

# Uniform Manifold Approximation with Two-phase Optimization

Hyeon Jeon<sup>1 \* †</sup> Hyung-Kwon Ko<sup>2 \* ‡</sup> Soohyun Lee<sup>1 §</sup> Jaemin Jo<sup>3 ¶</sup> Jinwook Seo<sup>1 ||</sup>

<sup>1</sup> Seoul National University

<sup>2</sup> NAVER Webtoon Corp.

<sup>3</sup> Sungkyunkwan University

## ABSTRACT

We introduce Uniform Manifold Approximation with Two-phase Optimization (UMATO), a dimensionality reduction (DR) technique that improves UMAP to capture the global structure of high-dimensional data more accurately. In UMATO, optimization is divided into two phases so that the resulting embeddings can depict the global structure reliably while preserving the local structure with sufficient accuracy. In the first phase, hub points are identified and projected to construct a skeletal layout for the global structure. In the second phase, the remaining points are added to the embedding preserving the regional characteristics of local areas. Through quantitative experiments, we found that UMATO (1) outperformed widely used DR techniques in preserving the global structure while (2) producing competitive accuracy in representing the local structure. We also verified that UMATO is preferable in terms of robustness over diverse initialization methods, numbers of epochs, and subsampling techniques.

**Index Terms:** Human-centered computing—Visualization—Visualization techniques; Computing methodologies—Machine learning—Machine learning algorithms

## 1 INTRODUCTION

Dimensionality reduction (DR) is one of the most useful tools for exploring high-dimensional (HD) data in visual analytics [5, 10, 15]. In various domains, DR provides an effective means of understanding HD data by enabling visual inspection of HD data [16, 23, 33]. Especially, nonlinear DR techniques such as  $t$ -SNE [24], NeRV [37], and UMAP [26] made huge success in capturing complex local manifolds [21].

However, despite the success, nonlinear DR techniques often work poor in preserving the global structure of data [9, 18]. For example, in the case of  $t$ -SNE, inaccuracy in preserving the global structure comes from the fact that its loss function, Kullback-Leibler (KL) divergence, assigns too little penalty for the points that are distant in the original space and stay close in the embedding space [26]. UMAP avoids such a problem by employing the cross-entropy function as a loss function. However, it still barely captures the global structure as it optimizes the layout based on a  $k$ -nearest neighbor ( $k$ NN) graph. UMAP's approximation for fast computation (Sect. 2.1) also leads to inaccuracy; for example, a recent study [17] showed that UMAP does not outperform  $t$ -SNE in preserving the global structure given the same initial layout. One way to alleviate this problem is to not use approximation, which would increase the computational cost greatly.

To overcome the aforementioned limitations, we present Uniform Manifold Approximation with Two-phase Optimization (UMATO), a novel DR technique that enhances UMAP to better preserve the global structure. In UMATO, optimization is first run for a small number of points that represent the data (i.e., hub points). As finding the optimal embedding for a small number of points is relatively easy and robust, this makes approximation unnecessary and thus leads to the better preservation of the global structure. Moreover, this procedure makes the embedding more stable and less sensitive to initialization methods. After capturing the overall skeleton of the HD structure with the hub points, we gradually append the rest of the points in subsequent phases. Although the same approximation techniques as UMAP are used for these points, the embedding becomes more accurate in preserving the global structure because we use already embedded hub points as anchors.

We quantitatively compared UMATO against five widely used DR techniques using four datasets. The embeddings were then analyzed with global and local quality metrics such as KL divergence and Trustworthiness & Continuity [36], respectively. As a result, we found that UMATO outperformed the competitors in preserving the global structure while showing competitive performance in preserving the local structure. We also found that UMATO is more stable against sampling fluctuations and changes in the number of epochs, and is less sensitive to the initialization methods, compared to the competitors (Appendix D, E, and F).

## 2 BACKGROUND AND RELATED WORKS

### 2.1 UMAP

Because UMATO shares the overall pipeline of UMAP, we provide a brief introduction to UMAP. Although UMAP is grounded in a sophisticated mathematical foundation, its computation can be simply divided into two steps, graph construction and layout optimization. **Graph Construction.** UMAP starts by generating a weighted  $k$ NN graph; given  $k$  (the number of NN to consider) and a distance metric  $d : X \times X \rightarrow [0, \infty)$ , UMAP first computes  $\mathcal{N}_i$ , the  $k$ NN of  $x_i$  with respect to  $d$ . Then, two parameters,  $\rho_i$  and  $\sigma_i$ , are computed for each data point  $x_i$  to identify its local metric space.  $\rho_i$  is a nonzero distance from  $x_i$  to its nearest neighbor:

$$\rho_i = \min_{j \in \mathcal{N}_i} \{d(x_i, x_j) \mid d(x_i, x_j) > 0\}. \quad (1)$$

