

Article

Variational Information Bottleneck for Unsupervised Clustering: Deep Gaussian Mixture Embedding

Yiğit Uğur ^{1,2,*} , George Arvanitakis ² and Abdellatif Zaidi ^{1,*} ¹ Laboratoire d'informatique Gaspard-Monge, Université Paris-Est, 77454 Champs-sur-Marne, France² Mathematical and Algorithmic Sciences Lab, Paris Research Center, Huawei Technologies, 92100 Boulogne-Billancourt, France; george.arvanitakis@huawei.com

* Correspondence: ygtugur@gmail.com (Y.U.); abdellatif.zaidi@u-pem.fr (A.Z.)

Received: 3 December 2019; Accepted: 9 February 2020; Published: 13 February 2020



Abstract: In this paper, we develop an unsupervised generative clustering framework that combines the variational information bottleneck and the Gaussian mixture model. Specifically, in our approach, we use the variational information bottleneck method and model the latent space as a mixture of Gaussians. We derive a bound on the cost function of our model that generalizes the Evidence Lower Bound (ELBO) and provide a variational inference type algorithm that allows computing it. In the algorithm, the coders' mappings are parametrized using neural networks, and the bound is approximated by Markov sampling and optimized with stochastic gradient descent. Numerical results on real datasets are provided to support the efficiency of our method.

Keywords: clustering; unsupervised learning; Gaussian mixture model; information bottleneck

1. Introduction

Clustering consists of partitioning a given dataset into various groups (clusters) based on some similarity metric, such as the Euclidean distance, L_1 norm, L_2 norm, L_∞ norm, the popular logarithmic loss measure, or others. The principle is that each cluster should contain elements of the data that are closer to each other than to any other element outside that cluster, in the sense of the defined similarity measure. If the joint distribution of the clusters and data is not known, one should operate blindly in doing so, i.e., using only the data elements at hand; and the approach is called unsupervised clustering [1,2]. Unsupervised clustering is perhaps one of the most important tasks of unsupervised machine learning algorithms currently, due to a variety of application needs and connections with other problems.

Clustering can be formulated as follows. Consider a dataset that is composed of N samples $\{\mathbf{x}_i\}_{i=1}^N$, which we wish to partition into $|\mathcal{C}| \geq 1$ clusters. Let $\mathcal{C} = \{1, \dots, |\mathcal{C}|\}$ be the set of all possible clusters and C designate a categorical random variable that lies in \mathcal{C} and stands for the index of the actual cluster. If \mathbf{X} is a random variable that models elements of the dataset, given that $\mathbf{X} = \mathbf{x}_i$ induces a probability distribution on \mathcal{C} , which the learner should learn, thus mathematically, the problem is that of estimating the values of the unknown conditional probability $P_{C|\mathbf{X}}(\cdot|\mathbf{x}_i)$ for all elements \mathbf{x}_i of the dataset. The estimates are sometimes referred to as the assignment probabilities.

Examples of unsupervised clustering algorithms include the very popular K -means [3] and Expectation Maximization (EM) [4]. The K -means algorithm partitions the data in a manner that the Euclidean distance among the members of each cluster is minimized. With the EM algorithm, the underlying assumption is that the data comprise a mixture of Gaussian samples, namely a Gaussian Mixture Model (GMM); and one estimates the parameters of each component of the GMM while simultaneously associating each data sample with one of those components. Although they offer some advantages in the context of clustering, these algorithms suffer from some strong limitations.

For example, it is well known that the K -means is highly sensitive to both the order of the data and scaling; and the obtained accuracy depends strongly on the initial seeds (in addition to that, it does not predict the number of clusters or K -value). The EM algorithm suffers mainly from slow convergence, especially for high-dimensional data.

Recently, a new approach has emerged that seeks to perform inference on a transformed domain (generally referred to as latent space), not the data itself. The rationale is that because the latent space often has fewer dimensions, it is more convenient computationally to perform inference (clustering) on it rather than on the high-dimensional data directly. A key aspect then is how to design a latent space that is amenable to accurate low-complexity unsupervised clustering, i.e., one that preserves only those features of the observed high-dimensional data that are useful for clustering while removing all redundant or non-relevant information. Along this line of work, we can mention [5], which utilized Principal Component Analysis (PCA) [6,7] for dimensionality reduction followed by K -means for clustering the obtained reduced dimension data; or [8], which used a combination of PCA and the EM algorithm. Other works that used alternatives for the linear PCA include kernel PCA [9], which employs PCA in a non-linear fashion to maximize variance in the data.

Thisby's Information Bottleneck (IB) method [10] formulates the problem of finding a good representation \mathbf{U} that strikes the right balance between capturing all information about the categorical variable C that is contained in the observation \mathbf{X} and using the most concise representation for it. The IB problem can be written as the following Lagrangian optimization:

$$\min_{P_{\mathbf{U}|\mathbf{X}}} I(\mathbf{X}; \mathbf{U}) - sI(C; \mathbf{U}), \quad (1)$$

where $I(\cdot; \cdot)$ denotes Shannon's mutual information and s is a Lagrange-type parameter, which controls the trade-off between accuracy and regularization. In [11,12], a text clustering algorithm is introduced for the case in which the joint probability distribution of the input data is known. This text clustering algorithm uses the IB method with an annealing procedure, where the parameter s is increased gradually. When $s \rightarrow 0$, the representation \mathbf{U} is designed with the most compact form, i.e., $|\mathcal{U}| = 1$, which corresponds to the maximum compression. By gradually increasing the parameter s , the emphasis on the relevance term $I(C; \mathbf{U})$ increases, and at a critical value of s , the optimization focuses on not only the compression, but also the relevance term. To fulfill the demand on the relevance term, this results in the cardinality of \mathbf{U} bifurcating. This is referred as a phase transition of the system. The further increases in the value of s will cause other phase transitions, hence additional splits of \mathbf{U} until it reaches the desired level, e.g., $|\mathcal{U}| = |C|$.

However, in the real-world applications of clustering with large-scale datasets, the joint probability distributions of the datasets are unknown. In practice, the usage of Deep Neural Networks (DNN) for unsupervised clustering of high-dimensional data on a lower dimensional latent space has attracted considerable attention, especially with the advent of Autoencoder (AE) learning and the development of powerful tools to train them using standard backpropagation techniques [13,14]. Advanced forms include Variational Autoencoders (VAE) [13,14], which are generative variants of AE that regularize the structure of the latent space, and the more general Variational Information Bottleneck (VIB) of [15], which is a technique that is based on the information bottleneck method and seeks a better trade-off between accuracy and regularization than VAE via the introduction of a Lagrange-type parameter s , which controls that trade-off and whose optimization is similar to deterministic annealing [12] or stochastic relaxation.

In this paper, we develop an unsupervised generative clustering framework that combines VIB and the Gaussian mixture model. Specifically, in our approach, we use the variational information bottleneck method and model the latent space as a mixture of Gaussians. We derive a bound on the cost function of our model that generalizes the Evidence Lower Bound (ELBO) and provide a variational inference type algorithm that allows computing it. In the algorithm, the coders' mappings are parametrized using Neural Networks (NN), and the bound is approximated by Markov sampling and

optimized with stochastic gradient descent. Furthermore, we show how tuning the hyper-parameter s appropriately by gradually increasing its value with iterations (number of epochs) results in a better accuracy. Furthermore, the application of our algorithm to the unsupervised clustering of various datasets, including the MNIST [16], REUTERS [17], and STL-10 [18], allows a better clustering accuracy than previous state-of-the-art algorithms. For instance, we show that our algorithm performs better than the Variational Deep Embedding (VaDE) algorithm of [19], which is based on VAE and performs clustering by maximizing the ELBO. Our algorithm can be seen as a generalization of the VaDE, whose ELBO can be recovered by setting $s = 1$ in our cost function. In addition, our algorithm also generalizes the VIB of [15], which models the latent space as an isotropic Gaussian, which is generally not expressive enough for the purpose of unsupervised clustering. Other related works, which are of lesser relevance to the contribution of this paper, are the Deep Embedded Clustering (DEC) of [20] and the Improved Deep Embedded Clustering (IDEC) of [21,22]. For a detailed survey of clustering with deep learning, the readers may refer to [23].

