
Continuously Parameterized Mixture Models

Christopher M. Bender¹ Yifeng Shi¹ Marc Niethammer¹ Junier B. Oliva¹

Abstract

Mixture models are universal approximators of smooth densities but are difficult to utilize in complicated datasets due to restrictions on typically available modes and challenges with initializations. We show that by continuously parameterizing a mixture of factor analyzers using a learned ordinary differential equation, we can improve the fit of mixture models over direct methods. Once trained, the mixture components can be extracted and the neural ODE can be discarded, leaving us with an effective, but low-resource model. We additionally explore the use of a training curriculum from an easy-to-model latent space extracted from a normalizing flow to the more complex input space and show that the smooth curriculum helps to stabilize and improve results with and without the continuous parameterization. Finally, we introduce a hierarchical version of the model to enable more flexible, robust classification and clustering, and show substantial improvements against traditional parameterizations of GMMs.

1. Introduction

Mixture models have long served as reliable tools for both modeling (density estimation) and summarizing (cluster analysis) datasets. Through their probabilistic nature, mixture models provide an estimate of the underlying data distribution; through a parsimonious collection of components, mixture models succinctly summarize the major patterns found in a dataset. Gaussian mixture models (GMMs), which provide a universal approximation for smooth densities (Goodfellow et al., 2015), have been effectively deployed for modeling and clustering in a wide range of applications (Böhning et al., 2007; Gao et al., 2010; Jiao et al., 2020).

¹Department of Computer Science, The University of North Carolina, Chapel Hill, North Carolina, USA. Correspondence to: Christopher M. Bender <bender@cs.unc.edu>.

There are two major challenges in applying GMMs to complicated, high dimensional data. First, the partitions (modes) that we wish to characterize in data are rarely spherical or simple enough in nature to be faithfully captured with Gaussian components. As an alternative, one may consider more complicated components (Holzmann et al., 2006), however, this worsens identifiability issues and yields partitions that do not adequately summarize populations of interest in the data (Böhning et al., 2007). Second, mixture models converge to local minima and thus their optimization is *extremely* sensitive to initialization. Some approaches propose to initialize via local fitting methods such as k-means (Richardson & Weiss, 2018). However, the distance metrics, for example Euclidean distance, implied in those local fitting methods is rarely appropriate in high dimensions, imposing initial partitioning that is arduous to escape.

We propose to offset the first difficulty by taking a mixture of mixtures in a hierarchical approach by considering a *mixture of distinct dynamical systems*. In effect, we propose to learn to evolve the components of one partition (see Fig. 1(a)) in a fashion that results in an infinite mixture model in the limit. This alleviates the burden of learning a large number of separate modes through a shared generator of components in an approach that is akin to hypernetworks (Ha et al., 2016). We can then combine multiple partitions to cover the entire support (see Fig. 1(b)). Specifically, each partition is itself defined by multiple components that are spanned through smooth dynamics to represent complex, multi-modal distributions. We sample discrete components from this continuous parameterization during inference to construct a cheap, lightweight model that dramatically reduces the computational cost, maintains improved performance over typical mixture methods, and simplifies model interpretation.

We alleviate the second difficulty by creating a learned curriculum (Bengio et al., 2009) over the data. We construct the curriculum by transforming the data into a standard Gaussian through a normalizing flow (Chen et al., 2019b; Dinh et al., 2017; Grathwohl et al., 2019; Kingma & Dhariwal, 2018) and slowly annealing back to the original space. This removes the difficulty with initialization since the distribution early in the training process is *known*. As we advance through the curriculum, the model only requires minor perturbations from its previous state. We find that this process results in considerably improved performance regardless of

the mixture parameterization.

Main Contributions We propose a parameterization of Gaussian mixture models through the output of a neural ordinary differential equation (NODE). This formulation induces a smoothly-varying trajectory through the data, in essence, learning a probabilistic sheath across the data manifold. Once trained, the components of the GMM can be extracted and the NODE discarded, rendering storage and inference notably cheaper than most other modern methods. We include a hierarchical component to the model by considering multiple trajectory starting points to allow us to utilize simple Bayesian methods for clustering or classification. Finally, we demonstrate a curriculum-based training method to significantly improve model performance. Code can be found at <https://github.com/lupalab/cpmm>.

2. Background

2.1. Gaussian Mixture Models

Traditional (finite) mixture models model the presence of sub-populations within a dataset through a convex combination of distinct components of a chosen distribution. Gaussian mixture models (GMMs), the most common class of mixture models, restrain the set of chosen distributions that constitutes the mixture to the Gaussian family:

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mu_k, \Sigma_k), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^M$ is the input data, π is the component weight, $\mu \in \mathbb{R}^M$ is the mean, $\Sigma \in \mathbb{R}^{M \times M}$ is the covariance matrix, and K is the number of components in the mixture. Despite the seemingly stringent simplification GMMs make, they encompass a broad set of distributions. In effect, they are universal approximators to any smooth probability distribution given enough mixture components (Goodfellow et al., 2015). This offers us an insight when deploying GMMs for data modeling: the more complex the data distribution is, the more components the GMM would likely require to yield a reasonable approximation. Motivated by this insight, analogous to defining an integral as an infinite sum, we take the limit $K \rightarrow \infty$ in Eq. 1 to arrive at the following continuous representation

$$p(\mathbf{x}) = (2\pi)^{-M/2} \times \int_0^1 \pi(s) |\Sigma(s)|^{-0.5} \exp\left(-(\mathbf{x}-\mu(s))^T \Sigma^{-1}(s) (\mathbf{x}-\mu(s))/2\right) ds \quad (2)$$

where the set of parameters, $\{\pi_k, \mu_k, \Sigma_k\}_k$, for the finite GMMs become the set of functions, $\{\pi(s), \mu(s), \Sigma(s)\}$, that is parametrized by the variable s whose meaning depends on the specific context of the problem. Eq. 1 can then be reversely seen as a discretized formulation of Eq. 2.