Using binary search, UMAP finds  $\sigma_i$  that satisfies:

$$\sum_{j \in \mathcal{N}_i} \exp(-\max(0, d(x_i, x_j) - \rho_i)/\sigma_i) = \log_2(k). \quad (2)$$

Next, UMAP computes:

$$v_{ji} = \exp(-\max(0, d(x_i, x_j) - \rho_i)/\sigma_i), \quad (3)$$

which is the weight of the edge from a point  $x_i$  to another point  $x_j$ . A combined weight of an edge is then defined as  $v_{ij} = v_{ji} + v_{ij} - v_{ji} \cdot v_{ij}$ . In the embedded space, the similarity between two points  $y_i$  and  $y_j$  is defined as  $w_{ij} = (1 + a\|y_i - y_j\|_2^{2b})^{-1}$ , where  $a$  and  $b$  are positive constants defined by users. Setting both  $a$  and  $b$  to 1 is identical to using Student's  $t$ -distribution to measure the similarity.

\*equal contribution

†e-mail: [hj@hcl.snu.ac.kr](mailto:hj@hcl.snu.ac.kr)

‡e-mail: [hkko@hcl.snu.ac.kr](mailto:hkko@hcl.snu.ac.kr)

§e-mail: [shlee@hcl.snu.ac.kr](mailto:shlee@hcl.snu.ac.kr)

¶e-mail: [jmjo@skku.edu](mailto:jmjo@skku.edu)

||e-mail: [jseo@snu.ac.kr](mailto:jseo@snu.ac.kr)

**Layout Optimization.** The goal of layout optimization is to find the  $y_i$  that minimizes the difference (or loss) between  $v_{ij}$  and  $w_{ij}$ . Unlike  $t$ -SNE, UMAP employs cross-entropy:

$$CE = \sum_{i \neq j} [v_{ij} \cdot \log(v_{ij}/w_{ij}) - (1 - v_{ij}) \cdot \log((1 - v_{ij})/(1 - w_{ij}))], \quad (4)$$

between  $v_{ij}$  and  $w_{ij}$  as the loss function. UMAP initializes  $y_i$  through spectral embedding [2] and iteratively optimizes its position to minimize  $CE$ . Given the output weight  $w_{ij}$  as  $1/(1 + ad_{ij}^{2b})$ , where  $d_{ij}^{2b} = \|y_i - y_j\|_2^{2b}$ , the attractive gradient is:

$$\frac{CE}{y_i}^+ = \frac{-2abd_{ij}^{2(b-1)}}{1 + ad_{ij}^{2b}} v_{ij}(y_i - y_j), \quad (5)$$

and the repulsive gradient is:

$$\frac{CE}{y_i}^- = \frac{2b}{(\epsilon + d_{ij}^2)(1 + ad_{ij}^{2b})} (1 - v_{ij})(y_i - y_j). \quad (6)$$

$\epsilon$  is a small constant added to prevent division by zero and  $d_{ij}$  is the Euclidean distance between  $y_i$  and  $y_j$ . For efficient optimization, UMAP leverages the negative sampling technique [27, 33, 34]. The technique first chooses a target edge  $(i, j)$  and  $M$  negative sample points. Then,  $i$  and  $j$  are used to compute attractive forces, while negative samples are used to compute repulsive forces; for each epoch, positions of  $i$ ,  $j$ , and negative samples are updated. Considering negative sampling, the modified objective function is:

$$O = \sum_{(i,j) \in E} v_{ij}(\log(w_{ij}) + \sum_{k=1}^M E_{j_k \sim P_n(j)} \gamma \log(1 - w_{ij_k})). \quad (7)$$

Here,  $v_{ij}$  and  $w_{ij}$  are the similarities in the high and low-dimensional spaces, respectively, and  $\gamma$  is a weight constant to apply to negative samples.  $E_{j_k \sim P_n(j)}$  means that  $j_k$  is sampled from a noisy distribution  $P_n(j) \propto \deg_j^{3/4}$  [27], where  $\deg_j$  denotes the degree of point  $j$ .

## 2.2 Dimensionality Reduction Techniques for Preserving Global Structure

Preserving global structures in DR has long been considered an important research topic [9, 14, 17, 18, 28]. One way is to design a loss function or an optimization pipeline that targets global aspects of data. For instance, Isomap [35] preserves global structure by approximating its geodesic distances in the embedding. Another example is TopoAE [28], a deep-learning approach that uses a generative model. To make the latent space resemble the HD space, it appends a topological loss to the original reconstruction loss of autoencoders [12]. However, these techniques leverage a single-step optimization; thus, their focus on global structure reduces their attention to local structure. In contrast, UMATO focuses on both global and local structures by dividing the optimization into two phases.

