DEEP CLUSTERING AND REPRESENTATION LEARNING THAT PRESERVES GEOMETRIC STRUCTURES

Anonymous authors
Paper under double-blind review

ABSTRACT

In this paper, we propose a novel framework for Deep Clustering and Representation Learning (DCRL) that preserves the geometric structure of data. In the proposed DCRL framework, the clustering of data points from different manifolds is done in the latent space guided by a clustering loss. To overcome the problem that clustering-oriented losses may deteriorate the geometric structure of embeddings in the latent space, two structure-oriented losses, namely an isometric loss and a ranking loss, are proposed to preserve intra-manifold structure locally and inter-manifold structure globally. Experimental results on various datasets show that the DCRL framework leads to comparable performances to current state-of-the-art deep clustering algorithms, yet exhibits superior performance for downstream tasks. Our results also demonstrate the importance and effectiveness of the proposed method in preserving geometric structure in terms of visualization and performance metrics.

1 INTRODUCTION

Clustering, a fundamental tool for data analysis and visualization, has been an essential research topic in data science and machine learning. Conventional clustering algorithms such as K-Means (MacQueen, 1965), Gaussian Mixture Models (GMM) (Bishop, 2006), and spectral clustering (Shi & Malik, 2000) perform clustering based on distance or similarity. However, handcrafted distance or similarity measures are rarely reliable for large-scale high-dimensional data, making it increasingly challenging to achieve effective clustering. An intuitive solution is to transform the data from the high-dimensional input space to the low-dimensional latent space and then to cluster the data in the latent space. This can be achieved by applying dimensionality reduction techniques such as PCA (Wold et al., 1987), t-SNE (Maaten & Hinton, 2008), and UMAP (McInnes et al., 2018). However, since these methods are not specifically designed for clustering tasks, some of their properties may be contrary to our expectations, e.g., two data points from different manifolds that are close in the input space will be closer in the latent space derived by UMAP. Therefore, the first question here is **how to learn the representation that favors clustering?**

The two main points of the multi-manifold representation learning are (1) preserving the local geometric structure within each manifold and (2) ensuring the discriminability between different manifolds. However, it is challenging to decouple complex cross-over relations and ensure the discriminability between different manifolds, especially in unsupervised situations. One natural strategy is to perform clustering in the input space to get pseudo-labels and then perform representation learning for each manifold. However, in that case, representation learning's performance depends heavily on the clustering effect, but commonly used clustering algorithms such as K-Means do not work well on high-dimensional data. Thus, the second question here is **how to cluster data that favors representation learning?**

To answer these two questions, some pioneering work has proposed integrating deep clustering and representation learning into a unified framework by defining a clustering-oriented loss. Though promising performance has been demonstrated on various datasets, we observe that a vital factor has been ignored by these work that the defined clustering-oriented loss may deteriorate the geometric structure of the latent space, which in turn hurts the performance of visualization, clustering generalization, and downstream tasks. In this paper, we propose to jointly perform deep clustering and representation learning with geometric structure preservation. Inspired by Xie et al. (2016),

the clustering centers are defined as a set of **learnable** parameters, and we use a clustering loss to simultaneously guide the separation of data points from different manifolds and the learning of the clustering centers. To prevent clustering loss from deteriorating the latent space, an isometric loss and a ranking loss are proposed to preserve the intra-manifold structure locally and inter-manifold structure globally. Our experimental results show that our method exhibits far superior performance to counterparts in terms of clustering and representation learining, which demonstrates the importance and effectiveness of preserving geometric structure.

The contributions of this work are summarized as below:

- Proposing to integrate deep clustering and representation learning into a unified framework with local and global structure preservation.
- Unlike conventional multi-manifold learning algorithms that deal with all point pair relationships between different manifolds simultaneously, we set the clustering centers as a set of learnable parameters and achieve global structure preservation in a faster, more efficient, and easier to optimize manner by applying ranking loss to the clustering centers.
- Analyzing the contradiction between two optimization goals of clustering and local structure preservation, and proposing an elegant training strategy to alleviate it.
- The proposed DCRL algorithm outperforms competing algorithms in terms of clustering effect, generalizability to out-of-sample, and performance in downstream tasks.