2.1.1. MIXTURE OF FACTOR ANALYZERS

General GMMs require $\mathcal{O}(M^2)$ parameters for each component. To reduce this parameter cost, we utilize a Mixture of Factor Analyzers (Ghahramani & Hinton, 1996; Richardson & Weiss, 2018) which parameterizes the covariance matrix of each component as the sum of a diagonal matrix, D , and the inner product of a low-rank matrix, $A \in \mathbb{R}^{M \times R}$, with itself where $R \ll M$, such that $\Sigma = D + AA^T$. This brings the number of parameters down to $\mathcal{O}(MR)$ and allows for $\mathcal{O}(MR^2)$ cost to estimate the determinant and inverse per component.

2.2. Neural Ordinary Differential Equations

Built upon the observation that a residual network (He et al., 2016) can be seen as an Euler discretization for solving an ODE, Chen et al. (Chen et al., 2018) introduced *neural ODEs* (NODEs), a class of continuous-depth neural networks that parameterize the class of ODEs that aim to solve the initial value problem:

$$\frac{dh(t)}{dt} = f(h(t), t; \theta), \quad h(t_0) = h_{t_0}, \quad (3)$$

with known initial value h_{t_0} . Instead of being restricted to the Euler discretization, Neural ODEs are able to incorporate any black-box ODE solver in a memory efficient manner to compute

$$h(t_{i+1}) = h(t_i) + \int_{t_i}^{t_{i+1}} f(h(t), t; \theta) dt. \quad (4)$$

In practice, for tabular data, we set the architecture of f to be a multilayer perceptron with an explicit dependence on t and use either \tanh or \sin as the activation function. We find that when using \sin , the ODE seems to be less prone to becoming stiff (Blanchard et al., 2006) and hence is easier to numerically integrate. When using image data, we utilize a convolutional neural network with the same activation functions. In practice, we found that, in conjunction with gradient check pointing, our models required only 1-2 GB of GPU RAM to train.

3. Methods

In this section we discuss how we utilize neural ordinary differential equations (NODE) to parameterize mixture models over continuous indices and how to improve the training process by instituting a curriculum to guide the model from an easy-to-learn space to the true data space. The NODE allows for a hyper-network (Ha et al., 2016) type approach, which shares weights of a network to output parameters over multiple components, and whose limit is an infinite GMM. We then extend a continuously parameterized mixture model by allowing for multiple trajectories in a hierarchical structure that allows us to perform clustering or classification.

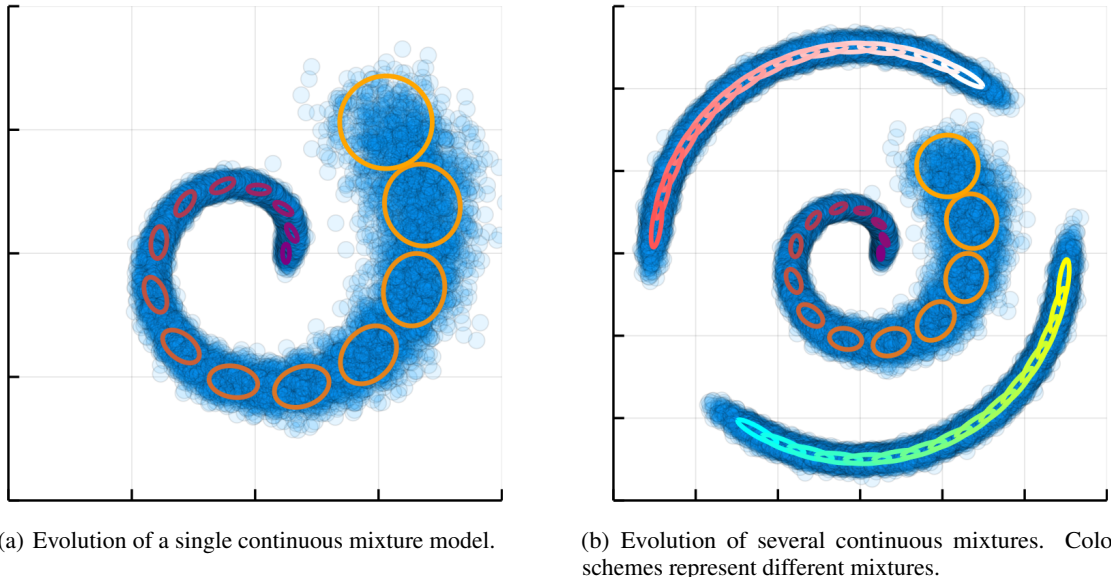


Figure 1. We illustrate how a continuously parameterized mixture model can evolve components to model the data distribution (samples in blue). Each ellipse corresponds to equal likelihood contours for a different component. Colors across components implicitly indicates the arrow of pseudotime. (a) A single *trajectory* through the space provides a smooth probabilistic model. (b) A hierarchy of three different partitions allows for different manifolds and enables flexible clustering or classification.

3.1. Continuously Parameterized Mixture Models

Inspired by the universal approximator properties of Eq. 2, where the parameters of each component depend on the latent variable, s , we chose to parameterize the MFA via a neural network. We begin by constructing a joint state across the MFA parameters. When the data is tabular, this amounts to flattening each parameter and concatenating them together:

$$h(t) = [\mu(t), \bar{A}(t), \log(d)(t), \log(\pi)(t)] \quad (5)$$

where the bar over A indicates the matrix has been flattened, d is the diagonal portion of D , and \log is applied element-wise. We choose to use a shared state across parameters rather than a separate ODE for each parameter so that the model can better coordinate and share relevant information. We learn the \log of d and π instead of the parameters directly to ensure that the covariance is positive definite and that the weights are non-negative. As a result, $h \in \mathbb{R}^J$ where $J = (R + 2)M + 1$. In most cases, we set the value of π to a constant and exclude it from the state to encourage contiguous modes.

There are two ways to proceed. If our goal was to best match Eq. 2, we would construct the NODE based on the joint state and solve the system of ODEs that includes the

continuous mixture

$$\begin{aligned} \frac{dh(t)}{dt} &= f(h(t), t; \theta) \\ \frac{dp(x)}{dt} &= (2\pi)^{-M/2} |\Sigma(t)|^{-0.5} \times \\ &\quad \exp\left(-(\mathbf{x} - \mu(t))^T \Sigma^{-1}(t) (\mathbf{x} - \mu(t)) / 2\right) \pi(t) \end{aligned}$$

where we evaluate each component in accordance with Eq. 2 on the instantaneous value of t within the ODE solver and set $p(x) = 0$ at $t = 0$ and take the integral over the range of t without producing intermediate values. Unfortunately, while this formulation is a better match to the continuous mixture model, integrating $p(x)$ is numerically unstable and often causes f to become stiff (Blanchard et al., 2006) or for the integrator to underflow, see Appendix for additional details. Additionally, this form would require that we keep the NODE once training is complete and resolve it for every new example encountered.

