent at f_0 in probability: i.e., for any weak neighborhood U of f_0 ,

$$\Pi(U | X_{1:N}) \xrightarrow{N \rightarrow \infty} 1 \quad \text{in } f_0\text{-probability.}$$

Assumption 5.1 also implies that for sufficiently large N , f_0 is a weak limit of finite mixtures in \mathbb{F} . The remainder of the proof is identical to that of Theorem 2.1. \square

6. Related work

In this work, we consider FMMs with a prior on the number of components. In the broader Bayesian mixture modeling literature, posterior consistency for the mixture density (Ghosal et al., 1999; Lijoi et al., 2004; Kruijer et al., 2010) and the mixing measure (Nguyen, 2013; Ho and Nguyen, 2016; Guha et al., 2019) is well established. But posterior consistency for the number of components is not as thoroughly characterized. There are several results establishing consistency for the number of components in well-specified FMMs. Nobile (1994, Proposition 3.5) and Guha et al. (2019, Theorem 3.1a) demonstrate that FMMs exhibit posterior consistency for the number of components when the model is well specified and Ψ is mixture-identifiable. The present work characterizes the behavior of the FMM posterior on the number of components under component misspecification. Under misspecification of the component family or the support of the true mixing measure, Guha et al. (2019, Theorem 4.1, Theorem 4.3) establish posterior rates of contraction for the mixing measure for Gaussian and Laplace location mixtures with compact parameter spaces. Our results, which rely on posterior density consistency results, assume weaker conditions on the prior and hold for more general classes of component families, such as multivariate Gaussians parameterized by a mean and covariance.

A related approach for handling a finite but unknown number of components is to specify a prior with a finite upper bound on the number of components (e.g. Ishwaran et al., 2001; Chambaz and Rousseau, 2008; Rousseau and Mengersen, 2011; Malsiner-Walli et al., 2016; Zhang et al., 2018; Frühwirth-Schnatter and Malsiner-Walli, 2019). In the setting of overfitted FMMs with well-specified component densities, Rousseau and Mengersen (2011, Theorem 1) show that under a stronger identifiability condition than mixture-identifiability and additional regularity assumptions on the model, the posterior will concentrate properly by emptying the extra components. Ishwaran et al. (2001, Theorem 1) consider the setting of estimating the number of components with the assumption of a known upper bound on the true number of components and well-specified components, and show that the posterior does not asymptotically underestimate the number of components when assuming a stronger identifiability condition than mixture-identifiability and a KL-support condition on the prior. Under a weaker notion of (second-order) strong identifiability (Chen, 1995) and a well-specified model, Chambaz and Rousseau (2008, Theorem 4) provide upper bounds on the underestimation and overestimation error of the number of components; fur-

thermore, they show that their conditions are satisfied by univariate Gaussians with bounded mean and variance (Cham-baz and Rousseau, 2008, Corollary 1). Notably, all of these methods with finite-support priors assume well-specified component families. By contrast, we show in Theorem B.1 that even for these finite-support priors, misspecified component families yield unreliable estimates of the number of components.

Frühwirth-Schnatter (2006) provides a wide-ranging review of methodology for finite mixture modeling. In (e.g.) Section 7.1, Frühwirth-Schnatter (2006) observes that, in practice, the learned number of mixture components will generally be higher than the true generating number of components when the likelihood is misspecified—but does not prove a result about the number of components under misspecification. Similarly, Miller and Harrison (2018, Section 7.1.5) discuss the issue of estimating the number of components in FMMs under model misspecification and state that the posterior number of components is expected to diverge to infinity as the number of samples increases, but no proof of this asymptotic behavior is provided.

Finally, a growing body of work is focused on developing more robust FMMs and related mixture models. In order to address the issue of component misspecification, a number of authors propose using finite mixture models with nonparametric component densities, e.g. Gaussian-mixture components (Bartolucci, 2005; Di Zio et al., 2007; Malsiner-Walli et al., 2017) or overfitted-mixture components (Aragam et al., 2020). However, for these finite mixture models that have mixtures as components, the posterior number of components and its asymptotic behavior have yet to be characterized.

