

Variational Clustering: Leveraging Variational Autoencoders for Image Clustering

Vignesh Prasad*

Technical University of Darmstadt
Germany
vignesh.prasad@tu-darmstadt.de

Dipanjan Das*

Embedded Systems and Robotics
TCS Research & Innovation, Kolkata, India
dipanjan.da@tcs.com

Brojeshwar Bhowmick

Embedded Systems and Robotics
TCS Research & Innovation, Kolkata, India
b.bhowmick@tcs.com

Abstract—Recent advances in deep learning have shown their ability to learn strong feature representations for images. The task of image clustering naturally requires good feature representations to capture the distribution of the data and subsequently differentiate data points from one another. Often these two aspects are dealt with independently and thus traditional feature learning alone does not suffice in partitioning the data meaningfully. Variational Autoencoders (VAEs) naturally lend themselves to learning data distributions in a latent space. Since we wish to efficiently discriminate between different clusters in the data, we propose a method based on VAEs where we use a Gaussian Mixture prior to help cluster the images accurately. We jointly learn the parameters of both the prior and the posterior distributions. Our method represents a true Gaussian Mixture VAE. This way, our method simultaneously learns a prior that captures the latent distribution of the images and a posterior to help discriminate well between data points. We also propose a novel reparametrization of the latent space consisting of a mixture of discrete and continuous variables. One key takeaway is that our method generalizes better across different datasets without using any pre-training or learnt models, unlike existing methods, allowing it to be trained from scratch in an end-to-end manner. We verify our efficacy and generalizability experimentally by achieving state-of-the-art results among unsupervised methods on a variety of datasets. To the best of our knowledge, we are the first to pursue image clustering using VAEs in a purely unsupervised manner on real image datasets.

Index Terms—Unsupervised Learning, Clustering, Variational Inference

I. INTRODUCTION

Image Clustering is a fundamental, challenging and widely studied problem in machine learning [3]–[8]. Some classical examples are K -means [9], Gaussian Mixture Models [10] and Spectral clustering [6] which are promising, but require a robust feature representation for good clustering. In recent years, Deep Learning has made huge progress in learning robust feature representations of images. These learned representations help cluster the data more accurately when used with traditional methods like K -means for example [11]. One way to use deep representations, off the shelf, is to extract the feature representation of an image from a pre-trained model and use them directly in any clustering algorithm [12].

This work was done when Vignesh Prasad worked at TCS Research & Innovation.

* - Equal Contribution

The problem with such approaches is that they don't fully exploit the power of deep neural networks. Song *et al.* [13] learn a representation to accurately cluster the images in the dataset by integrating K -means into the bottleneck layer of an Autoencoder. This association enables the model to learn a meaningful clustering-oriented representation.

With the motivation to pursue a robust and generalizable methodology in a principled way, we aim to make inferences in a latent space learned specifically for a clustering task. The idea is that it would be easier to group the data in this space, compared to an arbitrary space defined by pre-trained features. Off late, the use of generative methods for clustering has been on the rise as their expressive power helps efficiently capture, represent and recreate sampled data points. As we wish to experiment with data distributions in a latent space that can accurately represent the input data, the paradigm of Variational Autoencoders (VAEs) lends itself directly to the task at hand.

We build on the ideas of GMVAE [1] and VaDE [2] addressing their fallacies while maintaining the underlying motivation of using a Gaussian Mixture Model as the latent space distribution. Instead of deriving the prior from a random variable, as in GMVAE, our prior is deterministic. This is similar to VaDE however, we learn the parameters for the prior and posterior jointly, unlike in VaDE which uses a pre-training phase to initialize the parameters of the prior.

To illustrate the differences between our process, GMVAE [1] and VaDE [2] we visualize the graphical models in Fig. 1. In GMVAE, the Gaussian prior \mathbf{z}_2 depends on a noise variable \mathbf{z}_1 & varies for a given cluster, as shown in Fig. 1a. Ours is more intuitive as it depends only on the cluster, as shown in Fig. 1c. Secondly, GMVAE expresses the categorical posterior $q(k|\mathbf{z}_1, \mathbf{z}_2)$ with the prior $p_\beta(\mathbf{z}_2|\mathbf{z}_1, k)$ using Bayes' rule. This applies to VaDE too, along with fixing the GMM prior during pre-training. We learn it during training, giving more flexible learning, the effectiveness of which is seen in the results in Table I. This can also be seen on a toy dataset, compared to GMVAE, where our method learns a more compact cluster representation as compared to GMVAE, as shown in Fig. 3.

Our method is more principled as we directly learn the cluster assignment probabilities $q(k|\mathbf{z})$ instead of performing a Bayesian classification, as done in both GMVAE and VaDE. Once the cluster predictions $q(k|\mathbf{z})$ become closer to a one-hot vector, it, in turn, minimizes the loss for just a given

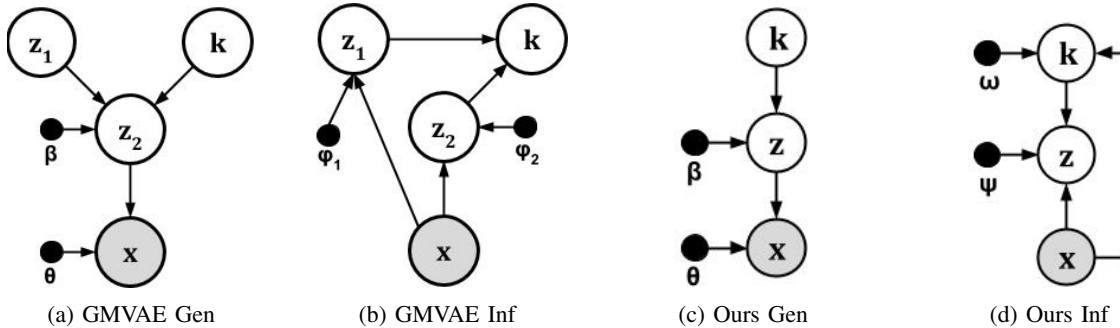


Fig. 1: Graphical representation of the inference and generative models used by us and GMVAE [1]. VaDE [2] also has a similar generative model as ours but differ in their inference model. Grey nodes represent data nodes. White ones represent latent nodes, where z , z_1 and z_2 are normally distributed variables where k is a discrete variable. Black nodes represent the network parameters. (Gen - Generative, Inf- Inference)

cluster instead of over all clusters leading to a more effective clustering. This process is shown in Fig. 2.

Along with this, we propose a novel augmentation loss as a mixture of L_2 distance between the cluster predictions and 2-Wasserstein distance [14] between the predicted posterior distributions. We do so to ensure consistency in the cluster predictions and also in the predicted posteriors thereby enforcing a similarity among samples from a given cluster. The idea behind this is to achieve invariance to marginal variations in the data and predict an accurate mapping for both real as well as augmented data. It would be counter-intuitive to map an augmented version of an input to a different distribution.

Therefore, our novelty is in modelling a GMM prior whose parameters are learnt during training, which is more flexible than fixed distributions like in other methods. This is too rigid a constraint mainly in complex datasets (CIFAR/STL). Due to our principled & effective modelling, coupled with our novel incorporation of an augmentation loss, our accuracy is significantly better than others and is very coherent on similar datasets. We perform well even in a cross-dataset transfer.