2 Related Work

Clustering analysis. As a fundamental tool in machine learning, it has been widely applied in various domains. One branch of classical clustering is K-Means (MacQueen, 1965) and Gaussian Mixture Models (GMM) (Bishop, 2006), which are fast, easy to understand, and can be applied to a large number of problems. However, limited by Euclidean measure, their performance on high-dimensional data is often unsatisfactory. Spectral clustering and its variants (such as SC-Ncut (Bishop, 2006)) extend clustering to high-dimensional data by allowing more flexible distance measures. However, limited by computational efficiency of the full Laplace matrix, spectral clustering is challenging to extend to large-scale datasets.

Deep clustering. The success of deep learning has contributed to the growth of *deep clustering*. One branch of deep clustering performs clustering after learning a representation through existing unsupervised techniques. For example, Tian et al. (2014) use autoencoder to learn low dimensional features and then run K-Means to get clustering results (AE+K-Means). Considering the geometric structure of the data, N2D applies UMAP to find the best clusterable manifold of the obtained embedding, and then run K-Means to discover higher-quality clusters (McConville et al., 2019). The other category of algorithms tries to optimize clustering and representation learning jointly. The closest work to us is Deep Embedding for Clustering (DEC) (Xie et al., 2016), which learns a mapping from the input space to a lower-dimensional latent space through iteratively optimizing a clustering objective. As a modified version of DEC, while IDEC claims to preserve the local structure of the data (Guo et al., 2017), in reality, their contribution is nothing more than adding a reconstruction loss. JULE proposes a recurrent framework for integrating clustering and representation learning into a single model with a weighted triplet loss and optimizing it end-to-end (Yang et al., 2016b). DSC devises a dual autoencoder to embed data into latent space, and then deep spectral clustering (Shaham et al., 2018) is applied to obtain label assignments (Yang et al., 2019).

Manifold Representation Learning. Isomap, as a representative algorithm of *single-manifold* learning, aims to capture global nonlinear features and seek an optimal subspace that best preserves the geodesic distance between data points (Tenenbaum et al., 2000). In contrast, some algorithms, such as the LLE (Roweis & Saul, 2000), are more concerned with the preservation of local neighborhood information. Combining DNN with manifold learning, the recently proposed MLDL algorithm achieves the preservation of local and global geometries by imposing LIS prior constraints (Li et al., 2020). Furthermore, *multi-manifold* learning is proposed to obtain intrinsic properties of different manifolds. Yang et al. (2016a) proposed a supervised MMD-Isomap where data points are partitioned into different manifolds according to label information. Similarly, Zhang et al. (2018) proposed a semi-supervised local multi-manifold learning framework, termed SSMM-Isomap, that applies the labeled and unlabeled training samples to perform the joint learning of local

neighborhood-preserving features. In most previous work on multi-manifold learning, the problem is considered from the perspective that the label is known or partially known, which significantly simplifies the problem. For unsupervised multi-manifold learning, it is still very challenging to decouple multiple overlapping manifolds, and that is exactly what this paper aims to explore.

3 Proposed Method

Consider a dataset X with N samples, and each sample $x_i \in \mathbb{R}^d$ is sampled from C different manifolds $\{M_c\}_{c=1}^{C}$. Assume that each category in the data set lies in a compact low-dimensional manifold, and the number of manifolds C is prior knowledge. Define two nonlinear mapping $z_i = f(x_i, \theta_f)$ and $y_i = g(z_i, \theta_g)$, where $z_i \in \mathbb{R}^m$ is the embedding of x_i in the latent space, y_i is the reconstruction of x_i. The j-th cluster center is denoted as $\mu_j \in \mathbb{R}^m$, where $\{\mu_j\}_{j=1}^{C}$ is defined as a set of **learnable** parameters. We aim to find optimal parameters θ_f so that the embedding features $\{z_i\}_{i=1}^{N}$ can achieve clustering with local and global structure preservation. To this end, a denoising autoencoder (Vincent et al., 2010) shown in Fig 1 is first pre-trained in an unsupervised manner to learn an initial latent space. Denoising autoencoder aims to optimize the self-reconstruction loss $L_{AE} = MSE(\hat{x}, y)$, where the $\hat{x}$ is a copy of x with Gaussian noise added, that is, $\hat{x} = x + N(0, \sigma^2)$. Then the autoencoder is finetuned by optimizing the following clustering-oriented loss $\{L_{cluster}(z, \mu)\}$ and structure-oriented losses $\{L_{rank}(x, \mu), L_{LIS}(x, z), L_{align}(z, \mu)\}$. Since the clustering should be performed on features of **clean** data, instead of noised data $\hat{x}$ that is used in denoising autoencoder, the clean data x is used for fine-tuning.