To the best of our knowledge, our algorithm performs the best in terms of clustering accuracy by using deep neural networks without any prior knowledge regarding the labels (except the usual assumption that the number of classes is known) compared to the state-of-the-art algorithms of the unsupervised learning category. In order to achieve the outperforming accuracy: (i) we derive a cost function that contains the IB hyperparameter s that controls optimal trade-offs between the accuracy and regularization of the model; (ii) we use a lower bound approximation for the KL term in the cost function, that does not depend on the clustering assignment probability (note that the clustering assignment is usually not accurate in the beginning of the training process); and (iii) we tune the hyperparameter s by following an annealing approach that improves both the convergence and the accuracy of the proposed algorithm.

Throughout this paper, we use the following notation. Uppercase letters are used to denote random variables, e.g., X ; lowercase letters are used to denote realizations of random variables, e.g., x ; and calligraphic letters denote sets, e.g., \mathcal{X} . The cardinality of a set \mathcal{X} is denoted by $|\mathcal{X}|$. Probability mass functions (pmfs) are denoted by $P_X(x) = \Pr\{X = x\}$ and, sometimes, for short, as $p(x)$. Boldface uppercase letters denote vectors or matrices, e.g., \mathbf{X} , where context should make the distinction clear. We denote the covariance of a zero mean, complex-valued, vector \mathbf{X} by $\Sigma_{\mathbf{X}} = \mathbb{E}[\mathbf{X}\mathbf{X}^\dagger]$, where $(\cdot)^\dagger$ indicates the conjugate transpose. For random variables X and Y , the entropy is denoted as $H(X)$, i.e., $H(X) = \mathbb{E}_{P_X}[-\log P_X]$, and the mutual information is denoted as $I(X;Y)$, i.e., $I(X;Y) = H(X) - H(X|Y) = H(Y) - H(Y|X) = \mathbb{E}_{P_{X,Y}}[\log \frac{P_{X,Y}}{P_X P_Y}]$. Finally, for two probability measures P_X and Q_X on a random variable $X \in \mathcal{X}$, the relative entropy or Kullback–Leibler divergence is denoted as $D_{\text{KL}}(P_X \| Q_X)$, i.e., $D_{\text{KL}}(P_X \| Q_X) = \mathbb{E}_{P_X}[\log \frac{P_X}{Q_X}]$.

2. Proposed Model

In this section, we explain the proposed model, the so-called Variational Information Bottleneck with Gaussian Mixture Model (VIB-GMM), in which we use the VIB framework and model the latent space as a GMM. The proposed model is depicted in Figure 1, where the parameters π_c, μ_c, Σ_c , for all values of $c \in \mathcal{C}$, are to be optimized jointly with those of the employed NNs as instantiation of the coders. Furthermore, the assignment probabilities are estimated based on the values of latent space vectors instead of the observations themselves, i.e., $P_{C|X} = Q_{C|U}$. In the rest of this section, we elaborate on the inference and generative network models for our method.

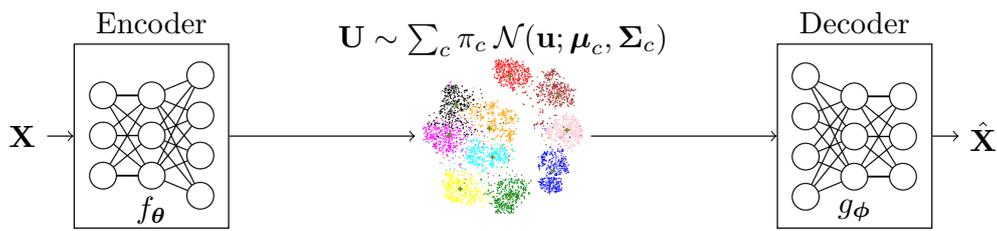


Figure 1. Variational Information Bottleneck with Gaussian mixtures.

2.1. Inference Network Model

We assume that observed data \mathbf{x} are generated from a GMM with $|\mathcal{C}|$ components. Then, the latent representation \mathbf{u} is inferred according to the following procedure:

1. One of the components of the GMM is chosen according to a categorical variable C .
2. The data \mathbf{x} are generated from the c th component of the GMM, i.e., $P_{\mathbf{X}|C} \sim \mathcal{N}(\mathbf{x}; \tilde{\mu}_c, \tilde{\Sigma}_c)$.
3. Encoder maps \mathbf{x} to a latent representation \mathbf{u} according to $P_{\mathbf{U}|\mathbf{X}} \sim \mathcal{N}(\mu_\theta, \Sigma_\theta)$.
 - 3.1. The encoder is modeled with a DNN f_θ , which maps \mathbf{x} to the parameters of a Gaussian distribution, i.e., $[\mu_\theta, \Sigma_\theta] = f_\theta(\mathbf{x})$.
 - 3.2. The representation \mathbf{u} is sampled from $\mathcal{N}(\mu_\theta, \Sigma_\theta)$.

For the inference network, shown in Figure 2, the following Markov chain holds:

$$C \ominus \mathbf{X} \ominus \mathbf{U} . \tag{2}$$

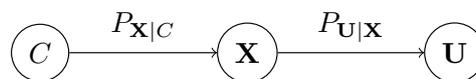


Figure 2. Inference network.

2.2. Generative Network Model

Since the encoder extracts useful representations of the dataset and we assume that the dataset is generated from a GMM, we model our latent space also with a mixture of Gaussians. To do so, the categorical variable C is embedded with the latent variable \mathbf{U} . The reconstruction of the dataset is generated according to the following procedure:

1. One of the components of the GMM is chosen according to a categorical variable C , with a prior distribution Q_C .
2. The representation \mathbf{u} is generated from the c th component, i.e., $Q_{\mathbf{U}|C} \sim \mathcal{N}(\mathbf{u}; \mu_c, \Sigma_c)$.
3. The decoder maps the latent representation \mathbf{u} to $\hat{\mathbf{x}}$, which is the reconstruction of the source \mathbf{x} by using the mapping $Q_{\mathbf{X}|\mathbf{U}}$.
 - 3.1 The decoder is modeled with a DNN g_ϕ that maps \mathbf{u} to the estimate $\hat{\mathbf{x}}$, i.e., $[\hat{\mathbf{x}}] = g_\phi(\mathbf{u})$.

For the generative network, shown in Figure 3, the following Markov chain holds:

$$C \ominus \mathbf{U} \ominus \mathbf{X} . \tag{3}$$

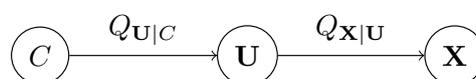


Figure 3. Generative network.

3. Proposed Method

In this section, we present our clustering method. First, we provide a general cost function for the problem of the unsupervised clustering that we study here based on the variational IB framework; and we show that it generalizes the ELBO bound developed in [19]. We then parametrize our model using NNs whose parameters are optimized jointly with those of the GMM. Furthermore, we discuss the influence of the hyperparameter s that controls optimal trade-offs between accuracy and regularization.

3.1. Brief Review of Variational Information Bottleneck for Unsupervised Learning

As described in Section 2, the stochastic encoder $P_{\mathbf{U}|\mathbf{X}}$ maps the observed data \mathbf{x} to a representation \mathbf{u} . Similarly, the stochastic decoder $Q_{\mathbf{X}|\mathbf{U}}$ assigns an estimate $\hat{\mathbf{x}}$ of \mathbf{x} based on the vector \mathbf{u} . As per the IB method [10], a suitable representation \mathbf{U} should strike the right balance between capturing all information about the categorical variable C that is contained in the observation \mathbf{X} and using the most concise representation for it. This leads to maximizing the following Lagrange problem:

$$\mathcal{L}_s(\mathbf{P}) = I(C; \mathbf{U}) - sI(\mathbf{X}; \mathbf{U}), \quad (4)$$

where $s \geq 0$ designates the Lagrange multiplier and, for convenience, \mathbf{P} denotes the conditional distribution $P_{\mathbf{U}|\mathbf{X}}$.