The main contributions of this paper are:

- A novel unsupervised image clustering algorithm by combining VAEs with a true Gaussian Mixture prior learnt without any pre-training in an end-to-end manner.
- More principled latent space priors that subsequently lead to a simpler inference model.
- A novel augmentation loss making our method robust and leading to strong intra-cluster relations.

In contrast to other methods, our method performs consistently well over all datasets using only raw pixel inputs, showcasing the generalization capability of our method. To the best of our knowledge, we are the first to do so using VAEs in a purely unsupervised manner, especially on real image datasets.

II. RELATED WORKS

Tian *et al.* [12] explore the use of deep methods in classical clustering by first learning a feature representation and subsequently performing clustering in the feature space. Nina *et al.* [15] follow a similar approach by using a triplet loss between positive and negative anchor images from the dataset and perform K-means on the learnt latent space. Similarly,

Das *et al.* [16] propose a novel method for selecting images for calculating a triplet loss, which however depends on the quality of the pre-training. Yang *et al.* [17] follow a similar approach for learning a latent space but follow a graph-based spectral clustering approach with the learnt representation.

Song *et al.* [13] take a step further by integrating K -means into the bottleneck of an Autoencoder. This enables the model to learn meaningful clustering-oriented representations. Deep Embedded Clustering (DEC) [18] builds on a similar idea by minimizing the KL Divergence between cluster assignment probabilities and an auxiliary target distribution derived from the current predictions. However, this causes the auxiliary distribution to change as the predictions do during training.

Deep Adaptive Clustering (DAC) [19] uses a pairwise classification approach to distinguish between image pairs. It learns features that correspond directly to predicted class labels and achieve reasonable accuracy. Similarly, Ji *et al.* [20] propose a method that tries to preserve pair-wise semantic mutual information in the data. Deep Embedded Regularized Clustering (DEPICT) [21] uses a convolutional autoencoder to jointly learn representations and cluster predictions by constraining clusters to have similar no. of samples. Their performance deteriorates when data is not distributed well across clusters. To this end, RDEC [22] uses adversarial training to improve performance on imbalanced datasets. Along similar lines, Joint Unsupervised Learning (JULE) [7] learns representations and cluster assignments simultaneously using agglomerative clustering. Though it gets good results, their method is computationally intensive and has a heavy memory usage [23]. Sarfaraz *et al.* [24] also develop an agglomerative clustering method but with lower overheads in computing distances by partitioning the data points more efficiently.

Haeusser *et al.* [25] attempt to enforce intra-class similarity with associative constraints between images and their augmented versions to ensure similarity in predictions. However, the robustness and generalization capability of their method are questionable since they show vastly differing results on CIFAR10 [26] and STL-10 [27], both of which are real-world datasets having 9 same classes out of 10.

Information-Maximizing Self-Augmented Training (IM-SAT) [28] learn a probabilistic classifier using Regularized In-

formation Maximization (RIM) [29]. For feature-rich datasets (CIFAR10 and CIFAR100 [26]), they use ResNet [30] features as input. Thus, it is difficult to properly gauge their performance as results without pre-trained models are unavailable.

Coming to Deep Generative methods, ClusterGAN [31] is a Generative Adversarial framework for unsupervised clustering using a discrete-continuous mixed approach for clustering. Ghasedi *et al.* [32] build on this idea by introducing a self-paced learning algorithm that helps guide the learning, similar to a curriculum learning framework, to get better results. Variational Deep Embedding (VaDE) [2] and Gaussian Mixture VAEs (GMVAE) [1] use VAEs with GMMs simultaneously to model the inference process. Yang *et al.* propose DGG [33] along similar lines with the additional constraint of minimizing the graph distances between embeddings of data points.

Both ClusterGAN and VaDE require pre-training to initialize cluster centroids, and DGG requires pre-training to initialize the graph embeddings. Their success relies on the success of the initial pre-training. In this context, GVMAE learns the prior and posterior parameters jointly. However, their prior representation seems counter-intuitive as they use a sample from a normal distribution to generate the parameters (mean and variance) of the prior for each of the clusters. This implies that the prior for each class is dependent on a random variable rather than the class itself and hence would vary each time one would want to sample a latent for a class.

III. VARIATIONAL AUTOENCODERS

Variational Autoencoders (VAEs) [34], [35] are an application of Autoencoders for performing Variational Bayesian Inference over a set of data points. The main idea of Variational Inference is to learn a distribution in a latent space that can accurately capture the true distribution of the dataset. In particular, we wish to represent the joint probability $p(x, z)$ for points in dataset x and their latent space representations z . This joint probability can be written as $p(x, z) = p(x|z)p(z)$ where $p(z)$ is a prior distribution from where latent variables are drawn and $p(x|z)$ is the conditional likelihood of a data point x conditioned on the drawn latent variable z . The goal of variational inference is to infer the latent distribution from observed samples i.e. to accurately calculate $p(z|x)$. Using Bayes theorem, we can write this posterior distribution as:

$$p(z|x) = \frac{p(x, z)}{p(x)} = \frac{p(x|z)p(z)}{p(x)} \quad (1)$$

The problem is in approximating the evidence $p(x)$, in the denominator, which requires an expensive marginalization over the latent variables. Therefore, Variational Inference seeks to approximate the posterior with a parameterized distribution $q_\theta(z|x)$ and minimize the KL-divergence between the true and approximated posterior w.r.t. to data distribution $p_{\mathcal{D}}(x)$.

$$\mathbb{E}_{p_{\mathcal{D}}} KL(q_\theta(z|x)||p(z|x)) = \mathbb{E}_{q, p_{\mathcal{D}}} \log \frac{q_\theta(z|x)}{p(x, z)} + \mathbb{E}_{p_{\mathcal{D}}} \log p(x) \quad (2)$$

This can be re-written as

$$\mathbb{E}_{p_{\mathcal{D}}} \log p(x) = \mathbb{E}_{p_{\mathcal{D}}} KL(q_\theta(z|x)||p(z|x)) + \mathbb{E}_{q, p_{\mathcal{D}}} \log \frac{p(x, z)}{q_\theta(z|x)} \quad (3)$$

By Jensen's inequality [36], the KL Divergence is always non-negative. Hence, the expectation in Eq. 3 acts as a lower bound for the evidence log-likelihood, and is hence called the Evidence Lower Bound (ELBO), given in Eq. 4.

$$ELBO = \mathbb{E}_{q, p_{\mathcal{D}}} \log \frac{p(x, z)}{q_\theta(z|x)} \quad (4)$$

This allows us to write Eq. 3 as:

$$\mathbb{E}_{p_{\mathcal{D}}} \log p(x) = \mathbb{E}_{p_{\mathcal{D}}} KL(q_\theta(z|x)||p(z|x)) + ELBO \quad (5)$$

Therefore, minimizing the KL Divergence and hence, maximizing the log-likelihood of the evidence can be done by maximizing the ELBO. Since deep networks have strong representational abilities, Autoencoders have shown to perform well in approximating the distributions, leading to the birth of Variational Autoencoders (VAEs). Further information on Variational Inference and VAEs can be found in [35].

IV. PROPOSED APPROACH

VAEs traditionally model the prior as a single multivariate Gaussian. Since our objective is to cluster the data accurately, we use a Gaussian mixture prior for the latent space representation such that each Gaussian represents a cluster in the data. We aim to learn these Gaussians and an effective representation of the probability models, thereby learning the distribution a given data point belongs to, which corresponds to learning the clustering of the data.

