



Universiteit
Leiden
The Netherlands

Deep recursive embedding for high-dimensional data

Zhou, Z.X.; , X.R. zu; Wang, Y.Y.; Lelieveldt, B.P.F.; Tao, Q.

Citation

Zhou, Z. X., Wang, Y. Y., Lelieveldt, B. P. F., & Tao, Q. (2022). Deep recursive embedding for high-dimensional data. *Ieee Transactions On Visualization And Computer Graphics*, 28(2), 1237-1248. doi:10.1109/TVCG.2021.3122388

Version: Publisher's Version
License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)
Downloaded from: <https://hdl.handle.net/1887/3567535>

Note: To cite this publication please use the final published version (if applicable).

Deep Recursive Embedding for High-Dimensional Data

Zixia Zhou¹, Xinrui Zu, Yuanyuan Wang², *Senior Member, IEEE*,
Boudewijn P. F. Lelieveldt³, and Qian Tao⁴

Abstract—Embedding high-dimensional data onto a low-dimensional manifold is of both theoretical and practical value. In this article, we propose to combine deep neural networks (DNN) with mathematics-guided embedding rules for high-dimensional data embedding. We introduce a generic deep embedding network (DEN) framework, which is able to learn a parametric mapping from high-dimensional space to low-dimensional space, guided by well-established objectives such as Kullback-Leibler (KL) divergence minimization. We further propose a recursive strategy, called deep recursive embedding (DRE), to make use of the latent data representations for boosted embedding performance. We exemplify the flexibility of DRE by different architectures and loss functions, and benchmarked our method against the two most popular embedding methods, namely, t-distributed stochastic neighbor embedding (t-SNE) and uniform manifold approximation and projection (UMAP). The proposed DRE method can map out-of-sample data and scale to extremely large datasets. Experiments on a range of public datasets demonstrated improved embedding performance in terms of local and global structure preservation, compared with other state-of-the-art embedding methods. Code is available at <https://github.com/tao-aimi/DeepRecursiveEmbedding>.

Index Terms—t-distributed stochastic neighbor embedding, uniform manifold approximation and projection, deep embedding network, deep recursive embedding, unsupervised learning

1 INTRODUCTION

EMBEDDING high-dimensional data onto a low-dimensional manifold is of both theoretical and practical value. It can be used for a variety of applications such as data visualization, representation learning, unsupervised clustering, and data exploration [1], [2], [3], [4]. t-distributed stochastic neighbor embedding (t-SNE) [5] is among the most well-known and widely-used methods for high-dimensional data visualization, which preserves local similarity of high-dimensional data in a drastically reduced low-dimensional map (typically 2). t-SNE has found widespread applications in life science research among others in the last decade, and established itself as an important visualization tool in the scientific community [6]. Nonetheless, it is so far largely recognized as a visualization tool, as some practical issues may have limited its wider use in learning-based tasks. First, t-SNE does not

immediately allow out-of-sample projection. Second, the computation of t-SNE is memory- and time-consuming, not well scalable to the extremely large datasets of today. Third, t-SNE (especially its fast implementations) focuses on local neighborhood, often losing global data structure of data.

A number of follow-up work have been proposed after the initial paper of t-SNE. For instance, Barnes-Hut-SNE (BH-SNE) [7], A-tSNE [8] and FIt-SNE [9] were proposed to accelerate the computation and reduce memory usage of t-SNE. Recently, another DR method called uniform manifold approximation and projection (UMAP) [10] was introduced and became another popular visualization tool. UMAP has demonstrated competitive visualization performance as t-SNE, with notably reduced runtime and arguably better preservation of global structure [10].

In the meanwhile, DR has also been studied in the deep learning field [11], [12]. For example, the auto-encoder (AE) is a classical DR method [11], which embeds high-dimensional data through a DNN with an encoder-decoder architecture. AE can be trained with mean square error (MSE) loss and standard optimization, and the intermediate tensors at the information bottleneck are seen as low-dimensional representations of the high-dimensional data. Although its embedding performance is generally inferior to that of t-SNE or UMAP, AE suggests great potential of DNN for DR through unsupervised learning.

We posit that the combination of DNN and dedicated, mathematics-guided embedding rules may further enrich the possibilities of DR and improve its performance. Embedding rules such as cross entropy and Kullback-Leibler (KL) divergence are more specific to the purpose of DR, compared to MSE which is intended for reconstructing data. Generally speaking, there exist two methodologies to combine DNN and

- Zixia Zhou is with the Department of Electronic Engineering, Fudan University, Shanghai 200433, China. E-mail: 16110720022@fudan.edu.cn.
- Xinrui Zu is with the Faculty of Electrical Engineering, Mathematics and Computer Science (EEMCS), University of Twente, 7522, NB, Enschede, The Netherlands. E-mail: x.zu@student.utwente.nl.
- Yuanyuan Wang is with the Department of Electronic Engineering, Fudan University, Shanghai 200433, China, and also with the Key Laboratory of Medical Imaging Computing and Computer Assisted Intervention of Shanghai, Shanghai 200032, China. E-mail: yywang@fudan.edu.cn.
- Boudewijn P. F. Lelieveldt is with the Department of Radiology, Division of Image Processing, Leiden University Medical Center, 2333, ZA, Leiden, The Netherlands. E-mail: B.P.F.Lelieveldt@lumc.nl.
- Qian Tao is with the Department of Imaging Physics, Delft University of Technology, 2628, CJ, Delft, The Netherlands. E-mail: q.tao@tudelft.nl.

Manuscript received 11 Feb. 2021; revised 14 Oct. 2021; accepted 15 Oct. 2021.

Date of publication 26 Oct. 2021; date of current version 31 Dec. 2021.

(Corresponding author: Qian Tao.)

Recommended for acceptance by P. Bremer.

Digital Object Identifier no. 10.1109/TVCG.2021.3122388

DR. Intuitively, we may teach the DNN to learn from a referenced DR method. The recent deep learning multidimensional projection method [13] follows this methodology, which first computes a t-SNE or UMAP map in conventional ways, then let the DNN learn the t-SNE or UMAP results by MSE. This method can achieve scalability and out-of-sample support, but its embedding performance has an upper limitation equal to that of the referenced method. The second methodology is to train a DNN with dedicated loss functions to directly capture the mathematical principles of embedding, therefore *mathematics-guided*. The parametric t-SNE (ptSNE) proposed in 2009 [14] is an early work in this direction, which employed the t-SNE loss function to train a restricted Boltzmann machine (RBM). ptSNE addresses the scalability issue, but still has some practical limitations. ptSNE involves a pre-training step and a fine-tuning step, which are complicated while not guaranteed to converge. Even with ideal convergence, ptSNE cannot exceed the performance of the original t-SNE.

We seek to develop a parametric (i.e., able to embed new data points) and scalable (i.e., able to embed infinite data points) DR method that can find an intrinsic low-dimensional representation of data. In this paper, we propose a generic *deep embedding network (DEN)* framework that is based on well-established embedding rules and optimized by modern DNN. In addition, to break the upper limit of the performance in classical DR methods, we propose a recursive training strategy called *deep recursive embedding (DRE)*, to make use of latent representations to boost the DR performance further.

Specifically, as we will show in later sections, the advantages of the proposed DRE method are four-fold: (1) It can readily map out-of-sample data points. (2) It is scalable and memory-efficient. (3) It is flexible in the design of its architecture, loss, and training strategy. (4) It breaks the upper limit of t-SNE or UMAP.

2 RELATED WORK

2.1 Classical Dimensionality Reduction Methods

Dimensionality reduction is a classical problem in machine learning. The most well-known linear DR method is principal component analysis (PCA) [15], which computes low-dimensional representation by linearly projecting high-dimensional data points onto the first few principal components that preserves the largest variance. Other classical DR methods include multidimensional scaling (MDS) [16], IsoMap [17], locally linear embedding (LLE) [18], and stochastic neighbor embedding (SNE) [19]. They usually produce better visualization results on nonlinear data compared with PCA, but with limitations in terms of scalability, speed, and stochasticity. In 2008, the t-SNE method [5], a variant of SNE, was proposed, which later became one of the most popular methods in the research community for high-dimensional data visualization. t-SNE emphasizes local neighborhood in data with a modified assumption (t-distribution in low-dimensional space) over SNE, and can produce higher quality embedding results compared with other nonlinear manifold learning methods.