We instead choose to solve for the joint state directly and return the parameters at a (possibly variable) number of pseudotimes. That is, for a (uniformly drawn) set of pseudotimes, $\{t_j\}_{j=1}^G$, we model the density as $p(\mathbf{x}) = \sum_{j=1}^G \pi(t_j) \mathcal{N}(\mathbf{x}; \mu(t_j), A(t_j)A(t_j)^T + D(t_j))$, where $\pi(t_j)$, $\mu(t_j)$, $A(t_j)$, $D(t_j)$ are the respective continuously indexed parameters and $\sum_j \pi(t_j) = 1$. This approach is essentially doing numerical integration to approximate Eq. 2, which is feasible given the $1d$ integral. This leaves us with a *mixture model* whose parameters are derived from

a *continuous process*. We will refer to this model as a continuously parameterized mixture model (CPMM). Figure 1 illustrates a dataset overlaid with the components extracted from a CPMM.

We have specifically chosen f so that the ODE is not autonomous (Blanchard et al., 2006). From a practical perspective, this means that our models are capable of learning loops and cycles. However, to further increase the flexibility of the model, we additionally augment the joint state

$$h(t) = [\mu(t), \bar{A}(t), \log(d)(t), \alpha(t)] \quad (6)$$

where $\alpha \in \mathbb{R}^L$ and we have excluded $\log(\pi)$ from the state. α is not directly used in evaluating any portion of the mixture but instead provides an unconstrained pathway that the network can use to pass information along. Without the augmented state, the network must encode information relevant to future times into the parameters of the current time, overloading those parameters and notably decreasing performance. We generally choose L to be $2-4 \times$ larger than J , so $h \in \mathbb{R}^{5J}$.

In this form, the NODE does not depend on the evaluation data point, \mathbf{x} , directly and can be solved independently. In fact, the NODE can be evaluated with only the initial value and the parameter times. We can therefore consider that our model is a type of hypernetwork (Ha et al., 2016). We train the model by evaluating/solving the hypernetwork against the *learnable* initial state and then evaluating the likelihood for each example in the batch. The initialization and hypernetwork dynamics are then updated to maximize the likelihood of the batch. During evaluation, we solve the hypernetwork *once* and cache the parameters (for a uniformly drawn set of pseudotimes $\{t_j\}_{j=1}^G$) against future executions. This places the total computational cost for inference at $\mathcal{O}(KMR)$.

When the input data is an image, we construct the joint (augmented) state by keeping the parameters in the image shape for the NODE and concatenating over channels. We then flatten the data and parameters to calculate the likelihood. This means that the unaugmented portion of the channel dimension is increased by a factor of $R + 2$ relative to the image channel count (when π is held fixed).

3.2. Hierarchical Mixture of Factor Analyzers

To facilitate clustering and classification, we extend the mixture of factor analyzers (MFAs) in a hierarchical manner. More specifically, we model the data with a *mixture of MFAs*

$$\begin{aligned} p(\mathbf{x}) &= \sum_{c=1}^C \eta_c p(\mathbf{x}|c) \\ &= \sum_{c=1}^C \eta_c \sum_{j=1}^G \left(\pi_c(t_j) \mathcal{N}(\mathbf{x}; \mu_c(t_j), A_c(t_j)A_c(t_j)^T + D_c(t_j)) \right), \end{aligned} \quad (7)$$

where C denotes the number of MFAs and, therefore the number of clusters/classes, $\sum_c \eta_c = 1$, and $\{t_j\}_{j=1}^G$ is a set of pseudotimes. We construct the mixture of the C MFAs by providing C different initial values to the NODE. Alternatively, we could learn a separate NODE and initial state for each MFA. This should allow for greater flexibility in each trajectory. However, we generally find that the improvement is not appreciable while the increase in compute is significant. Sharing the ODE means that the trajectories will learn from one another and share certain dynamic characteristics.

The hierarchical mixture allows the model to learn disparate data partitions without requiring that a single trajectory transition through low probability regions. Additionally, we can utilize Bayes rule to estimate $p(y = c | \mathbf{x})$ from Eq. 7 to perform clustering or classification.

3.3. Curriculum Through Spaces

Unfortunately, training large mixture models in high dimensions often gets stuck in local minima and is *extremely* sensitive to initialization. Direct mixtures (mixtures with disconnected parameters) are often initialized via a combination of kmeans and local fitting prior to gradient descent (Richardson & Weiss, 2018). However, for CPMMs, the parameters are the product of a learnable process and cannot be directly initialized. As a heuristic, one may initialize the ODE initial value based on kmeans or labeled examples, however this only provides a limited signal and one must resolve how to initialize other parameters such as A or $\log d$.

In order to circumvent this difficulty, we borrow ideas from curriculum learning (Bengio et al., 2009). We begin by defining a continuously indexible map, $\forall v \in [0, 1]$, $\mathcal{M}_v : \mathbb{R}^M \mapsto \mathbb{R}^M$ such that $\mathcal{M}_0(\mathbf{x}) = \mathbf{x}$, and \mathcal{M}_v progressively maps to a smoother, simpler space as v increases. We propose to leverage \mathcal{M} to construct a curriculum for learning the CPMM. Intuitively, we may begin training on the simplest space induced by \mathcal{M}_1 and slowly progress to training on our target input space \mathcal{M}_0 . That is, for a sequence $1 = v_1 > \dots > v_T = 0$, we progressively train on respective datasets $\mathcal{D}_t = \{\mathcal{M}_t(\mathbf{x}_i)\}_{i=1}^N$. An accessible choice is to define \mathcal{M}_v as a continuous normalizing flow (Grathwohl et al., 2019), $\mathcal{M}_v = \mathbf{u}_v$, where

$$\mathbf{u}_v(\mathbf{x}) = \int_0^v g(\mathbf{u}_s(\mathbf{x}), s; \phi) ds,$$

$$\log p_{\mathbf{u}_v}(\mathbf{u}_v(\mathbf{x})) = \log \mathcal{N}(\mathbf{u}_1(\mathbf{x}); 0, I) + \int_v^1 \text{tr} \left(\frac{\partial g}{\partial s} \right) ds, \quad (8)$$

$\mathbf{u}_0(\mathbf{x}) \equiv \mathbf{x}$ and $\mathbf{u}_v(\mathbf{x})$ is defined by a learnable function, g , which is trained via MLE so that $\mathbf{u}_1(\mathbf{x}) \sim \mathcal{N}(0, I)$.