An overview of our approach is shown in Fig. 2. Given a set of input images, we first predict their cluster probabilities. For each cluster, we infer a posterior latent distribution given the input and generate an independent prior distribution as well. We sample a latent from the posterior of each class from which we generate image samples which are used in computing the ELBO, along with the prior and the cluster probabilities.

A. Generative Process

We first select a cluster k , from a categorical distribution characterized by π . we draw a sample \mathbf{z} from the latent distribution $p(\mathbf{z}|k)$ of that cluster parameterized by a neural network β , to generate the conditional distribution $p(\mathbf{x}|\mathbf{z})$, which is parameterized by another neural network θ . In mathematical terms, our generative model can be seen as

$$p_{\beta, \theta}(\mathbf{x}, \mathbf{z}, k) = p_\theta(\mathbf{x}|\mathbf{z})p_\beta(\mathbf{z}|k)p(k) \quad (6)$$

$$p(k) = \text{Cat}(\pi) \quad (7)$$

$$p_\beta(\mathbf{z}|k) = \mathcal{N}(\mathbf{z}|\mu_\beta(\mathbf{e}_k), \text{diag}(\sigma_\beta^2(\mathbf{e}_k))) \quad (8)$$

$$p_\theta(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}|\mu_\theta(\mathbf{z}), \text{diag}(\sigma_\theta^2(\mathbf{z}))) \text{ or } \mathcal{B}(\mathbf{x}|\mu_\theta(\mathbf{z})) \quad (9)$$

where $\pi = \frac{1}{K}$ is the categorical probability, \mathbf{e}_k is a one-hot vector with a 1 at the k_{th} position, $\mathcal{N}(\cdot)$ refers to a Normal distribution & $\mathcal{B}(\cdot)$ refers to a Bernoulli distribution. $\mu_\beta, \sigma_\beta^2, \mu_\theta, \sigma_\theta^2$ are the means & variances, characterized by deep networks with parameters β & θ respectively.

Our process differs from GVMAE [1] & VaDE [2] in the following ways. In GMVAE, a neural network samples a standard normal distribution as an input (\mathbf{z}_1 in Fig. 1a) &

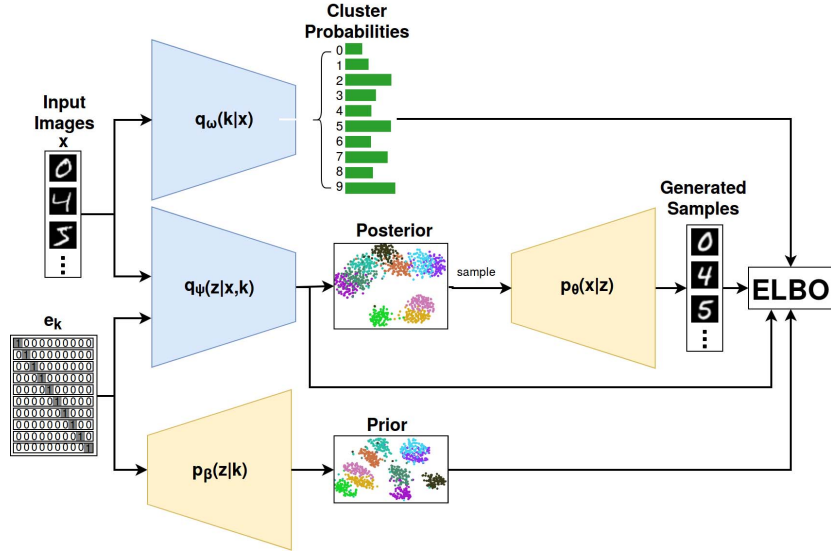


Fig. 2: A flow chart of our training process. Given a set of input images \mathbf{x} and its augments $\hat{\mathbf{x}}$, the posterior distribution $q_\omega(k|\mathbf{x})$ is first calculated, followed by $q_\psi(\mathbf{z}|\mathbf{x}, k)$ and prior $p_\beta(\mathbf{z}|k)$, where k is represented as \mathbf{e}_k , a one-hot vector having the k^{th} position is 1. We sample a latent from the posterior and generate a conditional prior $p_\theta(\mathbf{x}|\mathbf{z})$. We then calculate the ELBO according to Eq. 13. Given the predictions for the images and its augments, we calculate the augmentation loss \mathcal{L}_{aug} according to Eq. 14 and subsequently, calculate the final loss as given in Eq. 18. (Picture best viewed in colour)

generates a mean & variance for each cluster (\mathbf{z}_2 in Fig. 1a). Therefore, the latent distribution for any given cluster depends on a normally distributed variable & not just on the cluster it represents. Out of these, the k^{th} mean & variance is selected from which they sample their latent.

In VaDE, the process is similar to ours however, they first train a stacked autoencoder with the input data & then learn a GMM over the predicted latent representations from the encoder’s bottleneck. This learnt GMM is used as the prior from which the latent variables are sampled. The performance of VaDE depends on how well the pre-training & GMM initialization go, which would lead to a bad performance in case the initial learning doesn’t properly converge.

B. Inference Model

We perform variational inference by maximizing the Evidence Lower Bound (ELBO), which in turn leads to maximizing the log-likelihood of the evidence $\log p(x)$. With a slight abuse of notation, the ELBO can be written as:

$$ELBO = \mathbb{E}_{q, p_{\mathcal{D}}} \log \frac{p_{\beta, \theta}(\mathbf{x}, \mathbf{z}, k)}{q_{\omega, \psi}(\mathbf{z}, k|\mathbf{x})} \quad (10)$$

where $\mathbb{E}_{q, p_{\mathcal{D}}}$ is the expectation with respect to $q_{\omega, \psi}(\mathbf{z}, k|\mathbf{x})p_{\mathcal{D}}(\mathbf{x})$. Here, $q_{\omega, \psi}(\mathbf{z}, k|\mathbf{x})$ is the approximate posterior of the inference model, which we factorize as:

$$q_{\omega, \psi}(\mathbf{z}, k|\mathbf{x}) = q_\omega(k|\mathbf{x})q_\psi(\mathbf{z}|\mathbf{x}, k) \quad (11)$$

such that $\sum_{k=1}^K q_\omega(k|\mathbf{x}) = 1$ and $q_\psi(\mathbf{z}|\mathbf{x}, k) = \mathcal{N}(\mathbf{z}|\mu_\psi(\mathbf{x}, \mathbf{e}_k), \text{diag}(\sigma_\psi^2(\mathbf{x}, \mathbf{e}_k)))$ where μ_ψ, σ_ψ^2 refer to the mean and variance of the posterior parameterized by a deep network ψ .

Our inference model is more straight-forward than GMVAE [1] since they calculate their categorical posterior based on their latent’s prior which, as mentioned towards the ending of Sec. IV-A, depends on a sample from a standard normal distribution. This implies that the probability of a given category being assigned to a fixed category is stochastic in nature, which ideally shouldn’t be the case. A latent variable is a representation of a sample data point. Therefore, the categorical posterior that depends on the latent, it should directly be inferred from the sampled data point itself. Our method directly predicts the categorical posterior $q_\omega(k|\mathbf{x})$ directly from the sampled data point \mathbf{x} .