Many developments followed up the original t-SNE [20], [21]. An important work is Barnes-Hut t-SNE (BH t-SNE) [7], which approximates the high-dimensional space with sparse

distributions and uses the Barnes-Hut algorithm to speed up computation. The approximation largely reduces the algorithm complexity and memory consumption, enabling t-SNE to embed very large datasets that were previously impossible. A hierarchical stochastic neighbor embedding (HSNE) method was proposed to interactively render the visualization at different scales [22]. The hierarchical way of visualization allows small memory footprint by focusing on landmarks instead of all data. The authors also proposed an A-tSNE method [8], which trades off speed and accuracy, to enable interactive data exploration. More recently, Linderman *et al.* proposed a fast Fourier transform-accelerated interpolation-based t-SNE (Fit-SNE) [9] to accelerate the implementation of t-SNE. In Fit-SNE, an effective late exaggeration strategy was also proposed for better clustering. Lately, another important embedding method UMAP was introduced [10], which is rooted in mathematical foundations of Riemannian geometry and algebraic topology. UMAP is considered competitive with t-SNE in visualization, while faster and better scalable in implementation. It is also able to map out-of-sample data by graph embedding. However, it was reported that UMAP can be sensitive to the choice of hyper-parameters [13]. Although some work argued that the global data structure can be better preserved by UMAP, other studies suggested that UMAP preserves the global structure in a way similar to t-SNE with the same initialization [23].

2.2 Deep Learning Methods

AE is a classical DNN-based DR method, which maps the data to a low-dimensional representation with an encoder. Through unsupervised learning, the encoder can generate a nonlinear embedding in the low-dimensional space. While AE uses the MSE loss, the ptSNE method [14], proposed by van der Maaten in 2009, adopted the same KL loss function as t-SNE. ptSNE consists of three separate steps: training RBMs, stacking the RBMs to construct a pre-trained NN, and finetuning the pretrained NN. Such a training process is cumbersome, and its performance is frequently poorer than that of the standard t-SNE due to the difficulty of optimization. Another DNN method was later proposed to directly cluster data, using a loss function similar to that of t-SNE, but based on estimated centroids [24]. Very recently, a DL Projection method [13] was introduced to learn from any referenced DR methods: it first selects a reference method (e.g., t-SNE, UMAP) to generate a projection of a subset; then trains a DNN to learn the embedding. The method is effective and intuitive, but its performance is restricted by its teachers. Although the network itself does not demand complex parameter settings, the front-end process (i.e., the DR method to generate the ground truth for supervised learning) still requires elaborate tuning.

We believe there is still much potential in DNN for DR, given its strong capability of representation learning. In this work, we propose a dedicated DNN-based DR framework, named *deep embedding network (DEN)*, to embed high-dimensional data in a unsupervised manner by well-established mathematical principles. Within this framework, we further propose a novel *deep recursive embedding (DRE)* method, which can make use of the latent representations of the original data for further boosted DR performance. To demonstrate the

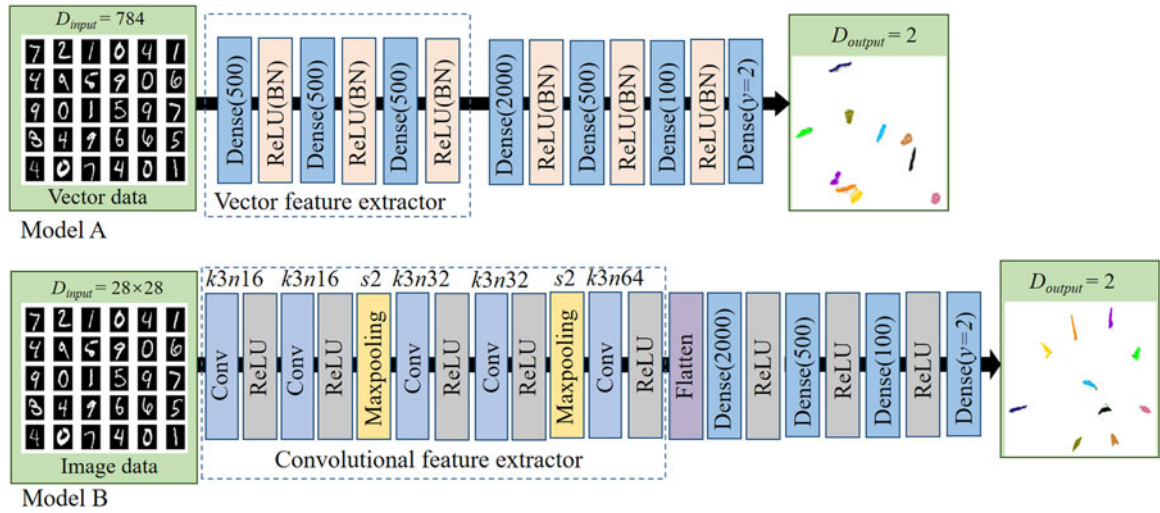


Fig. 1. Two exemplar architectures of the DEN model. Model A is a FCNN for vector input, and Model B is a CNN for image input. The CNN modules in Model B allow better extraction of features from images.

flexibility of our framework, we introduce a two-stage *t-SNE-UMAP loss*, which is able to combine the favorable properties of t-SNE and UMAP. We provide a further detailed comparison between our methods and other deep-learning-based methods in Discussion.

3 METHODS

3.1 Overview of DEN

We propose a generic framework, called Deep Embedding Network, to accomplish nonlinear DR through DNNs. In short, DEN is an unsupervised learning method which uses DNN to optimize a dedicated loss function such that the output low-dimensional representation optimally preserves the desirable geometrical properties of the input high-dimensional data. Once trained, it forms a parametric mapping from input to output, able to embed any out-of-sample data. DEN is flexible in its design: it can make use of any modern modules of DNN, take in data of different forms (i.e., not only vectors but also tensors), and allow different loss functions.

3.2 Flexible Architecture: FCNN and CNN

DEN allows very flexible design of the network architecture. Generally speaking, DEN can be divided into two parts: a feature extraction part and a dimensionality reduction part. In Fig. 1 we present two DEN architectures, named as Model A and Model B. Model A is a classical fully connected neural network (FCNN), while Model B includes convolutional neural network (CNN) modules for better extraction of features from images.

Each network consists of two parts: (1) Feature extraction: In this part, we aim to capture the high-level features from input data. In our experiments, we determined empirically the parameters of NN, and we noticed that the final performance is very robust to the setting of NN parameters. In Model A, the vector feature extractor contains 5 hidden layers, each of which followed by a ReLU and a batch normalization layer. In Model B, a standard design is followed: the first 2 convolutional layers (filter number=16) with ReLU and maxpooling capture shallow features, then 2 convolutional

layers with 32 filters followed by ReLU and maxpooling operation are used to extract deeper features related to the task.

(2) Dimension reduction: In this part, the dimensionality is gradually reduced, from 2000 to 500, then 100, and finally to the target dimension 2. We also followed a standard design in DNN: the first dense layer with 2000 neurons is used to lift the representations so that the capability to express data is increased. Then 2 dense layers with 500 and 100 neurons are used to reduce the dimension, as is often in NN design, i.e., a decreasing number of neurons at deeper layers. We note that the final DR results are not sensitive to the network hyperparameters.

3.3 Flexible Loss: Deep t-SNE and Deep UMAP

With the generic DEN architectures, we formulate dedicated loss functions to guide the training. DEN can take flexible loss functions, including that of t-SNE [5] and UMAP [10], both in an information-theoretic sense.

t-SNE [5] minimizes the Kullback-Leibler (KL) divergence between two estimated probability distributions, P and Q , where P is for x in the high-dimensional space, and Q for their representation y in the low-dimensional space

$$L_{t-SNE} = \text{KL}(P||Q) = \sum_i \sum_{j \neq i} p_{ij} \log \frac{p_{ij}}{q_{ij}}, \quad (1)$$

where p_{ij} and q_{ij} are normalized pairwise similarities: p_{ij} is the probability that a data point would choose the data point x_j as its neighbor under the Gaussian distribution in the high-dimensional space, and q_{ij} is the probability that a data point y_i would choose the data point y_j as its neighbor under the t-distribution in the low-dimensional space. p_{ij} and q_{ij} are calculated as follows:

$$p_{ji} = \frac{\exp\left(\frac{-\|x_i - x_j\|^2}{2\sigma_i^2}\right)}{\sum_{k \neq i} \exp\left(\frac{-\|x_i - x_k\|^2}{2\sigma_i^2}\right)} \quad (2)$$

$$p_{ij} = \frac{p_{ji} + p_{i|j}}{2N} \quad (3)$$

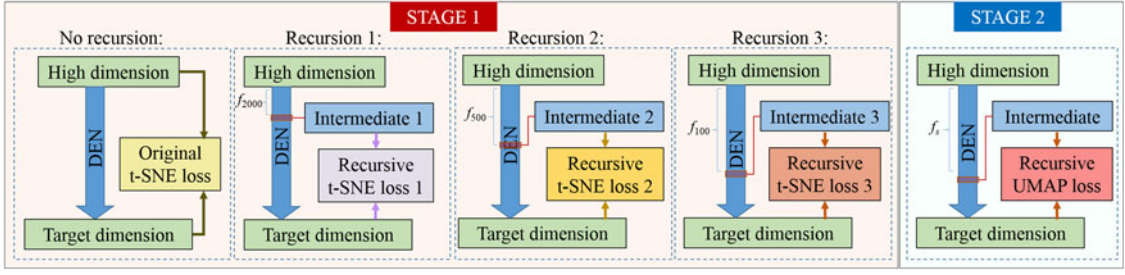


Fig. 2. The schematic of the proposed DRE method: Stage 1 shows the recursive training of DEN using the t-SNE loss, while Stage 2 further updates the network with the UMAP loss.

$$q_{ji} = \frac{(1 + \|y_i - y_j\|^2)^{-\frac{\alpha-1}{2}}}{\sum_k \sum_{l \neq k} (1 + \|y_k - y_l\|^2)^{-\frac{\alpha-1}{2}}}, \quad (4)$$

where N is the number of data points, σ_i is the Gaussian kernel at x_i , computed through a user defined perplexity [5], and α is the degree of freedom of t-distribution.