We construct the curriculum for the CPMM by first training the CPMM on \mathbf{u}_1 . Since the data is (ideally) a standard Gaussian in this space, it is trivial to learn a good CPMM




Figure 2. Evenly distributed transitions through different spaces. The data begins in a latent space as a standard Gaussian (top left) and ends in the true space (bottom right).

and the initialization of the CPMM parameters can be created as perturbations from the standard Gaussian. We then marginally perturb the space by taking a small step in v towards the input, e.g., we train the CPMM on $\mathbf{u}_{1-\epsilon}$ where, for sufficiently small ϵ , $p_{\mathbf{u}_{1-\epsilon}}(\mathbf{u}_{1-\epsilon}) \approx p_{\mathbf{u}_1}(\mathbf{u}_1)$. We repeat this process of perturbing the space and then updating the CPMM until we have arrived back at the input space. Figure 2 shows this smooth transition on a two-dimensional dataset as the map transitions between the latent space, through several intermediate spaces, before arriving back at the input space.

One interpretation of this curriculum is that we learn the initialization for one space based on training over a marginally simpler version of the data. This interpretation applies to both the parameters of the NODE hypernetwork and the initial values of each trajectory. Without this slow update procedure, even with a good initial value, the CPMM can be hard to train since the initial trajectories can point in the wrong directions and the model can struggle to shift the probability mass appropriately.

An important distinction between our use of a curriculum and the typical form of curriculum learning is that we do not want to remember earlier “tasks.” Doing so would mean that the CPMM would still capture the intermediate spaces between the Gaussian noise and the input space. Since these spaces are only used as a guide and have no intrinsic meaning, maintaining them has no direct value.

4. Related Work

Mixture Models Mixture models can be applied to solve an array of ML problems, e.g., clustering and density es-

timization (Hastie et al., 2001), and are widely deployed in modern ML methods. Viroli et al. (Viroli & McLachlan, 2019) model the variables at each layer of a deep network with a Gaussian mixture model (GMM), leading to a set of nested mixtures of linear models. Izmailov et al. (Izmailov et al., 2020) use a GMM as the base distribution in conjunction with normalizing flows for semi-supervised learning. Richardson et al. (Richardson & Weiss, 2018) demonstrate that GMMs better capture the statistical modes of the data distribution than generative adversarial models (GANs), while Eghbal-zadeh et al. (Eghbal-zadeh et al., 2019) incorporate a GMM into the discriminator of a GAN to encourage the generator to exploit different modes in the data.

Tractable Likelihood Models Normalizing flows and autoregressive models are two common types of extremely effective tractable likelihood estimators. Normalizing flows (NF) learn invertible transformations to a latent space where the data has a known, prechosen, distribution, typically the standard normal. There exist considerable variations between models based on the families of invertible functions allowed (Chen et al., 2019b; Dinh et al., 2017; Grathwohl et al., 2019; Kingma & Dhariwal, 2018). Autoregressive (AR) models (Oliva et al., 2018; Papamakarios et al., 2017; van den Oord et al., 2016) exploit the probabilistic chain rule and learn a distribution for each dimension conditioned on the previous features. Despite their impressive performance as likelihood estimators and as generative methods, these models are surprisingly brittle. In particular, they show higher likelihoods for completely different datasets than the ones they were trained on which prevents them from being utilized for outlier detection (Hendrycks et al.,

2019; Nalisnick et al., 2019).

Deep Clustering Several approaches to apply deep networks for clustering have been proposed in the past few years (Caron et al., 2018; Xie et al., 2016), centering around the concept that the input space in which traditional clustering algorithms operate is of importance. There have also been works on incorporating traditional clustering methods, such as spectral clustering or hierarchical clustering, directly into deep networks (Chen et al., 2019a; Law et al., 2017; Shaham et al., 2018). Mukherjee et al. (Mukherjee et al., 2019) extend GANs for clustering by using a combination of discrete and continuous latent variables.

Mixture Models + Deep Nets for Clustering Since mixture models are the traditional models of choice for clustering tasks, arming them further with the recent development in deep learning is only natural. Zhang et al. (Zhang et al., 2017) directly deployed a mixture of autoencoders with the assumption that different clusters are effectively different local data manifolds, and thus could be parametrized and learned by autoencoders. In the same vein, Pires et al. (Pires & Figueiredo, 2020) model the data using a mixture of normalizing flows, where the mixture weights are computed through optimizing the variational lower bound. Jiang et al. (Jiang et al., 2017) use a GMM as the prior distribution for the latent code in a variational auto-encoder (VAE), thus allowing for clustering of the input data in the latent space.

Non-parametric Bayesian Mixture Models Non-parametric Bayesian models (Müller & Quintana, 2004) assume the number of parameters of the underlying model grows with the number of data samples. In the case of the underlying model being a non-parametric Bayesian mixture model, this indicates that for every new data sample, the model can either group it with the already-discovered mixture components or initiate a new mixture component for that data sample, increasing the number of mixture components deployed and hence the number of parameters utilized (Görür & Rasmussen, 2010). This is accomplished by using a Dirichlet process for the prior distribution of the parameters of non-parametric Bayesian mixture models (Gelfand et al., 2005). Despite effectively deploying an infinite number of mixture components during the training stage, the number of components eventually learned is determined by the finite data samples available, resulting in a learned finite mixture distribution at test time.