7. Experiments

In this section, we demonstrate one of the primary practical implications of our theory: the inferred number of components can change drastically depending on the amount of observed data in misspecified finite mixture models. For all experiments below, we use a finite mixture model with a multivariate Gaussian component family having diagonal covariance matrices and a conjugate prior on each dimension. In particular, consider number of components k , mixture weights $p \in \mathbb{R}^k$, Gaussian component precisions $\tau \in \mathbb{R}_+^{k \times D}$ and means $\theta \in \mathbb{R}^{k \times D}$, labels $Z \in \{1, \dots, k\}^N$, and data $X \in \mathbb{R}^{N \times D}$.

Then the probabilistic generative model is

$$\begin{aligned} k &\sim \text{Geom}(r) & p &\sim \text{Dirichlet}_k(\gamma, \dots, \gamma) \\ \tau_{jd} &\stackrel{\text{i.i.d.}}{\sim} \text{Gam}(\alpha, \beta) & \theta_{jd} &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(m, \kappa_{jd}^{-1}) \\ Z_n &\stackrel{\text{i.i.d.}}{\sim} \text{Categorical}(p) & X_{nd} &\stackrel{\text{ind}}{\sim} \mathcal{N}(\theta_{z_n d}, \tau_{z_n d}^{-1}), \end{aligned}$$

where j ranges from $1, \dots, k$, d ranges from $1, \dots, D$, and n ranges from $1, \dots, N$. For posterior inference, we use a Julia implementation of split-merge collapsed Gibbs sampling (Neal, 2000; Jain and Neal, 2004) from Miller and Harrison (2018).^{*} The model and inference algorithm are described in more detail in Miller and Harrison (2018, Sec. 7.2.2, Algorithm 1). Note that we use this model primarily to illustrate the problem of posterior divergence under model misspecification; it should not be interpreted as a carefully-specified model for the data examples that we study. Also note that while the empirical examples below involve Gaussian FMMs, our theory applies to a more general class of component distributions.

7.1. Synthetic data

Gaussian and Laplace mixtures. Our first experiments on synthetic data are inspired by Figure 3 of Miller and Dunson (2019), which investigates the posterior of a mixture of perturbed Gaussians. Here we study the effects of varying data set sizes under both well-specified and misspecified models. We generated data sets of increasing size $N \in \{50, 200, 1000, 5000, 10000\}$ from 1- and 2-component univariate Gaussian and Laplace mixture models, where the 1-component distributions have mean 0 and scale 1, and the 2-component distributions have means $(-5, 5)$, scales $(1.5, 1)$, and mixing weights $(0.4, 0.6)$. We generated the sequence of data sets such that each was a subset of the next, larger data set in the sequence. Following Miller and Harrison (2018, Section 7.2.1), we set the hyperparameters of the Bayesian finite mixture model as follows: $m = \frac{1}{2}(\max_{n \in [\tilde{N}]} X_n + \min_{n \in [\tilde{N}]} X_n)$ where $\tilde{N} = 10,000$, $\kappa = (\max_{n \in [\tilde{N}]} X_n - \min_{n \in [\tilde{N}]} X_n)^{-2}$, $\alpha = 2$, $r = 0.1$, $\gamma = 1$, and $\beta \sim \text{Gam}(0.2, 10/\kappa)$. We refer to Miller and Harrison (2018, Section 7.2.1) for additional details on the choice of model hyperparameters and the sampling of β . We ran a total of 100,000 Markov chain Monte Carlo iterations per data set; we discarded the first 10,000 iterations as burn-in.

The results of the simulations are shown in Figure 2. For the data generated from the 1-component models, the posterior on the number of components concentrates around 1 in the case of Gaussian-generated data as the sample size increases (Figure 2a), whereas the posterior on the number of components diverges for the Laplace data (Figure 2d). We observe similar behavior in the 2-component case, where the posterior concentrates around the correct value in the Gaussian case (Figure 2b) but not the Laplace case (Figure 2e).