UMAP minimizes the fuzzy set cross entropy (CE) between two fuzzy membership functions [10], V and W , where V is for x in the high-dimensional space, and W for their representation y in the low-dimensional space. The loss is defined as

$$L_{\text{UMAP}} = \text{CE}(V||W) \quad (5)$$

$$= \sum_i \sum_{j \neq i} \left(v_{ij} \log \frac{v_{ij}}{w_{ij}} + (1 - v_{ij}) \log \frac{1 - v_{ij}}{1 - w_{ij}} \right). \quad (6)$$

In the high-dimensional space, v_{ij} is calculated as

$$v_{ji} = \exp\left(\frac{-d(x_i, x_j) + \rho_i}{\sigma_i}\right) \quad (7)$$

$$v_{ij} = (v_{ji} + v_{ij}) - v_{ji}v_{ij}, \quad (8)$$

where $d(x_i, x_j)$ is the distance between data point x_i and x_j . ρ_i is the minimum distance between x_i and its neighbors to ensure connectivity of the map. σ_i is derived as follows:

$$\sum_{j=1}^k \exp\left(\frac{-\max(0, d(x_i, x_j) - \rho_i)}{\sigma_i}\right) = \log_2(k), \quad (9)$$

where k is the number of nearest neighbors when constructing the UMAP graph.

In the low-dimensional space, w_{ij} is calculated as

$$w_{ij} = \left(1 + a\|y_i - y_j\|_2^{2b}\right)^{-1}, \quad (10)$$

where a and b are user-defined values [10].

We name the DEN trained with the t-SNE and UMAP loss *deep t-SNE* and *deep UMAP*, respectively. In the training process, mini-batches are used to reduce memory consumption for the expensive computation of P and V . We first randomly shuffle the entire dataset, then divide the dataset into mini-batches of a fixed size. After that, the batches are trained one-by-one. When all batches of the entire dataset have been trained, one epoch is accomplished. The loss calculation is as define in (1) and (6), and the two probability

distributions P and Q for t-SNE (or V and W for UMAP) are calculated from data within mini-batches.

3.4 Deep Recursive Embedding

With introduction of the modern DNN components such as ReLU activation and batch normalization, we can largely improve the ease and efficiency of training, compared with the RBM-based ptSNE. However, we note that the upper limit of the embedding performance remains the same as that of the original t-SNE or UMAP, given that the loss function is identical to its original definition. To push the limit, we further present a recursive training strategy, called deep recursive embedding (DRE).

DRE is proposed to make use of the latent representations extracted by DEN. In principle, the intermediate feature maps from different layers of DNN reflect different levels of abstraction of the original input. As discussed in [25], [26], the deeper layers of DNN provide more abstract representations specific to the learning task. In [27], the authors demonstrated that the deep layers of DNNs can extract a rich set of features, and leveraging these intermediate representations for knowledge distillation led to significant improvement of network efficiency. In [28], the high-level intermediate feature maps extracted from a pretrained DNN are used to compute the perceptual loss to improve the image reconstruction performance. For our purpose, the high-level representations obtained by Deep t-SNE (or Deep UMAP) also contain relevant information related to the task of DR. Although unsupervised, these representations are optimized with the information-theoretic loss that strives to match high-dimensional and low-dimensional statistics. Such representations therefore serves as proper representation of the original data to progressively improve the performance of DR.

We propose to make use of the intermediate output of DEN for computation of P or V . This process is therefore recursive, as illustrated in Fig. 2. In each recursion step, we invoke intermediate outputs of feature extraction. Then, based on these extracted feature maps, we calculate a new probability distributions \tilde{p}_{ji} and \tilde{v}_{ji} for \tilde{P} and \tilde{V}

$$\tilde{p}_{ji} = \frac{\exp\left(\frac{-\|f(x_i) - f(x_j)\|^2}{2\sigma_i^2}\right)}{\sum_{k \neq i} \exp\left(\frac{-\|f(x_i) - f(x_k)\|^2}{2\sigma_i^2}\right)} \quad (11)$$

$$\tilde{v}_{ji} = \exp\left(\frac{-d(f(x_i), f(x_j)) - \rho_i}{\sigma_i}\right), \quad (12)$$

where $f(\cdot)$ is high-level feature extraction function realized by the feature extraction part of DEN. Based on this new probability distributions, we can define the recursive t-SNE loss and recursive UMAP loss

$$L'_{t-SNE} = \text{KL}(\tilde{P}||Q) \quad (13)$$

$$L'_{UMAP} = \text{CE}(\tilde{V}||W). \quad (14)$$

At the beginning of recursive training, we calculate P and V once to estimate the loss in (1) and (6). In principle, we can also continuously update P and V during the training. However, this turned to be time-consuming and our experiments showed that the improvement was minor.

3.5 Multi-Stage Losses

Our design also allows flexible combination of loss functions at different stage of recursion. In theory, UMAP pushes the dispersed points between different clusters to the nearest ones, driven by the second term in its loss function (6). This results in a desirable visual effect promoting clean margins between classes (albeit not necessarily better classification performance or representation, as show in Results). We exemplify the flexible design of DEN by introducing a two-stage loss: first optimize DEN by the t-SNE loss, then refine the final visualization by the UMAP loss. The workflow is as follows:

In stage 1, the network is trained with the original t-SNE loss L_{t-SNE} between y and x ; then, the network is recursively updated by minimizing the recursive t-SNE loss L'_{t-SNE} between y and the latent representations, generated by the output of Dense 2000, Dense 500 and Dense 100 layers, where the feature extractors are defined as $f_{2000}(\cdot)$, $f_{500}(\cdot)$, and $f_{100}(\cdot)$, respectively. This is shown as Recursion 1, Recursion 2 and Recursion 3 in Fig. 2 (STAGE 1).

In stage 2, the DEN is further fine-tuned with the UMAP loss. This step is optional and for visualization purpose. This is shown in Fig. 2 (STAGE 2).

4 EXPERIMENTS

4.1 Datasets

We experimented with 5 public datasets, covering different data types, sizes, and characteristics: (1) MNIST [29], the handwritten digits dataset widely used to evaluate machine learning algorithms, which contains 60,000 training images and 10,000 testing images with size 28×28 , in 10 classes from digit 0 to dig-it 9. (2) Fashion-MNIST [30], the fashion product database by Zalando, developed in the same format as MNIST, containing 60,000 training images and 10,000 testing images of size 28×28 , including 10 classes: T-shirt, trousers, pullover, dress, coat, sandal, shirt, sneaker, bag, and ankle boot. Fashion-MNIST is another benchmark database for machine learning algorithms, more challenging than MNIST. (3) RNA-Seq dataset includes 23,822 single-cell transcriptomes. The cells were isolated from the primary visual cortex (VISp) and anterior lateral motor cortex (ALM) of adult mouse. The cluster label is defined in [31], which is used for visualizing different cell types. (4) IMDB dataset [32] includes 25,000 movie rating data used for sentiment analysis. All data are classified into positive and negative comments. Each

comment was preprocessed to transform the textual sequences into a 500-dimensional word-vector. (5) InfiMNIST [33]: an open source method to generate an infinitely large database of handwritten digits based on MNIST. This dataset is used to test the scalability of DR algorithms.

