

# MAPLE: Self-Supervised Learning-Enhanced Nonlinear Dimensionality Reduction for Visual Analysis

Zeyang Huang, Takanori Fujiwara, Angelos Chatzimparmpas, Wandrille Duchemin, and Andreas Kerren

**Abstract**—We present a new nonlinear dimensionality reduction method, *MAPLE*, that enhances UMAP by improving manifold modeling. *MAPLE* employs a self-supervised learning approach to more efficiently encode low-dimensional manifold geometry. Central to this approach are *maximum manifold capacity representations (MMCRs)*, which help untangle complex manifolds by compressing variances among locally similar data points while amplifying variance among dissimilar data points. This design is particularly effective for high-dimensional data with substantial intra-cluster variance and curved manifold structures, such as biological or image data. Our qualitative and quantitative evaluations demonstrate that *MAPLE* can produce clearer visual cluster separations and finer subcluster resolution than UMAP while maintaining a tractable computational cost.

**Index Terms**—Visualization, dimensionality reduction, manifold learning, embedding algorithms, UMAP

## I. INTRODUCTION

**H**IGH-DIMENSIONAL datasets are rapidly growing across scientific domains, from single-cell transcriptomics to large-scale image collections. Advances in genomics, machine learning, and related fields often produce measurements with tens of thousands of features, such as genes, pixels, or descriptors, that naturally place the data in very high dimensions. As humans cannot directly see high-dimensional spaces, effective visualization is essential to analyze these datasets. Dimensionality reduction (DR) methods make high-dimensional data visually explorable by projecting them into a low-dimensional space, typically 2D or 3D. Effective DR results can facilitate data annotation, reveal rare populations, and lead to the generation of new hypotheses.

A wide range of DR methods has been proposed for visualization [24], [46]. Linear DR methods, including principal component analysis (PCA) [39], remain popular for their interpretability [20], [29], but are limited in capturing complex, curved data structures. Nonlinear DR methods address this limitation by prioritizing local neighborhood relationships, allowing them to better represent data lying on complex manifolds [77]. Among these, uniform manifold approximation

and projection (UMAP) [37], [55] has become a method of choice across many domains, alongside alternatives such as *t*-SNE [76]. UMAP is valued for its scalability, graph-based formulation, and ability to capture nonlinear structure without strong parametric assumptions. UMAP constructs a weighted *k*-nearest neighbor (*k*NN) graph in the input space to model an underlying manifold and then optimizes a low-dimensional layout so that the structure of the layout aligns with the graph according to a cross-entropy objective. This procedure is computationally efficient, and its cross-entropy objective often yields layouts with clear cluster separation.

The quality of the resulting UMAP layout is strongly influenced by the quality of the weighted *k*NN graph. In practice, this graph may include connections that bring semantically distinct points close together in the low-dimensional layout, and vice versa. We argue that a key limitation of UMAP lies not in the *k*NN algorithm itself—which can serve as an effective proxy for nonlinear structure—but in how the graph is constructed, specifically the choice of data representation and distance space used to define neighborhood relationships. More specifically, standard *k*NN construction based on raw high-dimensional data and predefined distance metrics faces the following two issues. First, as dimensionality increases, pairwise distances in high-dimensional space (e.g., Euclidean distances) tend to concentrate toward similar values [28], making *k*NN graphs increasingly sensitive to noise. Second, many datasets lie on curved manifolds [50], where points that appear as neighbors in the *k*NN graph may actually belong to distinct semantic regions [19]. In 2D or 3D layouts, these effects result in local clusters with spurious overlaps, thereby mixing groups that are semantically distinct. Recent theoretical work has also analyzed these issues of UMAP [11], [22], emphasizing its limitations in capturing subtle geometric differences.

To address the above fundamental challenges, we introduce a novel nonlinear DR method, *MAPLE* (short for manifold-aware projection with learned edges). A central idea in *MAPLE* is that explicitly learning neighborhood relationships can substantially improve the quality of low-dimensional layouts. *MAPLE* refines neighborhood definitions by adapting multi-view self-supervised learning (MVSSL) [65], [75] to the context of nonlinear DR. Unlike standard methods that treat a weighted *k*NN graph directly constructed from raw high-dimensional data as a fixed representation of the underlying manifold, *MAPLE* treats the initial neighbor relationships as weak, tentative self-supervised signals. We employ a learning

Zeyang Huang is with Linköping University, Sweden. E-mail: zeyang.huang@liu.se.

Takanori Fujiwara is with the University of Arizona, United States. E-mail: tfujiwara@arizona.edu.

Angelos Chatzimparmpas is with Utrecht University, Netherlands. E-mail: a.chatzimparmpas@uu.nl.

Wandrille Duchemin is with University of Basel, Switzerland. E-mail: wandrille.duchemin@unibas.ch.

Andreas Kerren is with Linköping University and Linnaeus University, Sweden. E-mail: andreas.kerren@{liu/lnu}.se.

phase to iteratively refine these signals. Instead of directly operating in the input high-dimensional space, we employ encoder and projector networks [33], [48], [88] to map data into an embedding space where the distance relationships are learned rather than fixed by a predefined distance metric. A refined graph is then constructed based on these optimized distance relationships. The core of this learning phase is the *maximum manifold capacity representation (MMCR)* objective [87], which compresses variance among locally similar points (i.e., neighbors) while amplifying variance across dissimilar ones (i.e., non-neighbors). The outcome of this learning phase is a refined, weighted  $k$ NN graph that is subsequently laid out in a low-dimensional space using UMAP’s cross-entropy optimization.

We evaluate MAPLE against multiple DR methods using a set of quality measures. Additionally, we qualitatively evaluate MAPLE through visual analysis on image classification datasets. We further demonstrate MAPLE’s utility in case studies on single-cell transcriptomics in collaboration with a bioinformatics expert (a co-author of this work). Our results show that MAPLE reveals fine-grained, meaningful patterns that support deeper downstream analysis.

Our contributions are as follows:

- MAPLE, a new nonlinear DR method that learns neighborhood structure through MVSSL;
- the adaptation of MMCRs to graph-based DR to untangle a complex underlying manifold through compression and diversification of local manifolds; and
- a set of quantitative metrics and qualitative evaluations using multiple datasets across domains, showing that MAPLE produces high-quality layouts.

**Terminology.** MAPLE produces two types of low-dimensional representations. To avoid ambiguity, we use the following terminology throughout this paper:

*Embedding:* an intermediate, relatively low-dimensional representation produced before the layout optimization. In MAPLE, embeddings are the outputs of the encoder–projector networks, and the weighted  $k$ NN graph is constructed and refined in this embedding space.

*Layout:* the final low-dimensional representation (typically 2D or 3D) obtained after the layout optimization phase of DR.

The remainder of this paper is organized as follows: Sec. II discusses background and related work. Sec. III introduces MAPLE. Sec. IV presents our experimental results, while Sec. V presents case studies with single-cell datasets. Sec. VI presents the limitations of our approach and opportunities for future work. In Sec. VII, we conclude our paper.

## II. BACKGROUND AND RELATED WORK

DR aims to extract low-dimensional representations from high-dimensional data while preserving meaningful structure. Nonlinear DR methods [77] are designed to capture data patterns lying on nonlinear manifolds (e.g., curved surfaces). From the perspective of manifold modeling, nonlinear DR methods can be categorized by how they represent neighborhood structure. Graph-based modeling constructs a  $k$ NN

graph or a similar structure to capture local proximities along a manifold. This category includes early spectral and geodesic methods, such as ISOMAP [72] and Laplacian eigenmaps [6], as well as more recent methods, including LargeVis [71],  $t$ -SNE [76], and UMAP [55]. Diffusion-based modeling (e.g., spectral embeddings [56] and PHATE [58]) simulates diffusion processes (e.g., random walks) on the data and uses transition probabilities as similarities to preserve manifold structure. Kernel-based modeling [32], [57] employs kernel functions to implicitly map data into higher-dimensional feature spaces, enabling the modeling of complex shapes. Our new method, MAPLE, belongs to the graph-based family. We specifically aim to advance graph-based modeling, using UMAP (a widely used state-of-the-art method) as our primary baseline.