Instead of (4), which is not always computable in our unsupervised clustering setting, we find it convenient to maximize an upper bound of $\mathcal{L}_s(\mathbf{P})$ given by:

$$\tilde{\mathcal{L}}_s(\mathbf{P}) := I(\mathbf{X}; \mathbf{U}) - sI(\mathbf{X}; \mathbf{U}) \stackrel{(a)}{=} H(\mathbf{X}) - H(\mathbf{X}|\mathbf{U}) - s[H(\mathbf{U}) - H(\mathbf{U}|\mathbf{X})], \quad (5)$$

where (a) is due to the definition of mutual information (using the Markov chain $C \leftrightarrow \mathbf{X} \leftrightarrow \mathbf{U}$, it is easy to see that $\tilde{\mathcal{L}}_s(\mathbf{P}) \geq \mathcal{L}_s(\mathbf{P})$ for all values of \mathbf{P}). Noting that $H(\mathbf{X})$ is constant with respect to $P_{\mathbf{U}|\mathbf{X}}$, maximizing $\tilde{\mathcal{L}}_s(\mathbf{P})$ over \mathbf{P} is equivalent to maximizing:

$$\mathcal{L}'_s(\mathbf{P}) := -H(\mathbf{X}|\mathbf{U}) - s[H(\mathbf{U}) - H(\mathbf{U}|\mathbf{X})] \quad (6)$$

$$= \mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [\log P_{\mathbf{X}|\mathbf{U}} + s \log P_{\mathbf{U}} - s \log P_{\mathbf{U}|\mathbf{X}}] \right]. \quad (7)$$

For a variational distribution $Q_{\mathbf{U}}$ on \mathcal{U} (instead of the unknown $P_{\mathbf{U}}$) and a variational stochastic decoder $Q_{\mathbf{X}|\mathbf{U}}$ (instead of the unknown optimal decoder $P_{\mathbf{X}|\mathbf{U}}$), let $\mathbf{Q} := \{Q_{\mathbf{X}|\mathbf{U}}, Q_{\mathbf{U}}\}$. Furthermore, let:

$$\mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}) := \mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [\log Q_{\mathbf{X}|\mathbf{U}}] - s D_{\text{KL}}(P_{\mathbf{U}|\mathbf{X}} \| Q_{\mathbf{U}}) \right]. \quad (8)$$

Lemma 1. For given \mathbf{P} , we have:

$$\mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}) \leq \mathcal{L}'_s(\mathbf{P}), \quad \text{for all } \mathbf{Q}.$$

In addition, there exists a unique \mathbf{Q} that achieves the maximum $\max_{\mathbf{Q}} \mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}) = \mathcal{L}'_s(\mathbf{P})$ and is given by:

$$Q_{\mathbf{X}|\mathbf{U}}^* = P_{\mathbf{X}|\mathbf{U}}, \quad Q_{\mathbf{U}}^* = P_{\mathbf{U}}.$$

Proof. The proof of Lemma 1 is given in Appendix A. \square

Using Lemma 1, maximization of (6) can be written in term of the variational IB cost as follows:

$$\max_{\mathbf{P}} \mathcal{L}'_s(\mathbf{P}) = \max_{\mathbf{P}} \max_{\mathbf{Q}} \mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}). \quad (9)$$

Next, we develop an algorithm that solves the maximization problem (9), where the encoding mapping $P_{\mathbf{U}|\mathbf{X}}$, the decoding mapping $Q_{\mathbf{X}|\mathbf{U}}$, as well as the prior distribution of the latent space $Q_{\mathbf{U}}$ are optimized jointly.

Remark 1. As we already mentioned in the beginning of this section, the related work [19] performed unsupervised clustering by combining VAE with GMM. Specifically, it maximizes the following ELBO bound:

$$\mathcal{L}_1^{\text{VaDE}} := \mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [\log Q_{\mathbf{X}|\mathbf{U}}] - D_{\text{KL}}(P_{\mathbf{C}|\mathbf{X}} \| Q_{\mathbf{C}}) - \mathbb{E}_{P_{\mathbf{C}|\mathbf{X}}} [D_{\text{KL}}(P_{\mathbf{U}|\mathbf{X}} \| Q_{\mathbf{U}|\mathbf{C}})] \right]. \tag{10}$$

Let, for an arbitrary non-negative parameter s , $\mathcal{L}_s^{\text{VaDE}}$ be a generalization of the ELBO bound in (10) of [19] given by:

$$\mathcal{L}_s^{\text{VaDE}} := \mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [\log Q_{\mathbf{X}|\mathbf{U}}] - sD_{\text{KL}}(P_{\mathbf{C}|\mathbf{X}} \| Q_{\mathbf{C}}) - s\mathbb{E}_{P_{\mathbf{C}|\mathbf{X}}} [D_{\text{KL}}(P_{\mathbf{U}|\mathbf{X}} \| Q_{\mathbf{U}|\mathbf{C}})] \right]. \tag{11}$$

Investigating the right-hand side (RHS) of (11), we get:

$$\mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}) = \mathcal{L}_s^{\text{VaDE}} + s\mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [D_{\text{KL}}(P_{\mathbf{C}|\mathbf{X}} \| Q_{\mathbf{C}|\mathbf{U}})] \right], \tag{12}$$

where the equality holds since:

$$\mathcal{L}_s^{\text{VaDE}} = \mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [\log Q_{\mathbf{X}|\mathbf{U}}] - sD_{\text{KL}}(P_{\mathbf{C}|\mathbf{X}} \| Q_{\mathbf{C}}) - s\mathbb{E}_{P_{\mathbf{C}|\mathbf{X}}} [D_{\text{KL}}(P_{\mathbf{U}|\mathbf{X}} \| Q_{\mathbf{U}|\mathbf{C}})] \right] \tag{13}$$

$$\stackrel{(a)}{=} \mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [\log Q_{\mathbf{X}|\mathbf{U}}] - sD_{\text{KL}}(P_{\mathbf{U}|\mathbf{X}} \| Q_{\mathbf{U}}) - s\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [D_{\text{KL}}(P_{\mathbf{C}|\mathbf{X}} \| Q_{\mathbf{C}|\mathbf{U}})] \right] \tag{14}$$

$$\stackrel{(b)}{=} \mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}) - s\mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}} [D_{\text{KL}}(P_{\mathbf{C}|\mathbf{X}} \| Q_{\mathbf{C}|\mathbf{U}})] \right], \tag{15}$$

where (a) can be obtained by expanding and re-arranging terms under the Markov chain $\mathbf{C} \leftrightarrow \mathbf{X} \leftrightarrow \mathbf{U}$ (for a detailed treatment, please look at Appendix B); and (b) follows from the definition of $\mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q})$ in (8).

Thus, by the non-negativity of relative entropy, it is clear that $\mathcal{L}_s^{\text{VaDE}}$ is a lower bound on $\mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q})$. Furthermore, if the variational distribution \mathbf{Q} is such that the conditional marginal $Q_{\mathbf{C}|\mathbf{U}}$ is equal to $P_{\mathbf{C}|\mathbf{X}}$, the bound is tight since the relative entropy term is zero in this case.

3.2. Proposed Algorithm: VIB-GMM

In order to compute (9), we parametrize the distributions $P_{\mathbf{U}|\mathbf{X}}$ and $Q_{\mathbf{X}|\mathbf{U}}$ using DNNs. For instance, let the stochastic encoder $P_{\mathbf{U}|\mathbf{X}}$ be a DNN f_{θ} and the stochastic decoder $Q_{\mathbf{X}|\mathbf{U}}$ be a DNN g_{ϕ} . That is:

$$\begin{aligned} P_{\theta}(\mathbf{u}|\mathbf{x}) &= \mathcal{N}(\mathbf{u}; \boldsymbol{\mu}_{\theta}, \boldsymbol{\Sigma}_{\theta}), \quad \text{where } [\boldsymbol{\mu}_{\theta}, \boldsymbol{\Sigma}_{\theta}] = f_{\theta}(\mathbf{x}), \\ Q_{\phi}(\mathbf{x}|\mathbf{u}) &= g_{\phi}(\mathbf{u}) = [\hat{\mathbf{x}}], \end{aligned} \tag{16}$$

where θ and ϕ are the weight and bias parameters of the DNNs. Furthermore, the latent space is modeled as a GMM with $|\mathcal{C}|$ components with parameters $\psi := \{\pi_c, \boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c\}_{c=1}^{|\mathcal{C}|}$, i.e.,

$$Q_{\psi}(\mathbf{u}) = \sum_c \pi_c \mathcal{N}(\mathbf{u}; \boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c). \tag{17}$$