4.2 Implementation

Our environment is a Linux workstation with a NVIDIA Tesla V100 GPU, with 50GB system memory and 16GB GPU memory. The proposed DEN has a number of loss-related and network-related hyperparameters. In our experiments, we used Model A in Fig. 1 for vector input including RNA-Seq and IMDB, and Model B in Fig. 1 for image input including MNIST, Fashion-MNIST, and InfiMNIST. For the loss-related hyperparameters, we set the perplexity to 30 and the degree of freedom of t-distribution to 1. For the network-related hyperparameters, a mini-batch size of 2500 is chosen. Given one mini-batch of 2500 randomly selected data points from the input, P_{ij} is computed to derive the gradient for updating the network parameters. Then the next mini-batch follows to update the network parameters further. The Adam optimizer was used for DNN training, with an initial learning rate of 10^{-3} , β_1 (the exponential decay rate for the 1st moment estimate) of 0.9, β_2 (the exponential decay rate for the 2nd moment estimate) of 0.999 and an epsilon of 10^{-7} . The training epochs of Deep t-SNE is set as 100. The DRE method has the first 100 epochs trained with the original t-SNE loss, and 50 epochs for each subsequent recursion in Stage 1. In Stage 2, an additional 100 epochs were trained with the UMAP loss.

4.3 Evaluation Metrics

We used 6 metrics to evaluate the DR performance [34], [35]: (1) 1-nearest neighbor (1NN), (2) neighborhood hit, (3) trustworthiness, (4) continuity, (5) Shepard goodness, and (6) normalized stress. The definition of evaluation metrics are given as follows, where the dataset is denoted as $D = \{x_i\}$, $i = 1, \dots, N$, N is the number of sample points.

(1) The 1-nearest neighbor (1NN) classification accuracy is a standard measurement as reported in [5], [14], which can be used to estimate the quality of clustering. With its absolute simplicity, the 1NN classification accuracy can be an indication of the goodness of clustering in the low-dimensional space.

(2) The neighborhood hit measures how well separable the data is in the low-dimensional space, which helps gauge if a technique is good for data exploration. The neighborhood hit is defined as

$$\sum_{i=1}^N \frac{j \in N_i^{(K)} : l_j = l_i}{KN}, \quad (15)$$

which indicates the proportion of K neighbors $N_i^{(K)}$ of a point i in the low-dimensional space who have the same label l as point i itself, averaged over all points in the low-dimensional space ($K = 7$ as commonly used in literature).

(3) The trustworthiness measures the proportion of points in D that are also close in its mapping, suggesting how much one can trust the local patterns in a projection. The trustworthiness is defined as

$$1 - \frac{2}{NK(2n - 3K - 1)} \sum_{i=1}^N \sum_{j \in U_i^{(K)}} (r(i, j) - K), \quad (16)$$

where $r(i, j)$ represents the rank of the low-dimensional data point j according to the pairwise distances between the low-dimensional data points, and $U_i^{(K)}$ represents the set of points that are among the K nearest neighbors in the low-dimensional space but not in the high-dimensional space.

(4) The continuity measures the proportion of points in its mapping that are also close together in its original space, and is closely related to the missing neighbors of a projected point. The continuity is defined as

$$1 - \frac{2}{NK(2n - 3K - 1)} \sum_{i=1}^N \sum_{j \in V_i^{(K)}} (\hat{r}(i, j) - K), \quad (17)$$

where $\hat{r}(i, j)$ represents the rank of the high dimensional data point j according to the pairwise distances between the low-dimensional data points, and $V_i^{(K)}$ represents the set of points that are among the K nearest neighbors in the high-dimensional space but not in the low-dimensional space.

(5) The normalized stress measures the relative preservation of point-pairwise distances, which can be expressed as

$$\frac{\sum_{i,j} (\Delta^L(x_i, x_j) - \Delta^H(P(x_i), P(x_j)))^2}{\sum_{i,j} \Delta^L(x_i, x_j)}, \quad (18)$$

where Δ^L and Δ^H are distance metrics for data points in low- and high-dimensional space, respectively.

(6) The Shepard goodness measures the overall distance preservation by computing the Spearman rank correlation of Shepard diagram. The formula of generating a Shepard diagram is as follows:

$$\text{Scatterplot}(\|x_i - x_j\|, \|P(x_i) - P(x_j)\|), 1 \leq i \leq N, i \neq j. \quad (19)$$

5 RESULTS

5.1 Effect of Recursive Training

During recursive training, we observed that the clusters were gradually refined to better reflect the inter-cluster relationship, resulting in better separable clusters. This is reminiscent of the late exaggeration in Fit-SNE [9], which also forces clusters to be tighter and more apart. For an empirical comparison between Fit-SNE and DRE, we implemented Fit-SNE with different late exaggeration settings, as well as DRE with different number of recursions. Fig. 3 shows a comparison between the two different strategies on the MNIST and RNA-Seq datasets. We observed that both late exaggeration and recursive training resulted in similar effect of map optimization by shortening the distance between data points within the same class while prolonging the distance between data points in different classes, both in an unsupervised manner. While Fit-SNE realizes this effect by increasing the repulsive forces explicitly during optimization, our method recalculates the distribution P using an updated representation of the original high-dimensional data, obtained through DNN. The two strategies are not equivalent *per se*, but both lead to a better separation of the

inherent clusters in data. Closer observation shows that the DRE preserved a better visual balance between global and local structures. Within clusters we can better appreciate sub-structures in data, especially in the RNA-Seq data (lower panel b) with a natural hierarchical structure.

5.2 Comparison With Other Embedding Methods

We compared the proposed DEN methods with other reference methods on the MNIST, Fashion MNIST, RNA-Seq and IMDB datasets. For illustrations We included in total 7 methods for visual evaluation: (1) PCA, (2) AE, (3) t-SNE (Barnes-Hut implementation in Python), (4) UMAP, (5) DL Projection [13] (trained by t-SNE), (6) DRE (with t-SNE loss), (7) DRE (with t-SNE and UMAP loss). Quantitative evaluation metrics are reported in Table 1, which compared a more extensive set of embedding methods. The comparison further includes RBM-based ptSNE [14], Fit-SNE [9], DL Projection trained by UMAP [13], and Deep t-SNE (i.e., without recursion).

Fig. 4 shows the results of MNIST and Fashion-MNIST, both in training and testing. For the t-SNE method (column c), no testing results was plotted as there is no learning mechanism. For MNIST, it can be observed that the PCA method (column a) could hardly differentiate the 10 classes, with poor performance in terms of 1NN error and neighborhood hit. The AE method (column b) performed better than PCA in terms of 1NN error, neighborhood hit, and trustworthiness, but the embedding remained visually poor. The Barnes-Hut t-SNE (column c) and UMAP (column d) are the two currently most popular visualization methods, which generated largely improved embeddings over PCA and AE. The DL projection method (column e) method reproduced the results of t-SNE (column c) both in training and testing. However, we noticed a change of cluster distribution (column c and e), due to the different optimization of t-SNE in separate runs. The last two columns show the embedding results from the proposed DRE method, with t-SNE loss (column f) and t-SNE + UMAP loss (column g), respectively. Both exhibited visually improved 2D embeddings. Even without color coding, we can still observe 10 distinct clusters, better separated compared with t-SNE or UMAP. Quantitatively, our proposed DRE method also showed superior 1NN and neighborhood hit value, comparable trustworthiness, continuity and normalized stress compared with those from the t-SNE and UMAP methods. Nonetheless, we noticed that the Shepard goodness was reduced, as the recursive training strategy exaggerated the inter-cluster distances, while Shepard goodness is a metric based on the original data distribution.

The Fashion-MNIST dataset is generally more challenging to cluster or classify than the MNIST dataset. We observed more distinct clusters (e.g., tops, bottoms and shoes) by our method than UMAP or t-SNE, however, the confusion of certain classes (e.g., shirt, pullover and coat) remained similar.

The results of RNA-Seq are shown in the first and second rows of Fig. 5, in two color schemes reflecting the hierarchy of global and local data structure. The RNA-Seq can be divided into three major types: glutamatergic excitatory neurons (GABA), GABAergic inhibitory neurons (Gluta) and non-neuronal cells (Non-neu), color-coded in red, blue and black in the first row. It can be observed that the PCA

