Revisiting Dimensionality Reduction Techniques for Visual Cluster Analysis: An Empirical Study

Jiazhi Xia, Yuchen Zhang, Jie Song, Yang Chen, Yunhai Wang and Shixia Liu
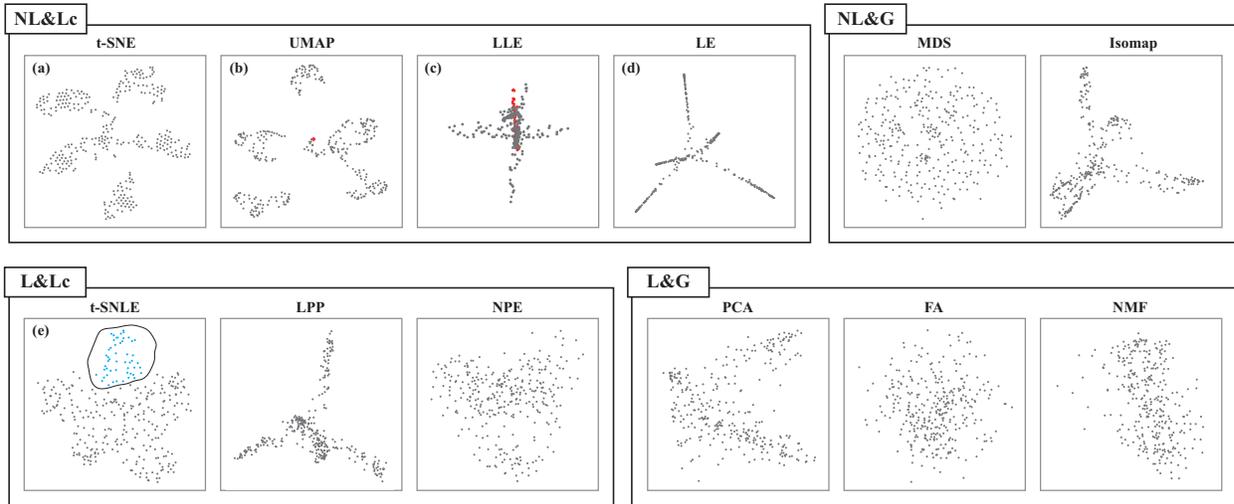
Fig. 1: Different dimensionality reduction results of the ExtYaleB dataset, which contains 320 images of five people. The twelve dimensionality reduction techniques fall into four categories: Non-linear & Local (NL&Lc), Non-linear & Global (NL&G), Linear & Local (L&Lc), and Linear & Global (L&G). Instances of task stimuli: (a) Cluster identification (**T1**), (b) Membership identification (**T2**), (c) Distance comparison (**T3**), and (d) Density comparison (**T4**). (e) lasso tool used in **T1–T4**.

Abstract—Dimensionality Reduction (DR) techniques can generate 2D projections and enable visual exploration of cluster structures of high-dimensional datasets. However, different DR techniques would yield various patterns, which significantly affect the performance of visual cluster analysis tasks. We present the results of a user study that investigates the influence of different DR techniques on visual cluster analysis. Our study focuses on the most concerned property types, namely the linearity and locality, and evaluates twelve representative DR techniques that cover the concerned properties. Four controlled experiments were conducted to evaluate how the DR techniques facilitate the tasks of 1) cluster identification, 2) membership identification, 3) distance comparison, and 4) density comparison, respectively. We also evaluated users' subjective preference of the DR techniques regarding the quality of projected clusters. The results show that: 1) Non-linear and Local techniques are preferred in cluster identification and membership identification; 2) Linear techniques perform better than non-linear techniques in density comparison; 3) UMAP (Uniform Manifold Approximation and Projection) and t-SNE (t-Distributed Stochastic Neighbor Embedding) perform the best in cluster identification and membership identification; 4) NMF (Nonnegative Matrix Factorization) has competitive performance in distance comparison; 5) t-SNLE (t-Distributed Stochastic Neighbor Linear Embedding) has competitive performance in density comparison.

Index Terms—Dimensionality reduction, visual cluster analysis, perception-based evaluation

---

◆

---

1 INTRODUCTION

Visual cluster analysis usually employs dimensionality reduction (DR) techniques to project high-dimensional data into 2D scatterplots, in which analysts can visually identify cluster patterns [53, 59]. The

- *Jiazhi Xia, Yuchen Zhang and Jie Song are with School of Computer Science and Engineering, Central South University, China. E-mail: {xiajiazhi,yuchenzhang,jiesong}@csu.edu.cn;*
- *Yang Chen is with I4 data. E-mail: chen1984yang@gmail.com.*
- *Yunhai Wang is with School of Computer Science and Technology, Shandong University, China. E-mail: cloudseawang@gmail.com.*
- *Shixia Liu is with School of Software, Tsinghua University, China. E-mail: shixia@tsinghua.edu.cn.*
- *Jiazhi Xia and Yang Chen are corresponding authors.*

cluster analysis is an inherent human-in-the-loop task that lacks a universal ground truth [6]. Therefore, visual cluster analysis, which assumes clusters in the projection are faithful representation matching with actual clusters, is widely used in evaluating automatic clustering algorithms [6] and clustering applications [7, 21, 33, 58, 60].

However, different DR techniques perform differently in a set of aspects in visual cluster analysis, e.g., cluster separation, membership preservation, distance preservation, and density preservation. As shown in Fig. 1, the twelve DR techniques generated various visual patterns out of the same dataset. Although each DR technique provides different insights into clusters, most of them are rarely used in real applications. For example, we have surveyed papers in the mainstream visualization journals and conferences from 2010 to 2020 with a focus on applying DR techniques in visual cluster analysis. Among the 51 resulted papers, we found that only PCA (Principal Component Analysis) [55], MDS (Multidimensional Scaling) [28], and t-SNE (t-Distributed Stochastic Neighbor Embedding) [47] are used for more than 2 times. A revisit to DR techniques is needed to guide the use of them for visual cluster

analysis.

Moreover, a high-level view of the performances of DR techniques on visual cluster analysis is a common concern of researchers. When choosing DR techniques from a large set, analysts usually consider the property that the DR technique preserves. Among various properties, linearity and locality have gained the most attention [16, 38]. For example, linear DR techniques are believed to outperform in preserving the density of clusters [11, 37]. On the other hand, non-linear DR techniques, such as t-SNE and UMAP (Uniform Manifold Approximation and Projection) [36], are widely used because of their performance in separating clusters. The locality property, i.e., if local or global pairwise relationships are preserved, also yields significant differences in the generated scatterplots [13, 16].

Nevertheless, many qualitative comparisons have been conducted towards understanding how different DR techniques perform in visual cluster analysis tasks [17, 49, 57, 61]. However, we argue that two significant gaps exist in current research. First, while the aforementioned high-level properties are of concern to both DR researchers and analysts, there lacks experimental studies that test and verify them against user's visual perception in visual cluster analysis tasks. Second, none of the existing experimental studies has evaluated a set of DR techniques that can be representative of all the properties or cover the most popular DR techniques, calling for a comprehensive and systematic comparison of these methods.

To fill these gaps, we conduct controlled experiments to evaluate DR techniques in visual cluster analysis. Given the same high-dimensional datasets and certain cluster analysis tasks, our study compares the effectiveness of graphical outcomes generated by different DR techniques and investigates how their essential properties affect user's visual perception and the task performance. The study starts with a comprehensive literature review of existing DR techniques, upon which twelve widely-used techniques are selected for evaluation (see Table 1). They are grouped based on two properties, namely locality and linearity, both of which are among the most concerned topics in the existing DR research. This grouping allows a high-level view of the DR techniques regarding their core features and thus the study results can be easily understood by data analysts. To ensure the study results are intuitive and generalized to various domains, we carefully define four analytical tasks that are essential for general visual cluster analysis. Each task is evaluated through a controlled experiment. The four corresponding controlled experiments are assessed based on ground truth established in the data space.

Regarding the effectiveness of the selected DR techniques, a total of twelve hypotheses are formulated before the study. They are derived based on the summarized advantages of the DR techniques from our literature review. In the experiment design, three key design choices are made to ensure the soundness of the experiment and obtain compelling results. First, we select a total of eight datasets whose dimension count, data points count, and cluster patterns are balanced in the design. Then, a pre-study is conducted to tune and optimize the parameters for the selected DR techniques, ensuring their robustness for each of the datasets. In addition to the four controlled experiments, a subjective experiment is conducted where the participants rank the clustering outcomes based on their personal preference. It allows us to evaluate the DR from a more subjective perspective.

Based on the experiment results, we perform a comprehensive statistical analysis. The analysis results of the objective metrics validate seven of the hypotheses but leave five rejected. We found that although non-linear and local techniques show significant advantages in tasks such as cluster identification and cluster membership identification, there is not a technique that can outperform the others in all the four visual cluster tasks. Moreover, some linear techniques that receive less citation, such as NMF (Nonnegative Matrix Factorization) [29] and t-SNLE (t-Distributed Stochastic Neighbor Linear Embedding) [9], catch our eyes because they yield better performance in tasks such as cluster distance comparison and density comparison, respectively. We also discuss the impact of the two DR properties on the effectiveness of the visual perception and the potential improvement of the experiment design. The DR techniques, datasets, and records of the study are available at https://github.com/DR-approach/DR-approaches.