Another common scheme to capture global structure is to utilize sample points to better model the original space; these points are usually called hubs, landmarks, or anchors. Silva et al. [32] proposed L-Isomap, a landmark version of classical multidimensional scaling (MDS), to alleviate its computation cost. However, in L-Isomap, landmarks are chosen randomly without considering their importance. HSNE [30] and HUMAP [25] instead chose landmarks based on  $k$ NN graph, which is similar to UMATO. However, in HSNE and HUMAP, landmarks' main role is constructing the early stage of hierarchical embeddings in which users can interactively drill down. Instead, UMATO utilizes hubs to capture the global structure and work as global anchors for the layout of non-hub points.

The most similar work to ours is At-SNE [9], which optimizes the anchor points and remaining points with two different loss functions.

Nonetheless, as the anchors wander during the optimization and the loss function (KL divergence) does not care about distant points, it barely captures the global structure (Table 1). UMATO avoids this problem by separating hubs and non-hub points; the hubs take their position in the first phase and barely move but guide other points in the second phase so that the global structure can be preserved robustly. Applying cross-entropy as a loss function also helps in preserving both structures [26].

## 3 UMATO

We present UMATO, which splits UMAP's optimization into two phases to preserve global structure better while maintaining the capability to represent local structure. For ease of understanding, we illustrated the UMATO pipeline in Fig. 1, presented the pseudo code in Appendix A, and publicized the source code [1].

### 3.1 Points Classification

In overall, UMATO follows the pipeline of UMAP; we first find the  $k$ NN, and by calculating  $\rho$  (Equation 1) and  $\sigma$  (Equation 2) for each point, we obtain the pairwise similarity for every pair of points in  $k$ NN indices. Once  $k$ NN indices are established, we unfold them and check the frequency of each point to sort them in descending order so that the index of the popular points can come to the front.

Then, we divide points into three disjoint sets: hubs, expanded nearest neighbors (eNN), and outliers. At first, we repeat the following steps until no points remain unconnected: (1) choose the point with the biggest frequency as a hub among points that are not yet selected; (2) remove the  $k$ NN of the chosen hub from the sorted list. Thanks to the sorted indices, the most popular point in each iteration—but not too densely located—becomes the hub point. Then, we recursively seek out hub points' NN and again look for the NN of those neighbors until there are no points to be newly appended. In other words, we find all connected points that are expanded from the original hub points, which, in turn, is called eNN. Any remaining point, neither a hub nor an eNN, becomes an outlier.

Our optimization takes distinctive approaches for different sets so that both global and local structures can be well preserved. That is, we first optimize hubs for the global structure, then perform local optimization for eNN. We do not optimize outliers to prevent them from disturbing the overall structure of the embedding. In the following section, we explain each step in detail.

### 3.2 Layout Optimization

**Global Optimization.** At first, we run the global optimization for hubs to construct the skeletal layout of the embedding. First, we initialize the positions of hub points using PCA, which makes the optimization process more stable than using random initial positions [18]. Next, we optimize the positions of hub points by minimizing the cross-entropy (Equation 4). Let  $f(X) = \{f(x_i, x_j) | x_i, x_j \in X\}$  and  $g(Y) = \{g(y_i, y_j) | y_i, y_j \in Y\}$  be two adjacency matrices in high- and low-dimensional spaces. If  $X_h$  represents a set of points selected as hubs in HD space and  $Y_h$  is a set of corresponding points in the embedding,  $CE(f(X_h) || g(Y_h))$  is minimized. Unlike UMAP, our global optimization excludes negative sampling approximation; therefore, the embedding becomes more robust and less biased in capturing global structure. As the optimization runs for a few hub points, it still does not require too much time to compute.

**Local Optimization.** In the second phase, UMATO additionally embeds eNN. For each eNN, we set its initial position as an average position of  $m$  (e.g., 10) NN with a small random perturbation. We then perform local optimization, similar to UMAP but slightly different. As explained in Sect. 2.1, UMAP starts by constructing a  $k$ NN graph; we perform the same task but only with the hubs and eNN. To do this, we need to update  $k$ NN indices constructed in advance (Sect. 3.1) to exclude outliers. We do this by simply replacing the outliers with new NN; thus, the computation is inexpensive.