Sum-Product Networks Sum-product networks (SPN) (Peharz et al., 2020a;b; Poon & Domingos, 2011) are a class of probabilistic circuits (Darwiche, 2003) that produce a tractable likelihood estimate akin to normalizing flows and autoregressive methods. SPNs are constructed as a directed acyclic graph composed of alternating “sum” and

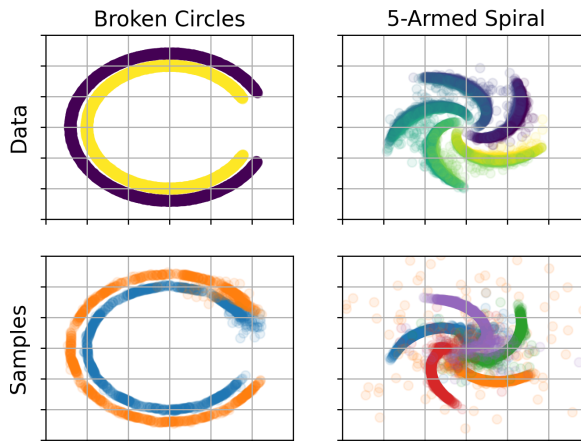


Figure 3. Two synthetic datasets designed to assess the limitations of the CPMM.

“product” operations over the evaluation of (typically) one-dimensional densities at each leaf node. Learning is often performed through the EM algorithm by updating the parameters of each leaf distribution and the weights at each sum node. This structure allows SPNs to cheaply evaluate a variety of statistical quantities (e.g., marginals) without approximation. When leaf densities are Normal distributions, the SPN can be expressed as a GMM with many components (Jaini et al., 2018).

5. Experiments

We construct our models in PyTorch (Paszke et al., 2019) and train using PyTorch Lightning (Falcon, 2019). We used Adam (Kingma & Ba, 2015) as the optimizer in all cases. Unless otherwise stated, we extract 25 components per trajectory for each hierarchical CPMM. Curriculum models were constructed using simple CNFs from FFJORD (Grathwohl et al., 2019). Since the execution of a CNF is computationally expensive, we extracted the smoothly varying datasets (see Sec. 3.3), \mathcal{D}_t , immediately after training and saved them to disk. In general, we choose to use 51 total spaces (including the latent space, \mathbf{u}_1 , and the input space, \mathbf{x}) and allocate one epoch per space. We then fine-tune on the input space. All models trained directly in the input space utilize standard data augmentations (see Appendix for further details). Models trained with the curriculum are initialized randomly; other models are initialized via kmeans or a combination of kmeans and fitting a local Factor Analyzer as in (Richardson & Weiss, 2018). When training a hierarchical model, we often find it helpful to include entropy regularizations across trajectories.

5.1. Synthetic Data

We demonstrate the effectiveness and weaknesses of CPMM on two toy datasets. The first dataset consists of two con-




Figure 4. Means from a hierarchical CPMM trained on MNIST with a curriculum. Each trajectory is evaluated at evenly spaced pseudotimes between zero and one and displayed from left to right. Despite the 3/8 and 4/9 confusion, the model does an excellent job of learning smoothly-varying transitions. The initial component (far left cases) are never the most likely component and could be discarded.

centric, broken circles. The radii of both circles are chosen to differ by 10% and have similar thicknesses. The second dataset consists of a noisy five-armed spiral that begins near the origin before curving outward. This dataset is intentionally constructed with “outliers” (points between the arms without a clear cluster identity) to assess how the CPMM will handle examples “far” from the manifold. Both CPMMs are trained using a curriculum.

Figure 3 contains true data points in the top row and samples drawn from the CPMM on the bottom row. Since the model is trained unsupervised, we use different colors between true labels and cluster labels (e.g., between the top and bottom row). We see that the CPMM does a reasonable job of capturing the general manifolds of both datasets. However, the model seems to struggle with the end points of the trajectory and the samples are less precise than during the other portions of the trajectory. We observe this trend on all datasets: the CPMM often struggles with endpoints, typically the early pseudotimes. In the spiral data, we see that the model has attempted to capture the outliers, though they have larger spread than the true outliers and it attributes them all to the same cluster.

5.2. Images

To process image datasets, we first rescale the input pixels to be between zero and one and then apply an element-wise logit transform, a standard preprocessing technique (e.g., (Grathwohl et al., 2019)). We perform these transformations so that the “input” data has support over \mathbb{R} . These transformations are applied before we train any model and the corresponding $\log \det \text{Jacobian}$ is accounted for

when estimating bits per dimension as in a normalizing flow.

We test training CPMMs on MNIST (Deng, 2012) and Fashion-MNIST (Xiao et al., 2017) and compare to two different GMM baselines along with several standard likelihood models and clustering models. CPMM NODEs are constructed using CNNs with a depth of 3, a hidden channel width of 64, and utilize \sin activations. We additionally augment the state space by a factor of four and explicitly condition the ODEs on pseudotime by concatenating it as an extra channel. To avoid degeneracies, we soft-clip $\log d$ so that the minimum value cannot be less than -6.

We consider two baseline GMMs for comparison. In the first model, we choose the total number of components equal to the number of clusters (e.g., 10). This simple procedure allows for easy cluster assignments; each component corresponds to exactly one cluster. In the second case, we train a GMM with 250 components since this corresponds exactly to the number of components extracted across trajectories when using a CPMM. However, attributing components to clusters is less obvious in this case. We choose to use spectral clustering (Filippone et al., 2008) over components based on the Fréchet distance over Gaussians (Dowson & Landau, 1982). Specifically, we attribute examples to components and components to clusters based on a spectral clustering model that constructs an affinity matrix through the Fréchet distance. Both models are initialized and trained as discussed in Sec. 5.

Table 1 summarizes the results for the different mixture models in juxtaposition to standard baselines (Agarap & Azcarraga, 2020; Ding et al., 2019; Mahon & Lukasiewicz, 2021; Peharz et al., 2020a). The matched GMM models

Table 1. Bits per dimension (BPD) and clustering accuracy (Acc.).