2 RELATED WORK

2.1 DR techniques for visual cluster analysis

Because our study focuses on the two essential properties of linearity and locality, we divide existing DR techniques into four groups as Espadoto et al. [16]: Non-Linear&Local, Linear&Local, Non-Linear&Global, and Linear&Global.

Non-Linear&Local utilizes non-linear functions and preserves local neighborhood in DR processes. Representatives include LLE (Locally Linear Embedding) [42], LE (Laplacian Eigenmaps) [8], LAMP (Local Affine Multidimensional Projection) [25], t-SNE [47], and UMAP [36]. They are usually considered as manifold learning techniques that can capture the clusters in manifolds. For example, LLE preserves the linear combination of local neighborhoods. To achieve a similar goal, LE embeds data points according to the eigenfunctions of the Laplace Beltrami operator on the manifold. LAMP uses orthogonal mapping theory to build accurate local transformations. Especially, t-SNE and UMAP are widely used in visual cluster analysis because of their robust performance on clustering separation [2, 32].

Linear&Local utilizes linear functions and preserves local neighborhood in DR processes. Representatives include LPP (Locality Preserving Projection) [23], NPE (Neighborhood Preserving Embedding) [22], and t-SNLE [9]. All the three techniques are linear variants of LE, LLE, and t-SNE, respectively. Compared to their non-linear counterparts, they retain the advantage in clustering separation but are limited by the linear projection [9].

Non-Linear&Global utilizes non-linear functions to preserve the pairwise distance among all the data points. Representatives include LSP (Least Square Projection) [39], MDS (non-linear version) [24, 28], and Isomap [46]. LSP uses least squares approximations. The non-linear version of MDS preserves distances of high-dimensional space in the embedded low-dimensional space. Different from MDS, Isomap attempts to preserve the geodesic distance between any two pairs of data points. The scheme enables a consistent overview of the data based on the globally optimal layouts [12].

Linear&Global uses linear functions to preserve the pairwise distance structure amongst all the data samples. Representatives include PCA [55], FA (Factor Analysis) [34], and NMF [29]. PCA aims to preserve data patterns in the aspect of variance. Ding et al. [14] prove that principal components are actually the continuous solution of the cluster membership indicators in the K-means clustering method, which indicates that PCA implicitly performs clustering. Compared to PCA, FA intends to identify latent variables underlying a higher-dimensional space of measurements. NMF aims to find two non-negative matrices W and H whose product can well approximate the original matrix V, namely $W \times H \approx V$. The matrix W is a low-dimensional representation of the original data [10].

2.2 Evaluations of DR techniques in the aspect of visual cluster/class analysis

Evaluation based on quantitative metrics. A variety of quantitative metrics has been used to measure the spatial aspects of projections generated by different DR techniques. The metrics can be used either standalone to gauge desirable aspects of a projection, or jointly to assess the overall quality of outcomes. For example, aggregate metrics such as trustworthiness, continuity, neighborhood hit, distance and class consistency can be used to assess the quality of data clusters in the projected space [16]. Local metrics, such as projection precision score [43] and average local errors [35], separately compare small neighborhoods in a projection, upon which the detailed levels of the result can be assessed. To further assess the effectiveness of human perception on the projection, Aupetit et al. [5] proposed a set of visual measures that can mimic the human notion of separability in data classes. By validating the measures with a machine learning framework, they found the average proportion of same-class neighbors has the best prediction accuracy on human perception. Unlike the above measures [5], which are specific to color-coded scatterplots, Clumpy scagnostic [54] and ClustMe [1] are visual quality measures to quantify clusters in monochrome scatterplots.

ClustMe was based on a quantitative study of cluster counting task. Wang et al. [51] used GONG [4] and SC (silhouette coefficient) [41] to compare the class separation performance of 10 DR techniques such as PCA and t-SNE on 93 datasets. Spathis et al. [45] used SC to evaluate the performance of interactive DR techniques such as LAMP, PLMP (Part-Linear Multidimensional Projection), and KELP (Kernel-based Linear Projection). Vernier et al [50] evaluated performance of eleven DR techniques on ten datasets using the neighborhood preservation.

Evaluation based on user perception. Lewis et al. [30] and Sedl-mair et al. [44] reported that quantitative metrics might perform differently from human perception in cluster recognition and visual class separation tasks, respectively. A set of researches focus on perception based evaluations. Lewis et al. [31] designed a study to investigate human agreement on the embedding quality. Their results show that expert users are reasonably consistent judges of embedding quality, whereas novice users are very inconsistent with one another. Etemadpour et al. [17] conducted a controlled experiment to evaluate the performance of DR techniques on five visual cluster analysis tasks. Ventocilla et al. [49] conducted a similar perception evaluation but with foci on PCA, t-SNE, Radviz, and SC. Xu et al. [57] conducted an experiment to compare advanced feature transformation technology (FT-high) and traditional DR techniques, such as PCA and MCML, regarding their performance in visual cluster analysis tasks. The results show that FT-high would yield better results than the DR techniques, especially in the accuracy of identifying clusters. Zhao et al. [61] conducted a controlled experiment to evaluate the performance of four DR techniques on fuzzy clusters analysis. Nevertheless, none of the above experimental studies has considered the two important properties, namely linearity and locality, in terms of their impact on the performance of the DR techniques. Another important difference lies in the task design and measure. Most existing studies are based on color-encoded scatterplots rather than monochrome scatterplots. In the task of identifying clusters [17], although monochrome scatterplots are used, their metric is based on the number of identified clusters only and lacks an accurate measure. Furthermore, most existing studies do not explore other cluster properties, such as density and membership.

3 EVALUATION LANDSCAPE

In order to establish a solid basis for our study, we have conducted a comprehensive literature review, upon which the DR techniques, the datasets, and the analytical tasks of the study are derived.

3.1 Selection of DR techniques

To select the most representative techniques from the wealth of existing DR research, we consider the following three metrics: (1) *influence* - the techniques should be influential in general research fields, and more importantly, in the visualization field as well. In practice, we consider the citation as a major indicator of the influence; (2) *task-suitability* - the techniques should have been widely applied in visual cluster analysis, and considered essential for accomplishing the related analytical tasks. Specifically, we focus on those that are the most cited techniques for certain visual cluster analysis tasks; (3) *comparison-balance* - since we intend to investigate the techniques based on their high-level categories, each category should have a consistent number of techniques for an effective comparison.

Based on the first two metrics, we conducted a comprehensive literature survey on recent research papers in the visualization community. Specifically, the source includes papers published in the journal of IEEE TVCG and three mainstream visualization conferences (IEEE VIS, EuroVis, and PacificVis) between 2010 and 2020. We distilled 937 papers that not only consist of the keywords "cluster" or "clustering" but also are relevant to DR techniques. Among them, we narrowed the keyword search related to the visual cluster analysis, which derived 51 papers as the finally refined source.

We conducted an in-depth survey on each of the papers and selected nine DR techniques based on our metrics. They were grouped based on the properties of linearity and locality. Table 1 summarizes the selected techniques with their categories, their google scholar citation counts, and the number of times it has been cited within the 51 papers. The

Techniques	Linearity	Locality	Citation	NO. of Applications
t-SNE [47]	non-linear	local	18655	30
Isomap [46]	non-linear	global	14394	2
UMAP [36]	non-linear	local	2244	2
MDS [28]	non-linear	global	7942	19
LE [8]	non-linear	local	4803	2
LLE [42]	non-linear	local	15811	2
PCA [55]	linear	global	8807	14
FA [34]	linear	global	14660	1
NMF [29]	linear	global	12365	1
LPP [23]	linear	local	4702	0
NPE [22]	linear	local	1829	0
t-SNLE [9]	linear	local	87	0

Table 1: DR techniques used in our evaluation.

Datasets	Data Items	Dimensions	Clusters
EcoliProteins [15]	336	7	8
Dermatology [15]	259	34	6
ExtYaleB [20]	320	30	5
World12d [44]	151	12	5
Boston [44]	155	13	3
Mnist64 [15]	1083	64	6
Weather [49]	366	194	7
Olive [18]	572	8	3

Table 2: Datasets used in our evaluation.

classical techniques, such as PCA, have multiple related papers that have high citation counts. For clarity, we selected the paper with the highest citation counts for a technique. While there was no DR technique in the group of Linear&Local, we selected three more techniques, namely LPP, NPE, and t-SNLE, which are linear variants of LE, LLE, and t-SNE, respectively. Similar to Espadoto et al. [16], these twelve techniques were distributed into four groups, including Linear & Local (L&Lc), Linear & Global (L&G), Non-linear&Local (NL&Lc), and Non-linear & Global (N&G) (see Table 1).