### A. UMAP

UMAP has been widely used in biology [5], computational social science [31], and other data-intensive fields [24], due to its scalability, visually appealing layouts, and robust behavior across disciplines. The key idea of UMAP is to approximate the manifold structure of high-dimensional data by first constructing a weighted  $k$ NN graph, called a fuzzy graph, in the input space and then computing a low-dimensional layout that preserves the graph’s structure. McInnes et al. [55] describe these two phases as (1) **graph construction** and (2) **graph layout**. MAPLE follows the same two-phase structure, while it refines the graph construction phase using MVSSL with MMCRs. For clarity, we next provide a concise formal description of UMAP’s two phases, which is necessary to highlight MAPLE’s unique contributions.

**Phase 1: Graph construction.** Let  $\mathbf{X} = [\mathbf{x}_1 \cdots \mathbf{x}_N]^T \in \mathbb{R}^{N \times D}$  denote the original input data, where  $N$  and  $D$  are the numbers of data points and dimensions, respectively. In this phase, UMAP first constructs a  $k$ NN graph,  $G$ , from  $\mathbf{X}$  based on a chosen distance function  $\text{dist}(\cdot)$  (Euclidean distance by default). UMAP then converts this graph into a fuzzy graph by assigning edge weights as follows:

$$w_{ij} = \exp(-\max(0, \text{dist}(\mathbf{x}_i, \mathbf{x}_j) - \rho_i) / \sigma_i) \quad (1)$$

where  $\rho_i$  and  $\sigma_i$  are local connectivity parameters automatically determined by UMAP to distribute data points uniformly over the manifold. The weight,  $w_{ij}$ , reflects the strength of neighborhood relationships between points  $i$  and  $j$  (precisely, the membership strength of a 1-simplex between them [55]).

**Phase 2: Graph layout.** Let  $\mathbf{Y} = [\mathbf{y}_1 \cdots \mathbf{y}_N]^T \in \mathbb{R}^{N \times M}$  denote the layout positions in an  $M$ -dimensional space (usually,  $M = 2$  or  $3$ ). UMAP initializes  $\mathbf{Y}$  using a certain strategy, such as random initialization, PCA of  $\mathbf{X}$ , or spectral embedding of  $\mathbf{X}$ —with spectral embedding as the default. Next, the edge weight between two points in the layout space is defined as:  $v_{ij} = (1 + a \|\mathbf{y}_i - \mathbf{y}_j\|^{2b})^{-1}$ , where  $a$  and  $b$  are user-specified hyperparameters controlling how tightly neighbors are placed together. In practice,  $a$  and  $b$  are derived from a single hyperparameter,  $\text{min\_dist}$ . UMAP then optimizes  $\mathbf{Y}$  to minimize the cross-entropy loss between  $w_{ij}$  and  $v_{ij}$  through:

$$\arg \min_{\mathbf{Y}} \sum_{i \neq j} \left[ w_{ij} \log \left( \frac{w_{ij}}{v_{ij}} \right) + (1 - w_{ij}) \log \left( \frac{1 - w_{ij}}{1 - v_{ij}} \right) \right]. \quad (2)$$

To perform this optimization, UMAP first symmetrizes the fuzzy graph as  $w_{ij} \leftarrow w_{ij} + w_{ji} - w_{ij}w_{ji}$  and then applies negative sampling with stochastic gradient descent.

**Limitations of the graph construction phase.** The cross-entropy loss encourages the similarity in the layout ( $v_{ij}$ ) to be high when the high-dimensional weight ( $w_{ij}$ ) is large (corresponding to an attractive force), and low when  $w_{ij}$  is small (a repulsive force) [11], [55]. Because of this formulation, the quality of the constructed graph directly determines the quality of the resulting layout. Dammrich and Hamprecht [22] further showed that UMAP’s stochastic negative sampling strongly down-weights the repulsive force, which increases the layout’s sensitivity to changes in neighborhood connections. In other words, the initial construction of the unweighted  $k$ NN graph is a critical factor in determining the quality of the final layout.

As described in Eq. (1), UMAP assumes that data points are *uniformly distributed* on the manifold under some Riemannian metric. Under this assumption, any ball of fixed volume should contain approximately the same number of points, regardless of its position on the manifold. This motivates the use of fixed-size  $k$ NNs to represent local relationships. In practice, however, real data are rarely uniform: they often vary in density or lie on curved manifolds. As a result,  $k$ NNs constructed with a conventional metric (e.g., Euclidean distance) and the connectivity parameters,  $\rho_i$  and  $\sigma_i$ , derived from them can misrepresent geodesic proximity. Since UMAP’s cross-entropy objective optimizes only over the weighted  $k$ NN graph, any mismatch introduced during graph construction directly propagates into the final layout. This observation highlights a research gap: improving neighborhood construction is essential for achieving low-dimensional layouts that faithfully reflect the underlying manifold geometry.

### B. Algorithmic Improvements of UMAP

Several works have refined UMAP’s pipeline, either by modifying its loss formulation to better capture data characteristics or by guiding its hyperparameter selection. For example, Narayan et al. [59] showed that UMAP tends to overly expand dense regions while overly contracting sparse regions. To mitigate this issue, they introduced DensMAP, which augments UMAP’s cross-entropy with a density-preservation term so that the layout better reflects underlying data density. Similarly, Islam and Fleischer [43] analyzed UMAP’s attraction–repulsion force dynamics and modified these forces to improve cluster consistency under random initialization. Other works have focused on guiding hyperparameter choices. For instance, Xia et al. [85] introduced scDEED, a statistical framework that evaluates layout reliability by comparing local neighborhoods before and after the graph layout phase, and then tunes hyperparameters such as the number of neighbors,  $k$ , to improve trustworthiness. In general, these refinements improve UMAP’s layout behavior, but unlike MAPLE, they do not address the critical limitation rooted in the graph construction.

### C. Self-supervised Learning and Integration into DR

Self-supervised learning (SSL) has become a dominant approach for learning representations from unlabeled data [52].

Its central idea is to let the data provide its own training signals. A large family of SSL methods is known as MVSSL. MVSSL’s key concept is to create multiple *views* of the same data that capture it from different perspectives. These views can be generated through data augmentations, such as cropping, masking, and rotating images, or they can come from different modalities, such as paired audio and video. A machine learning model (e.g., a neural network) is then trained to produce robust embeddings that remain consistent across different views of the same data [65]. Depending on how this alignment is achieved, MVSSL methods can be categorized [67] into contrastive methods (e.g., SimCLR [16]), clustering-based methods (e.g., DeepCluster [13]), distillation-based methods (e.g., BYOL [33]), and redundancy-reduction methods (e.g., Barlow Twins [88]).

Recently, MMCR-based methods have emerged as a new family within MVSSL. These methods are inspired by manifold capacity theory [17], which studies the capacity of a manifold to separate a random set of regions. Intuitively, this capacity reflects the linear separability of local manifolds that compose the global manifold. MMCRs were first introduced by Yerxa et al. [87]. MMCRs employ the nuclear norm [64] as a tractable surrogate for manifold capacity. Since the nuclear norm is defined as the sum of singular values of the data matrix, it provides an estimate of the effective rank of the corresponding manifold. Minimizing the nuclear norm compresses individual local manifolds into lower-rank structures, while maximizing the nuclear norm of manifold centroids promotes their separation. This geometric formulation is distinctive to MMCRs, compared to the other MVSSL families, as it directly ties embedding quality to manifold capacity rather than pairwise alignment [67].

To the best of our knowledge, MAPLE is the first method that adopts MMCRs into graph-based modeling for DR. Below, we discuss DR methods that adopt other SSL families.

**Integration of SSL into DR.** Several works have integrated SSL into DR to improve layout quality. These works share critiques of the graph construction phase, including unreliable distances in high-dimensional space [9], semantically unstable neighborhoods [40], and distortion caused by curvature or density differences [27], [48]. SSL objectives can help learn more semantically meaningful results. For example,  $t$ -SimCNE [9] combines the contrastive learning method, SimCLR, with a  $t$ -SNE-like loss to improve similarity estimation of data points. However, as it relies on SimCLR,  $t$ -SimCNE is only applicable to image datasets. TLDR [48] employs a single nearest neighbor as the self-supervised signal and Barlow Twins [88] as the SSL loss. While effective for data compression and preprocessing, TLDR is not designed for visualization purposes. Moreover, its reliance on a fixed 1-nearest neighbor graph prevents it from addressing misaligned neighborhood issues. Similarly, NCLDR [27] also uses nearest neighbors as signals, but its architecture is tailored to images.