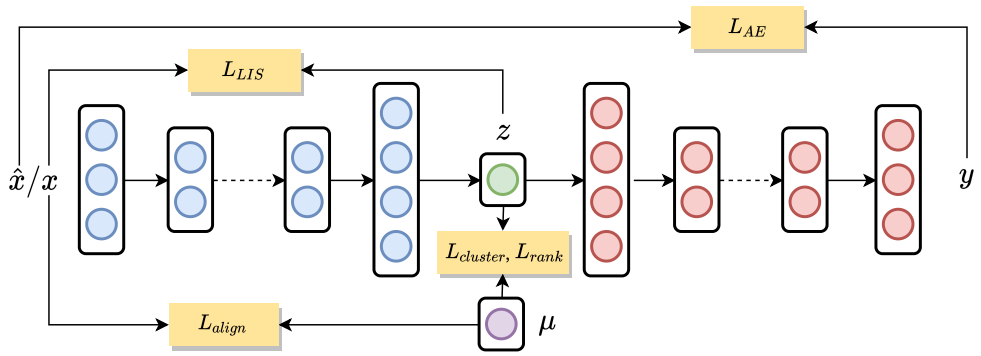

Figure 1: The framework of the proposed DCRL method. The encoder, decoder, latent space, and cluster centers are marked as blue, red, green, and purple, respectively.

3.1 Clustering-oriented Loss

First, the cluster centers $\{\mu_j\}_{j=1}^{C}$ in the latent space Z are initialized (the initialization method will be introduced in Sec 4.1). Then the similarity between the embedded point z_i and cluster centers $\{\mu_j\}_{j=1}^{C}$ is measured by Student's t-distribution:

$$q_{ij} = \frac{\left(1 + \|z_i - \mu_j\|^2\right)^{-1}}{\sum_{j'} \left(1 + \|z_i - \mu_{j'}\|^2\right)^{-1}} \tag{1}$$

The auxiliary target distribution is designed to help manipulate the latent space, defined as:

$$p_{ij} = \frac{q_{ij}^2 / f_j}{\sum_{j'} q_{ij'}^2 / f_{j'}}, \quad where \quad f_j = \sum_i q_{ij} \tag{2}$$

where f_j is the normalized cluster frequency, used to balance the size of different clusters. Then the encoder is optimized by the following objective:

$$L_{cluster} = \mathrm{KL}(P \| Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}} \tag{3}$$

The gradient of $L_{cluster}$ with respect to each learnable cluster center μ_j can be computed as:

$$\frac{\partial L_{cluster}}{\partial \mu_j} = -\sum_i \left(1 + \|z_i - \mu_j\|^2\right)^{-1} \cdot (p_{ij} - q_{ij})(z_i - \mu_j) \tag{4}$$

$L_{cluster}$ facilitates the aggregation of data points within the same manifold, while data points from different manifolds are kept away from each other. However, we find that the clustering-oriented loss may deteriorate the geometric structure of the latent space, which hurts the clustering accuracy and leads to meaningless representation. To prevent the deterioration of clustering loss, we introduce isometry loss L_{LIS} and ranking loss L_{rank} to preserve the local and global structure, respectively.

3.2 Structure-oriented Loss

Intra-manifold Isometry Loss. The intra-manifold local structure is preserved by optimizing the following objective:

$$L_{LIS} = \sum_{i=1}^{N} \sum_{j \in \mathcal{N}_i^Z} |d_X(x_i, x_j) - d_Z(z_i, z_j)| \cdot \pi(l(x_i) = l(x_j)) \tag{5}$$

where $\mathcal{N}_i^Z$ represents the neighborhood of data point z_i in the feature space Z, and the kNN is applied to determine the neighborhood. $\pi(\cdot) \in \{0, 1\}$ is an indicator function, and $l(x_i)$ is a manifold determination function that returns the manifold s_i where sample x_i is located, that is $s_i = l(x_i) = \arg\max_j p_{ij}$. Then we can derive C manifolds $\{M_c\}_{c=1}^{C}$: $M_c = \{x_i;\ s_i = c, i = 1, 2, ..., N\}$. In a nutshell, the loss L_{LIS} constrains the isometry within each manifold.