Priors that vary with N . Next, we considered the same finite Gaussian mixture model described above but with a prior that varies with the data. Specifically, for

^{*}Code available at <https://github.com/jwmi/BayesianMixtures.jl>.

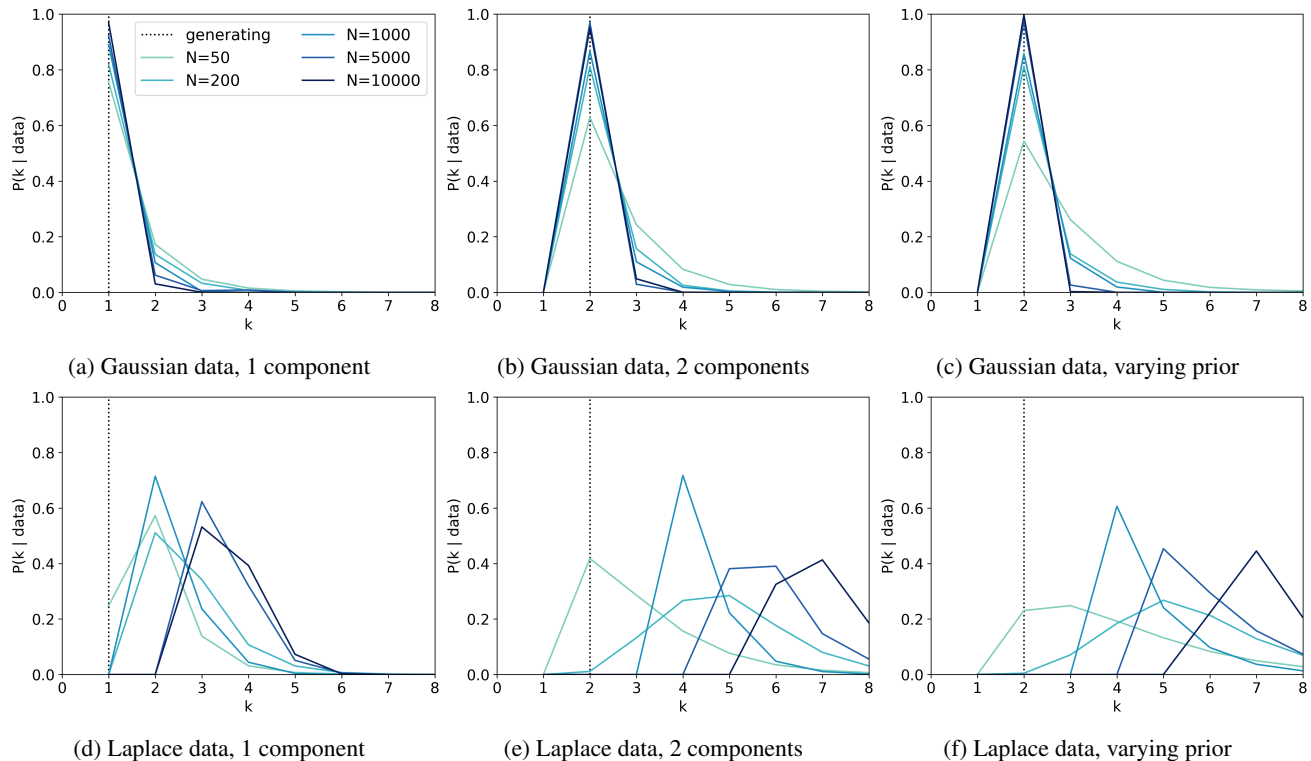


Figure 2. *Top row (well-specified)*: Posterior probability of the number of components k for Gaussian mixture models with a fixed prior fit to univariate data generated from (a, b) 1- and 2-component Gaussian mixture models, and (c) a varying prior with data generated from a 2-component Gaussian mixture model. *Bottom row (misspecified)*: Posterior probability of the number of components k for Gaussian mixture models with a fixed prior fit to univariate data generated from (c, d) 1- and 2-component Laplace mixture models, and (e) a varying prior with data generated from a 2-component Laplace mixture model.