In summary, existing SSL-integrated DR methods are either restricted to specific domains (e.g., image datasets) or designed for preprocessing rather than visualization. They do not provide a general solution for handling misaligned neighborhoods, diverse datasets, or visualization of high-dimensional data.




Fig. 1. Conceptual overview of MAPLE. Self-supervised neighborhood refinement turns a cross-connected, entangled  $k$ NN graph into a simplified geometry that supports clearer graph layouts.

### III. MANIFOLD-AWARE PROJECTION WITH LEARNED EDGES

In this section, we first outline UMAP’s critical limitations, which motivate our work. We then introduce *MAPLE* (manifold-aware projection with learned edges).<sup>1</sup>

#### A. Motivation

High-dimensional data introduce fundamental challenges for DR methods with graph-based modeling, such as  $t$ -SNE [76], UMAP [55], and PaCMAP [83]. First, as dimensionality increases, conventional distance metrics (e.g., Euclidean and cosine distances) become more prone to the curse of dimensionality—pairwise distances measured in the high-dimensional space become more sensitive to noise within a data distribution [8]. This makes it difficult to distinguish true neighbors from unrelated points when generating neighbor graphs from high-dimensional data. Second, many high-dimensional datasets (e.g., images or single-cell data in genomics) vary along curved low-dimensional manifolds defined by nonlinear generative factors such as pose or cell state. In these cases, close points in the high-dimensional space may be far apart along the underlying manifold and, therefore, may produce neighbor graphs where semantically distinct regions are connected. A recent DR method, LocalMAP [84], is designed to mitigate this issue in PaCMAP [83]. LocalMAP dynamically reweights neighbor edges and resamples

neighbors during optimization. However, this approach relies on heuristic adjustments and does not resolve the broader issue of noisy variance and manifold-level distortions when constructing neighborhood graphs.

UMAP models the underlying manifold as a fuzzy graph derived from a  $k$ NN graph that is constructed based on distances in high-dimensional space. Due to the above two-fold characteristics of high-dimensional data, the fuzzy graph can cause semantic mismatches, where points with no meaningful relationships are connected only because of distances in high-dimensional space rather than the underlying manifold geometry. As discussed in Sec. II-A, these semantic mismatches propagate into the final layout. This issue often manifests as loss of subcluster details in UMAP’s outcomes when intra-cluster variation is not adequately represented, and as noisy or irregular cluster boundaries when neighborhoods contain points belonging to different clusters. Therefore, with MAPLE, we aim to refine the graph construction phase of UMAP (see Phase 1 in Sec. II-A).

#### B. Overview of MAPLE

MAPLE enhances graph-based modeling of the underlying manifold through SSL. Fig. 1 illustrates a conceptual process by which MAPLE produces a concise graph representation of the manifold. Building on UMAP’s pipeline, MAPLE shifts graph construction from a fixed-distance space (e.g., defined by Euclidean distances) into a learned embedding space. This shift is enabled by incorporating MMCRs during graph construction. MMCRs introduce two complementary forces: (1) a compression force that reduces variance within local neighborhoods, flattening noisy or curved regions to make substantial structures (e.g., subclusters) more apparent; and (2) a diversification force that increases variance across dissimilar regions, reducing edges that connect dissimilar points and clarifying boundaries. The learned graph is then laid out in a low-dimensional space using UMAP’s cross-entropy optimization.

#### C. Phase 1: Self-supervised Graph Construction

MAPLE’s graph construction is performed through three steps, as illustrated in Phase 1 of Fig. 2. This phase represents the core contribution of MAPLE, where we integrate the MVSSL scheme into the DR process and incorporate MMCRs into the manifold modeling objective.

1) *Step 1: Data Embedding*: In the first step, we embed the original high-dimensional data into a lower-dimensional space (e.g., 128 dimensions) using simple neural networks. This reduces the influence of noisy data points, thereby improving both training stability and the quality of the learned graph.

Following prior work in self-supervised learning [33], [48], [88], MAPLE employs *encoder and projector networks*. The encoder network,  $f_\theta$ , consists of two fully connected hidden layers (default: 512 and 2,048 nodes), while the projector network,  $g_\theta$ , consists of two layers (512 and 128 nodes). Together, these networks transform the original high-dimensional data,  $\mathbf{X}$ , into embedded data,  $\mathbf{Z} \in \mathbb{R}^{N \times D'}$  where  $D' = 128$  under the default setting (i.e.,  $\mathbf{Z} = g_\theta(f_\theta(\mathbf{X}))$ ). During this transformation, each embedded data point is normalized to a

<sup>1</sup>The source code and supplementary materials can be found at <https://github.com/zeyh/MAPLE>.

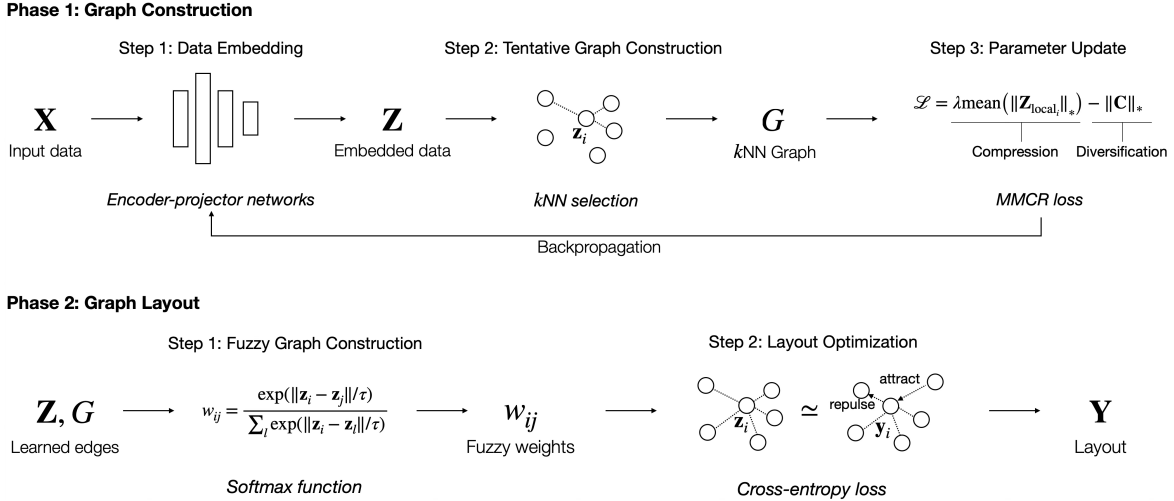


Fig. 2. The MAPLE pipeline, consisting of two phases: (1) graph construction and (2) graph layout. Phase 1 learns a graph representation of the input data through three steps that integrate SSL with MMCRs. Phase 2 performs the cross-entropy-based layout optimization, following UMAP. MAPLE’s core contribution lies in Phase 1.

unit vector to eliminate scaling differences when computing similarities. This unit-vector constraint is also consistent with the original implementation of MMCRs [87].

We apply this default architecture across datasets as a generic, lightweight feature extractor. For high-dimensional inputs ( $D > 512$ ), the first layer provides a moderate compression that helps filter out noise and stabilizes subsequent training. For lower-dimensional inputs ( $D < 512$ ), although the first layer could technically be reduced in size, we retain this expanded capacity to maintain a consistent and standardized architecture across datasets. Alternatively, the architecture could be scaled (e.g., linearly) with the input dimensionality; exploring such scaling strategies is a worthwhile direction for future work. Although more complex architectures such as ResNet [36] could be employed, we deliberately adopt this minimal encoder–projector design as an efficient and broadly applicable approach for high-dimensional data.

2) *Step 2: Tentative Graph Construction:* The second step constructs a tentative graph,  $\mathbf{G}$ , from the current embedded data,  $\mathbf{Z}$ . We first build a  $k$ NN graph as  $\mathbf{G}$  based on  $L^2$  distances in the embedding space. Unlike UMAP, MAPLE iteratively refines  $\mathbf{G}$  by updating the neural network parameters of  $f_\theta$  and  $g_\theta$  with MMCRs, as described in Step 3.