C. Evidence Lower Bound

Eq. 10 can be factorized using Eq. 6 & 11 as:

$$ELBO = \mathbb{E}_{q, p_{\mathcal{D}}} \left[\log \frac{p_\theta(\mathbf{x}|\mathbf{z})p_\beta(\mathbf{z}|k)p(k)}{q_\omega(k|\mathbf{x})q_\psi(\mathbf{z}|\mathbf{x}, k)} \right] \quad (12)$$

Since $\log p(k)$ is a constant, we drop it from Eq. 12 for ease of notation. Eq. 12 can then be re-written as:

$$ELBO = \mathbb{E}_{q, p_{\mathcal{D}}} [\log p_\theta(\mathbf{x}|\mathbf{z})] - \mathbb{E}_{q_\omega(k|\mathbf{x})p_{\mathcal{D}}(\mathbf{x})} [\log q_\omega(k|\mathbf{x})] - \mathbb{E}_{q_\omega(k|\mathbf{x})p_{\mathcal{D}}(\mathbf{x})} [KL(q_\psi(\mathbf{z}|\mathbf{x}, k)||p_\beta(\mathbf{z}|k))] \quad (13)$$

In Eq. 13, the first term $\mathbb{E}_{q, p_{\mathcal{D}}} [\log p_\theta(\mathbf{x}|\mathbf{z})]$ is the likelihood of a sampled data point \mathbf{x} w.r.t. the generated distribution. It is similar to ensuring a proper reconstruction of the sampled input. It is further written as $\frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K q_\omega(k|\mathbf{x}_i) \mathbb{E}_{\mathbf{z} \sim q_\psi(\mathbf{z}|\mathbf{x}_i, k)} [\log p_\theta(\mathbf{x}_i|\mathbf{z})]$, where N is the batch size. This expansion allows us to circumvent the issue of having to sample from a discrete distribution which in

turn would involve a discrete reparametrization. This expansion leaves us with having to sample a continuous latent \mathbf{z} for each cluster k , which can easily be reparametrized for each k . The second term $\mathbb{E}_{q_\omega(k|\mathbf{x})p_{\mathcal{D}}(\mathbf{x})}[\log q_\omega(k|\mathbf{x})]$ is the expected entropy of the posterior distribution of k w.r.t. the sampled data point \mathbf{x} , written as $\frac{-1}{N} \sum_{i=1}^N \sum_{k=1}^K q_\omega(k|\mathbf{x}_i) \log q_\omega(k|\mathbf{x}_i)$.

The third and final term is the \mathbf{z} KL term $\mathbb{E}_{q_\omega(k|\mathbf{x})p_{\mathcal{D}}(\mathbf{x})}[KL(q_\psi(\mathbf{z}|\mathbf{x}, k)||p_\beta(\mathbf{z}|k))]$, which refers to the Kullback-Leibler Divergence of the prior w.r.t. the posterior distribution of \mathbf{z} . Since they are both Gaussians, the KL Divergence has a closed form solution.

One interesting point to note in the formulation is the anti-clustering nature of the *ELBO*, caused by the maximization of the entropy of $q_\omega(k|\mathbf{x})$. This regularization forces the information stored in $q_\omega(k|\mathbf{x})$ to be distributed among the clusters, rather than learning a one-hot prediction for the cluster labels. This is however mitigated by the reconstruction term, which would force the data to belong to a particular cluster. One advantage of this entropy maximization is that it prevents the trivial solution where all inputs are mapped to a single cluster, also known as "mode collapse" which is a common problem associated with VAEs. In our implementation, we found that the scale of the reconstruction term is exponentially larger than the KL term on real image data. Hence, to ensure numerical stability, we multiply the reconstruction term in Eq. 13 with a scale-factor, λ_{recons} .

D. Augmentation Loss

To make our model robust to input variations, we add augmented images along with the original images. The advantage of this is two-fold. Firstly, it allows the network to learn from a larger number of samples. Secondly, it allows us to use predictions of the original images to guide the predictions of the augmented images, resulting in a form of self-supervision. Different augments used are given in Sec. V-A.

Let \mathbf{x} be an input image and $\hat{\mathbf{x}}$, its augmented image. Intuitively, both should belong to the same cluster and posterior distribution. This enforces a strong constraint on the posterior predicted for $\hat{\mathbf{x}}$ by providing a form of supervision. Therefore, to ensure similarity in the predictions, we minimize the L_2 distance between the predicted clusters and the expectation of the Wasserstein distance [14] between the posteriors.

$$\mathcal{L}_{aug} = \sum_{k=1}^K \|q_\phi(k|\mathbf{x}) - q_\phi(k|\hat{\mathbf{x}})\|_2^2 + \mathbb{E}_{q_\phi(k|\mathbf{x})}[\mathcal{W}_2(q_\psi(\mathbf{z}|\mathbf{x}, k), q_\psi(\hat{\mathbf{z}}|\hat{\mathbf{x}}, k))] \quad (14)$$

where \mathcal{W}_2 is the 2-Wasserstein distance. The expectation in the latter allows us to give a higher weight to the predicted cluster and a lower weight for clusters that are not associated with the given image. This has a closed form solution [37], due to their Gaussian form, given by

$$\mathcal{W}_2(q_\psi(\mathbf{z}|\mathbf{x}, k), q_\psi(\hat{\mathbf{z}}|\hat{\mathbf{x}}, k)) = \|\mu_\psi(\mathbf{x}, \mathbf{e}_k) - \mu_\psi(\hat{\mathbf{x}}, \mathbf{e}_k)\|_2^2 + \text{trace}(C_1 + C_2 - 2(C_2^{\frac{1}{2}} C_1 C_2^{\frac{1}{2}})^{\frac{1}{2}}) \quad (15)$$

the trace in Eq. 15 can be simplified as:

$$C_1 + C_2 - 2(C_2^{\frac{1}{2}} C_1 C_2^{\frac{1}{2}})^{\frac{1}{2}} = \text{diag}(\sigma_\psi^2(\mathbf{x}, \mathbf{e}_k)) + \text{diag}(\sigma_\psi^2(\hat{\mathbf{x}}, \mathbf{e}_k)) - 2\text{diag}(\sigma_\psi(\mathbf{x}, \mathbf{e}_k)\sigma_\psi(\hat{\mathbf{x}}, \mathbf{e}_k)) = \text{diag}(\sigma_\psi^2(\mathbf{x}, \mathbf{e}_k) + \sigma_\psi^2(\hat{\mathbf{x}}, \mathbf{e}_k) - 2\sigma_\psi(\mathbf{x}, \mathbf{e}_k)\sigma_\psi(\hat{\mathbf{x}}, \mathbf{e}_k)) = \text{diag}((\sigma_\psi(\mathbf{x}, \mathbf{e}_k) - \sigma_\psi(\hat{\mathbf{x}}, \mathbf{e}_k))^2) \quad (16)$$

Substituting Eq. 16 in Eq. 15, we get

$$\mathcal{W}_2(q_\psi(\mathbf{z}|\mathbf{x}, k), q_\psi(\hat{\mathbf{z}}|\hat{\mathbf{x}}, k)) = \|\mu_\psi(\mathbf{x}, \mathbf{e}_k) - \mu_\psi(\hat{\mathbf{x}}, \mathbf{e}_k)\|_2^2 + \text{trace}(\text{diag}(\sigma_\psi(\mathbf{x}, \mathbf{e}_k) - \sigma_\psi(\hat{\mathbf{x}}, \mathbf{e}_k))^2) = \|\mu_\psi(\mathbf{x}, \mathbf{e}_k) - \mu_\psi(\hat{\mathbf{x}}, \mathbf{e}_k)\|_2^2 + \|\sigma_\psi(\mathbf{x}, \mathbf{e}_k) - \sigma_\psi(\hat{\mathbf{x}}, \mathbf{e}_k)\|_2^2 \quad (17)$$