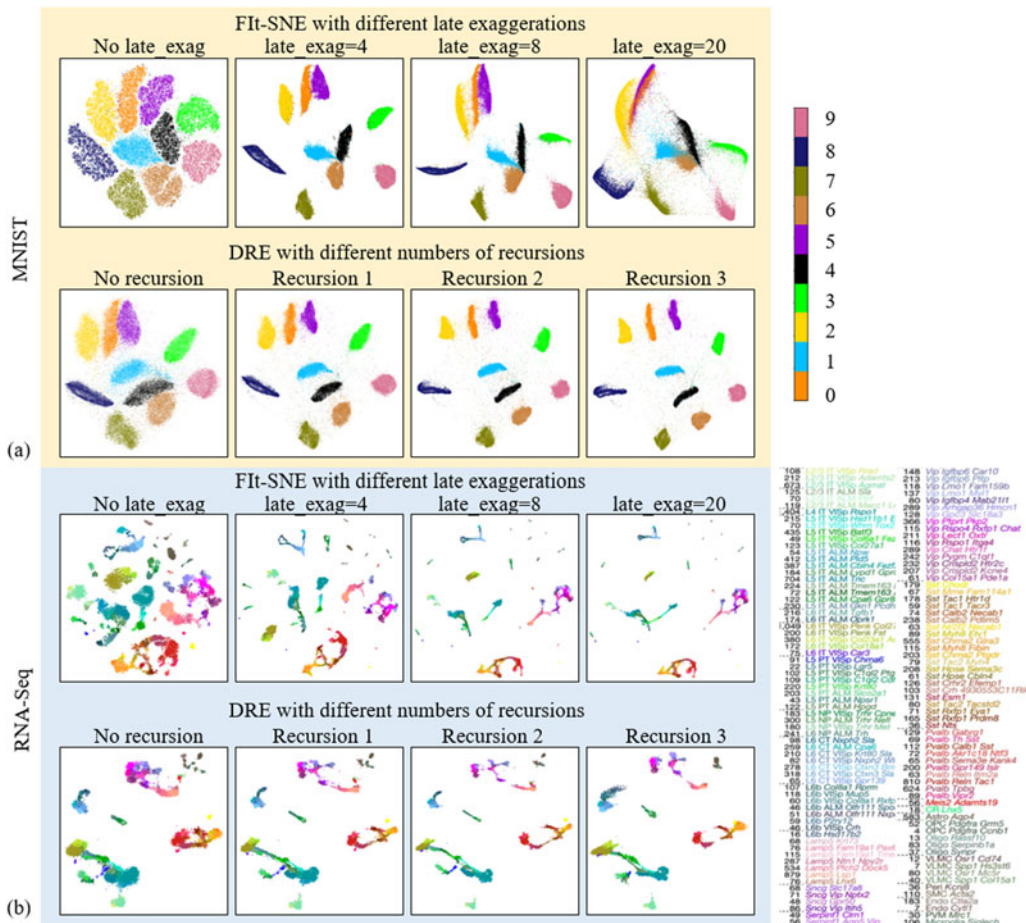


Fig. 3. Comparison of embedding performances of Fit-SNE with different late exaggeration settings and DRE with different recursions. (a) and (b) are results from the MNIST and RNA-Seq dataset, respectively. Different colors indicate different classes. Note that the color is for visualization, while the mapping was unsupervised in both Fit-SNE and DRE.

TABLE 1
Comparison of the Embedding Performance on MNIST, Fashion-MNIST, IMDB, and RNA-Seq Data by Different Methods

Experiments	PCA	Auto-encoder	t-SNE (Barnes-Hut)	Fit-SNE	UMAP	RBM ptSNE	DL projection (t-SNE)	DL projection (UMAP)	Deep t-SNE	DRE t-SNE	DRE t-SNE + UMAP
MNIST	M_{NH}	0.47	0.80	0.93	0.92	0.92	0.81	0.90	0.95	0.93	0.97
	M_t	0.74	0.94	0.98	0.95	0.96	0.92	0.92	0.93	0.94	0.92
	M_C	0.94	0.96	0.98	0.98	0.97	0.97	0.97	0.97	0.97	0.96
	$1 - M_\sigma$	0.42	0.43	0.54	0.54	0.50	0.24	0.53	0.54	0.53	0.52
	M_S	0.50	0.43	0.43	0.39	0.37	0.53	0.32	0.31	0.34	0.29
Fashion-MNIST	M_{NH}	0.53	0.71	0.77	0.77	0.73	0.70	0.71	0.73	0.73	0.77
	M_t	0.91	0.97	0.99	0.99	0.98	0.97	0.96	0.96	0.97	0.97
	M_C	0.98	0.98	0.99	0.99	0.99	0.99	0.98	0.99	0.99	0.98
	$1 - M_\sigma$	0.65	0.50	0.61	0.64	0.56	0.63	0.61	0.65	0.62	0.62
	M_S	0.88	0.72	0.58	0.65	0.58	0.76	0.50	0.63	0.63	0.60
IMDB	M_{NH}	0.59	0.59	0.65	0.62	0.65	0.59	0.61	0.62	0.59	0.76
	M_t	0.69	0.62	0.89	0.72	0.79	0.68	0.63	0.67	0.73	0.74
	M_C	0.66	0.63	0.84	0.80	0.85	0.70	0.75	0.76	0.79	0.79
	$1 - M_\sigma$	0.26	0.23	0.34	0.38	0.32	0.23	0.24	0.28	0.37	0.44
	M_S	0.16	0.08	0.29	0.27	0.26	0.37	0.34	0.29	0.39	0.38
RNA-Seq	M_{NH}	0.91	0.27	0.99	1.00	1.00	0.96	0.98	0.99	1.00	1.00
	M_t	0.85	0.82	1.00	0.99	0.99	0.98	0.98	0.98	0.99	0.99
	M_C	0.93	0.91	1.00	1.00	0.99	0.99	0.99	0.98	0.99	0.99
	$1 - M_\sigma$	0.53	0.51	0.59	0.61	0.61	0.56	0.55	0.61	0.62	0.61
	M_S	0.88	0.44	0.40	0.57	0.55	0.79	0.30	0.56	0.64	0.62

M_{NH} , M_t , M_C , M_σ , M_S indicate the values of neighborhood hit, trustworthiness, continuity, Shepard goodness, and normalized stress, respectively.

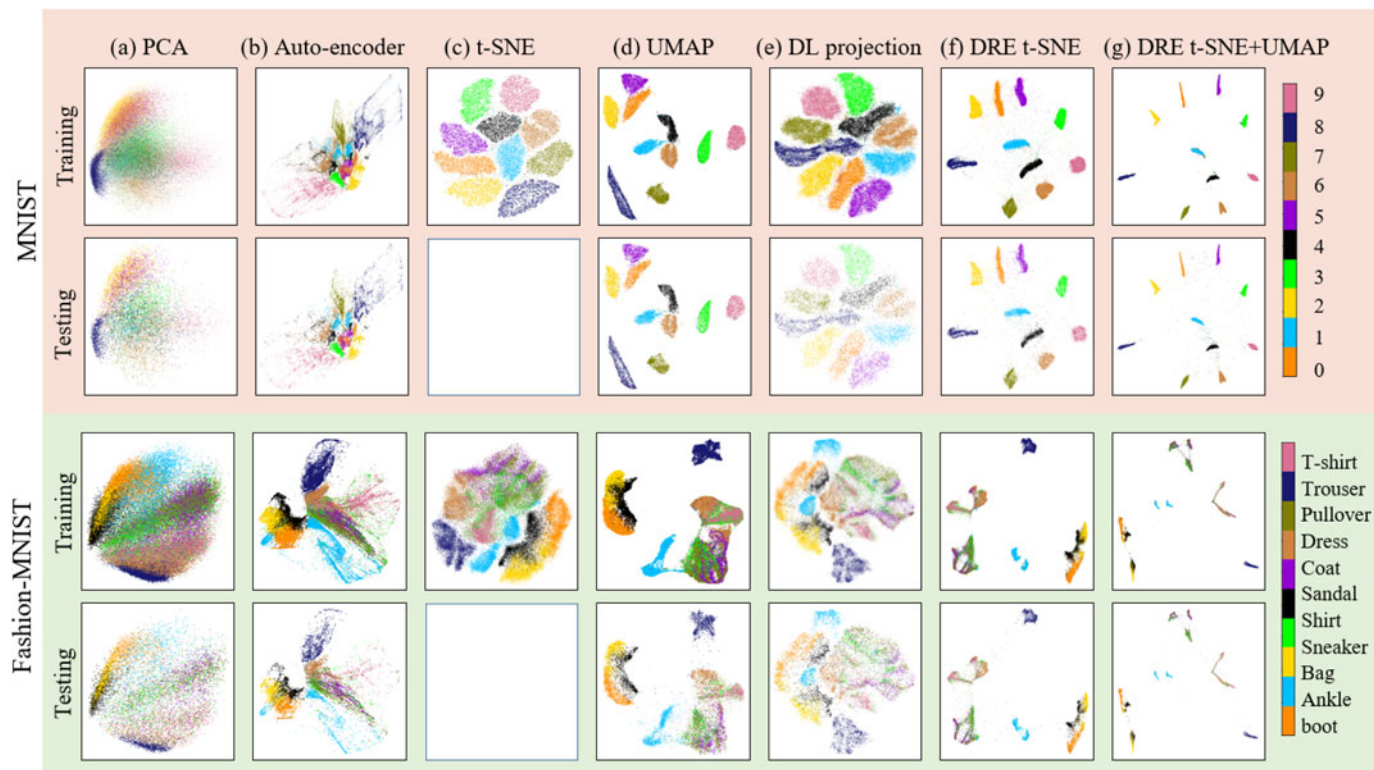


Fig. 4. Comparison of embedding performance on the MNIST and Fashion-MNIST datasets by 7 methods: (a) PCA, (b) AE, (c) t-SNE (Barnes-Hut implementation in Python), (d) UMAP, (e) DL Projection (trained by t-SNE), (f) DRE (with t-SNE loss), and (g) DRE (with two-stage t-SNE + UMAP loss). The results of both training and testing datasets are shown.