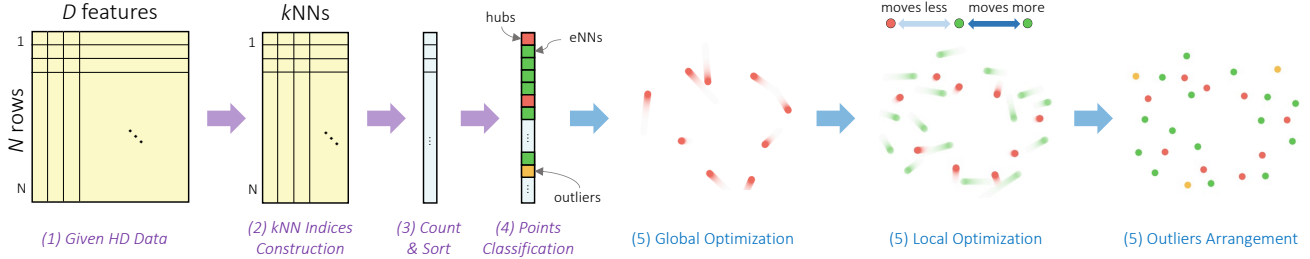


Figure 1: The illustration of the overall UMATO pipeline. Based on (1) given HD data, (2) we first find  $k$ NN for each point. (3) Then, we count the frequency of each point in the  $k$ NN index table and sort them in decreasing order (4) Next, points are categorized into hub points, expanded nearest neighbors (eNN), and outliers based on their connectivity to other points. (5) We initialize and optimize the positions of hubs (global optimization). (6) Leveraging the hubs' positions, the positions of eNN are optimized (local optimization). (7) Finally, outliers are arranged. The procedures described in Sect. 3.1 are depicted with purple arrow and captions, while the ones described in Sect. 3.2 are represented in blue.

Afterward, similar to UMAP, local optimizations of hubs and eNN are performed based on the cross-entropy loss function. The negative sampling technique (Equation 7) is also leveraged. However, rather than updating the positions of entire points equally, UMATO tries to keep the positions of hubs as much as possible because they have already formed the global structure. For this aim, while sampling a target edge  $(i, j)$ , we select  $i$  among eNN and choose  $j$  from both hubs and eNN. Then, if  $j$  is a hub, we penalize the attractive force for  $j$  by assigning a small weight (e.g., 0.1). This makes  $j$  not excessively affected by  $i$  if it is a hub point. In addition, as repulsive force can disperse the local attachment, making the point veer off for each epoch and eventually destroying the well-shaped global layout, we multiply a penalty (e.g., 0.1) when calculating the repulsive gradient (Equation 6) for the points selected as negative samples.

**Outliers Arrangement.** Isolated points (i.e., outliers) mostly have the same distance to all the other data points in HD space due to the curse of dimensionality [3]. Therefore, optimizing outliers can make them mingle with the already embedded points (i.e., hubs, eNN), thus sabotaging global and local structures. Therefore, we do not optimize outliers but simply append them using a previously embedded point that is located nearest to each outlier in the HD space. That is, for an outlier  $x_i \in C_n$  where  $C_n$  is the connected component to which  $x_i$  belongs, we first find  $x_j \in C_n$  that has already been embedded and is closest to  $x_i$ . Then, we set the low-dimensional position of  $x_i$  as the one of  $x_j$  with random noise. This helps us benefit from the comprehensive composition of the embedding that we have already optimized. Note that all outliers can find a previously embedded point as their neighbor. This is because we picked more than one hub from each connected component of the NN graph, and thus at least one point of each component has an optimized position (Sect. 3.1).

## 4 QUANTITATIVE EXPERIMENT

To evaluate UMATO in preserving global and local structures of HD data, we conducted experiments on one synthetic (Spheres) and three real-world datasets (MNIST [19], Fashion MNIST (FMNIST) [39], and Kuzushiji MNIST (KMNIST) [8]). We compared UMATO with six baseline DR techniques (PCA, Isomap,  $t$ -SNE, UMAP, TopoAE, and  $At$ -SNE) utilizing global (DTM, KL divergence) and local (T&C, MRREs) quality metrics.

### 4.1 Design

**Datasets.** For a synthetic dataset, we used Spheres [28]; it consists of 11 101-dimensional spheres, where ten spheres with relatively small radius of 5 and the number of points of 500 are enclosed by a larger sphere with a radius of 25 and the number of 5,000. A total of 10,000 points form the data. In the case of real-world datasets, we used MNIST, FMNIST, and KMNIST. The datasets represent images of digits, fashion items, and Japanese characters, respectively, and consist of 60,000 rows of 784-dimensional ( $28 \times 28$ ) vectors.