To prevent the augmentation from driving the initial learning phases, we anneal the effect of this loss so as to bring about its significance once a suitable amount of learning has progressed by modifying its weight λ_{aug} . Our final loss can be written as

$$\mathcal{L} = -ELBO + \lambda_{aug}\mathcal{L}_{aug} \quad (18)$$

V. EXPERIMENTAL DETAILS

Details about the datasets used are given in Sec. V-A and about the network architecture is given in the appendix in Sec. V-B. We implement the system using Tensorflow [38] using Adam [39] as our optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and the learning rate as 10^{-3} for MNIST & Fashion-MNIST and 10^{-4} for the rest. The value of λ_{aug} is 0.01 for the first 40 epochs, followed by 0.5 till 80 epochs and 1 thereafter. We use a batch-size of 100 for MNIST and Fashion-MNIST and 30 for the others. The value of the scale-factor for the reconstruction term, λ_{recons} , is kept at 0.5 for 40 epochs after which we keep it as 0.1 for CIFAR10 and STL-10.

A. Datasets

- **Toy Data** used in the GMVAE paper, which is a 2D dataset consisting of 10000 points from the arcs of 5 circles. No augmentation is used for this.
- **MNIST** [41] consists of 70000 black and white images of handwritten digits of size 28x28 split into 60000 training images and 10000 testing images. We use random rotations between 3° to -3° for the augmentation.
- **Fashion MNIST** [42] consists of grayscale images of fashion items like sandals, T-shirts, handbag etc. It has the same size and split as MNIST but a greater complexity, making it more challenging. We add random brightness variations within a delta factor of 0.2 followed by an image binarization by checking if the pixel value is greater than a uniform random value.
- **STL-10** [27] consists of 13000 labelled and 100000 unlabelled colour images of real world objects from

Method	Remarks	MNIST	Fashion-MNIST	STL-10	CIFAR10	CIFAR100	FRGCv2
k -means on pixels	-	53.49	47 [†]	22.0	20.4	-	-
IMSAT [28]	use ResNet [30]	98.4	-	94.1*	45.6*	27.5*	-
VaDE [2]	-	94.46	57.8 [‡]	84.5*	-	-	-
Sarfaraz <i>et al.</i> [24]	Generative, use ResNet [30]	91.89	-	95.24*	-	-	-
DDG [33]	-	97.58	-	90.59*	-	-	-
Das <i>et al.</i> [16]	Unsupervised, Pre-training	98.93	-	-	44.19	25.4	47.28
Yang <i>et al.</i> [17]		98	66.2	-	-	-	-
Nina <i>et al.</i> [15]		96.8	70.98	-	30.9	-	-
JULE [7]	Purely Unsupervised	96.1	56.3 [‡]	-	-	-	46
DEC [18]		84.3	61.8 [†]	35.9	-	14.3*	37.8
DEPICT [21]		96.3	39.2 [‡]	-	-	-	47
RDEC [22]		98.41	-	21.27	-	-	-
DEC-DA [40]		98.6	58.6	-	-	-	-
Haeusser <i>et al.</i> [25]		98.7	-	38.9	26.7	-	43.7
GMVAE [1]		96.92	59.56 [†]	25.36 [†]	24.74 [†]	10.10 [†]	21.24 [†]
ClusterGAN [31]		95	63	-	-	-	-
Ghasedi <i>et al.</i> [32]		96.4	-	42.3	41.2	-	47.6
Ours		Generative, Purely Unsupervised	98.2 95.4 ($k = 8$) 98.4 ($k = 20$)	71.72	43.9	44.5	25.6 ($k = 100$)

TABLE I: Clustering accuracy (%). * - Use pre-trained features as input. † - Results reported using source code. ‡ - Results reported in [40]. We put a "-" where results are unavailable. For our results, we use $k = 10$ except when specified. The methods above the red line, although unsupervised, use pre-trained features as input which itself boosts their performance by a big margin. Our main comparison is with unsupervised methods, shown below the red line. (Best viewed in colour)

ImageNet like cats, trucks, birds, ships etc. of size 96x96. The 13000 labelled images are split into 5000 training images and 8000 testing images. It is highly complex compared to MNIST and Fashion-MNIST. We augment the data using random shifts within a factor of 0.2 times the image dimensions, random rotations between 3° to -3°, random shears within a factor of 0.2 and random vertical and horizontal flips.

- **CIFAR10** and **CIFAR100** [26] are diverse datasets of real world images. 9 classes in CIFAR10 are similar to STL-10. Achieving good clustering accuracy on them in an unsupervised manner is truly challenging. They consist of 50000 training and 10000 testing images each with 10 classes in CIFAR10 and 100 in CIFAR100. We use the same augmentations as in STL-10.
- **FRGCv2** consists of 2462 images having 20 face samples. It's a subset of the larger FRGC dataset. We use the subset provided by the authors of [7], which is what other methods report their results on as well.

B. Network Architecture

The network architectures for the prior and posterior networks are given in Tables IIa and IIb respectively. The z prior consists of fully connected networks as there aren't any images involved. The convolutional part of the posterior networks for CIFAR10, STL-10, CIFAR100 entail complex representations which are thereby, captured better by weight-sharing.

VI. RESULTS

A. Clustering Results

Firstly, to see the representation capability of our method, we show the latent space learnt by our method & GMVAE on the synthetic dataset given in their paper. As seen in Fig. 3, our clusters are tightly packed individually & more spread

overall compared to GMVAE. Good latent space separation is key for clustering as it can partition the data in a better way.

Furthermore, to see how well the clusters are separated in the latent space, we visualize the means of the posterior and the prior using t-SNE [43]. The posterior representation of a subset of testing images can be seen in Fig. 4. For the prior, we randomly generate 500 samples from each cluster, to see how well the prior has been learnt. This can be seen in Fig. 5. Evidently, we can see a clear separation between clusters from the beginning itself (11 epochs), which gets refined as the iterations go on. This not only shows the discriminatory nature of the learnt representations, but also the speed with which it becomes a characteristic of the latent space.

For evaluating our method, we use the *unsupervised clustering accuracy (ACC)* [18], which can be given as

$$ACC = \max_{m \in \mathcal{M}} \frac{\sum_{i=1}^N 1(l_i == m(c_i))}{N} \quad (19)$$

where N is the total no. of images, l_i and c_i are the label and cluster for the i^{th} image respectively and \mathcal{M} is the set of possible one-to-one mappings between clusters and labels.

Our accuracy can be seen in Table I. It is important to note that we only lag behind methods which use pre-trained ResNet [30] features as input, which itself is an extremely informative representation, giving them a large advantage as it has already covered thousands of samples of the classes present in STL-10 and CIFAR10 in a supervised manner. In comparison with Das *et al.* [16] who use a lot of epochs for pre-training, we perform better especially on more complex datasets like Fashion-MNIST, STL-10, CIFAR10 and CIFAR100.