The neighbors in the tentative graph can be viewed as noisy “self-references”. We treat neighbors as similar data points and non-neighbors as dissimilar ones to enable self-supervised learning in Step 3. In the context of MVSSL, these neighbor data points can be regarded as *augmented views* of each data point, used to learn invariances. In our case, a data point and its neighbors act as multiple views reflecting the same local structure. However, unlike typical augmentation scenarios (e.g., creating augmented images via rotation), these views may not align perfectly with the same structure. Since the networks are initially untrained, early-epoch neighborhoods can still reflect noisy structure inherited from  $\mathbf{X}$ . We therefore refine these neighborhoods through regularization with MMCRs.

3) *Step 3: MMCR Loss Computation and Parameter Update:* In MAPLE, we reinterpret the objective of SSL through MMCRs [87], specifically in the context of neighbor graph learning for DR. MMCRs aim to produce *efficient* data embeddings for downstream classification tasks by maximizing the linear separability of local manifolds, which often correspond to different classes, in the embedding space. Within neighbor graph learning, MMCRs refine the embedding space to better preserve neighborhood relationships along the underlying manifold.

MMCRs consider two types of variance to improve the linear separability of local manifolds: (1) *between-variance*, the variance of samples across different local manifolds, and (2) *within-variance*, the variance of samples within the same local manifold. Linear separability improves when between-variance increases and within-variance decreases.

**Between-variance.** As discussed in Sec. III-C2, we treat each embedded data point and its neighborhoods as views of the same structure. In other words, we assume these points are samples from the same local manifold. Let  $\mathcal{P}_i$  be the set of indices of the  $i$ th data point and its  $k$  nearest neighbors. For the local manifold around a data point  $\mathbf{z}_i$ , the centroid of the manifold,  $\mathbf{c}_i$ , can be approximated as:  $\mathbf{c}_i = \sum_{j \in \mathcal{P}_i} \mathbf{z}_j / (k + 1)$ . The between-variance is then defined as:

$$\text{Var}_{\text{between}} := \|\mathbf{C}\|_* \quad (3)$$

where  $\mathbf{C} = [\mathbf{c}_1 \cdots \mathbf{c}_N]^T \in \mathbb{R}^{N \times D'}$  and  $\|\cdot\|_*$  denotes the nuclear norm operator, with  $\|\mathbf{C}\|_* = \text{trace}(\sqrt{\mathbf{C}\mathbf{C}^T})$ . Intuitively, the nuclear norm measures the total variance of  $\mathbf{C}$  across all principal directions.  $\|\mathbf{C}\|_*$  is large when the centroids spread across many orthogonal directions, and small when the centroids nearly lie in a low-dimensional subspace (e.g., a plane).

Maximizing the between-variance encourages the centroids of local manifolds to spread across multiple orthogonal directions. In other words, it promotes variance to be distributed more evenly across multiple principal components, rather than collapsing into a few dominant axes. Geometrically, this

pushes centroids to form large pairwise angles on the unit sphere, suppressing concentration in a single direction that would otherwise create hub nodes in the neighborhood graph.

**Within-variance.** Let  $\mathbf{Z}_{\text{local}_i}$  be the matrix consisting of the  $i$ -th data point in  $\mathbf{Z}$  and its  $k$  nearest neighbors (i.e.,  $\mathbf{Z}_{\text{local}_i} \in \mathbb{R}^{(k+1) \times D'}$ ). Then, similar to the between-variance, the within-variance of the local manifold around each point  $\mathbf{z}_i$  is defined as:

$$\text{Var}_{\text{within}} := \|\mathbf{Z}_{\text{local}_i}\|_* \quad (4)$$

Minimizing the within-variance encourages samples from each local manifold to align in a lower-dimensional subspace, ideally approximating a flat tangent plane. The nuclear norm is upper-bounded by the Frobenius norm [64] scaled by the square root of the rank, i.e.,  $\|\mathbf{Z}_{\text{local}_i}\|_* \leq \sqrt{\text{rank}(\mathbf{Z}_{\text{local}_i})} \|\mathbf{Z}_{\text{local}_i}\|_F$ . Thus, minimizing the within-variance suppresses variance orthogonal to the leading singular directions, making each neighborhood more compact and low-rank. As a result, semantic neighbors are pulled closer together, while irrelevant neighbors are naturally down-weighted.

**MMCR loss function.** Using the between- and within-variances of local manifolds, we define the MMCR loss as:

$$\mathcal{L} = \lambda \text{ mean}_{i \in \{1, \dots, N\}} (\|\mathbf{Z}_{\text{local}_i}\|_*) - \|\mathbf{C}\|_* \quad (5)$$

where  $\lambda > 0$  is a trade-off hyperparameter. This loss balances two complementary forces: the first term acts as a compression force, shrinking the size of local manifolds, while the second term acts as a diversification force, encouraging separation among different local manifolds.  $\lambda$  controls the relative strength of these forces. MAPLE optimizes this loss by updating the parameters  $\theta$  of the encoder and projector networks, ensuring that the resulting embedding  $\mathbf{Z}$  minimizes  $\mathcal{L}$  (i.e., solving  $\text{argmin}_{\theta} \mathcal{L}$ ).

**Efficient optimization.** Although Steps 1–3 could, in theory, be repeated at every update, we adopt a computationally efficient optimization strategy. First, we use mini-batch gradient descent, computing the loss over batches of data. Second, we reuse the constructed graph across updates in Steps 1 and 3, and only reconstruct it every fixed number of epochs (20 by default).

**Alternative designs.** MAPLE uses the encoder–projector networks and MMCRs. We also explored several alternative designs during development. For example, we tested multiple different self-supervised learning approaches such as Barlow Twins [88]. Our supplementary materials provide detailed discussions on tested designs.

#### D. Phase 2: Graph Layout

From the outcomes of Phase 1, MAPLE constructs a fuzzy graph and then performs layout optimization, as illustrated in Fig. 2.

1) *Step 1. Fuzzy Graph Construction:* Given the learned graph,  $G$ , and the embedded data,  $\mathbf{Z}$ , MAPLE computes fuzzy weights for each edge using a softmax function:

$$w_{ij} = \frac{\exp(-\|\mathbf{z}_i - \mathbf{z}_j\|/\tau)}{\sum_{l \in i\text{'s neighbors}} \exp(-\|\mathbf{z}_i - \mathbf{z}_l\|/\tau)} \quad (6)$$

where  $\tau$  is a softmax temperature parameter and  $\sum_j w_{ij} = 1$ .  $w_{ij}$  defines a probability distribution over  $i$ 's neighbors: closer neighbors (higher similarity) receive higher weights, while the softmax normalizes across the neighborhood, balancing the attractive forces from different neighbors. MAPLE's fuzzy weight computation differs from UMAP's original one (Eq. (1)), which is designed to eliminate weak neighbor connections. Since the learned graph in MAPLE has already removed such connections, applying UMAP's procedure would result in unnecessary elimination.

2) *Step 2: Layout Optimization:* For layout optimization, we adopt UMAP's original approach. The layout,  $\mathbf{Y}$ , is initialized using a spectral layout and then optimized based on the cross-entropy loss between pairwise similarities in the embedded space (128D by default) and the layout space (typically 2D or 3D). We preserve all aspects of UMAP's optimization procedure, including negative sampling, learning rate scheduling, and stochastic updates. This setting ensures that MAPLE's improvements primarily arise from the self-supervised graph construction rather than modifications to the layout optimization.

In many cases (e.g., Fig. 1, Fig. 7(a)), MAPLE produces locally linear, skeleton-like motifs in the final layout. This behavior can be interpreted as a geometric consequence of optimizing manifold capacity in Eq. (5). Since the nuclear norm jointly reflects manifold radius and intrinsic dimensionality [87], minimizing this term suppresses variance along minor directions within local neighborhoods. As a result, local structures are encouraged to align with their principal axes, approximating low-rank subspaces rather than high-dimensional ellipsoids. Compared to UMAP, which often yields a series of blob-like clusters, this process can produce layouts composed of concatenated, locally flattened segments.