Using the parametrizations above, the optimization of (9) can be rewritten as:

$$\max_{\theta, \phi, \psi} \mathcal{L}_s^{\text{NN}}(\theta, \phi, \psi) \tag{18}$$

where the cost function $\mathcal{L}_s^{\text{NN}}(\theta, \phi, \psi)$ is given by:

$$\mathcal{L}_s^{\text{NN}}(\theta, \phi, \psi) := \mathbb{E}_{P_{\mathbf{X}}} \left[\mathbb{E}_{P_{\theta}(\mathbf{U}|\mathbf{X})} [\log Q_{\phi}(\mathbf{X}|\mathbf{U})] - s D_{\text{KL}}(P_{\theta}(\mathbf{U}|\mathbf{X}) \| Q_{\psi}(\mathbf{U})) \right]. \quad (19)$$

Then, for given observations of N samples, i.e., $\{\mathbf{x}_i\}_{i=1}^N$, (18) can be approximated in terms of an empirical cost as follows:

$$\max_{\theta, \phi, \psi} \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{s,i}^{\text{emp}}(\theta, \phi, \psi), \quad (20)$$

where $\mathcal{L}_{s,i}^{\text{emp}}(\theta, \phi, \psi)$ is the empirical cost for the i th observation \mathbf{x}_i and given by:

$$\mathcal{L}_{s,i}^{\text{emp}}(\theta, \phi, \psi) = \mathbb{E}_{P_{\theta}(\mathbf{U}_i|\mathbf{x}_i)} [\log Q_{\phi}(\mathbf{x}_i|\mathbf{U}_i)] - s D_{\text{KL}}(P_{\theta}(\mathbf{U}_i|\mathbf{x}_i) \| Q_{\psi}(\mathbf{U}_i)). \quad (21)$$

Furthermore, the first term of the RHS of (21) can be computed using Monte Carlo sampling and the reparametrization trick [13]. In particular, $P_{\theta}(\mathbf{u}|\mathbf{x})$ can be sampled by first sampling a random variable \mathbf{Z} with distribution $P_{\mathbf{Z}}$, i.e., $P_{\mathbf{Z}} = \mathcal{N}(\mathbf{0}, \mathbf{I})$, then transforming the samples using some function $\tilde{f}_{\theta} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{U}$, i.e., $\mathbf{u} = \tilde{f}_{\theta}(\mathbf{x}, \mathbf{z})$. Thus,

$$\mathbb{E}_{P_{\theta}(\mathbf{U}_i|\mathbf{x}_i)} [\log Q_{\phi}(\mathbf{x}_i|\mathbf{U}_i)] = \frac{1}{M} \sum_{m=1}^M \log q(\mathbf{x}_i|\mathbf{u}_{i,m}), \quad \mathbf{u}_{i,m} = \boldsymbol{\mu}_{\theta,i} + \boldsymbol{\Sigma}_{\theta,i}^{\frac{1}{2}} \cdot \boldsymbol{\epsilon}_m, \quad \boldsymbol{\epsilon}_m \sim \mathcal{N}(\mathbf{0}, \mathbf{I}),$$

where M is the number of samples for the Monte Carlo sampling step.

The second term of the RHS of (21) is the KL divergence between a single component multivariate Gaussian and a Gaussian mixture model with $|\mathcal{C}|$ components. An exact closed-form solution for the calculation of this term does not exist. However, a variational lower bound approximation [24] of it can be obtained as:

$$D_{\text{KL}}(P_{\theta}(\mathbf{U}_i|\mathbf{x}_i) \| Q_{\psi}(\mathbf{U}_i)) = -\log \sum_{c=1}^{|\mathcal{C}|} \pi_c \exp \left(-D_{\text{KL}}(\mathcal{N}(\boldsymbol{\mu}_{\theta,i}, \boldsymbol{\Sigma}_{\theta,i}) \| \mathcal{N}(\boldsymbol{\mu}_c, \boldsymbol{\Sigma}_c)) \right). \quad (22)$$

In particular, in the specific case in which the covariance matrices are diagonal, i.e., $\boldsymbol{\Sigma}_{\theta,i} := \text{diag}(\{\sigma_{\theta,i,j}^2\}_{j=1}^{n_u})$ and $\boldsymbol{\Sigma}_c := \text{diag}(\{\sigma_{c,j}^2\}_{j=1}^{n_u})$, with n_u denoting the latent space dimension, (22) can be computed as follows:

$$\begin{aligned} & D_{\text{KL}}(P_{\theta}(\mathbf{U}_i|\mathbf{x}_i) \| Q_{\psi}(\mathbf{U}_i)) \\ &= -\log \sum_{c=1}^{|\mathcal{C}|} \pi_c \exp \left(-\frac{1}{2} \sum_{j=1}^{n_u} \left[\frac{(\mu_{\theta,i,j} - \mu_{c,j})^2}{\sigma_{c,j}^2} + \log \frac{\sigma_{c,j}^2}{\sigma_{\theta,i,j}^2} - 1 + \frac{\sigma_{\theta,i,j}^2}{\sigma_{c,j}^2} \right] \right), \end{aligned} \quad (23)$$

where $\mu_{\theta,i,j}$ and $\sigma_{\theta,i,j}^2$ are the mean and variance of the i th representation in the j th dimension of the latent space. Furthermore, $\mu_{c,j}$ and $\sigma_{c,j}^2$ represent the mean and variance of the c th component of the GMM in the j th dimension of the latent space.

Finally, we train NNs to maximize the cost function (19) over the parameters θ, ϕ , as well as those ψ of the GMM. For the training step, we use the ADAM optimization tool [25]. The training procedure is detailed in Algorithm 1.

Algorithm 1 VIB-GMM algorithm for unsupervised learning.

- 1: **input:** Dataset $\mathcal{D} := \{\mathbf{x}_i\}_{i=1}^N$, parameter $s \geq 0$.
 - 2: **output:** Optimal DNN weights θ^*, ϕ^* and GMM parameters $\psi^* = \{\pi_c^*, \mu_c^*, \Sigma_c^*\}_{c=1}^{|C|}$.
 - 3: **initialization** Initialize θ, ϕ, ψ .
 - 4: **repeat**
 - 5: Randomly select b mini-batch samples $\{\mathbf{x}_i\}_{i=1}^b$ from \mathcal{D} .
 - 6: Draw m random i.i.d samples $\{\mathbf{z}_j\}_{j=1}^m$ from $P_{\mathbf{Z}}$.
 - 7: Compute m samples $\mathbf{u}_{i,j} = \tilde{f}_\theta(\mathbf{x}_i, \mathbf{z}_j)$
 - 8: For the selected mini-batch, compute gradients of the empirical cost (20).
 - 9: Update θ, ϕ, ψ using the estimated gradient (e.g., with SGD or ADAM).
 - 10: **until** convergence of θ, ϕ, ψ .
-

Once our model is trained, we assign the given dataset into the clusters. As mentioned in Section 2, we do the assignment from the latent representations, i.e., $Q_{C|U} = P_{C|X}$. Hence, the probability that the observed data \mathbf{x}_i belongs to the c th cluster is computed as follows:

$$p(c|\mathbf{x}_i) = q(c|\mathbf{u}_i) = \frac{q_{\psi^*}(c)q_{\psi^*}(\mathbf{u}_i|c)}{q_{\psi^*}(\mathbf{u}_i)} = \frac{\pi_c^* \mathcal{N}(\mathbf{u}_i; \mu_c^*, \Sigma_c^*)}{\sum_c \pi_c^* \mathcal{N}(\mathbf{u}_i; \mu_c^*, \Sigma_c^*)}, \tag{24}$$

where $*$ indicates the optimal values of the parameters as found at the end of the training phase. Finally, the right cluster is picked based on the largest assignment probability value.