Our main comparison is with unsupervised methods, shown below the red line in Table I. Our performance significantly exceeds them in all the complex real world datasets achieving state-of-the-art results among unsupervised methods. Our performance exceeds that of generative methods [31], [32] on all

Dataset	Type of Layer	$p_{\beta}(\mathbf{z} k)$	$p_{\theta}(\mathbf{x} \mathbf{z})$
Toy Data	Input	Input \mathbf{e}_k [batch, K]	Input \mathbf{z} [batch, 2]
	Hidden	FCN 120	FCN 120 FCN 120
	Output	μ_{β} , FCN 2, No activation σ_{β}^2 , FCN 2, softplus activation	μ_{θ} , FCN 2, No activation σ_{θ}^2 , FCN 2, softplus activation
MNIST Fashion Mnist	Input	Input \mathbf{e}_k [batch, K]	Input \mathbf{z} [batch, 32]
	Hidden	-	FCN 512 FCN 2304, reshape [3,3,256] UpConv 128 [3x3] UpConv 128 [3x3], stride=1, padding=valid UpConv 128 [3x3], stride=1, padding=valid UpConv 64 [3x3], stride=1, padding=valid UpConv 64 [3x3] UpConv 64 [3x3], stride=1, padding=valid
	Output	μ_{β} , FCN 32, No activation σ_{β}^2 , FCN 32, softplus activation	μ_{θ} , UpConv 1 [3x3], stride=1, padding=valid, No activation
CIFAR10 STL-10 CIFAR100	Input	Input \mathbf{e}_k [batch, K]	Input \mathbf{z} [batch, 32]
	Hidden	FCN 64 FCN 64	FCN 3072, reshape [4,4,192] UpConv 192 [3x3], padding=valid UpConv 192 [3x3], stride=1, padding=valid UpConv 96 [3x3], stride=1, padding=valid UpConv 96 [3x3], padding=valid, UpConv 96 [3x3], stride=1, padding=valid UpConv 96 [3x3] stride=1, padding=valid
	Output	μ_{β} , FCN 32, No activation σ_{β}^2 , FCN 32, softplus activation	μ_{θ} , UpConv 3 [3x3], stride=1, padding=valid, No activation σ_{θ}^2 , UpConv 3 [3x3], stride=1, padding=valid, softplus activation
FRGCv2	Input	Input \mathbf{e}_k [batch, K]	Input \mathbf{z} [batch, 32]
	Hidden	FCN 64	FCN 512, reshape [4,4,32] UpConv 64 [3x3] UpConv 32 [3x3]
	Output	μ_{β} , FCN 32, No activation σ_{β}^2 , FCN 32, softplus activation	μ_{θ} , UpConv 1 [5x5], No activation σ_{θ}^2 , UpConv 1 [5x5], softplus activation

(a) Network architecture for the prior probabilities $p_{\beta}(\mathbf{z}|k)$ & $p_{\theta}(\mathbf{x}|\mathbf{z})$.

Dataset	Type of Layer	$q_{\omega}(k \mathbf{x})$	$q_{\psi}(\mathbf{z} \mathbf{x}, k)$
Toy Data	Input	Input \mathbf{x} [batch, 2]	Input \mathbf{x} [batch, 2], \mathbf{e}_k [batch, K]
	Hidden	FCN 120 FCN K No activation	FCN 10 concat \mathbf{e}_k FCN 10
	Output	$\hat{\mathbf{e}}_k$, softmax activation	μ_{ψ} , FCN 2, No activation σ_{ψ}^2 , FCN 2, softplus activation
MNIST Fashion Mnist	Input	Input \mathbf{x} [batch, 28, 28, 1]	Input \mathbf{x} [batch, 28, 28, 1], \mathbf{e}_k [batch, K]
	Hidden	Conv 64 [3x3], stride=1, padding=valid Conv 64 [3x3], stride=1, padding=valid Conv 64 [3x3], stride=1, padding=valid Conv 128 [3x3], padding=valid Conv 128 [3x3], stride=1, padding=valid Conv 128 [3x3], stride=1, padding=valid Conv 256 [3x3], padding=valid	Flatten, concat \mathbf{e}_k FCN 256 FCN 512
	Output	$\hat{\mathbf{e}}_k$, softmax activation	μ_{ψ} , FCN 32, No activation σ_{ψ}^2 , FCN 32, softplus activation
CIFAR10 STL-10 CIFAR100	Input	Input \mathbf{x} [batch, 32, 32, 3]	Input \mathbf{x} [batch, 32, 32, 3], \mathbf{e}_k [batch, K]
	Hidden	Conv 96 [3x3], stride=1, padding=valid Conv 96 [3x3], stride=1, padding=valid Conv 96 [3x3], padding=valid Conv 192 [3x3], stride=1, padding=valid Conv 192 [3x3], stride=1, padding=valid Conv 192 [3x3], padding=valid	Flatten, concat \mathbf{e}_k FCN 512 FCN 512
	Output	$\hat{\mathbf{e}}_k$, softmax activation	μ_{ψ} , FCN 32, No activation σ_{ψ}^2 , FCN 32, softplus activation
FRGCv2	Input	Input \mathbf{x} [batch, 32, 32, 3]	Input \mathbf{x} [batch, 32, 32, 3], \mathbf{e}_k [batch, K]
	Hidden	Conv 32 [5x5] Conv 64 [3x3] Conv 128 [3x3] Conv 32 [1x1]	Flatten, concat \mathbf{e}_k FCN 512 FCN 512
	Output	$\hat{\mathbf{e}}_k$, softmax activation	μ_{ψ} , FCN 32, No activation σ_{ψ}^2 , FCN 32, softplus activation

(b) Network architecture for the prior probabilities $p_{\beta}(\mathbf{z}|k)$ & $p_{\theta}(\mathbf{x}|\mathbf{z})$.

TABLE II: Network architectures for (a) posterior probabilities and (b) prior probabilities. All the Conv and UpConv layers have a stride of 2 with padding as 'same' unless specified otherwise. For MNIST and Fashion MNIST, all layers have ReLU activations unless specified. For others, all layers have leaky ReLU activations unless specified. K - number of clusters (Toy Data - 5, MNIST, Fashion MNIST, STL-10, CIFAR10 - 10, FRGCv2 - 20, CIFAR100 - 100)

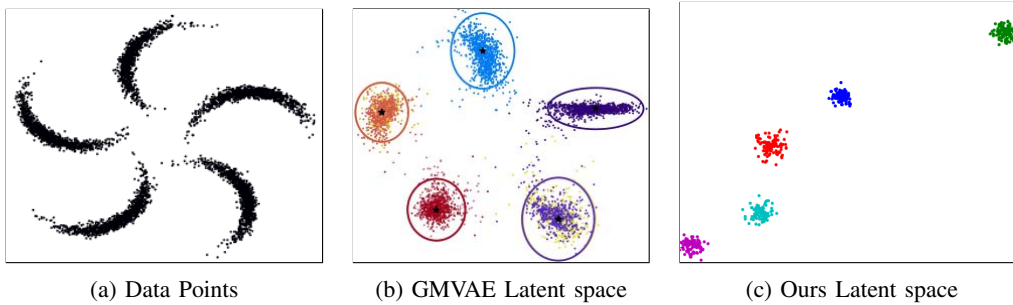


Fig. 3: Toy dataset & learnt latent spaces




Fig. 4: t-SNE Visualization of posterior latent representation of MNIST samples during training




Fig. 5: t-SNE Visualization of samples drawn from the learnt prior latent distribution of MNIST during training

datasets, showing that the inference power of VAEs is greater than GANs for clustering, where a strong representation of the data distribution is important. We perform better than GMVAE [1] and VaDE [2] on all datasets showing the effectiveness of our approach to using a Gaussian Mixture latent space.