the prior on the means, we set the hyperparameters to $m_N = \frac{1}{2}(\max_{n \in [N]} X_n + \min_{n \in [N]} X_n)$ and $\kappa_N = (\max_{n \in [N]} X_n - \min_{n \in [N]} X_n)^{-2}$, which is the setting considered by Miller and Harrison (2018, Section 7.2.1); the other hyperparameters were otherwise set to the same values above. We used the 2-component Gaussian and Laplace data sets constructed above for the fixed prior case. The bottom row of Figure 2 shows the results of the posterior number of components under this prior for the well-specified and misspecified cases; again we observe that the posterior diverges under model misspecification.

ϵ -contamination. Finally, in order to study the posterior number of components under a very slight amount of misspecification, we applied the fixed-prior Gaussian mixture model above to data generated with ϵ -contamination. That is, we generated the data according to the ϵ -contaminated distribution $f_0 = (1 - \epsilon)f + \epsilon q$, where f is a 2-component Gaussian mixture distribution with means $(5, 10)$, variances $(1, 1.5)$, and mixing weights $(0.4, 0.6)$, and q is a Laplace distribution with location 0 and scale 1. We generated two data sets: one with $\epsilon = 0.01$ and one with $\epsilon = 0.1$. In Figure 3, we observe that even under very small amounts

of misspecification, the posterior number of components diverges.

7.2. Gene expression data

Computational biologists are interested in classifying cell types by applying clustering techniques to gene expression data (Yeung et al., 2001; Medvedovic and Sivaganesan, 2002; McLachlan et al., 2002; Medvedovic et al., 2004; Rasmussen et al., 2008; de Souto et al., 2008; McNicholas and Murphy, 2010). In our next set of experiments, we apply the Gaussian finite mixture model to two gene expression data sets: (1) single-cell RNA sequencing data from mouse cortex and hippocampus cells (Zeisel et al., 2015) with the same feature selection as Prabhakaran et al. (2016) ($N = 3008$, $D = 558$, 11,000 Gibbs sampling steps with 1,000 of those as burn-in) and (2) mRNA expression data from human lung tissue (Bhattacharjee et al., 2001) ($N = 203$, $D = 1543$, and 10,000 Gibbs sampling steps with 1,000 of those burn-in). Our experiments here represent a simplified version of previous mixture model analyses for these and other related data sets (de Souto et al., 2008; Prabhakaran et al., 2016; Armstrong et al., 2001; Miller and




Figure 3. Posterior probability of the number of components k for Gaussian mixture models with a fixed prior fit to data generated from an ϵ -contaminated 2-component Gaussian mixture model, where ϵ is the proportion of data generated from a Laplace distribution.

Harrison, 2018).

As these gene expression data sets contain counts, we first transformed the data to real numerical values. In particular, we used a base-2 log transform followed by standardization—such that each dimension of the data had zero mean and unit variance—per standard practices (e.g., Miller and Harrison (2018)). Then to examine the effect of increasing data set size on inferential results, we randomly sampled subsets of increasing size without replacement; each smaller subset was contained in the next larger data set. For both data sets, we used hyperparameters $\alpha = 1$, $\beta = 1$, $m = 0$, $\kappa_{jd} = \tau_{jd}$, $r = 0.1$, and $\gamma = 1$.

For the single-cell RNAseq data set, the posterior on the number of components is shown in Figure 1a. Here the ground truth number of clusters is captured when the data set size is $N = 100$. But as predicted by our theory, as we increase the number of data points, the posterior number of components diverges.

The posterior on the number of components for the lung gene expression data is shown in Figure 1b. Again we find that on the smallest data subsets, the posterior appears to capture the ground truth number of clusters, but that as we examine more and more data, the posterior diverges.

The diagonal covariance Gaussian components are a particularly simple form of cluster shape. But no matter how complex the component model, one could wonder whether an even-more complex model might solve the issue that the number of components diverge. In the typical real-world situation that the component model cannot be specified in absolute perfection, our theory confirms that the divergence problem will remain. Thus, these examples suggest the need for more robust analyses.