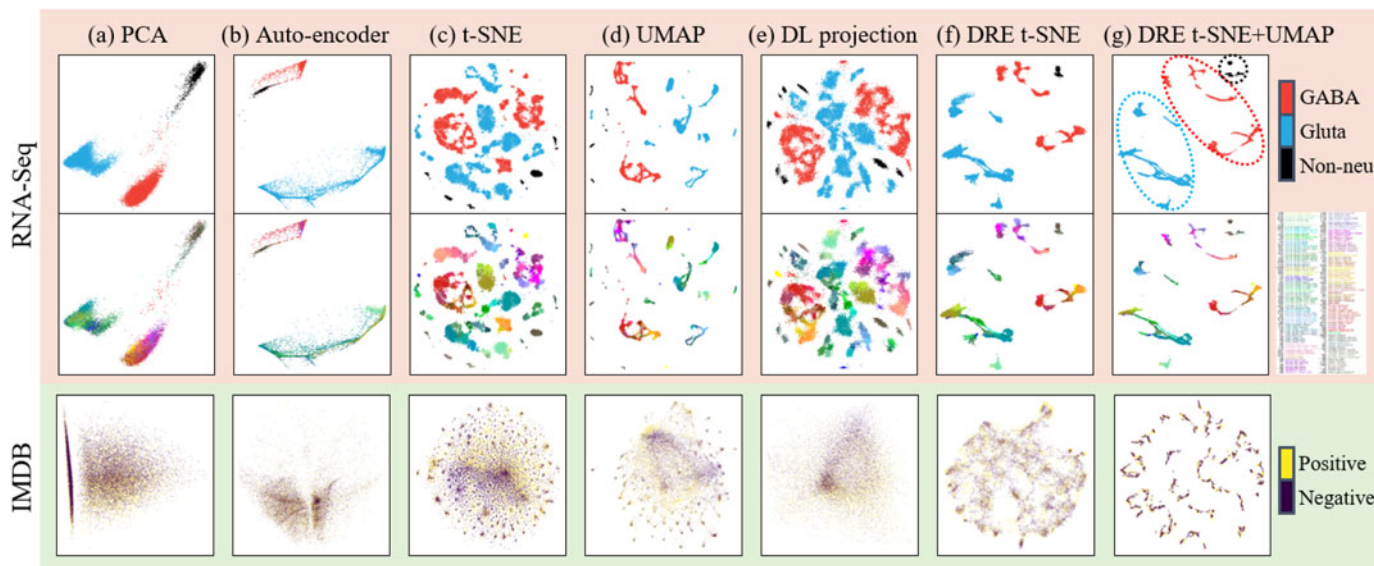


Fig. 5. Comparison of embedding performance on the RNA-Seq and IMDB datasets by 7 methods: (a) PCA, (b) AE, (c) t-SNE (Barnes-Hut implementation in Python), (d) UMAP, (e) DL Projection (trained by t-SNE), (f) DRE (with t-SNE loss), and (g) DRE (with two-stage t-SNE + UMAP loss). For readability of figures only the results of training datasets are shown. For RNA-seq, we showed two color schemes to highlight the hierarchy of global and local data structure.

method could differentiate the three major types of cells very well, presenting the global structure in data. In each major cell type, there are also a large number of subtypes, in total 133 (details can be find in [31]). The second row of Fig. 5 color-codes the cell subtypes in a refined color scale. As the local data structures are nonlinear, PCA can no longer differentiate these subtypes, in contrast to nonlinear methods such as t-SNE and UMAP, which exhibited clustering according to

cell subtype. However, we observed a loss of global structure in the results of t-SNE and UMAP, with the global distribution of three major types no longer consistent, i.e., scattered and mingled. The proposal DRE method (column f and g) showed a better balance between global and local data structures. The three major types are close as highlighted by the dotted ellipses in Fig. 5 column g, whereas the subtypes can also be appreciated.

TABLE 2
Comparison of the 1-Nearest Neighbor Classification Accuracy (%) of the 1 Million InfiMNIST Data, Embedded by Learning From Training Data of Different Sizes: 0.05, 0.1, and 0.3 Million

Training data size	PCA	Auto-encoder	UMAP	RBM ptSNE	DL projection (t-SNE)	DL projection (UMAP)	Deep t-SNE	DRE t-SNE	DRE t-SNE + UMAP
0.05M	35.80	64.14	87.00	64.84	83.68	86.52	89.96	93.33	93.52
0.1M	39.14	68.35	90.25	71.14	87.61	90.23	91.64	95.48	95.58
0.3M	52.36	78.50	93.60	82.38	90.39	93.07	94.76	97.14	97.17

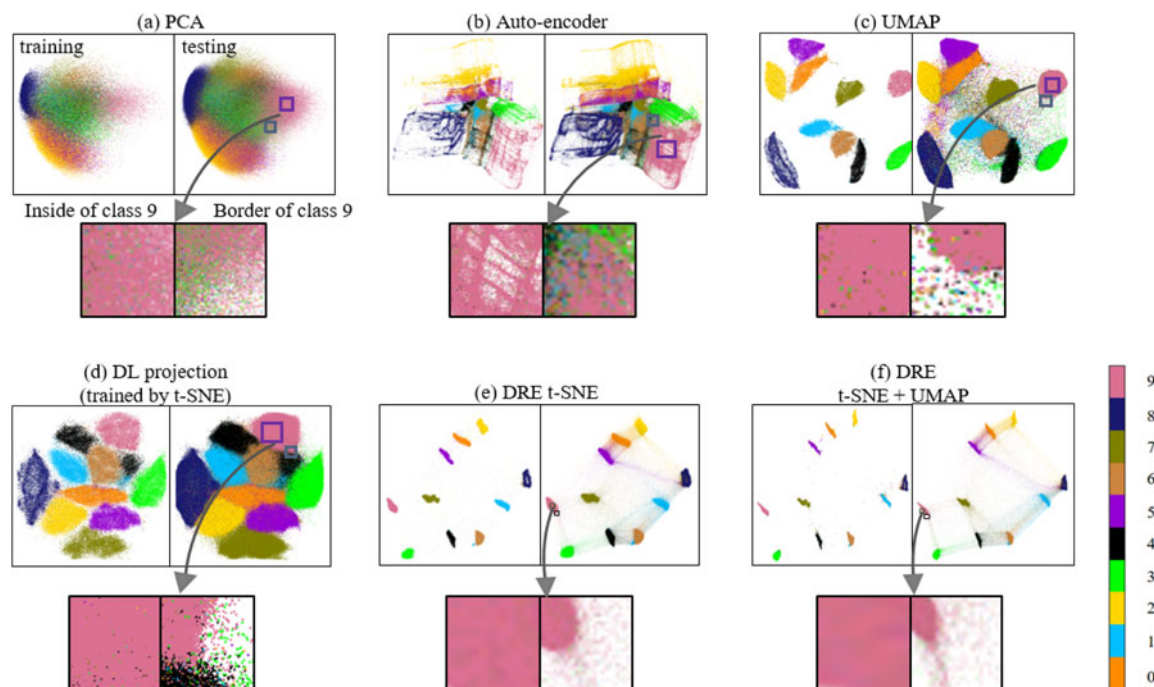


Fig. 6. Embedding 1 million InfiMNIST data by learning from 0.1 million data. The embeddings of both training and testing data are shown for (a) PCA, (b) AE, (c) UMAP, (d) DL projection trained with t-SNE, (e) DRE with t-SNE loss, and (f) DRE with t-SNE + UMAP loss. Local neighborhoods inside and at the border of class 9 are zoomed in for a closer look of embedding quality.

The results of IMDB dataset is shown in the third row of Fig. 5. The original data are only classified into positive and negative comments, however, the actual sentiments are much more complicated and subtle. We observed that all embedding methods mixed up the two categories of data, but nonlinear methods (column c-g) showed sub-clusters, which might reflect subtle semantics in comments. In Table 1, we note that the two DRE methods resulted in better quantitative metrics compared to other methods, with the highest neighborhood hit, trustworthiness and Shepard goodness, and lowest normalized stress. However, further evaluation demands more detailed labels of sub-clusters.