	MNIST		Fashion-MNIST	
	BPD	Acc.	BPD	Acc.
Matched GMM	6.61	61.75	6.58	55.00
GMM w/ S.Clust.	5.21	48.91	5.71	48.52
EiNet (Matched)	2.19	-	4.18	-
EiNet (Large)	1.95	-	3.98	-
GMM w/ Curr.	2.61	64.47	4.35	58.41
CPMM	2.21	80.07	3.91	65.01
Teacher	1.04	-	2.85	-
FFJORD	1.01	-	2.75	-
VADE	-	94.5	-	57.8
SPC	-	99.21	-	67.94
DLS	-	97.6	-	69.3

(one component per cluster) achieves surprisingly high clustering accuracies but subpar BPD. Unsurprisingly, introducing more components (GMM w/ S.Clust.) improves the BPD. However, spectral clustering across components is only marginally successful and the clustering accuracy drops relative to the baseline GMM. We additionally train two EiNets (Peharz et al., 2020a), a type of SPN. The first EiNet (Matched) is constrained to have the same total number of parameters (not components) as the CPMM. The second EiNet (Large) utilizes an order of magnitude more parameters. Finally, we train a standard GMM with 250 components that additionally utilizes the curriculum. The CPMM trained with the curriculum enjoys significantly improved BPD and clustering accuracy relative to the GMM baselines, better performance compared to the matched EiNet and standard GMM with the curriculum, and similar performance against the large EiNet. (See Appendix for additional ablation experiments.) Our results indicate that our proposed methodology closes the considerable gap between mixture models and modern neural baselines (listed in the bottom of Table 1). An especially notable case is the Fashion MNIST clustering accuracy, where the CPMM results in comparable-to-better than the standard baselines. Thus, CPMM allows one to retain the interpretability and inference speed of mixture models with improvements in modeling performance over tradition mixture approaches.

Interpretability Figure 4 shows the trajectories learned by a CPMM where we have manually ordered the clusters. The figure shows smooth transitions along the trajectory for all digits, in general transitioning from fatter, curvier digits to slimmer, straight digits. The early pseudotimes (far left) are never the most likely component and could be excluded for a slight increase in the BPD. We have not done so here for the sake of transparency and to avoid arbitrary post-hoc operations. This particular model does an excellent job of isolating different digits into different trajectories with

the exception of some 3/8 and 4/9 confusion that results in a clustering accuracy of 80%. This confusion is understandable given the similarities in the digit pairs but more importantly, this analysis highlights the *interpretability* of our model. That is, it is simple to resolve the model’s confusion from the trajectories since CPMM directly operates over input features. In contrast to more opaque clustering models, a quick visual inspection reveals CPMM’s partitioning of the space. Similar results and discussion can be found for Fashion-MNIST in the Appendix.

OOD Detection Finally, we test the CPMM’s ability to distinguish between different datasets based on likelihood. This is a known shortcoming of neural likelihood models (Hendrycks et al., 2019; Nalisnick et al., 2019) where the networks unfortunately predict that simpler datasets are more likely than the data the model was trained on; e.g., MNIST instances yield a higher likelihood than Fashion-MNIST instances for models trained on Fashion-MNIST. This surprising result limits the applicability of modern likelihood models for detecting out-of-distribution (OOD) and anomalous instances, despite their intuitive capabilities. Table 3 (appendix) illustrates the performance of our model against a CNF for out-of-distribution recognition when trained on a distinct inlier distribution. Results are given as areas under the receiver operating characteristic and precision/recall curves as a percent. Higher is better. Although both models perform nearly perfectly when MNIST is in-distribution and Fashion-MNIST is out-of-distribution, the CPMM performs quite well on the reverse problem where the CNF fails miserably. Notwithstanding a gap in the likelihood that is obtainable with CPMM as compared to CNF (Table 1), this experiment illustrates an advantage to directly modeling the input data through a mixture of components.

6. Conclusions

We have proposed a hierarchical method to parameterize a continuum of mixture models from neural ODEs to create a rich, multi-modal density over disparate partitions, essentially constructing a mixture of mixtures where the outer mixture encompasses different data regions and the inner mixture can be arbitrarily complex. We have innovated a curriculum that alleviates many of the difficulties associated with initialization via an annealing process from a prescribed latent space towards the true, nuanced data space. Since the initial distribution is known, it is trivial to initialize a model that is well-matched. Finally, we sample a finite number of components from the dynamic NODE to achieve an effective, light-weight model that significantly outperforms mixtures with a similar number of components, indicating that the combination of the curriculum and dynamic system allows for a more efficient use of the available components than traditional methods.

Acknowledgements

This research was partly funded by NSF grant IIS2133595 and by NIH 1R01AA02687901A1.