8. Discussion

We have shown that the posterior distribution for the number of components in finite mixtures diverges when the mixture component family is misspecified. Since misspecification is almost unavoidable in real applications, it follows that finite mixture models are typically unreliable for estimating the number of components. In practice, our conclusion implies that inferences on the number of components can change drastically depending on the size of the data set, calling into question the usefulness of these counts in application.

Since our analysis is inherently asymptotic, it is possible that the Bayesian component-count posterior may still provide useful inferences for a finite sample—for instance if care is taken to account for the dependence of inferential conclusions on data set size. Indeed, a number of authors have recently proposed robust Bayesian inference methods to mitigate likelihood misspecification (Woo and Sriram, 2006; 2007; Rodriguez and Dunson, 2011; Grünwald and van Ommen, 2017; Miller and Dunson, 2019; Bissiri et al., 2016; Wang et al., 2017; Holmes and Walker, 2017; Jewson et al., 2018; Huggins and Miller, 2019; Knoblauch et al., 2019; Rigon et al., 2020); it remains to better understand connections between our results and these methods.

Acknowledgements

We thank Jeff Miller for helpful conversations and comments on a draft. D. Cai was supported in part by a Google Ph.D. Fellowship in Machine Learning. T. Campbell was supported by a National Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and an NSERC Discovery Launch Supplement. T. Broderick was supported in part by ONR grant N00014-17-1-2072, an MIT Lincoln Laboratory Advanced Concepts Committee Award, a Google Faculty Research Award, the CSAIL-MSR Trustworthy AI Initiative, and an ARO YIP award.