**Implementation** We implemented MAPLE in Python 3, using NumPy/SciPy [81] for matrix calculations and scikit-learn [63] for  $k$ NN graph construction. The neural networks are built with PyTorch [61]. For the MMCR loss, we adopted the implementation by the original authors [87]. For layout optimization, we employed the official UMAP implementation [55]. Additional implementation details, including neural network settings, are provided in the supplementary materials.

**Summary of main hyperparameters** MAPLE introduces a small set of hyperparameters: (1) the embedding dimension  $D'$ , (2) the number of neighbors  $k$  for graph construction, (3) the trade-off parameter  $\lambda$  in the MMCR loss, and (4) the softmax temperature  $\tau$  for the fuzzy weight computation. Based on our experiments across multiple datasets, we observed that MAPLE is relatively robust to moderate variations in  $D'$  and  $\tau$ , due to the stability of the UMAP cross-entropy objective. Therefore, we consider  $k$  and  $\lambda$  to be the primary hyperparameters governing MAPLE's behavior. Both UMAP and MAPLE require the neighborhood size hyperparameter  $k$ . The remaining main parameter  $\lambda$  in MAPLE plays an analogous role to `min_dist` in UMAP, controlling the balance between diversification and compression forces. However, a key distinction is that while `min_dist` imposes a static constraint on inter-point distances, the MAPLE objective is

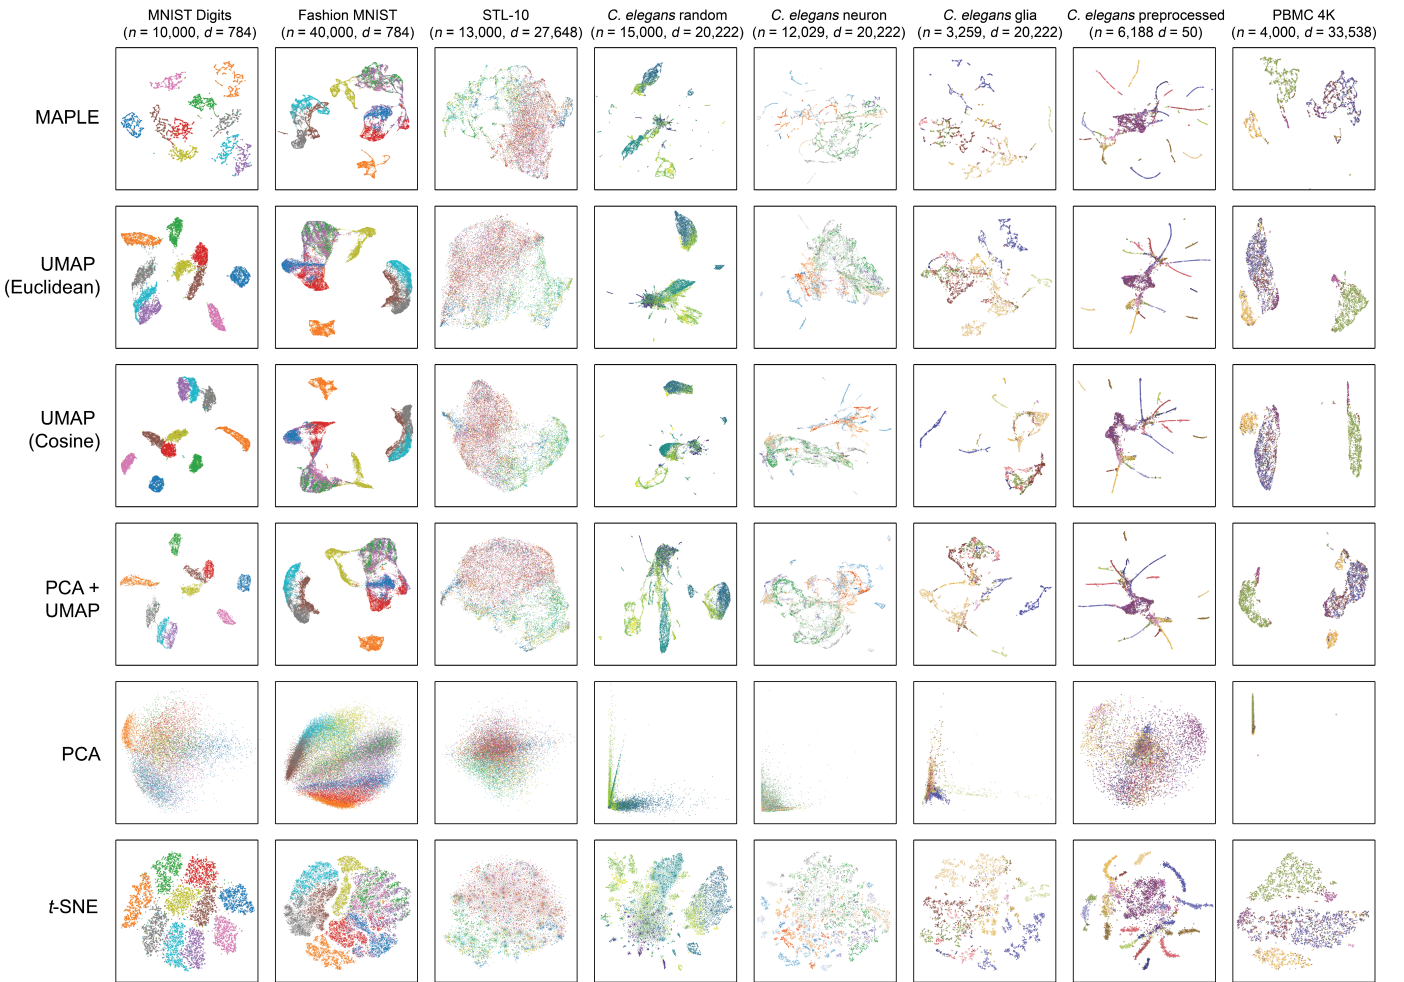


Fig. 3. Layouts of image (MNIST [23], Fashion-MNIST [86], STL-10 [18]) and single-cell datasets (*C. elegans* subsets [60] and PBMC 4K) using MAPLE and baseline methods. MAPLE preserves both global structure and local neighborhoods more consistently across datasets.

grounded in manifold capacity theory and adjusts the balance between compression and diversification during optimization. By optimizing the nuclear norm, the objective adaptively influences the spacing and density of local manifolds in the embedding space. Our parameter sensitivity analysis (detailed in Sec. IV-D) further shows that MAPLE remains robust across a wide range of  $\lambda$  and varying  $k$  values.

#### IV. RESULTS

We evaluate MAPLE on eight datasets with multiple classes or subtypes, representative of real-world high-dimensional analysis challenges. Fig. 3 shows layouts generated by MAPLE and five other DR methods. Tab. I presents quantitative results across 17 evaluation measures on the same datasets (in the same order as Fig. 3), comparing MAPLE with UMAP using a Euclidean input-space metric. The corresponding layouts can be found in the first two rows of Fig. 3.

**Datasets.** We use datasets from two representative domains that involve very high-dimensional data and commonly utilize DR for their analyses: image recognition and transcriptomics.

Image datasets:

- 1) MNIST handwritten digits test set [23]:  $N=10,000$ ,  $D=784$
  - 2) A subset of Fashion-MNIST [86]:  $N=40,000$ ,  $D=784$
  - 3) Labeled images from STL-10 [18]:  $N=13,000$ ,  $D=27,648$
- Single-cell RNA sequencing (scRNA-seq) datasets:
- 1) A random subset of *C. elegans* [60]:  $N=15,000$ ,  $D=20,222$
  - 2) Neuron subset from *C. elegans*:  $N=12,029$ ,  $D=20,222$
  - 3) Glia subset from *C. elegans*:  $N=3,259$ ,  $D=20,222$
  - 4) Preprocessed subset of *C. elegans*:  $N=6,188$ ,  $D=50$
  - 5) PBMC 4K from 10x Genomics [1], annotated with the Cellarium Cell Annotation Service [14]:  $N=4,000$ ,  $D=33,538$

The preprocessed subset of *C. elegans* follows the pipeline recommended in Monocle’s documentation [73]. We include this dataset to assess MAPLE’s performance when strong preprocessing has already been applied.