Remark 2. *It is worth mentioning that with the use of the KL approximation as given by (22), our algorithm does not require the assumption $P_{C|U} = Q_{C|U}$ to hold (which is different from [19]). Furthermore, the algorithm is guaranteed to converge. However, the convergence may be to (only) local minima; and this is due to the problem (18) being generally non-convex. Related to this aspect, we mention that while without a proper pre-training, the accuracy of the VaDE algorithm may not be satisfactory, in our case, the above assumption is only used in the final assignment after the training phase is completed.*

Remark 3. *In [26], it is stated that optimizing the original IB problem with the assumption of independent latent representations amounts to disentangled representations. It is noteworthy that with such an assumption, the computational complexity can be reduced from $\mathcal{O}(n_u^2)$ to $\mathcal{O}(n_u)$. Furthermore, as argued in [26], the assumption often results only in some marginal performance loss; and for this reason, it is adopted in many machine learning applications.*

3.3. Effect of the Hyperparameter

As we already mentioned, the hyperparameter s controls the trade-off between the relevance of the representation \mathbf{U} and its complexity. As can be seen from (19) for small values of s , it is the cross-entropy term that dominates, i.e., the algorithm trains the parameters so as to reproduce \mathbf{X} as accurately as possible. For large values of s , however, it is most important for the NN to produce an encoded version of \mathbf{X} whose distribution matches the prior distribution of the latent space, i.e., the term $D_{KL}(P_\theta(\mathbf{U}|\mathbf{X})\|Q_\psi(\mathbf{U}))$ is nearly zero.

In the beginning of the training process, the GMM components are randomly selected; and so, starting with a large value of the hyperparameter s is likely to steer the solution towards an irrelevant prior. Hence, for the tuning of the hyper-parameter s in practice, it is more efficient to start with a small value of s and gradually increase it with the number of epochs. This has the advantage of avoiding possible local minima, an aspect that is reminiscent of deterministic annealing [12], where s plays the role of the temperature parameter. The experiments that will be reported in the next section show that proceeding in the above-described manner for the selection of the parameter s helps in obtaining higher

clustering accuracy and better robustness to the initialization (i.e., no need for a strong pretraining). The pseudocode for annealing is given in Algorithm 2.

Algorithm 2 Annealing algorithm pseudocode.

- 1: **input:** Dataset $\mathcal{D} := \{x_i\}_{i=1}^n$, hyperparameter interval $[s_{\min}, s_{\max}]$.
 - 2: **output:** Optimal DNN weights θ^*, ϕ^* , GMM parameters $\psi^* = \{\pi_c^*, \mu_c^*, \Sigma_c^*\}_{c=1}^{|C|}$,
assignment probability $P_{C|X}$.
 - 3: **initialization** Initialize θ, ϕ, ψ .
 - 4: **repeat**
 - 5: Apply VIB-GMM algorithm.
 - 6: Update ψ, θ, ϕ .
 - 7: Update s , e.g., $s = (1 + \epsilon_s)s_{\text{old}}$.
 - 8: **until** s does not exceed s_{\max} .
-

Remark 4. As we mentioned before, a text clustering algorithm was introduced by Slonim et al. [11,12], which uses the IB method with an annealing procedure, where the parameter s is increased gradually. In [12], the critical values of s (so-called phase transitions) were observed such that if these values were missed during increasing s , the algorithm ended up with the wrong clusters. Therefore, how to choose the step size in the update of s is very important. We note that tuning s is also very critical in our algorithm, such that the step size ϵ_s in the update of s should be chosen carefully, otherwise phase transitions might be skipped that would cause a non-satisfactory clustering accuracy score. However, the choice of the appropriate step size (typically very small) is rather heuristic; and there exists no concrete method for choosing the right value. The choice of step size can be seen as a trade-off between the amount of computational resource spared for running the algorithm and the degree of confidence about scanning s values not to miss the phase transitions.

4. Experiments

4.1. Description of the Datasets Used

In our empirical experiments, we applied our algorithm to the clustering of the following datasets.

MNIST: A dataset of gray-scale images of 70,000 handwritten digits of dimensions 28×28 pixel.

STL-10: A dataset of color images collected from 10 categories. Each category consisted of 1300 images of size of 96×96 (pixels) $\times 3$ (RGB code). Hence, the original input dimension n_x was 27,648. For this dataset, we used a pretrained convolutional NN model, i.e., ResNet-50 [27], to reduce the dimensionality of the input. This preprocessing reduced the input dimension to 2048. Then, our algorithm and other baselines were used for clustering.

REUTERS10K: A dataset that was composed of 810,000 English stories labeled with a category tree. As in [20], 4 root categories (corporate/industrial, government/social, markets, economics) were selected as labels, and all documents with multiple labels were discarded. Then, tf-idf features were computed on the 2000 most frequently occurring words. Finally, 10,000 samples were taken randomly, which were referred to as the REUTERS10K dataset.

4.2. Network Settings and Other Parameters

We used the following network architecture: the encoder was modeled with NNs with 3 hidden layers with dimensions $n_x - 500 - 500 - 2000 - n_u$, where n_x is the input dimension and n_u is the dimension of the latent space. The decoder consisted of NNs with dimensions $n_u - 2000 - 500 - 500 - n_x$. All layers were fully connected. For comparison purposes, we chose the architecture of the hidden layers, as well as the dimension of the latent space $n_u = 10$ to coincide with those made for the DEC algorithm of [20] and the VaDE algorithm of [19]. All except the last layers of the encoder and

decoder were activated with the ReLU function. For the last (i.e., latent) layer of the encoder, we used a linear activation; and for the last (i.e., output) layer of the decoder, we used the sigmoid function for MNIST and linear activation for the remaining datasets. The batch size was 100, and the variational bound (20) was maximized by the ADAM optimizer of [25]. The learning rate was initialized with 0.002 and decreased gradually every 20 epochs with a decay rate of 0.9 until it reached a small value (0.0005 in our experiments). The reconstruction loss was calculated by using the cross-entropy criterion for MNIST and the mean squared error function for the other datasets.

4.3. Clustering Accuracy

We evaluated the performance of our algorithm in terms of the so-called unsupervised clustering Accuracy (ACC), which is a widely used metric in the context of unsupervised learning [23]. For comparison purposes, we also present those of algorithms from the previous state-of-the-art.

For each of the aforementioned datasets, we ran our VIB-GMM algorithm for various values of the hyper-parameter s inside an interval $[s_{\min}, s_{\max}]$, starting from the smaller value s_{\min} and gradually increasing the value of s every n_{epoch} epochs. For the MNIST dataset, we set $(s_{\min}, s_{\max}, n_{\text{epoch}}) = (1, 5, 500)$; and for the STL-10 dataset and the REUTERS10K dataset, we chose these parameters to be $(1, 20, 500)$ and $(1, 5, 100)$, respectively. The obtained ACC accuracy results are reported in Tables 1 and 2. It is important to note that the reported ACC results were obtained by running each algorithm ten times. For the case in which there was no pretraining (in [19,20], the DEC and VaDE algorithms were proposed to be used with pretraining; more specifically, the DNNs were initialized with a stacked autoencoder [28]), Table 1 states the accuracies of the best case run and average case run for the MNIST and STL-10 datasets. It was seen that our algorithm outperformed significantly the DEC algorithm of [20], as well as the VaDE algorithm of [19] and GMM for both the best case run and average case run. Besides, in Table 1, the values in parentheses correspond to the standard deviations of clustering accuracies. As seen, the standard deviation of our algorithm VIB-GMM was lower than the VaDE; which could be expounded by the robustness of VIB-GMM to non-pretraining. For the case in which there was pretraining, Table 2 states the accuracies of the best case run and average case run for the MNIST and REUTERS10K datasets. A stacked autoencoder was used to pretrain the DNNs of the encoder and decoder before running algorithms (DNNs were initialized with the same weights and biases of [19]). It was seen that our algorithm outperformed significantly the DEC algorithm of [20], as well as the VaDE algorithm of [19] and GMM for both the best case run and average case run. The effect of pretraining can be observed comparing Tables 1 and 2 for MNIST. Using a stacked autoencoder prior to running the VaDE and VIB-GMM algorithms resulted in a higher accuracy, as well as a lower standard deviation of accuracies; therefore, supporting the algorithms with a stacked autoencoder was beneficial for a more robust system. Finally, for the STL-10 dataset, Figure 4 depicts the evolution of the best case ACC with iterations (number of epochs) for the four compared algorithms.