References

- B. Aragam, C. Dan, E. P. Xing, and P. Ravikumar. Identifiability of nonparametric mixture models and Bayes optimal clustering. *Annals of Statistics*, 48(4):2277–2302, 2020.
- S. A. Armstrong, J. E. Staunton, L. B. Silverman, R. Pieters, M. L. d. Boer, M. D. Minden, S. E. Sallan, E. S. Lander, T. R. Golub, and S. J. Korsmeyer. MLL translocations specify a distinct gene expression profile that distinguishes a unique leukemia. *Nature Genetics*, 30(1):41, 2001.
- J. D. Banfield and A. E. Raftery. Model-based Gaussian and non-Gaussian clustering. *Biometrics*, 49(3):803–821, 1993.
- F. Bartolucci. Clustering univariate observations via mixtures of unimodal normal mixtures. *Journal of Classification*, 22(2):203–219, 2005.
- A. Bhattacharjee, W. G. Richards, J. Staunton, C. Li, S. Monti, P. Vasa, C. Ladd, J. Beheshti, R. Bueno, M. Gillette, M. Loda, G. Weber, E. Mark, E. Lander, W. Wong, B. Johnson, T. Golub, D. Sugarbaker, and M. Meyerson. Classification of human lung carcinomas by mRNA expression profiling reveals distinct adenocarcinoma subclasses. *Proceedings of the National Academy of Sciences*, 98(24):13790–13795, 2001.
- P. G. Bissiri, C. C. Holmes, and S. G. Walker. A general framework for updating belief distributions. *Journal of the Royal Statistical Society. Series B, Statistical methodology*, 78(5):1103, 2016.
- A. Chambaz and J. Rousseau. Bounds for Bayesian order identification with application to mixtures. *The Annals of Statistics*, 36(2):938–962, 2008.
- C. Chan, F. Feng, J. Ottinger, D. Foster, M. West, and T. B. Kepler. Statistical mixture modeling for cell subtype identification in flow cytometry. *Cytometry Part A*, 73(8):693–701, 2008.
- J. Chen. Optimal rate of convergence for finite mixture models. *The Annals of Statistics*, 23(1):221–233, 1995.
- M. C. de Souto, I. G. Costa, D. S. de Araujo, T. B. Ludermir, and A. Schliep. Clustering cancer gene expression data: a comparative study. *BMC Bioinformatics*, 9(1):497, 2008.
- M. Di Zio, U. Guarnera, and R. Rocci. A mixture of mixture models for a classification problem: The unity measure error. *Computational Statistics & Data Analysis*, 51(5):2573–2585, 2007.
- S. Frühwirth-Schnatter. *Finite mixture and Markov switching models*. Springer Series in Statistics, 2006.
- S. Frühwirth-Schnatter and G. Malsiner-Walli. From here to infinity: sparse finite versus Dirichlet process mixtures in model-based clustering. *Advances in Data Analysis and Classification*, 13(1):33–64, 2019.
- J. Geng, A. Bhattacharya, and D. Pati. Probabilistic community detection with unknown number of communities. *Journal of the American Statistical Association*, 114(526):893–905, 2019.
- S. Ghosal and A. van der Vaart. *Fundamentals of Nonparametric Bayesian Inference*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2017.
- S. Ghosal, J. Ghosh, and R. Ramamoorthi. Posterior consistency of Dirichlet mixtures in density estimation. *The Annals of Statistics*, 27(1):143–158, 1999.
- J. Ghosh and R. Ramamoorthi. *Bayesian Nonparametrics*. Springer Series in Statistics, 2003.
- J. Griffié, M. Shannon, C. L. Bromley, L. Boelen, G. L. Burn, D. J. Williamson, N. A. Heard, A. P. Cope, D. M. Owen, and P. Rubin-Delanchy. A Bayesian cluster analysis method for single-molecule localization microscopy data. *Nature Protocols*, 11(12):2499, 2016.
- P. Grünwald and T. v. Ommen. Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it. *Bayesian Analysis*, 12(4):1069–1103, 2017.
- A. Guha, N. Ho, and X. Nguyen. On posterior contraction of parameters and interpretability in Bayesian mixture modeling. *arXiv preprint arXiv:1901.05078*, 2019.
- P. Heinrich and J. Kahn. Strong identifiability and optimal minimax rates for finite mixture estimation. *The Annals of Statistics*, 46(6A):2844–2870, 2018.
- 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.
- C. Holmes and S. Walker. Assigning a value to a power likelihood in a general Bayesian model. *Biometrika*, 104(2):497–503, 2017.
- G. Hu, H.-C. Yang, and Y. Xue. Bayesian group learning for shot selection of professional basketball players. *Stat*, page e324, 2020.
- J. P. Huelsenbeck and P. Andolfatto. Inference of population structure under a Dirichlet process model. *Genetics*, 175(4):1787–1802, 2007.