3.2 Selection of Datasets

To reflect the common interests of visualization community, we select eight datasets out of 41 datasets that were used in the visualization applications in our survey. First, we considered the data size within a range from 100 to 1500 data instances and the cluster size from 3 to 10 to avoid scalability issue. Second, in order to make sure that the study can distinguish the performance of DR techniques, we select datasets for which cluster patterns can be visually separated in at least one projection result of a DR technique. Due to the huge decision space (41 datasets and 12 DR techniques), we assumed classes given by data labels formed clearly separable clusters. We used visual separation measures [5] GONG and KNNG to quantify class separation and select the best projections for our study. Noting that we have not fixed the DR technique that can visually separate clusters, selecting datasets with visually distinguishable clusters introduces few bias into the study, if there is. As a result, the selected eight datasets are shown in Table 2). We use PCA to reduce the dimensionality of ExtYaleB from 32,256 to 30 for a meaningful projection. All datasets are normalized by dimension.

3.3 Selection of Typical Tasks

The goal of the task formulation is to select the most common analytical tasks of the visual cluster analysis within general application domains. To this end, we extracted all the analytical tasks from the 51 surveyed cluster analysis papers and selected the top four most-cited tasks.

- **Cluster identification (T1)**: given a scatterplot generated by a DR technique, identify dense and well-separated clusters. Input: projected data. Output: color-coded points from each identified cluster. Action: lasso selection.

Techniques	Key Parameters	Settings	Other Parameters
t-SNE	*Perplexity*	5,15,30,40,50	*Early exaggeration*=6 *Iteration*=3000
t-SNLE	*Perplexity*	5,15,30,40,50	*Early exaggeration*=6 *Iteration*=3000
UMAP	*Neighbors*	4,7,10,13,16	*Initialization*=spectral *Minimum distance*=0.1 *Iteration*=500
LLE	*Neighbors*	4,7,10,13,16	*IReg*=0.001 *Iteration*=200
LPP	*Neighbors*	4,7,10,13,16	None
NPE	*Neighbors*	4,7,10,13,16	None
Isomap	*Neighbors*	4,7,10,13,16	None
LE	*Neighbors*	4,7,10,13,16	None
MDS	None		*Iteration*=500
FA	None		*Iteration*=2000
NMF	None		*Iteration*=400
PCA	None		No parameters

Table 3: The parameter settings of DR techniques in the pre-study. The formal study used the same settings except for *perplexity* and *neighbors*.

- **Membership identification (T2)**: given a point in a scatterplot, identify the cluster it belongs to. Input: projected data and one color-coded target point. Output: color-coded points of the cluster identified to contain the target point. Action: lasso selection.
- **Distance comparison (T3)**: given a cluster in a scatterplot, identify the nearest cluster to it. Input: projected data and color-coded points from one target cluster. Output: color-coded points from the cluster identified to be nearest to the target one. Action: lasso selection.
- **Density comparison (T4)**: given multiple clusters, identify the cluster with the largest density. Input: projected data. Output: color-coded points from the cluster identified to be of largest density. Action: lasso selection.

All tasks are completed in projection space. For each task, we created an experiment to compare the performance of different DR techniques. They are described in Section 5.2.

4 PRE-STUDY: PARAMETERS OF DR TECHNIQUES

The projection of the DR technique is usually significantly affected by their hyperparameters [19]. Improper parameters will lead to unreliable comparison results. In order to optimize the parameters of selected DR techniques for each dataset, we have conducted a preliminary experiment.

4.1 Parameters

In this study, we focus on hyperparameters that significantly affect the projection results and require trial-and-error tuning. Specifically, we considered two parameters, namely *neighborhood value* and *perplexity*, as they significantly affect the projection quality of LPP, NPE, Isomap, UMAP, LE, LLE, t-SNE, and t-SNLE in different degrees [16, 19, 48]. For the other parameters, such as iteration times, we adopted recommended values based on previous studies [44, 47]. Following the previous studies [16, 48], we formulated 5-level ranges for the parameter *neighborhood value* and *perplexity*, from which the optimal values can be selected for each DR method. The parameter range formulation ensures a moderate degree of discrimination for the projection results while keeping a reasonable range to fit the optimal values. The settings are presented in Table 3.

4.2 Task and Procedure

The same set of datasets was used in both formal study and pre-study. For a given dataset and a DR technique, participants were provided with 5 projections resulting from 5 parameter settings (see Table 3). The number and size of classes, which are taken as ground truth for clusters, are also expicitly mentioned to the participants of the pre-study

	1	2	3	4	5	6	7	8
t-SNE	40	5	15	5	15	40	5	40
t-SNLE	30	15	30	30	15	40	5	40
UMAP	13	16	10	7	7	16	4	10
Isomap	16	4	10	4	10	10	13	10
LLE	4	10	10	4	10	10	4	7
LE	4	4	7	10	10	16	10	13
LPP	10	7	10	13	7	4	10	16
NPE	4	10	7	4	4	16	13	7

Table 4: The optimal parameters, i.e., the *perplexity* and *neighbors* (see Table 3), of DR techniques on the 8 datasets. Each dataset is represented by a serial number ID from 1 to 8, which corresponds to the 8 datasets in an order of Weather, Dermatology, ExtYaleB, World12d, Boston, Mnist64, EcoliProteins, and Olive, respectively.

(see figure 5 in 7.1 of supplemental material). It is worth mentioning that this is only for the pre-study to select good projections, not for the formal study. They were asked to select the projection that best fits the description: **"Please select the scatterplot that presents the cluster structure of the given dataset most actually"**. For each participant, a total of $8(DR\ techniques) \times 8(datasets) = 64$ trials were evaluated.

Participants. We recruited 20 participants (16 males, 4 females) who are graduate students majoring in computer science. None of the participants reported color blindness or color weakness.

Procedure. The preliminary experiment was conducted on an online application that enables a selection of DR projections for the pre-processed datasets. Participants were instructed to log in to the application to complete the tasks remotely. The tasks were conducted on standard laptops with a 1,920×1,080 screen resolution using the Chrome browser. Before completing the tasks, each participant was provided with a 10-minute background introduction, including research purposes, related concepts, experimental procedures, and precautions. Then, the participants were instructed to finish a total of 64 trials in random order. There was no time limit for task completion, but all the participants finished within 50 minutes. To track the progress, the participants were asked to share the screen with a remote instructor during the experiments.

4.3 Results

For a given DR technique and a dataset, we retrieved the projection with the majority vote and considered the corresponding parameter values optimal for that dataset. The result is shown in Table 4. Based on the results, we utilized the optimal values of the DR techniques for the corresponding dataset in the formal study.

5 FORMAL STUDY

5.1 Hypotheses

This formal study aims to evaluate the performance of twelve selected DR techniques on four common visual cluster analysis tasks. Please refer to the tasks from section 3. Based on the general views from our literature review and practical experience, we put forward twelve hypotheses to verify the possible performance differences of different types of DR techniques on the four tasks.

H1.1–H1.3 Regarding **T1**, existing research pointed out that local DR techniques often ensure a clear visual separation of data entities in projections [23, 47]. It is also pointed out that non-linear DR techniques favor the optimization of neighborhood distances and mostly disregard large distances [11]. Again, this might have positive effects on differentiating clusters in the projections. Moreover, most applications employed t-SNE and UMAP for cluster analysis. Therefore, we have the following hypotheses:

- **H1.1** Local DR techniques perform better than global ones in **T1** if they have the same linearity type;
- **H1.2** Non-linear DR techniques perform better than linear ones in **T1** if they have the same locality type;
- **H1.3** t-SNE and UMAP perform better than other techniques in **T1** .

H2.1–H2.3 Regarding **T2**, although it is a different perception task from **T1**, it also requires the ability to separate clusters and preserve near neighbors. Therefore, we have hypotheses **H2.1**, **H2.2**, and **H2.3** in **T2** that are identical to **H1.1**, **H1.2**, and **H1.3** in **T1**, respectively.

H3.1–H3.3 Regarding **T3**, global DR techniques are designed to preserve global pairwise distances and thus could enable more faithful projections of global structures, e.g., cluster relationships, of high-dimensional spaces [26]. Compared to the non-linear embeddings, the linear embeddings are more reliable in terms of metric-preserving properties [27]. This might facilitate the representation of cluster relationships. Moreover, PCA is the most widely used global linear technique. Therefore, we have the following hypotheses:

- **H3.1** Global DR techniques perform better than local ones in **T3** when they have the same linearity;
- **H3.2** Linear DR techniques perform better than non-linear ones in **T3** when they have the same locality;
- **H3.3** PCA performs better than other techniques in **T3**.

H4.1–H4.3 When considering **T4**, there are many practical experiences [37, 52] pointing out that non-linear or local techniques neglect the density information. It is also considered that PCA performs better than non-linear techniques [11]. Therefore, we have hypotheses **H4.1**, **H4.2**, and **H4.3** in **T2** that are identical to **H3.1**, **H3.2**, and **H3.3** in **T3**, respectively.