**Metrics.** To analyze DR techniques in diverse perspectives, we utilized several global and local quality metrics. For global met-

rics, we utilized Distance to a Measure (DTM [6, 7]) and KL divergence. Both assesses how well embeddings capture global structure in terms of density estimation. We selected such density-based metrics to align our experiment with the one of TopoAE [28]. In two functions, the density of each HD point  $x$  is defined as  $f_{\sigma}^X(x) := \sum_{y \in X} \exp(-\text{dist}(x, y)^2 / \sigma)$ , where  $X$  denotes the dataset. The density of corresponding low-dimensional point  $y$  is defined as  $f_{\sigma}^Z(z)$ , where  $Z$  denotes the embedding. Then, DTM is computed as  $\sum_{(x, z) \in (X, Z)} f_{\sigma}^X(x) - f_{\sigma}^Z(z)$ , while KL divergence is defined as  $KL(f_{\sigma}^X(x) \| f_{\sigma}^Z(z))$ . In accordance with the experiment of TopoAE, we used three  $\sigma$  values, 0.01, 0.1, and 1, for both measures.

In the case of local metrics, we used mean relative rank errors (MRREs [20]) and Trustworthiness & Continuity (T&C [36]), which quantifies the quality of embedding based on the preservation of neighborhood structure. While Trustworthiness and MRRE<sub>Z</sub> evaluate the extent to which neighbors in the embedding are also neighbors in the HD space, Continuity and MRRE<sub>X</sub> evaluate the opposite. These two metrics are selected because they have been widely used in DR literature [11, 14, 28]. They require a hyperparameter  $k$ , the number of nearest neighbors; we used  $k = 5$  throughout our experiments, aligned with the experiment of TopoAE [28]. The results with  $k = 10$  and  $k = 15$  are depicted in Appendix C.

**Baseline DR techniques.** For competitors, we used the most widely used techniques plus the ones that focus on global structures (Sect. 2.2). For the former, PCA,  $t$ -SNE, and UMAP were picked, and for the latter, Isomap,  $At$ -SNE, and TopoAE, were picked. To initialize an embedding, we used PCA for  $t$ -SNE, following the recommendation in previous work [22], and spectral embedding for UMAP, which is the default. The embeddings are depicted in Fig. 2. The hyperparameters of the techniques are set to minimize  $KL_{0.1}$ . Appendix B describes our hyperparameter setting in more detail.

### 4.2 Results

Table 1 displays the quantitative evaluation results. In terms of global metrics, UMATO surpassed all other techniques for Spheres and exceeded all techniques except Isomap for MNIST and FMNIST. Although UMATO did not indisputably show the best performance for KMNIST, it still surpassed other techniques if  $\sigma = 0.1$ , while there was no winner for  $\sigma = 0.01$  and 1. TopoAE and Isomap also showed good performance in preserving global metrics overall. Although  $At$ -SNE utilizes landmark based optimization, which is similar to UMATO, it showed competitive scores to neither UMATO nor other global techniques. For local metrics,  $t$ -SNE and UMAP were common winners. However, UMATO showed competitive performance in terms of MRRE<sub>X</sub> and continuity for FMNIST and KMNIST and in terms of MRRE<sub>Z</sub> and Trustworthiness for Spheres.

### 4.3 Discussions

The results clearly show the benefit of UMATO in preserving the global structure. The shape of embeddings also supports such a