B. STL-CIFAR Transfer Learning

CIFAR10 and STL-10 are inherently similar, having 9 same classes out of 10 and has some scope for transfer learning. To show such a capability of our method, we test the model trained on CIFAR10 using STL-10 and vice-versa on these common classes. The results of this are shown in Table III.

Trained on	Tested on	Accuracy
STL-10	CIFAR10	38.8
CIFAR10	STL-10	41.03

TABLE III: Transfer learning between CIFAR10 and STL-10

C. Imbalanced Clusters

To study our performance on imbalanced data, we follow the imbalanced sampling mechanism as in DEC [18] for MNIST. We define which is the probability of retaining samples of

Method	r_{min}				
	0.1	0.3	0.5	0.7	0.9
k -means	47.14	49.93	53.65	54.16	54.39
AE+ k -means	66.82	74.91	77.93	80.04	81.31
DEC	70.10	80.92	82.68	84.69	85.41
Ours	89.6	90.3	96.1	96.9	97.1

TABLE IV: Clustering accuracy (%) for MNIST with imbalanced clusters using different min. retention rates r_{min} .

class 0 is defined as a minimum retention rate r_{min} , for class 9 the retention probability is 1 and the for other classes, it varies linearly from r_{min} to 1. The ratio of the sizes of the largest and smallest clusters is $1/r_{min}$. We perform significantly better than other methods as shown in Table IV.

D. Generated Samples

In order to see the effectiveness of our latent space representation, we generate images to see if captures the underlying data distribution in a meaningful manner. Since our approach is inherently a generative one, it should be able to extrapolate samples from the latent space, as is the case with any generative approaches. The ability to generate high quality samples is the true test for any generative approach.

As it can be seen in Fig. 6, our learned model is able to generate realistic samples from each cluster. Moreover, one more important inference that can be drawn is that our method learns to distinguish properly between classes since it is able to generate distinct samples from each of the clusters. This shows that our generative approach works well not only in distinguishing between different clusters but also in being able to generate realistic samples given latent data.




Fig. 6: Images generated from each learned cluster. Each row represents a cluster and each column is an independent Bernoulli distribution $p_{\theta}(\mathbf{x}|\mathbf{z})$ given a \mathbf{z} sampled from the Gaussian Distribution $p_{\beta}(\mathbf{z}|k)$ of the k^{th} cluster.

E. Latent Space Interpolation


To see how well the learning has progressed, we visualize how variations in the latent space correspond to variations in generated samples. We perform 2 types of interpolations, one in the continuous \mathbf{z} space and one in the k space. Though k is a discrete variable, we vary it continuously to see how data generation for unknown inputs looks like. In our model, we have two sets of variables on which the output of the generative model depends on. The first is the categorical distribution for k , followed by a Gaussian distribution corresponding to the k^{th} cluster, for obtaining \mathbf{z} . We generate samples by interpolating in both these spaces to see how the generated samples are, giving us further insight into the extent to which each of these variables influence the generation process.

1) *z-Interpolation*: We choose two clusters k_1 & k_2 and sample latents \mathbf{z}_1 & \mathbf{z}_2 from their respective Gaussians. We calculate our interpolated latent as $\tilde{\mathbf{z}} = \alpha\mathbf{z}_1 + (1 - \alpha)\mathbf{z}_2$ by varying α from 0 to 1, from which we generate our output.

As shown in Fig. 7a, smooth variations in the latent space generate smooth transitions in the image space. This shows that our distributions in the latent space capture the data distribution accurately. This is an important result showing the ability of our method to accurately capture changes in the latent space and reflect them in the input space.

2) *k-Interpolation*: We choose two clusters k_1 and k_2 and linearly interpolate between their one-hot vector representations \mathbf{e}_{k_1} and \mathbf{e}_{k_2} . We generate $\tilde{\mathbf{e}} = \alpha\mathbf{e}_{k_1} + (1 - \alpha)\mathbf{e}_{k_2}$ by varying α from 0 to 1. From $\tilde{\mathbf{e}}$, we generate the parameters of the prior for sampling \mathbf{z} from which we generate our output.

From Fig. 7b, it can be said that our method is able to extrapolate data not just from continuous variations in the latent space but more importantly from continuous variations



(a) \mathbf{z} -Interpolation (b) k -Interpolation

Fig. 7: Images generated by interpolating \mathbf{z} and k . The first column is for $\alpha = 0$ and the last for $\alpha = 1$.

Method	MNIST	Fashion-MNIST	CIFAR-10
VAE with GMM prior	96.8	65.3	37.1
With $\lambda_{recons} = 0.4$	-	-	38.5
With $\lambda_{recons} = 0.5$	-	-	39.1
With $\lambda_{recons} = 0.65$	-	-	41.8
With $\lambda_{recons} = 0.7$	-	-	39.2
With L_2 augmentation	97.2	66.4	42.7
With \mathcal{W}_2 augmentation	97.7	67.6	43.8
Ours-final $k = 8$	95.4	-	-
Ours-final $k = 10$	98.2	71.72	44.5
Ours-final $k = 20$	98.4	-	-

TABLE V: Ablation study of our proposed method done by incorporating each of our proposed additions individually. We report the clustering accuracy (%) of each variant.

in a discrete space. The main takeaway is that though we train our model using only one-hot vectors, our method is able to generate meaningful images representing the the interpolated vector $\tilde{\mathbf{e}}$. This result shows that our model is able to learn a truly effective representation for each cluster, where the data distribution is preserved not just in the latent space (\mathbf{z}) but in the categorical space (k) as well. The k -interpolation is not as smooth as the \mathbf{z} -interpolation mainly because the former is a discrete space whereas the latter is a smooth continuous one.

F. Ablation study

We perform an ablation study to understand the effect of different parts of our loss where we compare the following:

- A vanilla implementation of our proposed VAE.
- Adding a scale factor to the reconstruction term.
- Adding just the L_2 loss for the cluster predictions.
- Adding just the \mathcal{W}_2 loss for the predicted posteriors.
- The final proposed loss function.

The results for this study are given in Table V. As it can be seen, our vanilla method itself performs decently. Adding just the scale-factor for STL-10 helps improve the accuracy. Just adding the L_2 loss increases the accuracy by a small margin, while adding the \mathcal{W}_2 loss increases it by a larger margin. This shows that enforcing a consistent latent space representation helps make the learning more robust. Finally, we get the best performance by combining all our additions together.

VII. CONCLUSION AND FUTURE WORKS

We propose an image clustering method using VAEs with a GMM prior where each component represents a cluster. The prior is learned jointly with the posterior, which in turn learns a strong latent representation that leads to accurate clustering, as shown in extensive experiments. Our method doesn't require any pre-training and can be trained from scratch in an end-to-end manner. We show results on a variety of datasets ranging from simple handwritten digits to complex real world

objects, achieving state-of-the-art clustering accuracy among purely unsupervised methods. To the best of our knowledge, we are the first to achieve such significant results on real world datasets in a purely unsupervised manner. Moreover, we achieve comparable accuracy between similar datasets (CIFAR10/STL-10) in a transfer learning scenario. With the rise in the popularity of GANs, one way to move forward could be to replace the generator network in a GAN with our VAE model, thereby leading to a hybrid VAE-GAN model, that can has the strong latent representational powers of a VAE and the realistic generative powers of a GAN. Furthermore, adding explicit inference constraints on the GAN prior, somewhat as an increment along the lines of ClusterGAN [31], could be another way to approach the VAE-GAN hybrid architecture.