5.2 Experiments

Under the guidance of the four tasks (**T1-T4**) and the associated twelve hypotheses (**H1.1-H4.3**), we designed four controlled experiments: **E1–E4** for **T1–T4** to verify **H1–H4**, respectively. In addition to the four controlled experiments, a subjective evaluation **E5** was designed where we studied participants' preferences for different DR techniques based on their free exploration of datasets. All DR techniques and datasets for **E1–E5** are the same. There were $12(techniques) \times 8(datasets) = 96$ trials for each participant in **E1–E4** and $1(3 \times 4 \, layout \, for \, 12 \, techniques) \times 8(datasets) = 8$ trials for each participant in **E5**.

E1: **T1** was performed to test if projected clusters were separated cleanly and easily recognized by participants. In each trial, the participants were asked to **use a lasso tool to interactively select well-separated clusters, if any, from the visualization generated by a DR technique.** Compared with other evaluation methods such as counting the cluster number [1, 17], interactive selections enable more accurate reflection of participants' visual perceptions, especially for handling aggregated clusters that might lead to bias in counting the cluster number [3]. Since color would interfere with the judgments of the clusters, we rendered all data points in grey color and full opacity.

E2: **T2** was performed to test whether a meaningful belonging relationship between data points and clusters can be easily identified. In each trial, a pre-defined data point was highlighted in red color. Participants were asked to **use a lasso tool to interactively select a well-separated cluster, if any, that contains the highlighted data point.** When selecting the pre-defined data point, we prefer challenging ones that are close to the cluster boundary in the high-dimensional space. To this end, we extend the clustering metric of silhouette coefficient [41] to high-dimensional space, upon which the silhouette coefficient values of each data point are calculated. Then, the point with the lowest value was selected.

E3: **T3** was performed to evaluate the ability of different DR techniques to preserve the relative distance among clusters in terms of visual perception. In each trial, a pre-defined cluster was highlighted in red color. Participants were asked to **use a lasso tool to interactively select a well-separated cluster, if any, that is closest to the given cluster.**

E4: **T4** was performed to evaluate the ability of different DR techniques to preserve the density in the aspect of visual perception. In each trial, participants were asked to **use a lasso tool to interactively select a well-separated cluster, if any, that has the highest density among all clusters.**

E5: Subjective experiment was used to compare the overall performance of different DR techniques based on participants' personal
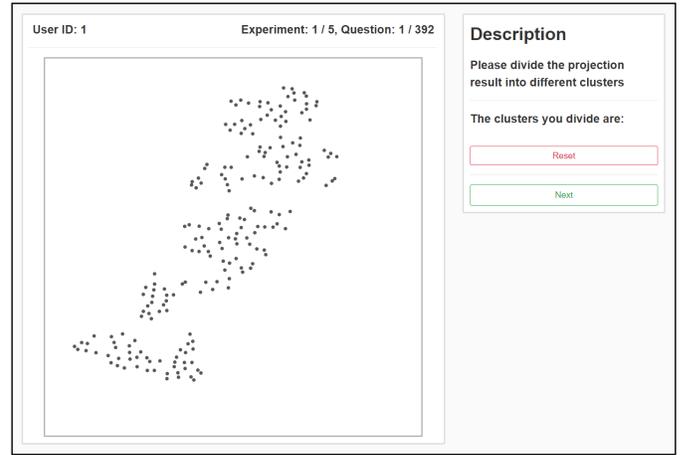
Fig. 2: Example interface of **E1** in the formal study. Participants were asked to identify clusters by lasso tool.
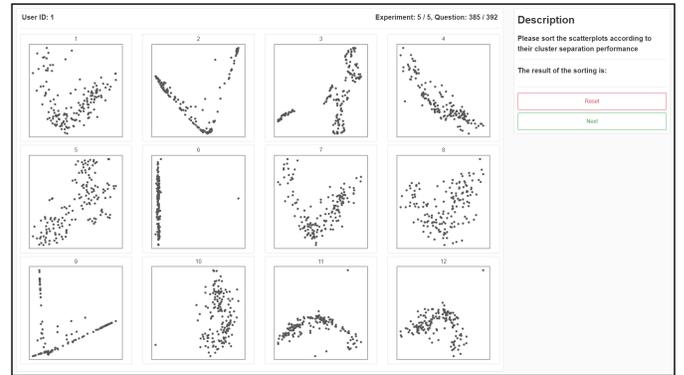
Fig. 3: Example interface of **E5** in the formal study. Participants were asked to rank the twelve visualizations according to their cluster separation performance.

preferences. In each trial, we assigned a participant twelve visualizations of a dataset generated by twelve DR techniques, respectively. They were asked to **rank the rank the visualization from 1 to 12 (1 is the best) by their personal preference in terms of cluster separation performance, identical ranking for multiple visualizations is allowed in case their qualities are hard to distinguish.** The visualizations were numbered and placed in a 3×4 layout with a random order (see Fig.3).

5.3 Participants, Apparatus, and Testing Data

Participants. We recruited 60 participants (40 males and 20 females) for the formal study. Their age was from 19 to 32 with an average of 24. 56 of them are graduate students and four of them are bachelor students with research experience. They come from majors including computer science (46), engineering (9), art (4), and mathematics (1). 19 participants have a data visualization background. 6 participants study clustering algorithms. 11 participants have performed cluster analysis in bioinformatics. Each participant was rewarded $20 per hour for completing the experiments. None of them reported color blindness or color weakness. Participants in the formal study are different from participants in the pre-study.

Apparatus. The formal study was conducted online on a pre-built web application in order to involve more participants. The application allows participants to complete all the controlled experiments and subjective studies (**E1–E5**), providing functionality such as visualizations, interactions, and time counter to facilitate the experiments. All the participants remotely took the experiments on standard laptops with a 1,920×1,080 screen resolution and chrome browser. The participants were asked to share their screens with the researcher during the experiments to enable

remote monitoring.

Testing Data. Before the formal study, we generated scatterplots based on the eight selected datasets. For each dataset, we created twelve scatterplots by each of the selected DR techniques, respectively. The parameter of each DR technique was determined based on the results of the pre-study. The same layout was used for a specific DR technique and a dataset for different experiments. Except for the given points in **E2** and the given clusters in **E3**, which were coded in red, the data points in the scatter plots were all coded in gray. The points were rendered with a radius of 3 pixels without transparency. The size of the scatterplots was 600×550 pixels in **E1–E4**, and 300×300 pixels in **E1**.

5.4 Procedure

A training session was conducted before the formal experiments. For the training session, we utilized the Iris dataset [15] to generate scatterplots for each DR technique. In the training session, the instructor explained the related concepts and the experimental procedure. Then, the participant was asked to complete a total of thirteen training trials, including three trials for each controlled experiment (**E1–E4**) and one trial for the subjective experiment (**E5**). The DR technique of the trial was randomly selected for each experiment. After the training session, the participant was asked to complete **E1–E5** in order. A 5-minutes break was allowed before each experiment. The completion order of the trials in each experiment was randomly assigned to each participant. The participants used the application to complete each trial, upon which the trial completion time and results were automatically stored by the application. There was no time limit for each experiment. All the participants finished the five experiments within 120 minutes.

Toward the end of **E5**, the participants were asked to complete a questionnaire about their background and subjective feedback of the experiments. The questionnaire includes three sections. The first section is about the participant background, such as the education level, visualization experience, and knowledge about the DR techniques. In the second section, the participant was asked to list the factors, including variance of density, cluster separation, variance of count, variance of size and variance of shape, if any, they relied on identifying the clusters in **E1**. In the last section, the participant was asked to rate the importance of the four analytical tasks (**T1-4**) in their data analysis experience using a five-point Likert scale. After the questionnaire, we conducted a subjective interview to understand how the scatterplots could affect participants' trial completion and their preference and/or experience on any DR techniques and visual clustering techniques in daily work.

6 STATISTICAL ANALYSIS

6.1 Result Measurement

To measure results of each trial for the five experiments, two conditions were considered: the trial completion time (the shorter the better), precision and recall (1 is best) were calculated for **E1–E4**; and the average ranking score of DR techniques was calculated for **E5**. The ground truth for computing the precision and recall are based on the class label of the datasets.

E1: While participants would identify different cluster structures, e.g., number of clusters, from the ground truth, matching the identification and the ground truth should be performed before the computation of precision and recall. We denote the identified clusters by participants as $\{C_i\}_{i=1}^{K}$, where K is the number of identified clusters. The ground-truth clusters are denoted as $\{G_j\}_{j=1}^{N}$, where N is the number of labels in the dataset. A matching function f matches an identified cluster C_i to $f(C_i)$, where $f(C_i) \in \{G_j\}_{j=1}^{N}$. Note that $f(C_i)$ could be an empty set when there is no matching for C_i. The match is based on two rules. First, if an identified cluster C_i contains data points from multiple ground-truth clusters, it is matched with the ground-truth cluster contributing the most data points to C_i. Second, if multiple identified clusters are matched to the same ground-truth cluster following the first rule, only the identified cluster containing the most data points from the ground-truth cluster is accepted. In each trial, we compute the precision as

$\frac{\sum_{i=1}^{K} |C_i \cap f(C_i)|}{\sum_{i=1}^{K} |C_i|}$, where $|C_i|$ is the size of C_i. The recall is computed as $\frac{\sum_{i=1}^{K} |C_i \cap f(C_i)|}{\sum_{j=1}^{N} |G_j|}$.