Inter-manifold Ranking Loss. The inter-manifold global structure is preserved by optimizing the following objective:

$$L_{rank} = \sum_{i=1}^{C} \sum_{j=1}^{C} \left| d_Z(\mu_i, \mu_j) - scale * d_X\left(v_i^X, v_j^X\right) \right| \tag{6}$$

where $\{v_j^X\}_{j=1}^{C}$ is defined as the centers of different manifolds in the original input space X with $v_j^X = \frac{1}{|M_j|} \sum_{i \in M_j} x_i$ ($j = 1, 2, ..., C$). The parameter $scale$ determines the extent to which different manifolds move away from each other. The larger $scale$ is, the further away the different manifolds are from each other. The derivation for the gradient of L_{rank} with respect to each learnable cluster center μ_j is placed in **Appendix A.1**. Additionally, contrary to us, the conventional methods for dealing with inter-manifold separation typically impose push-away constraints on all data points from different manifolds (Zhang et al., 2018; Yang et al., 2016a), defined as:

$$L_{sep} = -\sum_{i=1}^{N} \sum_{j=1}^{N} d_Z(z_i, z_j) \cdot \pi(l(x_i) \neq l(x_j)) \tag{7}$$

The main differences between L_{rank} and L_{sep} are as follows: (1) L_{sep} imposes constraints on embedding points $\{z_i\}_{i=1}^{N}$, which in turn indirectly affects the network parameters θ_f. In contrast, L_{rank} imposes rank-preservation constrains directly on learnable parameters $\{\mu_j\}_{j=1}^{C}$ in the form of *regularization* item to control the separation of the clustering centers. (2) L_{sep} involves $N \times N$ *point-to-point* relationships, while L_{rank} involves only $C \times C$ *cluster-to-cluster* relationships, so L_{rank} is easier to optimize, faster to process, and more accurate. (3) The parameter $scale$ introduced in L_{rank} allows us to control the extent of separation between manifolds for specific downstream tasks.

Alignment Loss. Note that the global ranking loss L_{rank} is imposed directly on the learnable parameter $\{\mu_j\}_{j=1}^{C}$, so optimizing L_{rank} will only update $\{\mu_j\}_{j=1}^{C}$ rather the encoder's parameter θ_f. Thus here we need to introduce an *auxiliary* item L_{align} to align learnable cluster centers $\{\mu_j\}_{j=1}^{C}$ with real cluster centers $\{v_j^Z\}_{j=1}^{C}$:

$$L_{align} = \sum_{j=1}^{C} ||\mu_j - v_j^Z|| \tag{8}$$

where $\{v_j^Z\}_{j=1}^{C}$ are defined as $v_j^Z = \frac{1}{|M_j|} \sum_{i \in M_j} z_i$ ($j = 1, 2, ..., C$). We place the derivation for the gradient of L_{align} with respect to each learnable cluster center μ_j in **Appendix A.1**.

3.3 TRAINING STRATEGY

3.3.1 CONTRADICTION

The contradiction between clustering and local structure preservation is analyzed from the *forces analysis* perspective. As shown in Fig 2, we assume that there exists a data point (red point) and its three nearest neighbors (blue points) around a cluster center (gray point). When clustering and local structure preserving are optimized simultaneously, it is very easy to fall into a local optimum, where the data point is in steady-state, and the resultant force from its three nearest neighbors is equal in magnitude and opposite to the gravitational forces of the cluster. Therefore, the following training strategy is applied to prevent such local optimal solutions.

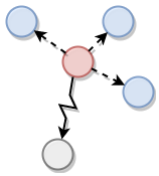

Figure 2: The force analysis of contradiction between clustering and local structure preservation.

3.3.2 ALTERNATING TRAINING AND WEIGHT GRADUALITY

Alternating Training. To solve the above problem and integrate the goals of clustering and structure preservation into a unified framework, we take an alternating training strategy. Within each epoch, we first jointly optimize $L_{cluster}$ and L_{rank} in a *mini-batch*, with joint loss defined as

$$L_1 = L_{AE} + \alpha L_{cluster} + L_{rank} \tag{9}$$

where α is the weighting factor that balances the effects of clustering and global rank-preservation. Then at each epoch, we optimize isometry loss L_{LIS} and L_{align} on *the whole dataset*, defined as

$$L_2 = \beta L_{LIS} + L_{align} \tag{10}$$

Weight Graduality. At different stages of training, we have different expectations for the clustering and structure-preserving. At the beginning of training, to successfully decouple the overlapping manifolds, we hope that the $L_{cluster}$ will dominate and L_{LIS} will be auxiliary. When the margin between different manifolds is sufficiently pronounced, the weight α for $L_{cluster}$ can be gradually reduced, while the weight β for L_{LIS} can be gradually increased, focusing on the preservation of the local isometry. The whole algorithm is summarized in Algorithm 1 in **Appendix A.2**.