Dataset	Algorithm	Global quality metrics						Local quality metrics ( $k = 5$ )			
		DTM <sub>0.01</sub>	DTM <sub>0.1</sub>	DTM <sub>1</sub>	KL <sub>0.01</sub>	KL <sub>0.1</sub>	KL <sub>1</sub>	Conti.	Trust.	MRRE <sub>X</sub>	MRRE <sub>Z</sub>
Spheres	PCA	1.0123	0.9950	0.1687	0.7568	0.6525	0.0153	0.7983	0.6088	0.7985	0.6078
	Isomap	<u>0.7020</u>	0.7784	0.1282	0.4492	0.4267	0.0095	<b>0.9041</b>	0.6266	<b>0.9039</b>	0.6268
	$t$ -SNE	0.9331	0.9144	0.1532	0.6091	0.5399	0.0130	<b>0.8916</b>	<b>0.7078</b>	<b>0.9045</b>	<b>0.7241</b>
	UMAP	0.9474	0.9209	0.1548	0.6100	0.5383	0.0134	0.8760	0.6499	0.8805	0.6494
	TopoAE	<b>0.4099</b>	<b>0.6890</b>	<b>0.1197</b>	<b>0.2063</b>	<b>0.3340</b>	<b>0.0076</b>	0.8317	0.6339	0.8317	0.6326
	Ar-SNE	0.9634	0.9448	0.1589	0.6584	0.5712	0.0138	0.8721	0.6433	0.8768	0.6424
	UMATO (ours)	<b>0.3271</b>	<b>0.3888</b>	<b>0.0529</b>	<b>0.1341</b>	<b>0.1434</b>	<b>0.0014</b>	0.7884	<b>0.6558</b>	0.7887	<b>0.6557</b>
Fashion MNIST	PCA	0.9373	0.2315	<b>0.0255</b>	0.6929	0.0454	<b>0.0006</b>	0.9843	0.9117	0.9853	0.9115
	Isomap	<b>0.9228</b>	<b>0.2272</b>	0.0352	<b>0.6668</b>	<b>0.0446</b>	0.0010	0.9865	0.9195	0.9872	0.9196
	$t$ -SNE	0.9987	0.2768	0.0442	0.8079	0.0663	0.0017	0.9899	<b>0.9949</b>	0.9919	<b>0.9955</b>
	UMAP	1.0125	0.2755	0.0438	0.8396	0.0641	0.0016	<b>0.9950</b>	0.9584	<b>0.9955</b>	0.9584
	TopoAE	0.9402	0.2329	<b>0.0311</b>	0.7301	0.0446	<b>0.0008</b>	0.9908	0.9591	0.9913	0.9590
	Ar-SNE	1.0187	0.2973	0.0454	0.8389	0.0702	0.0017	0.9826	<b>0.9847</b>	0.9849	<b>0.9848</b>
	UMATO (ours)	<b>0.9360</b>	<b>0.2035</b>	0.0314	<b>0.6852</b>	<b>0.0342</b>	0.0008	<b>0.9911</b>	0.9500	<b>0.9919</b>	0.9502
MNIST	PCA	1.3237	0.4104	0.0426	1.4981	0.1349	0.0014	0.9573	0.7340	0.9605	0.7342
	Isomap	<b>1.1936</b>	<b>0.3358</b>	<b>0.0382</b>	<b>1.0361</b>	<b>0.0857</b>	<b>0.0012</b>	0.9743	0.7527	0.976	0.7528
	$t$ -SNE	1.3193	0.4263	0.0553	1.4964	0.1523	0.0024	<b>0.9833</b>	<b>0.9954</b>	<b>0.9869</b>	<b>0.9963</b>
	UMAP	1.3428	0.4172	0.0588	1.5734	0.1430	0.0026	<b>0.9891</b>	0.9547	<b>0.9907</b>	0.9547
	TopoAE	1.3038	0.3686	<b>0.0366</b>	1.3818	0.1048	<b>0.0011</b>	0.9716	0.9429	0.9732	0.9429
	Ar-SNE	1.3312	0.4328	0.0466	1.5623	0.1482	0.0018	0.9768	<b>0.9765</b>	0.9830	<b>0.9777</b>
	UMATO (ours)	<b>1.2738</b>	<b>0.3525</b>	0.0414	<b>1.2785</b>	<b>0.1017</b>	0.0014	0.9792	0.8421	0.9813	0.8422
Kuzushiji MNIST	PCA	0.3756	0.4215	0.0440	0.1710	0.1317	0.0014	0.9380	0.7213	0.9420	0.7211
	Isomap	0.5267	<b>0.3458</b>	<b>0.0379</b>	0.2171	<b>0.0906</b>	<b>0.0012</b>	0.9573	0.7638	0.9589	0.7635
	$t$ -SNE	<b>0.2360</b>	0.4254	0.0571	<b>0.0483</b>	0.1369	0.0025	0.9843	<b>0.9688</b>	0.9871	<b>0.9693</b>
	UMAP	<b>0.2126</b>	0.3873	0.0566	<b>0.0417</b>	0.1148	0.0026	<b>0.9893</b>	0.9563	<b>0.9908</b>	0.9564
	TopoAE	0.3934	0.3730	<b>0.0373</b>	0.1495	0.1027	<b>0.0011</b>	0.9755	0.9442	0.9768	0.9440
	Ar-SNE	0.3029	0.3505	0.0384	0.0807	0.0978	0.0013	0.9786	<b>0.9671</b>	0.9824	<b>0.9676</b>
	UMATO (ours)	0.3993	<b>0.3231</b>	0.0435	0.1365	<b>0.0815</b>	0.0016	<b>0.9865</b>	0.8888	<b>0.9881</b>	0.8895

Table 1: Quantitative results of UMATO and six baseline algorithms. The hyperparameters of the DR algorithms are chosen to minimize KL<sub>0.1</sub>. The best one is bold and underlined, and the runner-up is bold. Only the first four digits are shown for conciseness.