E2–E4: The computation of precision and recall in **E2–E4** is simpler than **E1**, because there is only one identified cluster C_i. We denote the ground truth cluster as G_j, In each trial, the precision is the ratio of $|C_i \cap G_j|$ to $|C_i|$, and the recall is the ratio of $|C_i \cap G_j|$ to $|G_j|$. In **E3**, We calculated the distance among clusters in the high-dimensional space for the ground truth. As Etemadpour et al. [17], for each pair of clusters, we calculated the average pairwise Euclidean distances of all pairs of points belonging to the two clusters, respectively. Based on the distance among clusters, we randomly selected one cluster as the pre-defined cluster and its nearest cluster as the ground truth. In **E4**, the participants' selection was tested against the ground truth that has the highest density based on the minimum spanning tree-based measure in the high-dimensional space [17].

E5: For the subjective experiment **E5**, we compare the average ranking score of DR techniques in all trials. In each trial, the first one gets 12 points, the second one gets 11, ..., and the last one gets 1 point.

In **E1–E4**, For evaluating the performance of a DR technique in an experiment, we calculate the average precision and recall values on all participants and on all corresponding trials of the experiment. It is worth noting that both precision and recall are important for accurate measure of the quality. On

2D layout / HD data	Single cluster ●	Multiple clusters ●●
Single cluster ●	High precision High recall	High precision Low recall
Multiple clusters ●●	Low precision High recall	High precision High recall

one hand, when participants mistakenly identify unrelated points to the target cluster, the precision decreases. On the other hand, if they correctly identify partial target cluster, the recall decreases. As shown in the confusion matrix, a DR technique with high precision is capable of well separating ground truth clusters. Otherwise, multiple ground truth clusters that mixed in the projection would result in low precision. A low recall indicates that the DR technique cannot gather points together in one ground truth cluster, leading to a peak splitting event [56]. If a DR technique projects multiple ground truth clusters as one cluster in the 2D layout, participants would select all the connection points as one cluster, resulting in low precision and high recall.

6.2 Statistical Analysis Approaches

We utilize the average value and confidence interval of the result measures as indicators to test our hypotheses. Based on a Shapiro-Wilk test, we found that the results of **E1–E4** do not follow a normal distribution. Therefore, we employed a Friedman test to examine any significant differences among the various indicators for the DR techniques. For any resulted significant differences from the test, we further conducted detailed comparisons based on Tukey's HSD test. Here, we consider the standard significance level $\alpha = 0.05$ as a statistical significance for the results. When we take the effect size as 0.25 and the power as 0.95, the power analysis show that we need at least 18 samples in each experiment [40]. In our experiment, we have 60 samples for each trial that is sufficient for the p-value test.

6.3 Result Analysis

6.3.1 Objective Experiment Results

H1.1–H1.3 The results of **E1** are displayed in Fig. 4. Among all techniques, UMAP and t-SNE have the best precision (89.07% and 85.34%, respectively) and recall (88.98% and 84.99%, respectively). UMAP also has the shortest completion time (13.65s). The Friedman tests show statistical significance of precision ($\chi^2(11)$=524.90, p<0.05), recall($\chi^2(11)$=524.89, p<0.05) and average completion time ($\chi^2(11)$=261.82, p<0.05). Fig. 5 shows the pairwise significance relationships for all the DR techniques.

- **H1.1** Local DR techniques perform better than global ones in **T1** if they have the same linearity type. **H1.1** is partially confirmed.
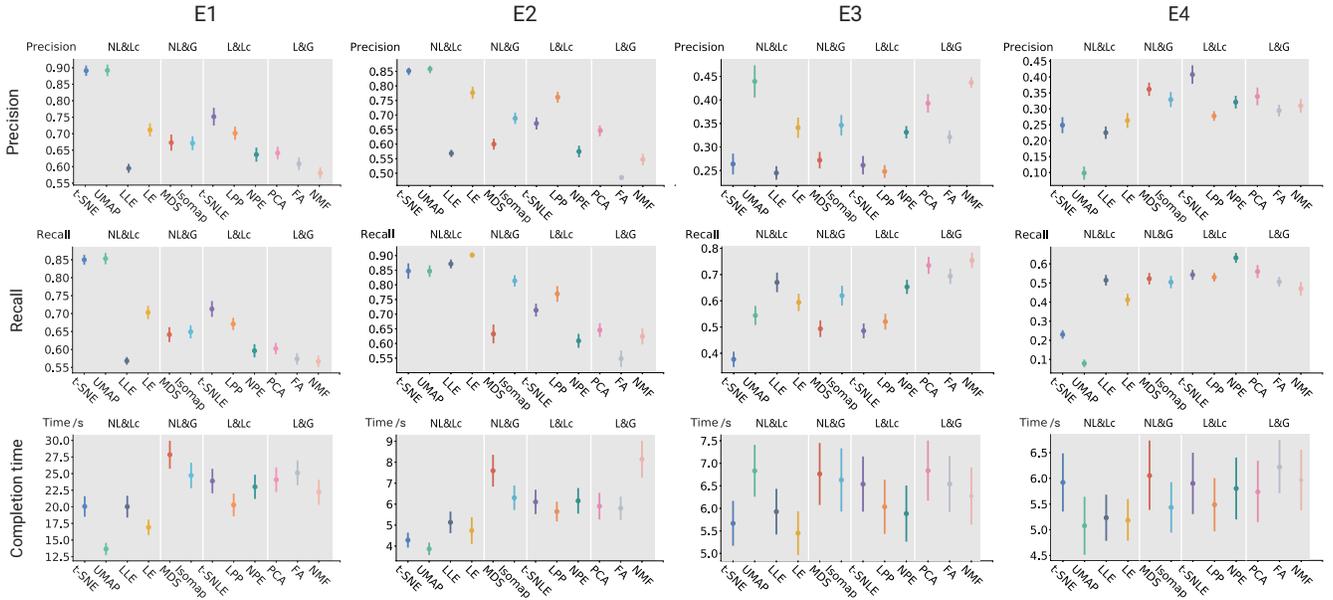
Fig. 4: Precision, Recall and Completion time of **E1–E4**. Error bars indicate 95% confidence intervals.

In the non-linear group (NL&Lc and NL&G), local techniques t-SNE and UMAP perform significantly better than all global techniques in terms of precision and recall. However, the local technique LLE performs significantly worse than global techniques MDS and Isomap. In the linear group (L&Lc and L&G), the local techniques t-SNLE and LPP performs significantly better than all the global techniques in terms of precision and recall. In terms of completion time, local techniques UMAP and LE perform significantly better than all global techniques.

- **H1.2** Non-linear DR techniques perform better than linear ones in **T1** if they have the same locality type. **H1.2** is partially confirmed.
In terms of precision and recall, in the local group (NL&Lc and L&Lc), non-linear techniques t-SNE and UMAP perform significantly better than all linear techniques. However, the non-linear technique LLE performs significantly poorer than linear techniques t-SNLE and LPP. In the global group (NL&G and L&G), non-linear techniques MDS and Isomap perform significantly better than all the linear techniques except for PCA.

- **H1.3** t-SNE and UMAP perform better than other techniques in **T1**. **H1.3** is confirmed.
In **E1**, UMAP and t-SNE perform significantly better than all other techniques in precision and recall. They also perform well in completion time. In terms of completion time, UMAP performs significantly better than other techniques except LE.

H2.1–H2.3 As shown in Fig. 4, UMAP has the highest precision (85.81%) and shortest completion time (3.87s) in **E2**. t-SNE is the second in precision(85.14%) and completion time (4.29s). The Friedman tests show significance among different DR techniques in terms of precision ($\chi^2(11)$=534.79,p<0.05), recall ($\chi^2(11)$=506.43,p<0.05) and average completion time($\chi^2(11)$=297.76,p<0.05). The pairwise significance relationships between each pair of DR techniques can be seen in Fig. 5.

- **H2.1** Local DR techniques perform better than global ones in **T2** if they have the same linearity type. **H2.1** is partially confirmed.
In the non-linear group (NL&Lc and NL&G), local techniques t-SNE and UMAP perform significantly better than all global techniques in terms of precision. However, the local technique LLE performs significantly poorer than the global technique Isomap. In terms of recall, local techniques perform significantly better than all global techniques except Isomap. In the linear group (L&Lc and L&G), in terms of precision and recall, local techniques t-SNLE and LPP perform significantly better than all global techniques except PCA. In terms of completion time, local techniques t-SNE and UMAP also perform significantly better than all global techniques.