- J. H. Huggins and J. W. Miller. Using bagged posteriors for robust inference and model criticism. *arXiv preprint 1912.07104*, 2019.
- H. Ishwaran, L. F. James, and J. Sun. Bayesian model selection in finite mixtures by marginal density decompositions. *Journal of the American Statistical Association*, 96(456):1316–1332, 2001.
- S. Jain and R. M. Neal. A split-merge Markov chain Monte Carlo procedure for the Dirichlet process mixture model. *Journal of Computational and Graphical Statistics*, 13(1):158–182, 2004.
- J. Jewson, J. Q. Smith, and C. Holmes. Principles of Bayesian inference using general divergence criteria. *Entropy*, 20(6):442, 2018.
- J. Knoblauch, J. Jewson, and T. Damoulas. Generalized variational inference: Three arguments for deriving new posteriors. *arXiv preprint arXiv:1904.02063*, 2019.
- W. Kruijjer, J. Rousseau, and A. van der Vaart. Adaptive Bayesian density estimation with location-scale mixtures. *Electronic Journal of Statistics*, 4:1225–1257, 2010.
- S. Legramanti, T. Rigon, D. Durante, and D. B. Dunson. Extended stochastic block models with application to criminal networks. *arXiv preprint arXiv:2007.08569*, 2020.
- A. Lijoi, I. Prünster, and S. Walker. Extending Doob’s consistency theorem to nonparametric densities. *Bernoulli*, 10(4):651–663, 2004.
- E. D. Lorenzen, P. Arctander, and H. R. Siegismund. Regional genetic structuring and evolutionary history of the impala *aepyceros melampus*. *Journal of Heredity*, 97(2):119–132, 2006.
- G. Malsiner-Walli, S. Frühwirth-Schnatter, and B. Grün. Model-based clustering based on sparse finite Gaussian mixtures. *Statistics and Computing*, 26(1-2):303–324, 2016.
- G. Malsiner-Walli, S. Frühwirth-Schnatter, and B. Grün. Identifying mixtures of mixtures using Bayesian estimation. *Journal of Computational and Graphical Statistics*, 26(2):285–295, 2017.
- G. J. McLachlan, R. Bean, and D. Peel. A mixture model-based approach to the clustering of microarray expression data. *Bioinformatics*, 18(3):413–422, 2002.
- P. D. McNicholas and T. B. Murphy. Model-based clustering of microarray expression data via latent gaussian mixture models. *Bioinformatics*, 26(21):2705–2712, 2010.
- M. Medvedovic and S. Sivaganesan. Bayesian infinite mixture model based clustering of gene expression profiles. *Bioinformatics*, 18(9):1194–1206, 2002.
- M. Medvedovic, K. Y. Yeung, and R. E. Bumgarner. Bayesian mixture model based clustering of replicated microarray data. *Bioinformatics*, 20(8):1222–1232, 2004.
- J. W. Miller and D. B. Dunson. Robust Bayesian inference via coarsening. *Journal of the American Statistical Association*, 114(527):1113–1125, 2019.
- J. W. Miller and M. T. Harrison. A simple example of Dirichlet process mixture inconsistency for the number of components. In *Advances in Neural Information Processing Systems*, pages 199–206, 2013.
- J. W. Miller and M. T. Harrison. Inconsistency of Pitman-Yor process mixtures for the number of components. *Journal of Machine Learning Research*, 15(1):3333–3370, 2014.
- J. W. Miller and M. T. Harrison. Mixture models with a prior on the number of components. *Journal of the American Statistical Association*, 113(521):340–356, 2018.
- S. Mukherjee, E. D. Feigelson, G. J. Babu, F. Murtagh, C. Fraley, and A. Raftery. Three types of gamma-ray bursts. *The Astrophysical Journal*, 508(1):314, 1998.
- R. M. Neal. Markov chain sampling methods for Dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, 9(2):249–265, 2000.
- X. Nguyen. Convergence of latent mixing measures in finite and infinite mixture models. *The Annals of Statistics*, 41(1):370–400, 2013.
- A. Nobile. *Bayesian analysis of finite mixture distributions*. PhD thesis, Carnegie Mellon University, 1994.
- E. Otranto and G. M. Gallo. A nonparametric Bayesian approach to detect the number of regimes in Markov switching models. *Econometric Reviews*, 21(4):477–496, 2002.
- J. Pella and M. Masuda. The Gibbs and split merge sampler for population mixture analysis from genetic data with incomplete baselines. *Canadian Journal of Fisheries and Aquatic Sciences*, 63(3):576–596, 2006.
- F. Petralia, V. Rao, and D. B. Dunson. Repulsive mixtures. In *Advances in Neural Information Processing Systems*, pages 1889–1897, 2012.
- S. Prabhakaran, E. Azizi, A. Carr, and D. Pe’er. Dirichlet process mixture model for correcting technical variation in single-cell gene expression data. In *International Conference on Machine Learning*, pages 1070–1079, 2016.