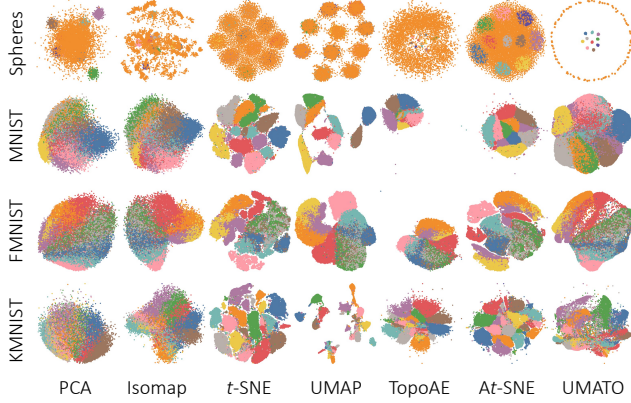


Figure 2: 2D embeddings of UMATO and six competitors. Overall, UMATO surpassed other techniques in preserving global structure while showing comparable performance in capturing local structure. A high-resolution version of this figure is provided in Appendix J.

claim. If we look at the first row of Fig. 2, we can see that the outer sphere encircles the inner spheres in a circular form in UMATO’s embedding, which is the most intuitive to understand the global relationship among different classes. However, all other embeddings did not show a clear separation between the outer and inner spheres.

Combining results from both global and local metrics, we can conclude that UMATO achieved high performance in capturing global structure at a slight loss for local structure preservation. Such a result aligns well with the design of UMATO; hubs help UMATO capture global structure in the first phase but work as a constraint for local optimization of eNN. Still, the fact that UMATO showed competitive performance in preserving local structures to  $t$ -SNE and UMAP clearly demonstrates its merit in analyzing HD data.

Refer to Appendix I for the results using synthetic datasets (e.g., Swiss Roll) that support our claim about UMATO’s capability.

## 5 ADDITIONAL EXPERIMENTS

We conducted additional experiments to understand UMATO’s features better and verify its benefit. At first, we compared the runtime of UMATO against previous nonlinear DR techniques (Appendix G). We also verified that UMATO avoids UMAP’s problem of clusters’ becoming dispersed as the number of epochs increases (Appendix D), and revealed that UMATO outperforms UMAP and  $t$ -SNE in terms of stability over subsampling (Appendix E) and robustness over diverse initializations (Appendix F). Finally, we discovered that utilizing more than two phases failed to improve the embedding results while consuming more computation time (Appendix H).

## 6 CONCLUSION

We present a two-phase DR algorithm called UMATO that can effectively preserve the global and local properties of HD data. In our experiments with diverse datasets, we have proven that UMATO outperforms previous widely used baselines (e.g.,  $t$ -SNE and UMAP) in capturing global structures while showing competitive performance in preserving local structures. In the future, we plan to optimize and accelerate UMATO in a heterogeneous system (e.g., GPU), as in previous attempts with other DR techniques [29, 31]. Qualitatively analyzing UMATO’s effect in analyzing HD data, such as visual cluster analysis [4, 13, 38], will also be an interesting direction.

## ACKNOWLEDGMENTS

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. NRF2019R1A2C208906213)