Table 1. Comparison of the clustering accuracy of various algorithms. The algorithms are run without pretraining. Each algorithm is run ten times. The values in (·) correspond to the standard deviations of clustering accuracies. DEC, Deep Embedded Clustering; VaDE, Variational Deep Embedding; VIB, Variational Information Bottleneck.

	MNIST		STL-10	
	Best Run	Average Run	Best Run	Average Run
GMM	44.1	40.5 (1.5)	78.9	73.3 (5.1)
DEC			80.6 [†]	
VaDE	91.8	78.8 (9.1)	85.3	74.1 (6.4)
VIB-GMM	95.1	83.5 (5.9)	93.2	82.1 (5.6)

[†] Values are taken from VaDE [19].

Table 2. Comparison of the clustering accuracy of various algorithms. A stacked autoencoder is used to pretrain the DNNs of the encoder and decoder before running algorithms (DNNs are initialized with the same weights and biases of [19]). Each algorithm is run ten times. The values in (·) correspond to the standard deviations of clustering accuracies.

	MNIST		REUTERS10K	
	Best Run	Average Run	Best Run	Average Run
DEC	84.3 [‡]		72.2 [‡]	
VaDE	94.2	93.2 (1.5)	79.8	79.1 (0.6)
VIB-GMM	96.1	95.8 (0.1)	81.6	81.2 (0.4)

[‡] Values are taken from DEC [20].

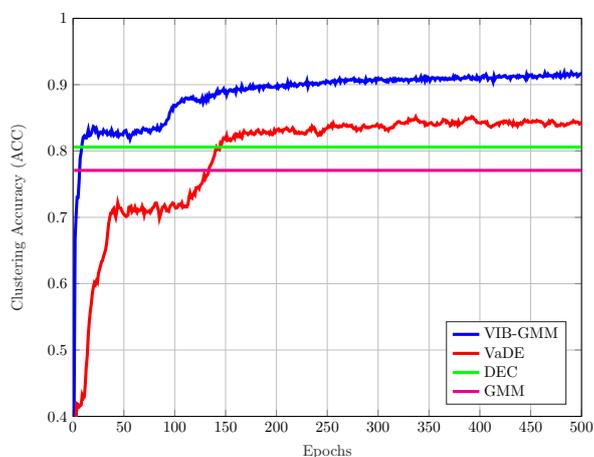


Figure 4. Accuracy vs. the number of epochs for the STL-10 dataset.

Figure 5 shows the evolution of the reconstruction loss of our VIB-GMM algorithm for the STL-10 dataset, as a function of simultaneously varying the values of the hyperparameter s and the number of epochs (recall that, as per the described methodology, we started with $s = s_{min}$, and we increased its value gradually every $n_{epoch} = 500$ epochs). As can be seen from the figure, the few first epochs were spent almost entirely on reducing the reconstruction loss (i.e., a fitting phase), and most of the remaining epochs were spent making the found representation more concise (i.e., smaller KL divergence). This was reminiscent of the two-phase (fitting vs. compression) that was observed for supervised learning using VIB in [29].

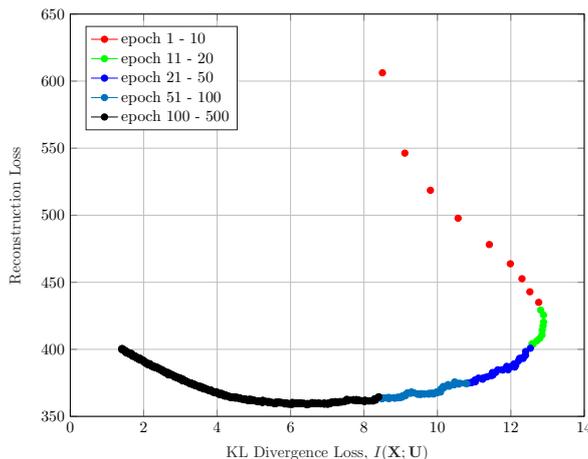


Figure 5. Information plane for the STL-10 dataset.

Remark 5. For a fair comparison, our algorithm VIB-GMM and the VaDE of [19] were run for the same number of epochs, e.g., n_{epoch} . In the VaDE algorithm, the cost function (11) was optimized for a particular value of hyperparameter s . Instead of running n_{epoch} epochs for $s = 1$ as in VaDE, we ran n_{epoch} epochs by gradually increasing s to optimize the cost (21). In other words, the computational resources were distributed over a range of s values. Therefore, the computational complexity of our algorithm and the VaDE was equivalent.

4.4. Visualization on the Latent Space

In this section, we investigate the evolution of the unsupervised clustering of the STL-10 dataset on the latent space using our VIB-GMM algorithm. For this purpose, we found it convenient to visualize the latent space through application of the t-SNE algorithm of [30] in order to generate meaningful representations in a two-dimensional space. Figure 6 shows 4000 randomly chosen latent representations before the start of the training process and respectively after 1, 5, and 500 epochs. The shown points (with a \cdot marker in the figure) represent latent representations of data samples whose labels are identical. Colors are used to distinguish between clusters. Crosses (with an \times marker in the figure) correspond to the centroids of the clusters. More specifically, Figure 6a shows the initial latent space before the training process. If the clustering is performed on the initial representations, it allows ACC as small as 10%, i.e., as bad as a random assignment. Figure 6b shows the latent space after one epoch, from which a partition of some of the points starts to be already visible. With five epochs, that partitioning is significantly sharper, and the associated clusters can be recognized easily. Observe, however, that the cluster centers seem not to have converged yet. With 500 epochs, the ACC accuracy of our algorithm reached %91.6, and the clusters and their centroids were neater, as is visible from Figure 6d.

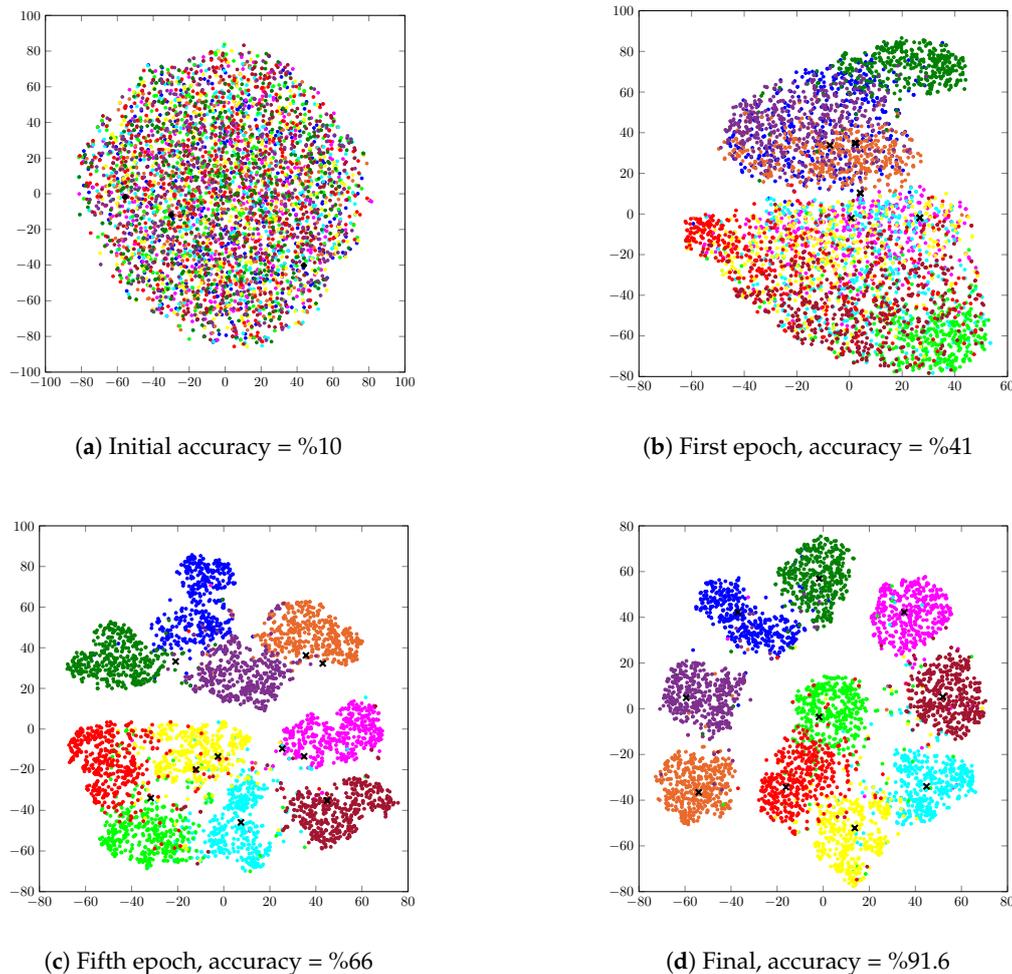


Figure 6. Visualization of the latent space before training; and after 1, 5, and 500 epochs.