##### A. Quantitative Evaluation

Quantitative evaluation of DR layouts is inherently challenging, as there is no universally accepted “ground truth” for high-dimensional structure, and different measures emphasize different aspects of layout quality [24], [45], [47], [53]. To perform quantitative evaluation, it is therefore necessary to select or design quality measures for DR layouts [44], [46].

TABLE I

MAPLE VS UMAP ACROSS LOCAL, GLOBAL, CLUSTERING, BOUNDARY, AND PERCEPTUAL MEASURES (HIGHER IS BETTER  $\uparrow$ , LOWER IS BETTER  $\downarrow$ ).

Measure	MNIST Digits		Fashion MNIST		STL-10		C. <i>el</i> Random		C. <i>el</i> Neuron		C. <i>el</i> Glia		C. <i>el</i> Preprocessed		PBMC 4K	
	MAPLE	UMAP	MAPLE	UMAP	MAPLE	UMAP	MAPLE	UMAP	MAPLE	UMAP	MAPLE	UMAP	MAPLE	UMAP	MAPLE	UMAP
<i>Local</i>																
Neighborhood hit (core) $\uparrow$	<b>0.922</b>	0.911	<b>0.761</b>	0.715	<b>0.244</b>	0.212	<b>0.554</b>	0.450	<b>0.626</b>	0.561	<b>0.750</b>	0.718	<b>0.756</b>	0.751	<b>0.598</b>	0.576
kNN classification accuracy (core) $\uparrow$	<b>0.813</b>	0.794	<b>0.818</b>	0.778	<b>0.302</b>	0.260	<b>0.650</b>	0.554	<b>0.660</b>	0.555	<b>0.825</b>	0.790	<b>0.859</b>	0.855	<b>0.678</b>	0.601
Distance consistency $\uparrow$	<b>0.913</b>	0.911	0.660	<b>0.670</b>	0.203	<b>0.224</b>	<b>0.341</b>	0.247	<b>0.307</b>	0.227	<b>0.424</b>	0.418	0.519	<b>0.606</b>	<b>0.255</b>	0.245
<i>Global</i>																
CH $\uparrow$	9282.101	<b>11480.073</b>	25828.332	<b>36460.773</b>	<b>643.450</b>	522.385	<b>521.646</b>	471.087	<b>557.774</b>	517.456	<b>463.124</b>	457.781	678.806	<b>1205.347</b>	560.435	<b>1262.673</b>
DB $\downarrow$	<b>0.926</b>	0.971	2.055	<b>1.947</b>	<b>16.232</b>	34.250	<b>15.664</b>	19.263	<b>8.131</b>	9.393	3.719	<b>3.170</b>	5.360	<b>4.245</b>	<b>5.119</b>	6.050
<i>Clustering (k-means)</i>																
ARI $\uparrow$	<b>0.810</b>	0.752	<b>0.493</b>	0.431	<b>0.072</b>	0.063	<b>0.237</b>	0.183	<b>0.174</b>	0.159	<b>0.365</b>	0.351	0.343	<b>0.436</b>	<b>0.337</b>	0.294
AMI $\uparrow$	0.792	<b>0.813</b>	<b>0.651</b>	0.623	<b>0.158</b>	0.138	<b>0.484</b>	0.469	<b>0.434</b>	0.397	<b>0.540</b>	0.531	0.547	<b>0.571</b>	<b>0.524</b>	0.502
NMI $\uparrow$	<b>0.793</b>	0.779	<b>0.655</b>	0.607	<b>0.164</b>	0.141	<b>0.483</b>	0.477	<b>0.424</b>	0.412	<b>0.545</b>	0.532	0.567	<b>0.614</b>	<b>0.530</b>	0.507
V-measure $\uparrow$	<b>0.832</b>	0.779	<b>0.649</b>	0.607	<b>0.159</b>	0.139	<b>0.515</b>	0.476	<b>0.428</b>	0.408	<b>0.548</b>	0.522	0.595	<b>0.615</b>	<b>0.543</b>	0.510
<i>Clustering (DBSCAN)</i>																
ARI $\uparrow$	<b>0.711</b>	0.405	<b>0.334</b>	0.219	-	-	<b>0.203</b>	0.061	<b>0.157</b>	0.074	<b>0.345</b>	0.078	<b>0.166</b>	0.056	<b>0.447</b>	0.400
AMI $\uparrow$	<b>0.794</b>	0.689	<b>0.596</b>	0.532	-	-	<b>0.514</b>	0.296	<b>0.499</b>	0.348	<b>0.552</b>	0.379	<b>0.525</b>	0.322	<b>0.584</b>	0.568
NMI $\uparrow$	<b>0.795</b>	0.690	<b>0.596</b>	0.532	-	-	<b>0.528</b>	0.302	<b>0.511</b>	0.355	<b>0.563</b>	0.384	<b>0.534</b>	0.328	<b>0.588</b>	0.571
V-measure $\uparrow$	<b>0.795</b>	0.690	<b>0.596</b>	0.532	-	-	<b>0.528</b>	0.302	<b>0.511</b>	0.355	<b>0.563</b>	0.384	<b>0.534</b>	0.328	<b>0.588</b>	0.571
<i>Boundary</i>																
Label trustworthiness $\uparrow$	<b>0.846</b>	0.823	<b>0.523</b>	0.431	<b>-0.511</b>	-0.576	<b>0.109</b>	-0.099	<b>0.255</b>	0.124	<b>0.504</b>	0.438	<b>0.513</b>	0.503	<b>0.198</b>	0.152
Hausdorff distance $\uparrow$	<b>5.112</b>	2.938	<b>5.211</b>	3.077	<b>0.792</b>	0.574	<b>8.898</b>	2.512	<b>6.583</b>	4.074	<b>7.761</b>	4.800	<b>7.994</b>	5.674	<b>11.049</b>	7.430
<i>Perceptual</i>																
2AFC accuracy $\uparrow$	<b>0.953</b>	0.949	0.898	<b>0.904</b>	<b>0.612</b>	0.602	<b>0.861</b>	0.848	<b>0.828</b>	0.809	0.840	<b>0.844</b>	0.853	<b>0.893</b>	0.869	<b>0.870</b>
Pairwise AUC (ROC-AUC) $\uparrow$	0.953	<b>0.954</b>	0.908	<b>0.918</b>	<b>0.610</b>	0.603	<b>0.795</b>	0.747	<b>0.796</b>	0.792	<b>0.845</b>	0.827	0.868	<b>0.895</b>	<b>0.881</b>	0.879

Note: Best scores per dataset are in **bold**. The *C. elegans* (*C. el*) datasets (Random, Neuron, Glia) are evaluated against fine-grained biological subtype annotations rather than broad cell classes. Clustering scores for STL-10 are not included because the density-based clustering algorithm DBSCAN identified only a single cluster for both MAPLE and UMAP.

While existing measures provide valuable insights, selecting appropriate ones remains challenging due to inherent trade-offs in the structural aspects they capture. A large class of existing measures heuristically quantifies how well neighborhood relationships in high-dimensional space are preserved in the resulting layout. For instance, the Kullback–Leibler divergence [38] compares neighborhood probability distributions, while trustworthiness and continuity [78] penalize neighbors that are gained or lost compared to the input *k*NNs. ZADU [44], a library compiling various quality measures, includes 20 measures, 15 of which fall into this class.

Although widely adopted as standard evaluation criteria for new DR methods [2], [47], these measures capture only limited aspects of layout quality and face two critical issues. First, many rely on *k*NN graphs constructed with distance metrics; however, as discussed in Sec. II-A and Sec. III-A, conventional *k*NN graphs do not necessarily reflect intrinsic manifold structure [7]–[9], [84]. Second, these measures emphasize statistical interpretability rather than geometric or informational expressiveness. They quantify whether certain local statistics (e.g., distances among neighbors) are preserved. However, as highlighted in SSL research [4], [65], raw data often contain noise, and such statistics may not reflect whether embeddings and layouts provide faithful geometric representations of the manifold (e.g., curvature, geodesics [58]). MAPLE is designed to address data with noisy local manifolds, and the layout quality should be judged by the ability to reveal semantically meaningful structures.

Given these limitations, we use class labels as a proxy for ground truth semantic structure in high-dimensional space. Accordingly, we use label-based measures to quantify layout quality. Jeon et al. [45] argue that label-based measures can be unreliable as classes do not always correspond to clusters in the original space. This concern is mitigated in our evaluation by our choice of datasets, where classes are well-defined.