- J. K. Pritchard, M. Stephens, and P. Donnelly. Inference of population structure using multilocus genotype data. *Genetics*, 155(2):945–959, 2000.
- C. Rasmussen, B. de la Cruz, Z. Ghahramani, and D. Wild. Modeling and visualizing uncertainty in gene expression clusters using Dirichlet process mixtures. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 6(4):615–628, 2008.
- S. Richardson and P. J. Green. On Bayesian analysis of mixtures with an unknown number of components (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(4):731–792, 1997.
- T. Rigon, A. H. Herring, and D. B. Dunson. A generalized Bayes framework for probabilistic clustering. *arXiv preprint arXiv:2006.05451*, 2020.
- A. Rodriguez and D. B. Dunson. Nonparametric Bayesian models through probit stick-breaking processes. *Bayesian Analysis*, 6(1):145–178, 2011.
- K. Roeder and L. Wasserman. Practical Bayesian density estimation using mixtures of normals. *Journal of the American Statistical Association*, 92(439):894–902, 1997.
- J. Rousseau and K. Mengersen. Asymptotic behaviour of the posterior distribution in overfitted mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(5):689–710, 2011.
- P. Rubin-Delanchy, G. L. Burn, J. Grifflé, D. J. Williamson, N. A. Heard, A. P. Cope, and D. M. Owen. Bayesian cluster identification in single-molecule localization microscopy data. *Nature methods*, 12(11):1072–1076, 2015.
- W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 1976.
- L. Schwartz. On Bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie*, 4:10–26, 1965.
- H. Teicher. Identifiability of mixtures. *The Annals of Mathematical Statistics*, 32(1):244–248, 1961.
- H. Teicher. Identifiability of finite mixtures. *The Annals of Mathematical Statistics*, pages 1265–1269, 1963.
- F. K. Teklehaymanot, A.-K. Seifert, M. Muma, M. G. Amin, and A. M. Zoubir. Bayesian target enumeration and labeling using radar data of human gait. In *26th European Signal Processing Conference*, pages 1341–1346. IEEE, 2018.
- S. T. Tokdar. Posterior consistency of Dirichlet location-scale mixture of normals in density estimation and regression. *Sankhyā: The Indian Journal of Statistics*, pages 90–110, 2006.
- G. Tonkin-Hill, J. A. Lees, S. D. Bentley, S. D. Frost, and J. Corander. Fast hierarchical Bayesian analysis of population structure. *Nucleic Acids Research*, 47(11):5539–5549, 2019.
- Y. Wang, A. Kucukelbir, and D. M. Blei. Reweighted data for robust probabilistic models. In *International Conference on Machine Learning*, page 3646–3655, 2017.
- M.-J. Woo and T. Sriram. Robust estimation of mixture complexity. *Journal of the American Statistical Association*, 101(476):1475–1486, 2006.
- M.-J. Woo and T. Sriram. Robust estimation of mixture complexity for count data. *Computational Statistics & Data Analysis*, 51(9):4379–4392, 2007.
- Y. Wu and S. Ghosal. Kullback Leibler property of kernel mixture priors in Bayesian density estimation. *Electronic Journal of Statistics*, 2:298–331, 2008.
- E. P. Xing, K.-A. Sohn, M. I. Jordan, and Y.-W. Teh. Bayesian multi-population haplotype inference via a hierarchical Dirichlet process mixture. In *International Conference on Machine Learning*, pages 1049–1056, 2006.
- S. J. Yakowitz and J. D. Spragins. On the identifiability of finite mixtures. *The Annals of Mathematical Statistics*, 39(1):209–214, 1968.
- K. Y. Yeung, D. R. Haynor, and W. L. Ruzzo. Validating clustering for gene expression data. *Bioinformatics*, 17(4):309–318, 2001.
- A. Zeisel, A. B. Muñoz-Manchado, S. Codeluppi, P. Lönnerberg, G. La Manno, A. Juréus, S. Marques, H. Munguba, L. He, C. Betscholtz, C. Rolny, G. Castelo-Branco, J. Hjerling-Leffler, and S. Linnarsson. Cell types in the mouse cortex and hippocampus revealed by single-cell RNA-seq. *Science*, 347(6226):1138–1142, 2015.
- Z. Zhang, G. Zhang, H. Goyal, L. Mo, and Y. Hong. Identification of subclasses of sepsis that showed different clinical outcomes and responses to amount of fluid resuscitation: a latent profile analysis. *Critical Care*, 22(1): 1–11, 2018.