- **H2.2** Non-linear DR techniques perform better than linear ones in **T2** if they have the same locality. **H2.2** is partially confirmed.
In the local group (NL&Lc and L&lc), non-linear techniques t-SNE and UMAP perform significantly better than all linear techniques in terms of precision and recall. However, the non-linear technique LLE performs significantly poorer than the linear technique t-SNLE in precision. In the global group (NL&G and L&G), the non-linear technique Isomap performs significantly better than all linear techniques. In terms of completion time, non-linear techniques t-SNE and UMAP perform significantly better than all techniques except LPP.

- **H2.3** t-SNE and UMAP perform better than other techniques in **T2**. **H2.3** is partially confirmed.
In **E2**, UMAP and t-SNE perform significantly better than all linear techniques in precision. Though they have lower recall than LE and LLE, there is no significant difference between the four techniques. In terms of completion time, UMAP and t-SNE also perform significantly better than the other techniques except LLE, LE and LPP.

H3.1–H3.3 The results of **E3** are displayed in Fig. 4. UMAP has the highest precision (43.94%) but ranked 8th in recall(54.42%). NMF has the second-highest precision(43.69%) and highest recall(75.42%). The Friedman tests show statistical significance among different DR techniques in terms of precision ($\chi^2(11)$=534.79,p<0.05) and recall ($\chi^2(11)$=506.43,p<0.05). However, Tukey's HSD test shows there are no pairwise significance relationships exist between each pair of DR techniques in terms of completion time. The pairwise significance relationships between each pair of DR techniques can be seen in Fig. 6.

- **H3.1** Global DR techniques perform better than local ones in **T3** when they have the same linearity. **H3.1** is rejected.
In the non-linear group (NL&Lc and NL&G), the local technique UMAP performs significantly better than all global techniques in terms of precision. In terms of recall, global techniques MDS and Isomap only significantly better than local techniques t-SNE. In the linear group (L&Lc and L&G), all global techniques perform significantly better than all local techniques except NPE in terms of precision and recall.

- **H3.2** Linear DR techniques perform better than non-linear ones in **T3** when they have the same locality. **H3.2** is rejected.
In the local group (NL&Lc and L&Lc), the non-linear technique UMAP performs significantly better than all linear techniques in terms of precision. The linear technique NPE performs significantly better than non-linear techniques t-SNE and LLE in precision. In the global group (NL&G and L&G), linear techniques PCA, FA, and NMF perform significantly better than the non-linear technique MDS in terms of precision and recall.
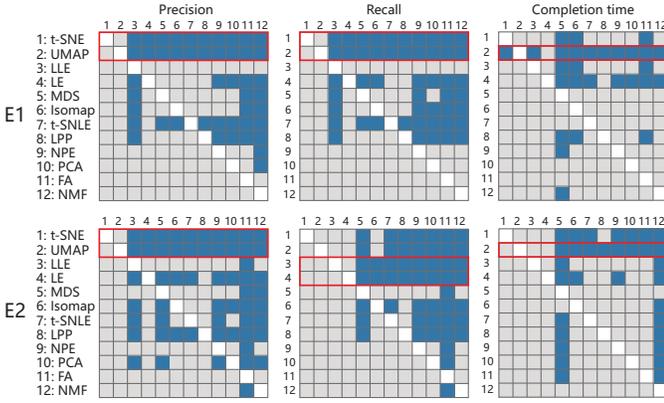
Fig. 5: Matrix depiction of the pairwise significance relationships of the precision, recall and completion time differences of the projection techniques in **E1–E2**. A blue cell in row i and column j denotes that the technique with ID i performs significantly better than the technique with ID j. The red rectangles denote the best DR techniques.

- **H3.3** PCA performs better than other techniques in **T3**. **H3.3** is rejected.

In **E3**, PCA ranks as the third technique in precision (39.27%) and the second technique in recall (73.5%). UMAP performs significantly better than PCA in precision.

H4.1–**H4.3** As shown in Fig. 4, t-SNLE has the highest precision (40.77%) and the third recall(54.29%) in **E4**. NPE has the fifth precision(32.11%) but the highest recall(63.15%). The Friedman tests show statistical significance among different DR techniques in terms of precision ($\chi^2(11)$=247.55,p<0.05) and recall ($\chi^2(11)$=352.34.43,p<0.05). However, Tukey's HSD test shows there are no pairwise significance relationships exist between each pair of DR techniques in terms of completion time. The pairwise significance relationships between each pair of DR techniques can be seen in Fig. 6.

- **H4.1** Global techniques perform better than local ones in **T4** if they have the same linearity. **H4.1** is rejected.

In the non-linear group (NL&Lc and NL&G), global techniques MDS and Isomap perform significantly better than all local techniques in terms of precision and recall. In the linear group (L&Lc and L&G), the local technique t-SNLE performs significantly better than all global techniques in terms of precision. In terms of recall, the local technique NPE performs significantly better than all global techniques.

- **H4.2** Linear techniques perform better than non-linear ones in **T4** if they have the same locality. **H4.2** is partially confirmed.

For the local group (NL&Lc and L&lc), in terms of precision, linear techniques t-SNLE and NPE perform significantly better than all non-linear techniques. In terms of recall, t-SNLE, LPP, and NPE also perform significantly better than all non-linear techniques except LLE. However, in the global group (NL&G and L&G), the non-linear technique MDS performs significantly better than the linear technique FA in terms of precision.

- **H4.3** PCA performs better than other techniques in **T4**. **H4.3** is rejected.

In **E4**, PCA ranks as the third technique in precision (33.91%) and the second technique in recall (55.96%). t-SNLE and NPE perform significantly better than PCA in precision and recall, respectively.

6.3.2 Subjective Experiment Results

The results of **E5** are shown in Fig. 7. The rankings of the most preferred techniques are: UMAP (9.26), t-SNE (8.60), LE (7.07), LPP (6.74), LLE (6.38), t-SNLE (6.26), Isomap (6.15), NPE (6.00), FA (5.95), PCA (5.86), MDS (5.16), NMF (4.53). Local techniques obtain the top four preferred techniques with UMAP and t-SNE ranked the
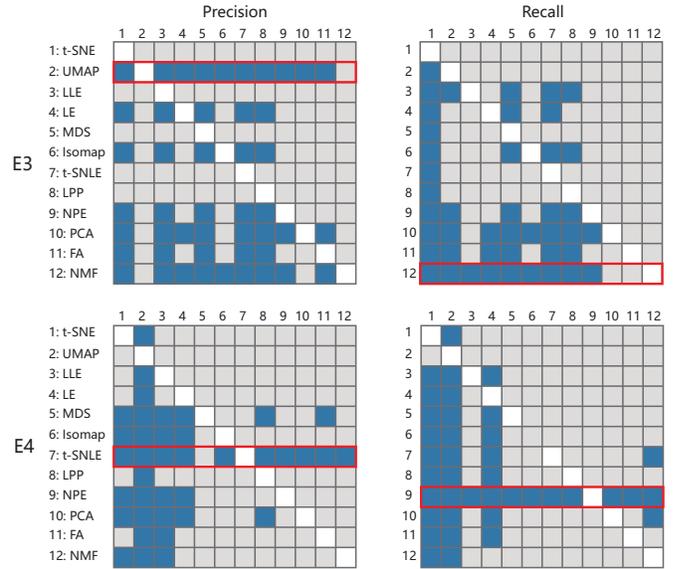
Fig. 6: Matrix depiction of the pairwise significance relationships of the precision and recall differences of the projection techniques in **E3–E4**. A blue cell in row i and column j denotes that the technique with ID i performs significantly better than the technique with ID j. The red rectangles denote the best DR techniques.
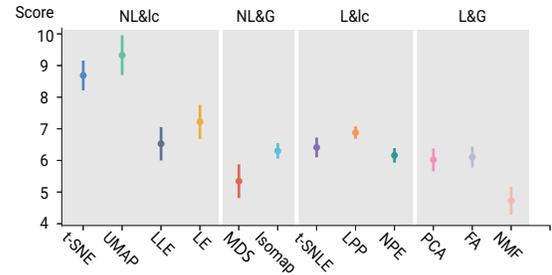
Fig. 7: Average ranking scores of **E5** (the higher the better). Error bars indicate 95% confidence intervals.

highest positions. It suggests a strong preference for the local techniques from the participants. Regarding linearity, LPP is ranked as the 4th preferred techniques. It indicates that there is no strong tendency toward a particular linearity type by the participants. It is worth noting that there could be misunderstandings in interpreting the projection results. For example, if a single cluster in the data space was scattered into two clusters by the projection, the participants could view it as a correct visual cluster. Similar to Lewis et al. [31], we made a correlation analysis between the results of the subjective experiment and the results of objective experiments. The correlation coefficients between the rankings of E5 and E1 are 0.73 (precision) and 0.76 (recall). The correlation confirms the consistency and differences between the subjective experiment and objective experiments.