As a result, we assembled a set of 17 label-based quality

measures suitable for comprehensive evaluation. Specifically, this set consists of five groups:

- *Local measures*, including neighborhood hit, *k*NN classification accuracy [55], and distance consistency, which measure label agreement at the neighborhood and centroid level [44], [62];
- *Global measures*, including the Calinski–Harabasz (CH) [12] and Davies–Bouldin (DB) [79] scores, both of which assess the quality of class margins (with higher CH and lower DB indicating clearer separation) [44];
- *Clustering measures*, obtained by applying *k*-means [35] and DBSCAN [25] on the 2D layouts and reporting four scores—adjusted Rand index (ARI), adjusted mutual info score (AMI), normalized mutual info score (NMI), and V-measure—to quantify alignment with ground-truth classes [44], [80];
- *Boundary measures*, which evaluate boundary quality at both local and global levels; and
- *Perceptual measures*, adapted from perceptual similarity measures in computer vision [89] to mimic human judgments.

For consistency, the local, global, and clustering measures in Tab. I were reimplemented from ZADU [44] in PyTorch. For all quantitative comparisons, UMAP was evaluated using neighborhood sizes matched to MAPLE, while all other UMAP parameters followed default settings. Additional implementation details and parameter settings are provided in the supplementary materials.

For the *local boundary* measure, we modify the classical trustworthiness score [78] into a local label trustworthiness measure. For each point, neighbors with mismatched labels are penalized more strongly when they appear among the closest neighbors. The final score is the average across all points, reflecting the cleanliness of class boundaries in the layout. To measure *boundary globally*, we compute a quantile Hausdorff distance [41], [70] between clusters. For each pair of classes,

TABLE II  
PER-CLASS  $k$ NN CLASSIFICATION ACCURACY SCORES (WITH  $k = 10$ ).

Fashion MNIST										
	T-shirt	Trouser	Pullover	Dress	Coat	Sandal	Shirt	Sneaker	Bag	Ankle Boot
MAPLE	<b>0.838</b>	<b>0.960</b>	<b>0.697</b>	<b>0.869</b>	<b>0.676</b>	<b>0.913</b>	<b>0.442</b>	<b>0.907</b>	<b>0.960</b>	<b>0.923</b>
UMAP	0.830	0.951	0.643	0.839	0.576	0.865	0.372	0.854	0.939	0.921

<i>C. elegans</i> Neuron										
	ADF	ADF_AWB	AFD	ASE	ASE_p	ASK	ASK_p	OLQ	OLQ_gp	OLQ_p
MAPLE	<b>0.718</b>	<b>0.588</b>	<b>0.825</b>	<b>0.745</b>	<b>0.772</b>	<b>0.764</b>	<b>0.567</b>	0.717	0.514	<b>0.671</b>
UMAP	0.588	0.343	0.696	0.585	0.651	0.506	0.440	<b>0.720</b>	<b>0.590</b>	0.578

Note: Additional results with varying  $k$  values and full per-class scores are provided in the supplementary materials.

we measure how close points from one class are to the other, with the focus on the nearest cross-class neighbors. Instead of taking only the single closest pair, which can be noisy, we use the percentile of these nearest-neighbor distances and then obtain the median of these values across all class pairs.

We design *perceptual* measures inspired by perceptual similarity tests in computer vision [89], which mimic how people judge whether two images look alike. We adopt the same sampling-based idea for evaluating the layouts perceptually. We randomly sample reference points from the layout produced by MAPLE or UMAP, and compare distances to same-class and different-class neighbors. Based on Zhang et al. [89], we define a two-alternative forced choice (2AFC) accuracy measure, which counts how often a reference is closer to a same-class point than to a different-class point. Similar to the Just Noticeable Difference (JND) test [89], we also compute the area under the receiver operating characteristic curve (ROC–AUC) score [10], which corresponds to sampling many pairs of points from the layout, recording their pairwise distances, and labeling them as same-class or different-class. The ROC curve captures how well pairwise distance discriminates between same-class and different-class pairs across thresholds, and the resulting AUC provides a single reliability score. Unlike the original perceptual measure that relies on deep neural network features (e.g., from AlexNet [51]) to approximate human visual perception, we sample points on the DR results and check them with ground-truth labels.

**Observations on core quantitative measures.** We focus here on *neighborhood hit* and *kNN classification accuracy*, as these measures are directly related to assessing the quality of the learned neighborhood structure in MAPLE. As shown in Tab. I, MAPLE outperforms UMAP across both measures. Both measures assume that the provided labels approximate ground truth semantics; consequently, higher scores in the low-dimensional space indicate better preservation of local decision boundaries and neighborhood relationships.

Neighborhood hit is a widely used measure in DR evaluation. It quantifies the fraction of a point’s  $k$ NNs that share the same label, with higher values indicating that local neighborhoods are more semantically consistent. We set  $k = 15$ . On the MNIST handwritten digits dataset [23], where the classes are already well separated, MAPLE achieved a neighborhood hit [62] of 0.922 compared to 0.911 for UMAP. This means that about 1% more neighbors per point were correctly recovered as belonging to the same class. The improvement is 5% more correct neighbors on the Fashion-MNIST dataset [86].

On STL-10, which has higher input dimensionality and greater intra-class variation than MNIST/Fashion-MNIST, MAPLE still improves neighborhood hit by roughly 4% even with our minimal MLP encoder–projector. For single-cell data [60], MAPLE achieved 0.554 vs. 0.450 on a *Random* subset of *C. elegans* cells, 0.626 vs. 0.561 on the *Neuron* subset, and 0.750 vs. 0.718 on the *Glia* subset. Those improvements correspond to 5%–10% more correctly grouped neighbors in biologically meaningful subtypes when using MAPLE instead of UMAP. This improvement is partly due to the high-resolution of the ground truth labels in this dataset. Unlike broader categories, our evaluation uses 54 identified cell subtypes, which makes the measures more sensitive to MAPLE’s ability to preserve fine-grained structure that other methods might blur.

$k$ NN classification accuracy evaluates how well local neighborhoods preserve class identity and enables per-class analysis. For each data point, we predict its label via a weighted majority vote among its  $k$  neighbors in the 2D layout ( $k = 10$ ). This approach yields per-class accuracy scores defined as the mean prediction accuracy within each class (Tab. II), while the overall score corresponds to the global mean accuracy across all data points (Tab. I). MAPLE outperforms UMAP on overall per-dataset scores in Tab. I. To better understand these gains, we examine class-level performance in Tab. II. On Fashion-MNIST, improvements primarily arise from resolving overlaps among visually similar classes such as ‘Shirt’ and ‘Coat’, with gains of 5–10%. A similar trend appears in the *C. elegans* neuron dataset, where MAPLE achieves a global improvement exceeding 11%. For specific chemosensory lineages (e.g., ADF and ASE), MAPLE more clearly separates distinct cell subtypes (see Fig. 7). In contrast, for the mechanosensory OLQ lineage, improved separation of specific subtypes increases fine-grained accuracy while reducing scores for broader groupings. This trade-off reflects MAPLE’s tendency to resolve substructure rather than blur distinct cell types to preserve coarse lineage groupings.

**Other observations.** Regarding global class separability, MAPLE generally produces layouts with clearer margins between classes compared to UMAP. On global measures, both the DB [79] and CH [12] scores compare within-cluster scatter against between-cluster separation, and they favor compact, well-separated clusters. MAPLE exhibits mixed performance compared to UMAP, depending on the data geometry. By revealing fine-grained classes, MAPLE often increases inter-class separation (thereby improving these scores); in some other cases, however, this increased resolution resulted in elongated structures that are penalized by these compactness-based measures. MAPLE also achieves higher Hausdorff distances [41], [70] across all experiments.