References

- Agarap, A. F. and Azcarraga, A. P. Improving k-means clustering performance with disentangled internal representations. *2020 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, 2020.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, pp. 41–48, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585161. doi: 10.1145/1553374.1553380. URL <https://doi.org/10.1145/1553374.1553380>.
- Blanchard, P., Devaney, R., and Hall, G. *Differential Equations*. Thomson Brooks/Cole, 2006. ISBN 9780495012658. URL <https://books.google.com/books?id=mwxX2pv9UvYC>.
- Böhning, D., Seidel, W., Alfö, M., Garel, B., Patilea, V., and Walther, G. Advances in mixture models. *Comput. Stat. Data Anal.*, 51:5205–5210, 2007.
- Caron, M., Bojanowski, P., Joulin, A., and Douze, M. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018.
- Chen, C., Li, G., Xu, R., Chen, T., Wang, M., and Lin, L. Clusternet: Deep hierarchical cluster network with rigorously rotation-invariant representation for point cloud analysis. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4989–4997, 2019a.
- Chen, R. T. Q., Behrmann, J., Duvenaud, D. K., and Jacobsen, J.-H. Residual flows for invertible generative modeling. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019b. URL <https://proceedings.neurips.cc/paper/2019/file/5d0d5594d24f0f955548f0fc0ff83d10-Paper.pdf>.
- Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. Neural ordinary differential equations. In *NeurIPS*, 2018.
- Darwiche, A. A differential approach to inference in bayesian networks. 50(3), 2003. ISSN 0004-5411. doi: 10.1145/765568.765570. URL <https://doi.org/10.1145/765568.765570>.
- Deng, L. The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- Ding, F., Luo, F., and Yang, Y. Double cycle-consistent generative adversarial network for unsupervised conditional generation, 2019. URL <https://arxiv.org/abs/1911.05210>.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real NVP. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=HkpbmH9lx>.
- Dowson, D. and Landau, B. The fréchet distance between multivariate normal distributions. *Journal of multivariate analysis*, 12(3):450–455, 1982.
- Eghbal-zadeh, H., Zellinger, W., and Widmer, G. Mixture density generative adversarial networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5813–5822, 2019.
- Falcon, WA, e. Pytorch lightning. *GitHub. Note: https://github.com/PyTorchLightning/pytorch-lightning*, 3, 2019.
- Filippone, M., Camastra, F., Masulli, F., and Rovetta, S. A survey of kernel and spectral methods for clustering. *Pattern Recognition*, 41(1):176–190, 2008. ISSN 0031-3203. doi: <https://doi.org/10.1016/j.patcog.2007.05.018>. URL <https://www.sciencedirect.com/science/article/pii/S0031320307002580>.
- Gao, M. M., Tai-hua, C., and xiang Gao, X. Application of gaussian mixture model genetic algorithm in data stream clustering analysis. *2010 IEEE International Conference on Intelligent Computing and Intelligent Systems*, 3:786–790, 2010.
- Gelfand, A. E., Kottas, A., and MacEachern, S. N. Bayesian nonparametric spatial modeling with dirichlet process mixing. *Journal of the American Statistical Association*, 100(471):1021–1035, 2005. doi: 10.1198/016214504000002078. URL <https://doi.org/10.1198/016214504000002078>.
- Ghahramani, Z. and Hinton, G. E. The EM algorithm for mixtures of factor analyzers. 1996.
- Goodfellow, I. J., Bengio, Y., and Courville, A. C. Deep learning. *Nature*, 521:436–444, 2015.

- Görür, D. and Rasmussen, C. E. Dirichlet process gaussian mixture models: Choice of the base distribution. *Journal of Computer Science and Technology*, 25:653–664, 2010.
- Grathwohl, W., Chen, R. T. Q., Bettencourt, J., and Duvenaud, D. Scalable reversible generative models with free-form continuous dynamics. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJxgknCcK7>.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks, 2016. URL <https://arxiv.org/abs/1609.09106>.
- Hastie, T., Tibshirani, R., and Friedman, J. The elements of statistical learning. 2001.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- Hendrycks, D., Mazeika, M., and Dietterich, T. Deep anomaly detection with outlier exposure. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=HyxCxhRcY7>.
- Holzmann, H., Munk, A., and Gneiting, T. Identifiability of finite mixtures of elliptical distributions. *Scandinavian Journal of Statistics*, 33, 2006.
- Izmailov, P., Kirichenko, P., Finzi, M., and Wilson, A. Semi-supervised learning with normalizing flows. *ArXiv*, abs/1912.13025, 2020.
- Jaini, P., Poupart, P., and Yu, Y. Deep homogeneous mixture models: Representation, separation, and approximation. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/c5f5c23belb71adb51ea9dc8e9d444a8-Paper.pdf>.
- Jiang, Z., Zheng, Y., Tan, H., Tang, B., and Zhou, H. Variational deep embedding: An unsupervised and generative approach to clustering. In *IJCAI*, 2017.
- Jiao, L., Denooux, T., Liu, Z., and Pan, Q. Egmm: an evidential version of the gaussian mixture model for clustering. *ArXiv*, abs/2010.01333, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pp. 10236–10245, 2018.
- Law, M. T., Urtasun, R., and Zemel, R. S. Deep spectral clustering learning. In *ICML*, 2017.
- Mahon, L. and Lukasiewicz, T. Selective pseudo-label clustering. In Edelkamp, S., Möller, R., and Rueckert, E. (eds.), *KI 2021: Advances in Artificial Intelligence*, pp. 158–178, Cham, 2021. Springer International Publishing. ISBN 978-3-030-87626-5.
- Mukherjee, S., Asnani, H., Lin, E., and Kannan, S. Clustergan: Latent space clustering in generative adversarial networks. *ArXiv*, abs/1809.03627, 2019.
- Müller, P. and Quintana, F. A. Nonparametric Bayesian Data Analysis. *Statistical Science*, 19(1):95 – 110, 2004. doi: 10.1214/088342304000000017. URL <https://doi.org/10.1214/088342304000000017>.
- Nalisnick, E., Matsukawa, A., Teh, Y. W., and Lakshminarayanan, B. Detecting out-of-distribution inputs to deep generative models using typicality. *arXiv preprint arXiv:1906.02994*, 2019.
- Oliva, J., Dubey, A., Zaheer, M., Póczos, B., Salakhutdinov, R., Xing, E., and Schneider, J. Transformation autoregressive networks. In *International Conference on Machine Learning*, pp. 3898–3907, 2018.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/6c1da886822c67822bcf3679d04369fa-Paper.pdf>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance.pdf>.

- Peharz, R., Lang, S., Vergari, A., Stelzner, K., Molina, A., Trapp, M., Van Den Broeck, G., Kersting, K., and Ghahramani, Z. Einsum networks: Fast and scalable learning of tractable probabilistic circuits. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7563–7574. PMLR, 13–18 Jul 2020a. URL <https://proceedings.mlr.press/v119/peharz20a.html>.
- Peharz, R., Vergari, A., Stelzner, K., Molina, A., Shao, X., Trapp, M., Kersting, K., and Ghahramani, Z. Random sum-product networks: A simple and effective approach to probabilistic deep learning. In Adams, R. P. and Gogate, V. (eds.), *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pp. 334–344. PMLR, 22–25 Jul 2020b. URL <https://proceedings.mlr.press/v115/peharz20a.html>.
- Pires, G. G. P. F. and Figueiredo, M. A. T. Variational mixture of normalizing flows. *ArXiv*, abs/2009.00585, 2020.
- Poon, H. and Domingos, P. Sum-product networks: A new deep architecture. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pp. 689–690, 2011. doi: 10.1109/ICCVW.2011.6130310.
- Richardson, E. and Weiss, Y. On gans and gmms. In *NeurIPS*, 2018.
- Shaham, U., Stanton, K., Li, H., Nadler, B., Basri, R., and Kluger, Y. Spectralnet: Spectral clustering using deep neural networks. *ArXiv*, abs/1801.01587, 2018.
- van den Oord, A., Kalchbrenner, N., Espeholt, L., Kavukcuoglu, k., Vinyals, O., and Graves, A. Conditional image generation with pixelcnn decoders. In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL <https://proceedings.neurips.cc/paper/2016/file/b1301141feffabac455e1f90a7de2054-Paper.pdf>.
- Viroli, C. and McLachlan, G. Deep gaussian mixture models. *Statistics and Computing*, 29:43–51, 2019.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Xie, J., Girshick, R. B., and Farhadi, A. Unsupervised deep embedding for clustering analysis. In *ICML*, 2016.
- Zhang, D., Sun, Y., Eriksson, B., and Balzano, L. Deep unsupervised clustering using mixture of autoencoders. *ArXiv*, abs/1712.07788, 2017.