Three-stage explanation. The entire training process can be roughly divided into three stages, as shown in Fig 3, to explain the training strategy more vividly. At first, four different manifolds overlap each other. At Stage 1, $L_{cluster}$ dominates, thus data points within each manifold are converging towards the clustering center to form a sphere, and the local structure of manifolds is destroyed. At Stage 2, L_{rank} dominates, thus different manifolds in the latent space move away from each other to increase the manifold margin and enhance the discriminability. At stage 3, the manifolds gradually recover their original local structure from the spherical shape with L_{LIS} dominating. It is worth noting that **all of the above losses coexist rather than independently at different stages**, but that the role played by different losses varies due to the alternating training and weight graduality.

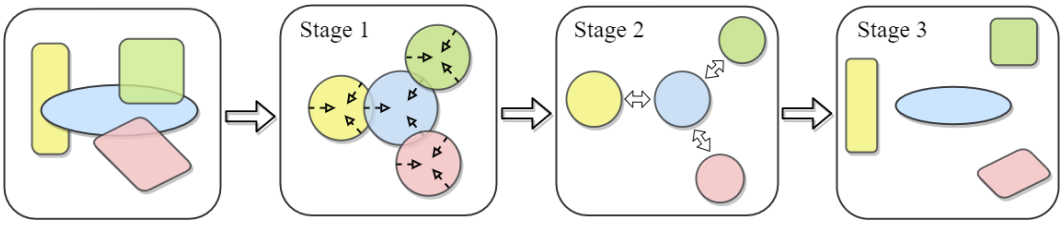

Figure 3: Schematic of training strategy. Four different colors and shapes represent four intersecting manifolds, and three stages involve the clustering, separation, and structure recovery of manifolds.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUPS

In this section, the effectiveness of the proposed framework is evaluated in 5 benchmark datasets: USPS[1], MNIST-full, MNIST-test (LeCun et al., 1998), Fashion-MNIST (Xiao et al., 2017) and REUTERS-10K (Lewis et al., 2004), on which our method is compared with 8 other methods mentioned in Sec 2 in 8 evaluation metrics including metrics designed specifically for clustering and representation learning. The brief descriptions of the datasets are given in **Appendix A.3**.

Parameters settings. The encoder is a multilayer perceptron (MLP) with dimensions d-500-500-2000-10 where d is the dimension of the input data, and the decoder is its mirror. After pretraining, in order to initialize the learnable clustering centers, the t-SNE is applied to transform the latent space Z to 2 dimensions further, and then the K-Means algorithm is run to obtain the label assignments for each data point. The centers of each category in the latent space Z are set as initial clustering centers $\{\mu_j\}_{j=1}^{C}$. The batch size is set to 256, the epoch is set to 300, the parameter k for nearest neighbor is set to 3, and the parameter $scale$ is set to 3 for all datasets. Besides, Adam optimizer (Kingma & Ba, 2014) with learning rate λ=0.001 is used. As described in Sec 3.3.2, the weight graduality is applied to train the model. The weight parameter α for $L_{cluster}$ decreases linearly from 0.1 to 0 within epoch 0-150. In contrast, the weight parameter β for loss L_{LIS} increases linearly from 0 to 1.0 within epoch 0-150. The implementation is based on the PyTorch running on NVIDIA v100 GPU.

Evaluation Metrics. Two standard evaluation metrics: Accuracy (ACC) and Normalized Mutual Information (NMI) (Xu et al., 2003) are used to evaluate clustering performance. Besides, six evaluation metrics are adopted in this paper to evaluate the performance of representation learning, including Relative Rank Error (RRE), Trustworthiness (Trust), Continuity (Cont), Root Mean Reconstruction Error (RMRE), Locally Geometric Distortion (LGD) and Cluster Rank Accuracy (CRA). Limited by space, their precise definitions are available in **Appendix A.4**.