5.3 Scalability and Generalization

To evaluate the scalability of the embedding methods, we generated huge training and testing datasets from InfiMNIST [33], and evaluated if the proposed DRE could scale to such extremely large dataset. We trained DRE with 0.05 million, 0.1 million and 0.3 million InfiMNIST data, and embedded an independently generated 1 million InfiMNIST testing dataset. Then we measured the 1NN classification accuracy on the 2D embedding, as reported in Table 2. The proposed DRE method out-performed all other embedding methods, with a 93.52%, 95.58% and 97.17% testing accuracy when

trained on 0.05 million, 0.1 million and 0.3 million data, respectively. Fig. 6 shows the embedding results of the 1 million testing data (trained on 0.1 million data), for 6 methods in comparison: (a) PCA, (b) AE, (c) UMAP, (d) DL projection trained with t-SNE, (e) DRE with t-SNE loss, and (f) DRE with t-SNE + UMAP loss, respectively. The embedding of both the training and testing datasets are shown. On the testing results, we zoomed in the local neighborhood at the inside and border of class 9 to inspect embedding details. As the learned embedding was extrapolated from 0.1 million training data onto 1.0 million testing data, we observed scattering of classes at the border. Overall, the proposed DRE methods resulted in the cleanest neighborhood compared to other embedding methods, which explains the high classification accuracy in Table 2.

In a second experiment, we used a training set of 10K data points to train the embedding methods, and tested them on an increasing testing set of 10K, 30K, and 60K. The 1NN classification accuracy is reported in Table 3. It can be seen that all embedding methods showed stable generalization, suggesting the manifold of MNIST can be adequately learned and represented in a dimensionality as low as 2. In both Tables 2 and 3, we observed that the two DRE methods had the best quantitative performance, followed by deep t-SNE

TABLE 3

Comparison of the 1-Nearest Neighbor Classification Accuracy (%) on InfiMNIST Testing Datasets of Different Sizes: 10K, 30K, and 60K

Testing data size	PCA	Auto-encoder	UMAP	RBM ptSNE	DL projection (t-SNE)	DL projection (UMAP)	Deep t-SNE	DRE t-SNE	DRE t-SNE + UMAP
10K	37.32	67.43	78.98	68.58	82.31	85.08	87.09	89.83	90.07
30K	37.85	66.89	87.76	67.79	81.66	84.45	87.01	90.09	90.08
60K	38.10	67.47	88.01	68.47	82.06	84.75	87.26	90.49	90.67

The training data size is 10K.

TABLE 4

Comparison of the Runtime of Different Embedding Methods (MNIST: 60,000 Training and 10,000 Testing)

Type	Parametric DR methods								Non-parametric DR Methods				
Method	PCA	Auto-encoder	RBM ptSNE	UMAP	DL projection (t-SNE)	DL projection (UMAP)	Deep t-SNE	DR t-SNE	DRE	Eigenmaps	BH t-SNE	Open t-SNE	Fit-SNE
Runtime Train (second)	0.029	31.78	7455.39	83.88	2122.01	1085.72	250.81	339.51	759.88	42026.70	5730.79	1149.47	304.17
Test	0.67	0.46	0.49	13.68	0.22	0.40	0.84	0.82	0.73	/	/	/	/

and the DL projection method. The results can be intuitively explained as follows: both deep t-SNE and DL projection methods are approaching the limit of t-SNE, while the DRE methods are pushing the limit further by recursive training. Comparing the two DRE methods using t-SNE and composite t-SNE+UMAP losses, however, we noticed little difference in 1NN classification performance or other quantitative measures, although the latter produces visually better separation and clearer margin.

5.4 Runtime

We tested the runtime of a large range of existing embedding methods in literature, categorized into parametric and non-parametric DR methods. The runtime are reported in Table 4. The PCA, Eigenmaps and BH t-SNE methods were implemented single threaded on CPU. UMAP used numba’s parallel implementation to do multithreaded processing with multiple cores on CPU. The DL projection methods adopted the MulticoreTSNE with 8 NVIDIA Tesla V100 GPU to parallelly calculate the network references, and fitted the model using 1 GPU. The t-SNE and Fit-SNE method utilized the multithreading and C/C++ compiler on CPU. The proposed methods single threaded all computations, and used 1 GPU to train the network. As shown in Table 4, the runtime of the proposed Deep t-SNE, DR t-SNE and DRE increased gradually with more recursions. Compared with non-parametric methods, the DRE method needed less runtime than the publicly available Open t-SNE [36], but more runtime than Fit-SNE [9]. Among the parametric methods, the proposed methods were faster than the RBM-based ptSNE. Although our methods are slower than AE and UMAP during training, they are very fast when applied on new data points (second row “test”).

6 DISCUSSION

In this work, we propose to integrate modern DNNs and mathematics principles to embed high-dimensional data, and we name this framework “deep embedding network” - DEN. DEN is a generic framework, which can be designed

in a flexible manner in terms of NN architectures, training strategies, and loss functions. Based on DEN, we showed how to design Deep t-SNE and Deep UMAP, and further introduced a recursive training strategy to boost the embedding performance over the original t-SNE or UMAP. We exemplified the flexibility of DEN by the recursive training with two-stage loss functions combining the popular t-SNE and UMAP. We tested the proposed DRE on a variety of public datasets, and evaluated its performance both visually and quantitatively, against a comprehensive set of classical and state-of-the-art embedding methods.

While inspired by t-SNE, the proposed DRE has made the following improvement: (1) the DRE method is parametric and can easily embed new, out-of-sample data; (2) the proposed recursive training strategy can boost the embedding performance over the original t-SNE; (3) empirical experiments show that the DRE method can better preserve the global data structure simultaneously; (4) it can integrate any mathematics rules beyond t-SNE.

Both DL projection method and the proposed DRE method can map large, out-of-sample datasets. This is an inherent advantage of learning-based methods. However, the training mechanisms are very different: DL Projection solves the DR problem by learning from the results of t-SNE or UMAP, i.e., the network is taught to generate an output similar to that of the reference methods. In contrast, the proposed DRE method provides the network a principled loss function, and then let the network explore the embedding itself.

There are a number of hyperparameters in the proposed method. The mini-batch used in DNN is critical for scalability and has an influence on the embedding results: a small mini-batch cannot fully sample the data distribution, while a big mini-batch demands too much memory as the square matrix of P is dependent on mini-batch size. We empirically selected 2500, but it can also be set larger if memory allows. Other network settings only had minor influence on the final performance according to our experiments, given the capability of DNNs to optimize highly complex loss functions. Another hyperparameter is the number of the recursions, for which we can use a simple rule of thumb: if we

are keen in visualizing the global clustering of data, we can increase the number of recursions (e.g., 3) and obtain better separated clusters in a global view; otherwise 1 or 2 recursions are sufficient as our experiments indicated. We argue that the improved global structure preservation by DRE may arise from the use of latent representations, which better capture the global structure in data (for example, PCA is a linear latent representation). Meanwhile, the local structures are captured by DRE as it is initiated from the t-SNE embedding rule, which emphasizes local neighborhood.

There are a few practical weaknesses of the proposed methods. First, our method obtains the parametric mapping by learning from a dataset, so it is more suitable for big data datasets than for small datasets. Second, the proposed methods still need more runtime than PCA, AE and UMAP. Nevertheless, we expect that the runtime can be reduced by multi-threaded implementation.

7 CONCLUSION

In this paper, we introduced a generic DEN framework, which is flexible in its NN architectures, training strategies, and loss functions. Based on DEN, we proposed a novel DRE strategy to further boost the embedding performance. The proposed framework can combine modern DNNs and any effective embedding rules such as those from t-SNE and UMAP. The proposed DRE method can map out-of-sample data and scale to extremely large datasets. Experiments on a range of datasets demonstrated its improved embedding performance in terms of local and global structure preservation, compared with other state-of-the-art embedding methods.