A. Numerical Integration

The typical strategy to train a likelihood model is to optimize based on the predicted *log-likelihood* and not on the likelihood directly. In fact, when training a Gaussian mixture model, we typically calculate the log-likelihood of each component separately prior to utilizing the numerically-stable $\log \sum \exp$ operator instead of performing the \exp , followed by the \sum , followed by the \log . This joint version stably handles difficulties resulting from over- or under-flow when the log-likelihood of any single component is very (large) small.

When attempting to evaluate the log-likelihood for a continuum of components we should utilize a similar $\log \int \exp$:

$$\log p(\mathbf{x}) = \frac{-M}{2} \log(2\pi) + \log \int_0^1 \exp \left(-\log |\mathbb{X}(s)| / 2 - (\mathbf{x} - \mu(s))^T \mathbb{X}^{-1}(s) (\mathbf{x} - \mu(s)) / 2 + \log \pi(s) \right) ds.$$

Unfortunately, none of the available black-box integrators match this form. This requires that we utilize a naive implementation of the $\log \int \exp$ operator. In practice, we find this sufficient to train the model for several epochs without issue but the ODE becomes increasingly stiff before ultimately underflowing and “NaNing” out. A better strategy to estimate a continuous mixture model will require a new black-box integrator which we consider an excellent avenue for future work.

B. Augmentations

For all image datasets we added uniform random noise between zero and one *prior to rescaling*, e.g., a pixel with an 8bit value of 34 would then be distributed uniformly between 34 and 35. For color images, we additionally perform standard random translations and horizontal flips. For tabular data, we add Gaussian noise to each feature such that the standard deviation of the noise is 1-10% of the spread of the data itself, e.g., we normalize each feature to be zero mean and unit standard deviation, add Gaussian noise scaled by 0.01 to 0.1 and unnormalize that data.

When training with the curriculum (see Sec. 3.3), we utilize similar augmentations prior to transforming into the various spaces such that the augmentation is shared across spaces. When training the CPMM student, we then add a small amount of Gaussian noise (standard deviation of 0.01) regardless of the space, including when fine-tuning on the input space.

C. MNIST Ablations

We experiment with different types of initializations and modifications of the CPMM on MNIST. In particular, we

test random initializations and kmeans-based initializations. We also explore restricting the diagonal portion of the covariance within each MFA to be constant (though still learnable and still with the low rank updates, A) and attempt to utilize a separate ODE for each cluster.

Table 2. MNIST Ablation

Init.	Diagonal	BPD	Acc.
Random	Constant	7.67	48.83
kmeans	Constant	7.56	64.93
kmeans	Variable	4.49	46.42
kmeans*	Constant	7.55	71.22

Table 2 contains some representative runs attempting to train the CPMM without guidance from the teacher. We find that training the models with a variable diagonal often results in the model degenerating and failing to train. In fact, training with a variable diagonal and a random initialization tends to fail after a few epochs and is therefore excluded from the table. Utilizing a kmeans initialization produces moderately good BPD (better than standard GMM training methods) but slightly worsened accuracies. Restricting the model to use a constant diagonal greatly stabilizes training and can produce reasonably good clustering accuracies. However, this strong assumption on the covariance results in very poor BPD. This is not surprising, as the model must raise the variance of pixels that would normally be very small (background pixels) to sufficiently capture the spread elsewhere. Finally, the bottom row in Table 2 contains results using a kmeans initialization with a constant diagonal but utilizes a separate ODE for each trajectory. This added flexibility results in virtually no improvement to the BPD but does provide a noticeable increase in accuracy. We find that training separate ODEs per cluster results in a significant increase to run time.

The results in Table 2 were chosen as representative values across many different runs. In general, it is possible to trade accuracy for BPD in any given row by adjust other hyperparameters or seeds. Utilizing the curriculum learning process helps to stabilize training and allows us to simultaneously improve the generative and discriminative portions of the model.

D. Fashion MNIST Results

Figure 5 contains the trajectories from a curriculum-based CPMM trained on Fashion-MNIST. Similar to MNIST, the trajectories show smooth variations between means. Many of the trajectories contain visible evolutions from start to finish, i.e., bags begin without handles but end with handles or shoes evolve from boots to stilletos. However, some of the fine detail is missing from the trajectories, in particular for the shirts, pullovers, and coats classes.




Figure 5. Means from a hierarchical CPMM trained on Fashion-MNIST with a space-based curriculum. Each trajectory is evaluated at evenly spaced pseudotimes between zero and one and displayed from left to right. The initial component (far left cases) are never the most likely component and could be discarded.

As in MNIST, we can clearly see why the model achieves 65.74% clustering accuracy. The shirts, pullovers, and coats classes all produce very similar trajectories. Similarly, the sandals, sneakers, and ankle boots classes see considerable overlap. Abstractly, this information can be extracted from any clustering model based on the confusion matrix. But, thanks to the interpretability of the mixture models, the decision process and resulting confusion is clear.

E. Out-of-Distribution Performance

Table 3. OOD Detection via Likelihood Thresholding

In	Out	CPMM		FFJORD	
		AUROC	AUPR	AUROC	AUPR
MNIST	Fashion	99.95	99.95	99.95	99.95
Fashion	MNIST	93.38	93.28	8.98	31.65