4.2 EVALUATION OF CLUSTERING

4.2.1 QUANTITATIVE COMPARISON

The metrics ACC/NMI of different methods on various datasets are reported in Tab 1. For those comparison methods whose results are not reported on some datasets, we run the released code using the hyperparameters provided in their paper and label them with (*). We find that our method outperforms K-Means and SC-Ncut with a significant margin and surpasses the other six competing DNN-based algorithms on all datasets except MNIST-test. With even the MNIST-test dataset, we still rank second, outperforming the third by 1.1%. In particular, we obtained the best performance on the Fashion-MNIST dataset and, more notably, our clustering accuracy exceeds the current best method (N2D) by 3.8%. Despite $L_{cluster}$ is inspired by and highly consistent with the design of DEC, our method achieves much better clustering results than them. With MNIST-full, for example, our clustering accuracy is 11.7% and 9.9% higher than DEC and IDEC, respectively.

Table 1: Clustering performance (ACC/NMI) of different algorithms on five datasets

Algorithms	MNIST-full	MNIST-test	USPS	Fashion-MNIST	REUTERS-10K
k-means	0.532/0.500	0.546/0.501	0.668/0.601	0.474/0.512	0.599/0.375*
SC-Ncut	0.656/0.731	0.660/0.704	0.649/0.794	0.508/0.575	0.658/0.401*
AE+k-means	0.818/0.747	0.815/0.784*	0.662/0.693	0.566/0.585*	0.721/0.432*
DEC	0.863/0.834	0.856/0.830	0.762/0.767	0.518/0.546	0.755/0.503*
IDEC	0.881/0.867	0.846/0.802	0.761/0.785	0.529/0.557	0.778/0.527*
JULE	0.964/0.913	0.961/0.915	0.950/0.913	0.563/0.608	0.797/0.551*
DSC	0.978/0.941	**0.980/0.946**	0.869/0.857	0.662/0.645	0.743/0.484*
N2D	0.979/0.942	0.948/0.882	0.958/0.901	0.672/0.684	0.808/0.548
DCRL (ours)	**0.980/0.946**	0.972/0.930	**0.960/0.902**	**0.710/0.685**	**0.836/0.590**

[1]https://cs.nyu.edu/roweis/data.html

4.2.2 GENERALIZABILITY EVALUATION

Tab 2 demonstrates that a learned DCRL can generalize well to unseen data with high clustering accuracy. Taking MNIST-full as an example, DCRL was trained using 50,000 training samples and then tested on the remaining 20,000 testing samples using the learned model. In terms of the metrics ACC and MNI, our method is optimal for both training and test samples. More importantly, there is hardly any degradation in the performance of our method on the test samples compared to the training samples, while all other methods showed a significant drop in performance, e.g., DEC from 84.1% to 74.8%. This demonstrates the importance of geometric structure maintenance for good generalizability. The testing visualization available in **Appendix A.5** shows that DCRL still maintains clear inter-cluster boundaries even on the test samples, which demonstrates the great generalizability of our method.

Table 2: Generalizability evaluated by ACC/NMI

Algorithms	**training samples**	**testing samples**
AE+k-means	0.815/0.736	0.751/0.711
DEC	0.841/0.773	0.748/0.704
IDEC	0.845/0.860	0.826/0.842
JULE	0.958/0.907	0.921/0.895
DSC	0.975/0.939	0.969/0.921
N2D	0.974/0.930	0.965/0.911
DCRL (ours)	**0.978/0.941**	**0.978/0.941**

4.2.3 CLUSTERING VISUALIZATION

The visualization of DCRL with several comparison methods is shown in Fig 4 (visualized using UMAP). From the perspective of clustering, our method is much better than the other methods. Among all methods, only DEC, IDEC and DCRL can hold clear boundaries between different clusters, while the cluster boundaries of the other methods are indistinguishable. Although DEC and IDEC can successfully separate different clusters, they group many data points from different classes into the same cluster. Most importantly, due to the use of the clustering-oriented loss, the embedding learned by algorithms such as DEC, IDEC, JULE, and DSC (especially DSC) tend to form spheres and disrupt the original topological structure. Instead, our method overcomes these problems and achieves almost perfect separation between different clusters while preserving the local and global structure. Additionally, the embedding of latent space during training process is visualized in **Appendix A.6**, which is highly consistent with the three-stage explanation mentioned in Sec 3.3.2.