REFERENCES

- [1] B. Wang, J. Zhu, E. Pierson, D. Ramazzotti, and S. Batzoglou, "Visualization and analysis of single-cell RNA-seq data by kernel-based similarity learning," *Nat. Methods*, vol. 14, no. 4, pp. 414–416, Mar. 2017.
- [2] Y. Bengio, "Deep learning of representations: Looking forward," in *Proc. 1st Int. Conf. Statist. Lang. Speech Process.*, 2013, pp. 1–37.
- [3] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2014.
- [4] S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci, "Visualizing high-dimensional data: Advances in the past decade," *IEEE Trans. Vis. Comput. Graphics*, vol. 23, no. 3, pp. 1249–1268, Mar. 2017.
- [5] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, no. 86, pp. 2579–2605, 2008. [Online]. Available: <http://jmlr.org/papers/v9/vandermaaten08a.html>
- [6] D. Kobak and P. Berens, "The art of using t-SNE for single-cell transcriptomics," *bioRxiv*, 2019. [Online]. Available: <https://www.biorxiv.org/content/early/2019/05/20/453449>
- [7] L. van der Maaten, "Accelerating t-SNE using tree-based algorithms," *J. Mach. Learn. Res.*, vol. 15, no. 93, pp. 3221–3245, 2014. [Online]. Available: <http://jmlr.org/papers/v15/vandermaaten14a.html>
- [8] N. Pezzotti, B. P. F. Lelieveldt, L. van der Maaten, T. Hollt, E. Eiseemann, and A. Vilanova, "Approximated and user steerable tSNE for progressive visual analytics," *IEEE Trans. Vis. Comput. Graphics*, vol. 23, no. 7, pp. 1739–1752, Jul. 2017.
- [9] G. C. Linderman, M. Rachh, J. G. Hoskins, S. Steinerberger, and Y. Kluger, "Fast interpolation-based t-SNE for improved visualization of single-cell RNA-seq data," *Nat. Methods*, vol. 16, no. 3, pp. 243–245, Feb. 2019.
- [10] L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform manifold approximation and projection," *J. Open Source Softw.*, vol. 3, no. 29, p. 861, 2018.
- [11] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006. [Online]. Available: <http://arxiv.org/abs/1511.05644>
- [12] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey, "Adversarial autoencoders," in *Proc. Int. Conf. Learn. Representations*, 2016. [Online]. Available: <http://arxiv.org/abs/1511.05644>
- [13] M. Espadoto, N. S. T. Hirata, and A. C. Telea, "Deep learning multidimensional projections," *Inf. Vis.*, vol. 19, no. 3, pp. 247–269, May 2020.
- [14] L. van der Maaten, "Learning a parametric embedding by preserving local structure," in *Proc. 12th Int. Conf. Artif. Intell. Statist.*, 2009, pp. 384–391. [Online]. Available: <http://proceedings.mlr.press/v5/maaten09a.html>
- [15] I. Jolliffe, "Principal component analysis: a review and recent developments," *Philos. Trans. Royal Soc. A*, vol. 374, 2016, doi: 10.1098/rsta.2015.0202.
- [16] A. Buja, D. F. Swayne, M. L. Littman, N. Dean, H. Hofmann, and L. Chen, "Data visualization with multidimensional scaling," *J. Comput. Graphical Statist.*, vol. 17, no. 2, pp. 444–472, Jun. 2008.
- [17] M. Balasubramanian, E. L. Schwartz, J. B. Tenenbaum, V. de Silva, and J. C. Langford, "The Isomap algorithm and topological stability," *Science*, vol. 295, no. 5552, pp. 7–7, 2002.
- [18] S. T. Roweis, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, no. 5500, pp. 2323–2326, Dec. 2000.
- [19] G. Hinton and S. T. Roweis, "Stochastic neighbor embedding," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2002, pp. 833–840.
- [20] D. M. Chan, R. Rao, F. Huang, and J. F. Canny, "T-SNE-CUDA: GPU-accelerated t-SNE and its applications to modern data," in *Proc. 30th Int. Symp. Comput. Architect. High Perform. Comput.*, 2018, pp. 330–338.
- [21] N. Pezzotti et al., "GPGPU linear complexity t-SNE optimization," *IEEE Trans. Vis. Comput. Graphics*, vol. 26, no. 1, pp. 1172–1181, Jan. 2020.
- [22] N. Pezzotti, T. Höllt, B. Lelieveldt, E. Eiseemann, and A. Vilanova, "Hierarchical stochastic neighbor embedding," *Comput. Graph. Forum*, vol. 35, no. 3, pp. 21–30, Jun. 2016.
- [23] D. Kobak and G. C. Linderman, "UMAP does not preserve global structure any better than t-SNE when using the same initialization," *BioRxiv*, 2019. [Online]. Available: <https://www.biorxiv.org/content/early/2019/12/19/2019.12.19.877522>
- [24] J. Xie, R. B. Girshick, and A. Farhadi, "Unsupervised deep embedding for clustering analysis," in *Proc. 33rd Int. Conf. Mach. Learn., ser. Mach. Learn. Research*, vol. 48, M. F. Balcan and K. Q. Weinberger, Eds., New York, USA: PMLR, Jun. 2016, pp. 478–487. [Online]. Available: <https://proceedings.mlr.press/v48/xieb16.html>
- [25] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson, "Understanding neural networks through deep visualization," *CoRR*, vol. abs/1506.06579, 2015. [Online]. Available: <http://arxiv.org/abs/1506.06579>
- [26] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 336–359, Oct. 2019.
- [27] A. T. Elthakeb, P. Pilligundla, F. Mireshghallah, A. Cloninger, and H. Esmailzadeh, "Divide and conquer: Leveraging intermediate feature representations for quantized training of neural networks," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 2880–2891. [Online]. Available: <http://proceedings.mlr.press/v119/elthakeb20a.html>
- [28] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 694–711.
- [29] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [30] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms," *CoRR*, vol. abs/1708.07747, 2017. [Online]. Available: <http://arxiv.org/abs/1708.07747>
- [31] B. Tasic, Z. Yao, L. T. Graybuck, K. A. Smith, and T. N. Nguyen, "Shared and distinct transcriptomic cell types across neocortical areas," *Nature*, vol. 563, no. 7729, pp. 72–78, Oct. 2018.

- [32] A. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts, "Learning word vectors for sentiment analysis," in *Proc. 49th Annu. Meeting Assoc. Comput. Linguistics, Hum. Lang. Technol.*, 2011, pp. 142–150.
- [33] P. Simard, B. Victorri, Y. Le Cun, and J. Denker, "Tangent prop: A formalism for specifying selected invariances in an adaptive network," in *Proc. 4th Int. Conf. Neural Inf. Process. Syst.*, 1991, pp. 895–903.
- [34] M. Espadoto, R. M. Martins, A. Kerren, N. S. T. Hirata, and A. C. Telea, "Toward a quantitative survey of dimension reduction techniques," *IEEE Trans. Vis. Comput. Graphics*, vol. 27, no. 3, pp. 2153–2173, Mar. 2021.
- [35] L. G. Nonato and M. Aupetit, "Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment," *IEEE Trans. Vis. Comput. Graphics*, vol. 25, no. 8, pp. 2650–2673, Aug. 2019.
- [36] P. G. Poličar, M. Stražar, and B. Zupan, "openTSNE: A modular python library for t-SNE dimensionality reduction and embedding," *BioRxiv*, 2019, doi: [10.1101/731877](https://doi.org/10.1101/731877).



Zixia Zhou received the PhD degree in biomedical engineering from Fudan University, Shanghai, China, in 2021. She is currently a postdoctoral researcher with the Department of Radiation Oncology, Stanford University. She is interested in improving medical imaging quality in cost-effective manners with artificial intelligence. Broadly, her research focuses on medical image reconstruction and high-dimensional data visualization.



Xinrui Zu received the BSc degree in theoretical and applied mechanics from Fudan University, Shanghai, China, in 2019. He is currently working toward the MSc degree in electrical engineering, mathematics and computer science at the University of Twente, Enschede, The Netherlands. His research mainly focuses on computer vision.



Yuanyuan Wang (Senior Member, IEEE) received the BSc, MSc, and PhD degrees in electronic engineering from Fudan University, Shanghai, China, in 1990, 1992, and 1994, respectively. From 1994 to 1996, he was a postdoctoral research fellow with the School of Electronic Engineering and Computer Science, University of Wales, Bangor, U.K. In 1996, he joined the Department of Electronic Engineering, Fudan University, as an associate professor, and was then promoted to a full professor in 1998. He is currently the director of Biomedical Engineering Center, Fudan University. He has authored or coauthored six books and 500 research papers. His research interests include medical ultrasound techniques and medical image processing.



Boudewijn P. F. Lelieveldt is professor of biomedical imaging with the Department of Radiology, Leiden University Medical Center, Leiden, The Netherlands, where he is heading the Division of Image Processing (LKEB). He is also appointed as Medical Delta professor with the Pattern Recognition and Bioinformatics Group, Delft University of Technology in the context of the Medical Delta consortium. His main research interest include data analytics for spatially resolved omics data, with applications in clinical imaging and biology. He is an affiliated investigator with the Allen Institute for Brain Science in Seattle, with a focus on 1) the development of fast and interactive data analysis methods for spatial transcriptomics and single cell data (BrainScope, Cytosplore), and 2) deploying the Allen Brain Atlases in the LUMC Translational Neuroscience Research Program. He serves as a member of the editorial board of the *Medical Image Analysis*, and is an associate editor of the *IEEE Transactions on Medical Imaging*.



Qian Tao received the BSc degree in electrical engineering and the MSc degree in biomedical engineering, both from Fudan University, Shanghai, China, and the PhD degree in biometrics authentication from the University of Twente, Enschede, The Netherlands. She is assistant professor of medical imaging with the Department of Imaging Physics, Delft University of Technology, The Netherlands. After PhD, she joined the Division of Image Processing, Department of Radiology, Leiden University Medical Center, and her research focused on cardiac image analysis (MRI and CT) and its applications in cardiac interventions. Her current research interests include medical image and data analysis, and their applications in clinical and biological research.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.**