6.4 Summary

Our experiment results indicate that there is no universally better DR technique and high-level property. Therefore, it is essential to choose appropriate techniques for enhancing particular tasks involving visual cluster analysis. (1) In cluster identification, local DR techniques perform better than global DR techniques in cluster identification if they have the same linearity type, except LLE; non-linear DR techniques perform better than linear DR techniques if they have the same locality type, except LLE; t-SNE and UMAP have an outstanding performance. (2) In membership identification, local DR techniques perform better than global DR techniques if they have the same linearity type, except

	E1P	E1R	E1T	E2P	E2R	E2T	E3P	E3R	E3T	E4P	E4R	E4T
t-SNE	0.05	0.12	0.33	0.14	0.18	0.15	0.39	0.36	0.36	0.46	0.31	0.36
UMAP	0.03	0.13	0.15	0.14	0.17	0.1	0.35	0.32	0.34	0.12	0.05	0.39
LE	0.26	0.27	0.26	0.23	0.15	0.28	0.41	0.38	0.19	0.36	0.31	0.29
LLE	0.31	0.29	0.38	0.41	0.35	0.28	0.34	0.38	0.16	0.25	0.37	0.32
MDS	0.28	0.26	0.16	0.31	0.34	0.27	0.27	0.31	0.24	0.37	0.35	0.35
Isomap	0.26	0.26	0.28	0.32	0.37	0.34	0.35	0.36	0.31	0.42	0.39	0.24
t-SNLE	0.35	0.34	0.32	0.42	0.34	0.27	0.35	0.39	0.39	0.35	0.36	0.37
LPP	0.15	0.12	0.26	0.24	0.35	0.22	0.27	0.36	0.19	0.35	0.44	0.36
NPE	0.32	0.33	0.15	0.35	0.44	0.27	0.38	0.35	0.24	0.3	0.36	0.27
PCA	0.23	0.25	0.34	0.32	0.37	0.26	0.33	0.35	0.37	0.34	0.32	0.35
FA	0.3	0.3	0.34	0.42	0.43	0.26	0.3	0.35	0.32	0.36	0.36	0.2
NMF	0.27	0.25	0.19	0.26	0.32	0.25	0.37	0.26	0.23	0.34	0.27	0.31

Table 5: The homogeneity of eight datasets. E#P, E#R, and E#T represent the normalized standard deviations of precision, recall, and completion time in E# in terms of eight datasets, respectively.
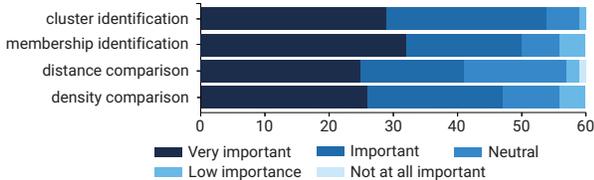
Fig. 8: The importance rating of the four visual cluster analysis tasks evaluated in our experiments.

LLE; non-linear DR techniques perform better than linear DR techniques if they have the same locality type, except LLE; t-SNE and UMAP have an outstanding performance. (3) In density comparison, linear DR techniques perform better than non-linear DR techniques if they have the same locality type, except MDS; t-SNLE and NPE are the preffered techniques. (4) In distance comparison, NMF and UMAP are the preffered techniques. The experiment results shed light on DR techniques that received less attentions, such as NPE, t-SNLE.

7 DISCUSSION

Visual cluster analysis tasks. The participants answered the subjective questionnaire of the importance of the four analysis tasks after completing the experiments. As Fig. 8 shows, the average scores of the four tasks are cluster identification (4.36), membership identification (4.30), density comparison (4.15), and cluster distance comparison (4.03), respectively. These scores once again confirm that these four tasks are common concerns when using DR techniques for visual cluster analysis. The experiment results also show that overall DR techniques perform better in E1 and E2 than in E3 and E4 in terms of precision and recall. The average precision of all DR techniques in E1–E4 are 0.69, 0.67, 0.32, and 0.29, respectively.The average recall of them are 0.67, 0.74, 0.60, and 0.46, respectively. It suggests a research opportunity to design DR techniques that support E3 and E4 better.

The homogeneity of datasets. We check the homogeneity of eight datasets by comparing the normalized standard deviations of precision, recall, and time for each task and each DR technique. As Table 5 shows, the normalized standard deviations are ranged from 0.03 to 0.46 and present a reasonable homogeneity of datasets. Another observation is that t-SNE and UMAP that perform better than other DR techniques in most tasks, also present low normalized standard deviations. Therefore, t-SNE and UMAP are recommended in the aspects of performance and stability on datasets with different characteristics.

Cluster identification and class separation. In this study, our evaluation is based on the perception of cluster identification. Participants were provided with monochrome scatterplots without color-coded class information. It is worth noting that we have also used the class label information with the datasets as the ground truth of clusters. Therefore, in selecting the datasets, we check the visual separation of clusters with the visual class separation measure. However, when a class is scattered into several clusters, the class separation measure also presents a good but misleading result. This issue can be addressed by additional user controlled selection with colored scatterplot. More discussion on this issue is referred to Aupetit [3].

Parameter robustness of DR techniques. Our pre-study provides an opportunity to analyze the DR techniques in the aspect of parameter robustness. In our pre-study, participants were asked to select the result that best fits the structure of datasets among five settings. If the selections are uniformly distributed in all five settings, we can infer that the tuning of parameters has few effects on the quality of projection. Otherwise, a big variation of the setting distribution indicates that the DR technique is sensitive to the parameter. We calculated the average standard deviation across datasets of each tested eight DR techniques in the pre-study. The rankings are: t-SNLE (3.25), t-SNE (3.15), LLE (3.11), UMAP (3.04), LE (2.57), NPE (2.39), Isomap (2.35), and LPP (1.85). The results show that although t-SNE is one of the most popular techniques, its sensitivity to the *perplexity* parameter should be considered. This observation is the same as Espadoto et al. [16]. It is also worth noting that UMAP is less sensitive than t-SNE while having similar performance in most tasks.

Conducting remote studies. Remote study lacks a face-to-face communication environment and yields several challenges for a control study. First, the test should be performed independently. In our study, participants are required to set an independent environment to avoid possible disturbance. Second, the test could be interrupted due to network issues. We ask participants to test their network to ensure a stable connection. During the test, if a participant is disconnected for more than ten minutes, we discard this sample to ensure the validity of tests. Third, the remote communication would be limited in a remote study. To address this issue, we monitor the experimental screen of participants through screen sharing throughout the test. We also communicate with participant during the test in voice.

Order of experiments and trials. A limitation of our study is the order of experiments and trials. Because E3 is conducted before E4, it may reveal a part of ground truth clusters and make a potential threat to validity of E4. If the given cluster in E3 happens to be the cluster with the highest density in E4: 1) when several classes overlap in a single cluster, then an increase of precision would occur; 2) when a single class is scattered in several clusters, then an increase of recall would occur. In addition, learning effects should be considered. A Latin square design could address this issue.

Limitations and future work. Due to the different characteristics of datasets, DR techniques will inevitably produce different performance in different types of datasets. Therefore, the experimental results may have a certain correlation with the datasets. In this study, we focus on the characteristics of locality and linearity and four objective visual clustering tasks. In the future, it is valuable to perform further studies on other characteristics (for example, false and missed neighbors could have an impact on clustering tasks that is worth evaluating [38]) and tasks. The linear-nonlinear pairs of techniques, such as t-SNE and t-SNLE, also shed light on the development of DR techniques and design of comparative analysis.

8 CONCLUSION

In this paper, we present an empirical evaluation of DR techniques from the perspective of perception in visual cluster analysis tasks. Twelve representative DR techniques are identified from a literature review. They are grouped into groups based on linearity and locality. We first conduct a pre-study to determine the proper input parameters for each DR technique. Then, the techniques are formally evaluated in four controlled experiments each of which focuses on hypotheses on a particular analytical task. A subjective task is also conducted to collect users' preferences on each technique. Several guidelines and interesting insights are provided by the result analysis.

REFERENCES

[1] M. M. Abbas, M. Aupetit, M. Sedlmair, and H. Bensmail. Clustme: A visual quality measure for ranking monochrome scatterplots based on cluster patterns. *Computer Graphics Forum*, 38(3):225–236, 2019.

[2] S. Arora, W. Hu, and P. K. Kothari. An analysis of the t-sne algorithm for data visualization. In *Proceedings of Conference on Learning Theory*, pp. 1455–1462, 2018.

[3] M. Aupetit. Sanity check for class-coloring-based evaluation of dimension reduction techniques. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, pp. 134–141, 2014.

[4] M. Aupetit, P. Couturier, and P. Massotte. γ-observable neighbours for vector quantization. *Neural networks*, 15(8-9):1017–1027, 2002.

[5] M. Aupetit and M. Sedlmair. Sepme: 2002 new visual separation measures. In *Proceedings of IEEE Pacific Visualization Symposium*, pp. 1–8, 2016.

[6] M. Aupetit, M. Sedlmair, M. M. Abbas, A. Baggag, and H. Bensmail. Toward perception-based evaluation of clustering techniques for visual analytics. In *Proceedings of IEEE Visualization Conference*, pp. 141–145, 2019.