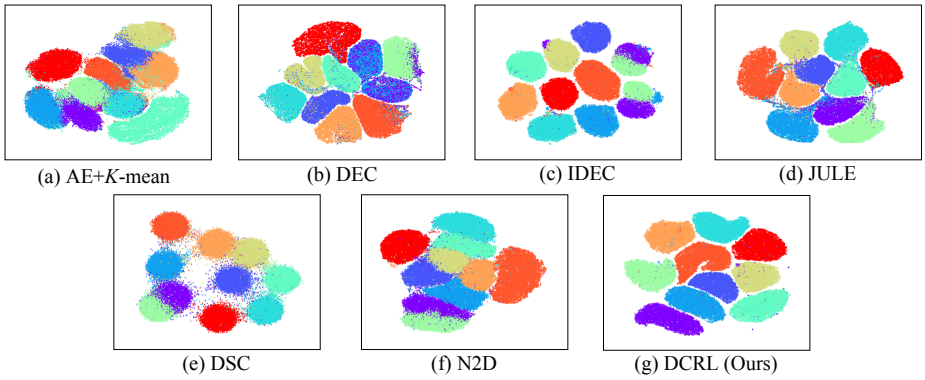

Figure 4: Visualization of the embedding learned by different algorithms on MNIST-full dataset.

4.3 EVALUATION OF REPRESENTATION LEARNING

4.3.1 QUANTITATIVE COMPARISON

Although numerous previous work has claimed that they brought clustering and representation learning into a unified framework, they all, unfortunately, lack an analysis of the effectiveness of the learned representations. In this paper, we compare DCRL with the other five methods in six evaluation metrics on five datasets. (Limited by space, only MNIST-full results are provided in the Tab 3 and the complete results are in **Appendix A.7**). The results show that DCRL outperforms all other methods, especially in the CRA metric, which is not only the best on all datasets but also reaches 1.0. This means that the "rank" between different manifolds in the latent space is completely preserved and undamaged, which proves the effectiveness of our global ranking loss L_{rank}. Moreover,

statistical analysis is performed in this paper to show the extent to which local and global structure is preserved in the latent space for each algorithm. Limited by space, they are placed in **Appendix A.8**.

Table 3: Performance for representation learning

Methods	RRE	Trust	Cont	RMSE	LGD	CRA
DEC	0.099	0.844	0.948	44.85	4.379	0.28
IDEC	0.009	0.998	0.979	24.58	1.714	0.33
JULE	0.026	0.936	0.983	28.34	2.129	0.27
DSC	0.097	0.873	0.925	6.98	1.198	0.23
N2D	0.010	0.992	0.984	5.71	0.699	0.21
DCRL	**0.005**	**0.999**	**0.987**	**5.49**	**0.691**	**1.00**

Table 4: Performance for downstream tasks

Methods	MLP	RFC	SVM	LR
AE	0.974	0.965	0.985	0.956
DEC	0.864	0.870	0.870	0.856
IDEC	0.979	0.973	0.985	0.965
JULE	0.980	0.982	0.978	0.974
DSC	0.962	0.950	0.983	0.975
N2D	0.979	0.980	0.979	0.979
DCRL	**0.985**	**0.987**	**0.986**	**0.984**

4.3.2 DOWNSTREAM TASKS

Recently, numerous deep clustering algorithms have claimed to obtain meaningful representations, however, they do not analyze and experiment with the so-called "meaningful". Therefore, we are interested to see whether these proposed methods can indeed learn representations that are useful for downstream tasks. Tab 4 compares DCRL with the other six methods on five datasets (Limited by space, only MNIST-full results are provided in the paper, and the complete results are in **Appendix A.9**). Four different classifiers, including a linear classifier (Logistic Regression; LR), two nonlinear classifiers (MLP, SVM), and a tree-based classifier (Random Forest Classifier; RFC) are used as downstream tasks, all of which use default parameters and default implementations in sklearn (Pedregosa et al., 2011) for a fair comparison. The learned representations are frozen and used as input for training. The classification accuracy evaluated on the test set serves as a metric to evaluate the effectiveness of learned representations. On the MNIST-full dataset, our method outperforms all the other methods. Moreover, we surprisingly find that with MLP and RFC as downstream tasks, all methods except DCRL could not even match the accuracy of AE. Significantly, the performance of DEC on downstream tasks deteriorates sharply and even shows a large gap with the simplest AEs, which once again shows that the clustering-oriented loss may damage the data geometric structure.