5. Conclusions and Future Work

In this paper, we proposed and analyzed the performance of an unsupervised algorithm for data clustering. The algorithm used the variational information bottleneck approach and modeled the latent space as a mixture of Gaussians. It was shown to outperform state-of-the-art algorithms such as the VaDE of [19] and the DEC of [20]. We note that although it was assumed that the number of classes was known beforehand (as was the case for almost all competing algorithms in its category), that number could be found (or estimated to within a certain accuracy) through inspection of the resulting bifurcations on the associated information-plane, as was observed for the standard information bottleneck method. Finally, we mention that among the interesting research directions in this line of work, one important question pertains to the distributed learning setting, i.e., along the counterpart, to the unsupervised setting, of the recent work [31–33], which contained distributed IB algorithms for both discrete and vector Gaussian data models.

Author Contributions: Software, Y.U.; Methodology, A.Z., Y.U. and G.A.; Supervision, A.Z.; Writing—original draft, A.Z., Y.U. and G.A. Writing—Review & Editing, Y.U. and A.Z. All authors have read and agreed to the published version of the manuscript.

Funding: The authors received no specific funding for this work.

Acknowledgments: The authors thank the reviewers for their comments that helped to improve the paper.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. The Proof of Lemma 1

First, we expand $\mathcal{L}'_s(\mathbf{P})$ as follows:

$$\begin{aligned} \mathcal{L}'_s(\mathbf{P}) &= -H(\mathbf{X}|\mathbf{U}) - sI(\mathbf{X}; \mathbf{U}) = -H(\mathbf{X}|\mathbf{U}) - s[H(\mathbf{U}) - H(\mathbf{U}|\mathbf{X})] \\ &= \iint_{\mathbf{u}\mathbf{x}} p(\mathbf{u}, \mathbf{x}) \log p(\mathbf{x}|\mathbf{u}) \, d\mathbf{u} \, d\mathbf{x} + s \int_{\mathbf{u}} p(\mathbf{u}) \log p(\mathbf{u}) \, d\mathbf{u} - s \iint_{\mathbf{u}\mathbf{x}} p(\mathbf{u}, \mathbf{x}) \log p(\mathbf{u}|\mathbf{x}) \, d\mathbf{u} \, d\mathbf{x}. \end{aligned}$$

Then, $\mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q})$ is defined as follows:

$$\mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}) := \iint_{\mathbf{u}\mathbf{x}} p(\mathbf{u}, \mathbf{x}) \log q(\mathbf{x}|\mathbf{u}) \, d\mathbf{u} \, d\mathbf{x} + s \int_{\mathbf{u}} p(\mathbf{u}) \log q(\mathbf{u}) \, d\mathbf{u} - s \iint_{\mathbf{u}\mathbf{x}} p(\mathbf{u}, \mathbf{x}) \log p(\mathbf{u}|\mathbf{x}) \, d\mathbf{u} \, d\mathbf{x}. \quad (\text{A1})$$

Hence, we have the following relation:

$$\mathcal{L}'_s(\mathbf{P}) - \mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}) = \mathbb{E}_{P_{\mathbf{U}}}[D_{\text{KL}}(P_{\mathbf{X}|\mathbf{U}}\|Q_{\mathbf{X}|\mathbf{U}})] + sD_{\text{KL}}(P_{\mathbf{U}}\|Q_{\mathbf{U}}) \geq 0,$$

where equality holds under equalities $Q_{\mathbf{X}|\mathbf{U}} = P_{\mathbf{X}|\mathbf{U}}$ and $Q_{\mathbf{U}} = P_{\mathbf{U}}$. We note that $s \geq 0$.

Now, we complete the proof by showing that (A1) is equal to (8). To do so, we proceed (A1) as follows:

$$\begin{aligned} \mathcal{L}_s^{\text{VB}}(\mathbf{P}, \mathbf{Q}) &= \int_{\mathbf{x}} p(\mathbf{x}) \int_{\mathbf{u}} p(\mathbf{u}|\mathbf{x}) \log q(\mathbf{x}|\mathbf{u}) \, d\mathbf{u} \, d\mathbf{x} \\ &\quad + s \int_{\mathbf{x}} p(\mathbf{x}) \int_{\mathbf{u}} p(\mathbf{u}|\mathbf{x}) \log q(\mathbf{u}) \, d\mathbf{u} - s \int_{\mathbf{x}} p(\mathbf{x}) \int_{\mathbf{u}} p(\mathbf{u}|\mathbf{x}) \log p(\mathbf{u}|\mathbf{x}) \, d\mathbf{u} \, d\mathbf{x} \\ &= \mathbb{E}_{P_{\mathbf{X}}}\left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}}[\log Q_{\mathbf{X}|\mathbf{U}}] - sD_{\text{KL}}(P_{\mathbf{U}|\mathbf{X}}\|Q_{\mathbf{U}})\right]. \end{aligned}$$

Appendix B. Alternative Expression $\mathcal{L}_s^{\text{VaDE}}$

Here, we show that (13) is equal to (14).

To do so, we start with (14) and proceed as follows:

$$\begin{aligned} \mathcal{L}_s^{\text{VaDE}} &= \mathbb{E}_{P_{\mathbf{X}}}\left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}}[\log Q_{\mathbf{X}|\mathbf{U}}] - sD_{\text{KL}}(P_{\mathbf{U}|\mathbf{X}}\|Q_{\mathbf{U}}) - s\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}}[D_{\text{KL}}(P_{C|\mathbf{X}}\|Q_{C|\mathbf{U}})]\right] \\ &= \mathbb{E}_{P_{\mathbf{X}}}\left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}}[\log Q_{\mathbf{X}|\mathbf{U}}] - s \int_{\mathbf{x}} p(\mathbf{x}) \int_{\mathbf{u}} p(\mathbf{u}|\mathbf{x}) \log \frac{p(\mathbf{u}|\mathbf{x})}{q(\mathbf{u})} \, d\mathbf{u} \, d\mathbf{x} \right. \\ &\quad \left. - s \int_{\mathbf{x}} p(\mathbf{x}) \int_{\mathbf{u}} p(\mathbf{u}|\mathbf{x}) \sum_c p(c|\mathbf{x}) \log \frac{p(c|\mathbf{x})}{q(c|\mathbf{u})} \, d\mathbf{u} \, d\mathbf{x} \right] \\ &\stackrel{(a)}{=} \mathbb{E}_{P_{\mathbf{X}}}\left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}}[\log Q_{\mathbf{X}|\mathbf{U}}] - s \iint_{\mathbf{u}\mathbf{x}} p(\mathbf{x})p(\mathbf{u}|\mathbf{x}) \log \frac{p(\mathbf{u}|\mathbf{x})}{q(\mathbf{u})} \, d\mathbf{u} \, d\mathbf{x} \right. \\ &\quad \left. - s \iint_{\mathbf{u}\mathbf{x}} \sum_c p(\mathbf{x})p(\mathbf{u}|\mathbf{x})p(c|\mathbf{x}) \log \frac{p(c|\mathbf{x})}{q(c|\mathbf{u})} \, d\mathbf{u} \, d\mathbf{x} \right] \\ &= \mathbb{E}_{P_{\mathbf{X}}}\left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}}[\log Q_{\mathbf{X}|\mathbf{U}}] - s \iint_{\mathbf{u}\mathbf{x}} \sum_c p(\mathbf{u}, c, \mathbf{x}) \log \frac{p(\mathbf{u}|\mathbf{x})p(c|\mathbf{x})}{q(\mathbf{u})q(c|\mathbf{u})} \, d\mathbf{u} \, d\mathbf{x} \right] \\ &= \mathbb{E}_{P_{\mathbf{X}}}\left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}}[\log Q_{\mathbf{X}|\mathbf{U}}] - s \iint_{\mathbf{u}\mathbf{x}} \sum_c p(\mathbf{u}, c, \mathbf{x}) \log \frac{p(c|\mathbf{x})}{q(c)} \frac{p(\mathbf{u}|\mathbf{x})}{q(\mathbf{u}|\mathbf{c})} \, d\mathbf{u} \, d\mathbf{x} \right] \\ &= \mathbb{E}_{P_{\mathbf{X}}}\left[\mathbb{E}_{P_{\mathbf{U}|\mathbf{X}}}[\log Q_{\mathbf{X}|\mathbf{U}}] - s \int_{\mathbf{x}} \sum_c p(c, \mathbf{x}) \log \frac{p(c|\mathbf{x})}{q(c)} \, d\mathbf{x} \right. \\ &\quad \left. - s \iint_{\mathbf{u}\mathbf{x}} \sum_c p(\mathbf{x})p(c|\mathbf{x})p(\mathbf{u}|\mathbf{c}, \mathbf{x}) \log \frac{p(\mathbf{u}|\mathbf{x})}{q(\mathbf{u}|\mathbf{c})} \, d\mathbf{u} \, d\mathbf{x} \right] \end{aligned}$$