REFERENCES

- [1] N. Dilokthanakul, P. A. Mediano, M. Garnelo, M. C. Lee, H. Salimbeni, K. Arulkumaran, and M. Shanahan, "Deep unsupervised clustering with gaussian mixture variational autoencoders," *arXiv preprint arXiv:1611.02648*, 2016. **1, 2, 3, 4, 6, 8**
- [2] Z. Jiang, Y. Zheng, H. Tan, B. Tang, and H. Zhou, "Variational deep embedding: an unsupervised and generative approach to clustering," in *International Joint Conference on Artificial Intelligence*, 2017. **1, 2, 3, 6, 8**
- [3] A. Y. Ng, M. I. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Advances in Neural Information Processing Systems*, 2002. **1**
- [4] J. Ye, Z. Zhao, and M. Wu, "Discriminative k-means for clustering," in *Advances in Neural Information Processing Systems*, 2008. **1**
- [5] Y. Yang, D. Xu, F. Nie, S. Yan, and Y. Zhuang, "Image clustering using local discriminant models and global integration," *IEEE Transactions on Image Processing*, 2010. **1**
- [6] U. Von Luxburg, "A tutorial on spectral clustering," *Statistics and computing*, 2007. **1**
- [7] J. Yang, D. Parikh, and D. Batra, "Joint unsupervised learning of deep representations and image clusters," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. **1, 2, 6**
- [8] Y. Gdalyahu, D. Weinshall, and M. Werman, "Self-organization in vision: stochastic clustering for image segmentation, perceptual grouping, and image database organization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2001. **1**
- [9] J. MacQueen *et al.*, "Some methods for classification and analysis of multivariate observations," *Berkeley Symposium on Mathematical Statistics and Probability*, 1967. **1**
- [10] C. M. Bishop, *Pattern recognition and machine learning*. Springer, 2006. **1**
- [11] B. Yang, X. Fu, N. D. Sidiropoulos, and M. Hong, "Towards k-means-friendly spaces: Simultaneous deep learning and clustering," in *International Conference on Machine Learning*, 2017. **1**
- [12] F. Tian, B. Gao, Q. Cui, E. Chen, and T.-Y. Liu, "Learning deep representations for graph clustering," in *AAAI Conference on Artificial Intelligence*, 2014. **1, 2**
- [13] C. Song, F. Liu, Y. Huang, L. Wang, and T. Tan, "Auto-encoder based data clustering," in *Iberoamerican Congress on Pattern Recognition*, 2013. **1, 2**
- [14] L. N. Vaserstein, "Markov processes over denumerable products of spaces, describing large systems of automata," *Problemy Peredachi Informatsii*, 1969. **2, 5**
- [15] O. Nina, J. Moody, and C. Milligan, "A decoder-free approach for unsupervised clustering and manifold learning with random triplet mining," in *IEEE International Conference on Computer Vision Workshops*, 2019. **2, 6**
- [16] D. Das, R. Ghosh, and B. Bhowmick, "Deep representation learning characterized by inter-class separation for image clustering," in *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019. **2, 6**
- [17] X. Yang, C. Deng, F. Zheng, J. Yan, and W. Liu, "Deep spectral clustering using dual autoencoder network," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. **2, 6**
- [18] J. Xie, R. Girshick, and A. Farhadi, "Unsupervised deep embedding for clustering analysis," in *International Conference on Machine Learning*, 2016. **2, 6, 8**
- [19] J. Chang, L. Wang, G. Meng, S. Xiang, and C. Pan, "Deep adaptive image clustering," in *IEEE International Conference on Computer Vision*, 2017. **2**
- [20] X. Ji, J. F. Henriques, and A. Vedaldi, "Invariant information clustering for unsupervised image classification and segmentation," in *IEEE International Conference on Computer Vision*, 2019. **2**
- [21] K. G. Dizaji, A. Herandi, C. Deng, W. Cai, and H. Huang, "Deep clustering via joint convolutional autoencoder embedding and relative entropy minimization," in *IEEE International Conference on Computer Vision*, 2017. **2, 6**
- [22] Y. Tao, K. Takagi, and K. Nakata, "Rdec: Integrating regularization into deep embedded clustering for imbalanced datasets," in *Asian Conference on Machine Learning*, 2018. **2, 6**
- [23] C.-C. Hsu and C.-W. Lin, "Cnn-based joint clustering and representation learning with feature drift compensation for large-scale image data," *IEEE Transactions on Multimedia*, 2018. **2**
- [24] S. Sarfraz, V. Sharma, and R. Stiefelhagen, "Efficient parameter-free clustering using first neighbor relations," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. **2, 6**
- [25] P. Haeusser, J. Plapp, V. Golkov, E. Aljalbout, and D. Cremers, "Associative deep clustering: Training a classification network with no labels," in *German Conference on Pattern Recognition (GCPR)*, 2018. **2, 6**
- [26] A. Krizhevsky, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009. **2, 3, 6**
- [27] A. Coates, A. Ng, and H. Lee, "An analysis of single-layer networks in unsupervised feature learning," in *International Conference on Artificial Intelligence and Statistics*, 2011. **2, 5**
- [28] W. Hu, T. Miyato, S. Tokui, E. Matsumoto, and M. Sugiyama, "Learning discrete representations via information maximizing self-augmented training," in *International Conference on Machine Learning*, 2017. **2, 6**
- [29] A. Krause, P. Perona, and R. G. Gomes, "Discriminative clustering by regularized information maximization," in *Advances in Neural Information Processing Systems*, 2010. **3**
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. **3, 6**
- [31] S. Mukherjee, H. Asnani, E. Lin, and S. Kannan, "Clustergan: Latent space clustering in generative adversarial networks," in *AAAI Conference on Artificial Intelligence*, 2019. **3, 6, 10**
- [32] K. Ghasedi, X. Wang, C. Deng, and H. Huang, "Balanced self-paced learning for generative adversarial clustering network," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2019. **3, 6**
- [33] L. Yang, N.-M. Cheung, J. Li, and J. Fang, "Deep clustering by gaussian mixture variational autoencoders with graph embedding," in *IEEE International Conference on Computer Vision*, 2019. **3, 6**
- [34] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," *arXiv preprint arXiv:1401.4082*, 2014. **3**
- [35] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013. **3**
- [36] J. L. W. V. Jensen, "Sur les fonctions convexes et les inegalités entre les valeurs moyennes," *Acta mathematica*, 1906. **3**
- [37] D. Dowson and B. Landau, "The fréchet distance between multivariate normal distributions," *Journal of Multivariate Analysis*, 1982. **5**
- [38] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, "Tensorflow: a system for large-scale machine learning," in *OSDI*, 2016. **5**
- [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014. **5**
- [40] X. Guo, E. Zhu, X. Liu, and J. Yin, "Deep embedded clustering with data augmentation," in *Asian Conference on Machine Learning*, 2018. **6**
- [41] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, 1998. **5**
- [42] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017. **5**
- [43] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of Machine Learning Research (JMLR)*, 2008. **6**