4.4 ABLATION STUDY

This evaluates the effects of the loss terms and training strategies in the DCRL with five sets of experiments: the model without (A) Structure-oriented Loss (SL); (B) Clustering-oriented Loss (CL); (C) Weight Graduality (WG); (D) Alternating Training (AT), and (E) the full model. Limited by space, only MNIST-full results are provided in the paper, and results for the other four datasets are in **Appendix A.10**. After analyzing the results, we can conclude: (1) CL is the most important factor for obtaining good clustering, the lack of which leads to unsuccessful clustering, hence the numbers in the table are not very meaningful and are shown in gray color. (2) SL not only brings subtle improvements in clustering but also greatly improves the performance of representation learning. (3) Our elegant training strategies (WG and AT) both improve the performance of clustering and representation learning to some extent, especially on metrics such as RRE, Trust, Cont, and CRA.

Table 5: Ablation study of loss items and training strategies on MNIST-full dataset

Datasets	Methods	ACC/NMI	RRE	Trust	Cont	RMSE	LGD	CRA
MNIST-full	w/o SL	0.976/0.939	0.0093	0.9967	0.9816	24.589	1.6747	0.32
	w/o CL	0.814/0.736	0.0004	0.9998	0.9990	7.458	0.0487	1.00
	w/o WG	0.977/0.943	0.0065	0.9987	0.9860	5.576	0.6968	0.98
	w/o AT	0.978/0.944	0.0069	0.9986	0.9851	5.617	0.7037	0.96
	full model	**0.980/0.946**	**0.0056**	**0.9997**	**0.9871**	**5.498**	**0.6916**	**1.00**

5 CONCLUSION

The proposed DCRL framework imposes clustering-oriented and structure-oriented constraints to optimize the latent space for simultaneously performing clustering and representation learning with local and global structure preservation. Extensive experiments on image and text datasets demonstrate that DCRL is not only comparable to the state-of-the-art deep clustering algorithms but also able to learn effective and robust representation, which is beyond the capability of those clustering methods that only care about *clustering accuracy*. Future work will focus on the adaptive determination of manifolds (clusters) number and extend our work to datasets with larger scale.

REFERENCES

Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.

Xifeng Guo, Long Gao, Xinwang Liu, and Jianping Yin. Improved deep embedded clustering with local structure preservation. In *IJCAI*, pp. 1753–1759, 2017.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5(Apr):361–397, 2004.

Stan Z Li, Zelin Zhang, and Lirong Wu. Markov-lipschitz deep learning. *arXiv preprint arXiv:2006.08256*, 2020.

Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.

J MacQueen. Some methods for classification and analysis of multivariate observations. In *Proc. 5th Berkeley Symposium on Math., Stat., and Prob*, pp. 281, 1965.

Ryan McConville, Raul Santos-Rodriguez, Robert J Piechocki, and Ian Craddock. N2d:(not too) deep clustering via clustering the local manifold of an autoencoded embedding. *arXiv preprint arXiv:1908.05968*, 2019.

Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.

Uri Shaham, Kelly Stanton, Henry Li, Boaz Nadler, Ronen Basri, and Yuval Kluger. Spectralnet: Spectral clustering using deep neural networks. *arXiv preprint arXiv:1801.01587*, 2018.

Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000.

Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.

Fei Tian, Bin Gao, Qing Cui, Enhong Chen, and Tie-Yan Liu. Learning deep representations for graph clustering. In *Aaai*, volume 14, pp. 1293–1299. Citeseer, 2014.

Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(12), 2010.

Svante Wold, Kim Esbensen, and Paul Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, pp. 478–487, 2016.

Wei Xu, Xin Liu, and Yihong Gong. Document clustering based on non-negative matrix factorization. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pp. 267–273, 2003.

Bo Yang, Ming Xiang, and Yupei Zhang. Multi-manifold discriminant isomap for visualization and classification. *Pattern Recognition*, 55:215–230, 2016a.

Jianwei Yang, Devi Parikh, and Dhruv Batra. Joint unsupervised learning of deep representations and image clusters. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5147–5156, 2016b.

Xu Yang, Cheng Deng, Feng Zheng, Junchi Yan, and Wei Liu. Deep spectral clustering using dual autoencoder network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4066–4075, 2019.

Yan Zhang, Zhao Zhang, Jie Qin, Li Zhang, Bing Li, and Fanzhang Li. Semi-supervised local multi-manifold isomap by linear embedding for feature extraction. *Pattern Recognition*, 76:662–678, 2018.