$$\stackrel{(b)}{=} \mathbb{E}_{P_X} \left[\mathbb{E}_{P_{U|X}} [\log Q_{X|U}] - sD_{\text{KL}}(P_{C|X} \| Q_C) - s\mathbb{E}_{P_{C|X}} [D_{\text{KL}}(P_{U|X} \| Q_{U|C})] \right],$$

where (a) and (b) follow due to the Markov chain $C \rightarrow X \rightarrow U$.

References

1. Sculley, D. Web-scale K -means clustering. In Proceedings of the 19th International Conference on World Wide Web, Raleigh, NC, USA, 26–30 April 2010; pp. 1177–1178.
2. Huang, Z. Extensions to the k -means algorithm for clustering large datasets with categorical values. *Data Min. Knowl. Disc.* **1998**, *2*, 283–304. [[CrossRef](#)]
3. Hartigan, J.A.; Wong, M.A. Algorithm AS 136: A K -means clustering algorithm. *J. R. Stat. Soc.* **1979**, *28*, 100–108. [[CrossRef](#)]
4. Dempster, A.P.; Laird, N.M.; Rubin, D.B. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc.* **1977**, *39*, 1–38.
5. Ding, C.; He, X. K -means clustering via principal component analysis. In Proceedings of the 21st International Conference on Machine Learning, Banff, AB, Canada, 4–8 July 2004.
6. Pearson, K. On lines and planes of closest fit to systems of points in space. *Philos. Mag.* **1901**, *2*, 559–572. [[CrossRef](#)]
7. Wold, S.; Esbensen, K.; Geladi, P. Principal component analysis. *Chemom. Intell. Lab. Syst.* **1987**, *2*, 37–52. [[CrossRef](#)]
8. Roweis, S. EM algorithms for PCA and SPCA. In Proceedings of the Advances in Neural Information Processing Systems 10, Denver, CO, USA, 1–6 December 1997; pp. 626–632.
9. Hofmann, T.; Scholkopf, B.; Smola, A.J. Kernel methods in machine learning. *Ann. Stat.* **2008**, *36*, 1171–1220. [[CrossRef](#)]
10. Tishby, N.; Pereira, F.C.; Bialek, W. The information bottleneck method. In Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing, Monticello, IL, USA, 22–24 September 1999; pp. 368–377.
11. Slonim, N.; Tishby, N. Document clustering using word clusters via the information bottleneck method. In Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Athens, Greece, 24–28 July 2000; pp. 208–215.
12. Slonim, N. The Information Bottleneck: Theory and Applications. Ph.D. Thesis, Hebrew University, Jerusalem, Israel, 2002.
13. Kingma, D.P.; Welling, M. Auto-encoding variational bayes. In Proceedings of the 2nd International Conference on Learning Representations, Banff, AB, Canada, 14–16 April 2014.
14. Rezende, D.J.; Mohamed, S.; Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In Proceedings of the 31st International Conference on Machine Learning, Beijing, China, 21–26 June 2014; pp. 1278–1286.
15. Alemi, A.A.; Fischer, I.; Dillon, J.V.; Murphy, K. Deep variational information bottleneck. In Proceedings of the 5th International Conference on Learning Representations, Toulon, France, 24–26 April 2017.
16. Lecun, Y.; Bottou, L.; Bengio, Y.; Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **1998**, *86*, 2278–2324. [[CrossRef](#)]
17. Lewis, D.D.; Yang, Y.; Rose, T.G.; Li, F. A new benchmark collection for text categorization research. *J. Mach. Learn. Res.* **2004**, *5*, 361–397.
18. Coates, A.; Ng, A.; Lee, H. An analysis of single-layer networks in unsupervised feature learning. In Proceedings of the 14th International Conference on Artificial Intelligence and Statistics, Fort Lauderdale, FL, USA, 11–13 April 2011; pp. 215–223.
19. Jiang, Z.; Zheng, Y.; Tan, H.; Tang, B.; Zhou, H. Variational deep embedding: An unsupervised and generative approach to clustering. In Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia, 19–25 August 2017; pp. 1965–1972.
20. Xie, J.; Girshick, R.; Farhadi, A. Unsupervised deep embedding for clustering analysis. In Proceedings of the 33rd International Conference on Machine Learning, New York, NY, USA, 19–24 June 2016; pp. 478–487.

21. Guo, X.; Gao, L.; Liu, X.; Yin, J. Improved deep embedded clustering with local structure preservation. In Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia, 19–25 August 2017; pp. 1753–1759.
22. Dilokthanakul, N.; Mediano, P.A.M.; Garnelo, M.; Lee, M.C.H.; Salimbeni, H.; Arulkumaran, K.; Shanahani, M. Deep unsupervised clustering with Gaussian mixture variational autoencoders. *arXiv* **2017**, arXiv:1611.02648.
23. Min, E.; Guo, X.; Liu, Q.; Zhang, G.; Cui, J.; Long, J. A survey of clustering with deep learning: From the perspective of network architecture. *IEEE Access* **2018**, *6*, 39501–39514. [[CrossRef](#)]
24. Hershey, J.R.; Olsen, P.A. Approximating the Kullback Leibler divergence between Gaussian mixture models. In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, Honolulu, HI, USA, 15–20 April 2007; pp. 317–320.
25. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. In Proceedings of the 3rd International Conference on Learning Representations, San Diego, CA, USA, 7–9 May 2015.
26. Achille, A.; Soatto, S. Information dropout: Learning optimal representations through noisy computation. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *40*, 2897–2905. [[CrossRef](#)] [[PubMed](#)]
27. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
28. Vincent, P.; Larochelle, H.; Lajoie, I.; Bengio, Y.; Manzagol, P.-A. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *J. Mach. Learn. Res.* **2010**, *11*, 3371–3408.
29. Schwartz-Ziv, R.; Tishby, N. Opening the black box of deep neural networks via information. *arXiv* **2017**, arXiv:1703.00810.
30. van der Maaten, L.; Hinton, G. Visualizing Data using t-SNE. *J. Mach. Learn. Res.* **2008**, *9*, 2579–2605.
31. Estella-Aguerri, I.; Zaidi, A. Distributed variational representation learning. *IEEE Trans. Pattern Anal. Mach. Intell.* **2020**, in press
32. Estella-Aguerri, I.; Zaidi, A. Distributed information bottleneck method for discrete and Gaussian sources. In Proceedings of the International Zurich Seminar on Information and Communication, Zürich, Switzerland, 21–23 February 2018.
33. Zaidi, A.; Estella-Aguerri, I.; Shamaï (Shitz), S. On the information bottleneck problems: Models, connections, applications and information theoretic views. *Entropy* **2020**, *22*, 151. [[CrossRef](#)]



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).