## REFERENCES

- [1] <https://github.com/hyungkwonko/umato>.
- [2] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems*, pp. 585–591, 2002.
- [3] R. Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966.
- [4] M. Cavallo and Ç. Demiralp. Clustrophile 2: Guided visual clustering analysis. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):267–276, 2018.
- [5] A. Chatzimpampas, R. M. Martins, and A. Kerren. t-visne: Interactive assessment and interpretation of t-sne projections. *IEEE Transactions on Visualization and Computer Graphics*, 26(8):2696–2714, 2020.
- [6] F. Chazal, D. Cohen-Steiner, and Q. Mérigot. Geometric inference for probability measures. *Foundations of Computational Mathematics*, 11(6):733–751, 2011.
- [7] F. Chazal, B. Fasy, F. Lecci, B. Michel, A. Rinaldo, A. Rinaldo, and L. Wasserman. Robust topological inference: Distance to a measure and kernel distance. *The Journal of Machine Learning Research*, 18(1):5845–5884, 2017.
- [8] T. Clanuwat, M. Bober-Irizar, A. Kitamoto, A. Lamb, K. Yamamoto, and D. Ha. Deep learning for classical japanese literature, 2018.
- [9] C. Fu, Y. Zhang, D. Cai, and X. Ren. Atsne: Efficient and robust visualization on gpu through hierarchical optimization. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 176–186, 2019.
- [10] T. Fujiwara, O.-H. Kwon, and K.-L. Ma. Supporting analysis of dimensionality reduction results with contrastive learning. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):45–55, 2019.
- [11] A. Gracia, S. González, V. Robles, and E. Menasalvas. A methodology to compare dimensionality reduction algorithms in terms of loss of quality. *Information Sciences*, 270:1–27, 2014.
- [12] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [13] H. Jeon, M. Aupetit, S. Lee, H.-K. Ko, Y. Kim, and J. Seo. Distortion-aware brushing for interactive cluster analysis in multidimensional projections. *arXiv preprint arXiv:2201.06379*, 2022.
- [14] H. Jeon, H.-K. Ko, J. Jo, Y. Kim, and J. Seo. Measuring and explaining the inter-cluster reliability of multidimensional projections. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):551–561, 2022. doi: 10.1109/TVCG.2021.3114833
- [15] J. Jo, J. Seo, and J.-D. Fekete. Panene: A progressive algorithm for indexing and querying approximate k-nearest neighbors. *IEEE Transactions on Visualization and Computer Graphics*, 26(2):1347–1360, 2018.
- [16] J. Kehler and H. Hauser. Visualization and visual analysis of multi-faceted scientific data: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 19(3):495–513, 2012.
- [17] D. Kobak and G. C. Linderman. Umap does not preserve global structure any better than t-sne when using the same initialization. *BioRxiv*, 2019.
- [18] D. Kobak and G. C. Linderman. Initialization is critical for preserving global data structure in both t-sne and umap. *Nature Biotechnology*, 39(2):156–157, 2021.
- [19] Y. LeCun and C. Cortes. MNIST handwritten digit database. 2010.
- [20] J. A. Lee and M. Verleysen. *Nonlinear dimensionality reduction*. Springer Science & Business Media, 2007.
- [21] J. A. Lee and M. Verleysen. Two key properties of dimensionality reduction methods. In *2014 IEEE Symposium on Computational Intelligence and Data Mining*, pp. 163–170. IEEE, 2014.
- [22] 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. *Nature methods*, 16(3):243–245, 2019.
- [23] S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci. Visualizing high-dimensional data: Advances in the past decade. *IEEE Transactions on Visualization and Computer Graphics*, 23(3):1249–1268, 2016.
- [24] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [25] W. E. Marcilio-Jr, D. M. Eler, F. V. Paulovich, and R. M. Martins. Humap: Hierarchical uniform manifold approximation and projection. *arXiv preprint arXiv:2106.07718*, 2021.
- [26] L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [27] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pp. 3111–3119, 2013.
- [28] M. Moor, M. Horn, B. Rieck, and K. Borgwardt. Topological autoencoders. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, Proceedings of Machine Learning Research. PMLR, 2020.
- [29] C. J. Nolet, V. Lafargue, E. Raff, T. Nanditale, T. Oates, J. Zedlewski, and J. Patterson. Bringing umap closer to the speed of light with gpu acceleration. *arXiv preprint arXiv:2008.00325*, 2020.
- [30] N. Pezzotti, T. Höllt, B. Lelieveldt, E. Eisemann, and A. Vilanova. Hierarchical stochastic neighbor embedding. In *Computer Graphics Forum*, vol. 35, pp. 21–30. Wiley Online Library, 2016.
- [31] N. Pezzotti, J. Thijssen, A. Mordvintsev, T. Höllt, B. Van Lew, B. P. Lelieveldt, E. Eisemann, and A. Vilanova. Gpgpu linear complexity t-sne optimization. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1172–1181, 2019.
- [32] V. D. Silva and J. B. Tenenbaum. Global versus local methods in nonlinear dimensionality reduction. In *Advances in Neural Information Processing Systems*, pp. 721–728, 2003.
- [33] J. Tang, J. Liu, M. Zhang, and Q. Mei. Visualizing large-scale and high-dimensional data. In *Proceedings of the 25th International Conference on World Wide Web*, pp. 287–297, 2016.
- [34] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web*, pp. 1067–1077, 2015.
- [35] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [36] J. Venna and S. Kaski. Neighborhood preservation in nonlinear projection methods: An experimental study. In *International Conference on Artificial Neural Networks*, pp. 485–491. Springer, 2001.
- [37] J. Venna, J. Peltonen, K. Nybo, H. Aidos, and S. Kaski. Information retrieval perspective to nonlinear dimensionality reduction for data visualization. *Journal of Machine Learning Research*, 11(2), 2010.
- [38] J. Xia, Y. Zhang, J. Song, Y. Chen, Y. Wang, and S. Liu. Revisiting dimensionality reduction techniques for visual cluster analysis: An empirical study. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):529–539, 2021.
- [39] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.