[7] E. Becht, L. McInnes, J. Healy, C.-A. Dutertre, I. W. Kwok, L. G. Ng, F. Ginhoux, and E. W. Newell. Dimensionality reduction for visualizing single-cell data using umap. *Nature biotechnology*, 37(1):38–44, 2019.

[8] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Proceedings of Conference on Neural Information Processing Systems*, pp. 585–591, 2002.

[9] K. Bunte, M. Biehl, and B. Hammer. A general framework for dimensionality-reducing data visualization mapping. *Neural Computation*, 24(3):771–804, 2012.

[10] D. Cai, X. He, X. Wu, and J. Han. Non-negative matrix factorization on manifold. In *Proceedings of IEEE International Conference on Data Mining*, pp. 63–72, 2008.

[11] A. Chatzimparmpas, 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.

[12] S. Cheng and K. Mueller. The data context map: Fusing data and attributes into a unified display. *IEEE transactions on visualization and computer graphics*, 22(1):121–130, 2015.

[13] V. De Silva and J. B. Tenenbaum. Global versus local methods in non-linear dimensionality reduction. In *Proceedings of Neural Information Processing Systems*, pp. 721–728, 2003.

[14] C. Ding and X. He. K-means clustering via principal component analysis. In *Proceedings of International Conference on Machine Learning*, p. 29, 2004.

[15] D. Dua and C. Graff. UCI machine learning repository, 2017.

[16] M. Espadoto, R. M. Martins, A. Kerren, N. S. Hirata, and A. C. Telea. Towards a quantitative survey of dimension reduction techniques. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2153–2173, 2021.

[17] R. Etemadpour, R. Motta, J. G. de Souza Paiva, R. Minghim, M. C. F. de Oliveira, and L. Linsen. Perception-based evaluation of projection methods for multidimensional data visualization. *IEEE Transactions on Visualization and Computer Graphics*, 21(1):81–94, 2014.

[18] M. Forina, C. Armanino, S. Lanteri, and E. Tiscornia. Classification of olive oils from their fatty acid composition. In *Proceedings of Food research and data analysis*, pp. 189–214, 1983.

[19] F. J. García Fernández, M. Verleysen, J. A. Lee, and I. Díaz Blanco. Stability comparison of dimensionality reduction techniques attending to data and parameter variations. In *Proceedings of Eurographics Conference on Visualization*, 2013.

[20] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE transactions on pattern analysis and machine intelligence*, 23(6):643–660, 2001.

[21] Q. Han, D. Thom, M. John, S. Koch, F. Heimerl, and T. Ertl. Visual quality guidance for document exploration with focus+ context techniques. *IEEE transactions on visualization and computer graphics*, 26(8):2715–2731, 2019.

[22] X. He, D. Cai, S. Yan, and H.-J. Zhang. Neighborhood preserving embedding. In *Proceedings of International Conference on Computer Vision*, pp. 1208–1213, 2005.

[23] X. He and P. Niyogi. Locality preserving projections. In *Proceedings of Neural Information Processing Systems*, pp. 153–160, 2004.

[24] S. Ingram, T. Munzner, and M. Olano. Glimmer: Multilevel mds on the gpu. *IEEE Transactions on Visualization and Computer Graphics*, 15(2):249–261, 2008.

[25] P. Joia, D. Coimbra, J. A. Cuminato, F. V. Paulovich, and L. G. Nonato. Local affine multidimensional projection. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2563–2571, 2011.

[26] I. T. Jolliffe. Principal components in regression analysis. In *Principal component analysis*, pp. 129–155. Springer, 1986.

[27] Y. Koren and L. Carmel. Visualization of labeled data using linear transformations. In *Proceedings of IEEE Symposium on Information Visualization*, pp. 121–128, 2003.

[28] J. B. Kruskal and M. Wish. *Multidimensional scaling*, vol. 11. Sage, 1978.

[29] D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

[30] J. Lewis, M. Ackerman, and V. de Sa. Human cluster evaluation and formal quality measures: A comparative study. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, pp. 1870–1875, 2012.

[31] J. Lewis, L. Van der Maaten, and V. de Sa. A behavioral investigation of dimensionality reduction. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, pp. 671—676, 2012.

[32] G. C. Linderman and S. Steinerberger. Clustering with t-sne, provably. *SIAM Journal on Mathematics of Data Science*, 1(2):313–332, 2019.

[33] M. Liu, J. Shi, Z. Li, C. Li, J. Zhu, and S. Liu. Towards better analysis of deep convolutional neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):91–100, 2017.

[34] P. Mair. Factor analysis. In *Modern Psychometrics with R*, pp. 17–61. Springer, 2018.

[35] R. M. Martins, D. B. Coimbra, R. Minghim, and A. C. Telea. Visual analysis of dimensionality reduction quality for parameterized projections. *Computers & Graphics*, 41:26–42, 2014.

[36] L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

[37] A. Narayan, B. Berger, and H. Cho. Assessing single-cell transcriptomic variability through density-preserving data visualization. *Nature Biotechnology*, pp. 1–10, 2021.

[38] L. G. Nonato and M. Aupetit. Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE Transactions on Visualization and Computer Graphics*, 25(8):2650–2673, 2018.

[39] F. V. Paulovich, L. G. Nonato, R. Minghim, and H. Levkowitz. Least square projection: A fast high-precision multidimensional projection technique and its application to document mapping. *IEEE Transactions on Visualization and Computer Graphics*, 14(3):564–575, 2008.

[40] B. Prajapati, M. Dunne, and R. Armstrong. Sample size estimation and statistical power analyses. *Optometry today*, 16(7):10–18, 2010.

[41] P. J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

[42] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000.

[43] T. Schreck, T. Von Landesberger, and S. Bremm. Techniques for precision-based visual analysis of projected data. *Information Visualization*, 9(3):181–193, 2010.

[44] M. Sedlmair, A. Tatu, T. Munzner, and M. Tory. A taxonomy of visual cluster separation factors. *Computer Graphics Forum*, 31(3pt4):1335–1344, 2012.

[45] D. Spathis, N. Passalis, and A. Tefas. Fast, visual and interactive semi-supervised dimensionality reduction. In *Proceedings of the European Conference on Computer Vision*, pp. 550–563, 2018.

[46] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323, 2000.

[47] L. Van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(86):2579–2605, 2008.

[48] L. Van Der Maaten, E. Postma, and J. Van den Herik. Dimensionality reduction: a comparative review. *J Mach Learn Res*, 10(66-71):13, 2009.

[49] E. Ventocilla and M. Riveiro. A comparative user study of visualization techniques for cluster analysis of multidimensional data sets. *Information Visualization*, 19(4):318–338, 2020.

[50] E. F. Vernier, R. Garcia, I. d. Silva, J. L. D. Comba, and A. C. Telea. Quantitative evaluation of time-dependent multidimensional projection techniques. In *Computer Graphics Forum*, pp. 241–252, 2020.

[51] Y. Wang, K. Feng, X. Chu, J. Zhang, C.-W. Fu, M. Sedlmair, X. Yu, and B. Chen. A perception-driven approach to supervised dimensionality reduction for visualization. *IEEE transactions on visualization and computer graphics*, 24(5):1828–1840, 2017.

[52] M. Wattenberg, F. Viégas, and I. Johnson. How to use t-sne effectively. *Distill*, 1(10):e2, 2016.

[53] J. Wenskovitch, I. Crandell, N. Ramakrishnan, L. House, and C. North. Towards a systematic combination of dimension reduction and clustering in visual analytics. *IEEE transactions on visualization and computer graphics*, 24(1):131–141, 2017.

[54] L. Wilkinson and G. Wills. Scagnostics distributions. *Journal of Computational and Graphical Statistics*, 17(2):473–491, 2008.

[55] S. Wold, K. Esbensen, and P. Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.

[56] S. Xiang, X. Ye, J. Xia, J. Wu, Y. Chen, and S. Liu. Interactive correction of mislabeled training data. In *Proceedings of IEEE Conference on Visual Analytics Science and Technology*, pp. 57–68, 2019.

[57] K. Xu, L. Zhang, D. Pérez, P. H. Nguyen, and A. Ogilvie-Smith. Evaluating interactive visualization of multidimensional data projection with feature transformation. *Multimodal Technologies and Interaction*, 1(3):13, 2017.

[58] J. Yuan, C. Chen, W. Yang, M. Liu, J. Xia, and S. Liu. A survey of visual analytics techniques for machine learning. *Computational Visual Media*, 7(1):3–36, 2021.

[59] J. Yuan, S. Xiang, J. Xia, L. Yu, and S. Liu. Evaluation of sampling methods for scatterplots. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1720–1730, 2020.

[60] X. Yue, J. Bai, Q. Liu, Y. Tang, A. Puri, K. Li, and H. Qu. sPortfolio: Stratified visual analysis of stock portfolios. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):601–610, 2019.

[61] Y. Zhao, F. Luo, M. Chen, Y. Wang, J. Xia, F. Zhou, Y. Wang, Y. Chen, and W. Chen. Evaluating multi-dimensional visualizations for understanding fuzzy clusters. *IEEE transactions on visualization and computer graphics*, 25(1):12–21, 2018.