Clustering and perceptual measures can assess MAPLE’s quality for downstream analytical workflows such as annotation and cluster analysis. For clustering measures, MAPLE outperforms UMAP on almost all datasets. This confirms that the layouts produced by MAPLE are easier to partition and consistent with the improvement in local measures. For perceptual measures (2AFC, ROC–AUC) [89], MAPLE and UMAP are comparable. The two perceptual measures are ranking-based tests that ask whether same-class distances are




Fig. 4. Closer visual inspection of MAPLE layouts on the Fashion-MNIST dataset (same layout as in Fig. 3). (a) Underlying average images of the trouser class. (b) Underlying average images from the coat/dress region.

consistently smaller than different-class distances, so they measure binary rankings of pairwise distances in the layout. Unless a method achieves substantially better class separation, layouts with rounded and compact clusters that UMAP often produces tend to score higher because their within-class distances are more uniformly small and stable. In contrast, MAPLE produces higher-resolution structures that can introduce greater variance within classes, which can lead to lower perceptual scores for datasets like Fashion-MNIST [86].

In summary, MAPLE generally outperforms UMAP across the selected quantitative measures. The magnitude of MAPLE’s improvements depends on dataset complexity [42] and the granularity of the ground truth labels. In this context, complexity refers to the difficulty of separating overlapping classes in high dimensions, while granularity refers to the level of detail in the annotations (e.g., broad categories vs. specific subtypes). On well-separated datasets like MNIST handwritten digits, there is limited room for improvement; MAPLE shows modest quantitative improvements, whereas for single-cell data, the improvements are more pronounced.

### B. Qualitative Evaluation

We qualitatively compare the layouts produced by MAPLE with other DR methods. Since part of our pipeline is adapted from UMAP, we treat UMAP as the main baseline and provide several of its variants for comparison. In Fig. 3, each row represents one DR method, including MAPLE, UMAP [55] with Euclidean distance for input space graph construction, UMAP using cosine similarity, UMAP (Euclidean) with PCA [54] pre-processing that reduces high-dimensional data to 50 principal components before applying UMAP, PCA, and  $t$ -SNE [76].

Across the datasets shown in Fig. 3, MAPLE produces layouts with noticeably finer internal structures. These finer internal structures contribute to clearer separation between broader categories, consistent with our quantitative results showing improved 2D neighborhood fidelity and boundary measures. To examine whether these structures correspond to meaningful variation in the input data, we focus on the Fashion-MNIST dataset, as the underlying images require no additional domain knowledge to interpret. In Fig. 3, UMAP layouts under various configurations typically render this dataset as a set of largely uniform clusters where internal

variations are difficult to identify. In contrast, MAPLE reveals a distinct geometric pattern across multiple runs. For example, within the ‘Trouser’ class (orange), MAPLE exposes two elongated structures in the upper region, a central compact group, and two smaller connected peaks at the bottom.

**Visual validation.** To examine the patterns revealed by MAPLE and relate them to the underlying images, we designed a grid-based aggregation visualization (see Fig. 4). For an interactively selected region of interest, we discretize the 2D space into a grid. In each grid cell containing data points, we visualize the pixel-wise average of the corresponding images. This analysis indicates that the structures produced by MAPLE correspond to visually meaningful fashion subtypes disentangled by the self-supervised objective. As shown in Fig. 4(a), the upper-left branch in the ‘Trouser’ class primarily contains wide-leg shorts, while the upper-right region captures thin, straight-leg cuts. The central cluster consists of standard jeans, and the bottom peaks correspond to trousers with one bent leg. Similarly, in the region containing the ‘Coat’ and ‘Dress’ classes (Fig. 4(b)), MAPLE shows roughly three distinct subareas: coats at the top, dark clothing to the left, and skirts or dresses (identified by longer lower hems) at the bottom. Additional examples on the MNIST handwritten digits dataset are provided in the supplementary materials.

This ability to resolve intra-class variation is a key strength of MAPLE. By encouraging local neighborhoods to align with lower-rank structures through a self-supervised objective, MAPLE introduces perceptible gaps in the final layout between visually distinct sub-groups. From an information visualization perspective, these gaps align with the Gestalt law of proximity [3], [26], supporting the perception of distinct groups and facilitating the identification of subtle data subtypes. While we do not claim that all geometric structures in the layout directly correspond to visually identifiable semantic components, the observed consistency between local layout structures and input-level visual patterns suggests that a substantial portion of the internal structures reflects coherent variation in the data and arises from the objective’s suppression of redundant variance. Overall, these results indicate that MAPLE enables higher-resolution visual exploration and supports the discovery of fine-grained intra-class patterns that are often obscured in other DR methods.

### C. Scalability and Runtime Analysis

We evaluate the performance of MAPLE on a laptop equipped with an Apple M1 Pro chip and 16 GB of unified memory. The training framework is implemented in PyTorch using the Metal Performance Shaders (MPS) backend to execute tensor operations on the integrated GPU. To improve throughput, we use automatic mixed precision for gradient computations. Due to current limitations in the MPS backend for certain singular-value decomposition operators, the nuclear norm required by the MMCR loss is computed on the CPU. To accommodate large datasets under limited memory, we use a batch-loading strategy that avoids memory overflow.

**Data generation.** To systematically evaluate scalability, we generate synthetic variations of the MNIST handwritten

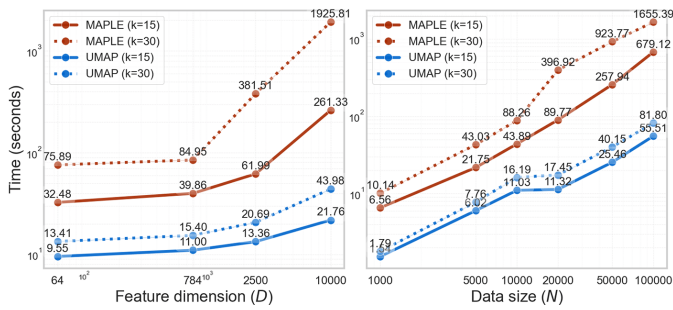


Fig. 5. Completion time of MAPLE and UMAP in seconds. Left: runtime with varying feature dimensions ( $D$ ) at fixed  $N = 10,000$ . Right: runtime with varying data sizes ( $N$ ) at fixed  $D = 784$ . Solid and dotted lines indicate neighborhood sizes of  $k = 15$  and  $k = 30$ , respectively. Both axes are shown on logarithmic scales, with tick labels  $10^2$  and  $10^3$  indicating orders of magnitude.

digits dataset [23] along two axes: feature dimensionality ( $D$ ) and data size ( $N$ ). To vary dimensionality, we spatially interpolate the original  $28 \times 28$  pixel images to resolutions ranging from  $8 \times 8$  up to  $100 \times 100$ , yielding feature vectors from  $D = 64$  to  $D = 10,000$ . To assess scalability with respect to sample size, we expand the dataset from the original 10,000 images up to  $N = 100,000$  via oversampling. Each replicated sample is augmented using one of four basic transformations: Gaussian noise addition, rotation, random flips, or brightness adjustment. All experiments use neighborhood sizes of  $k = 15$  and  $k = 30$ , a batch size of 4,096, and a training duration of 20 epochs.

**Observation.** The results are shown in Fig. 5. MAPLE exhibits approximately linear runtime scaling with respect to both data size  $N$  and feature dimensionality  $D$ , closely matching the scaling behavior of standard UMAP. For instance, increasing the data size from  $N = 50,000$  to  $N = 100,000$  increases runtime by approximately 2.7 times (from roughly 4 minutes to 11 minutes). Runtime also scales linearly with  $D$ . However, we observe increased sensitivity to neighborhood size for high-dimensional inputs. Beginning at  $D = 2,500$ , increasing  $k$  from 15 to 30 leads to a significant runtime increase (e.g., from 261 seconds to 1,925 seconds at  $D = 10,000$ ). This behavior arises from two factors. First, the encoder processes entire local neighborhoods, causing the dominant computational cost to scale as  $O(kND')$ . Second, at high dimensionality, the resulting dense tensors exceed CPU cache capacity, making system memory bandwidth a limiting factor. While batching prevents memory overflow, the runtime bottleneck shifts to data loading and computation. The bottleneck can be partially mitigated by reducing the batch size, although this introduces a trade-off with training efficiency [75].

The runtime difference between MAPLE and UMAP reflects an inherent trade-off in learning-based DR methods. MAPLE incurs additional computational cost due to the neural network training phase, resulting in a runtime of the order of minutes rather than seconds. However, this cost is predictable and scales linearly with dataset size. In practice, MAPLE’s runtime is comparable to other parametric DR methods [66]. We advise practitioners to be mindful of feature dimensionality

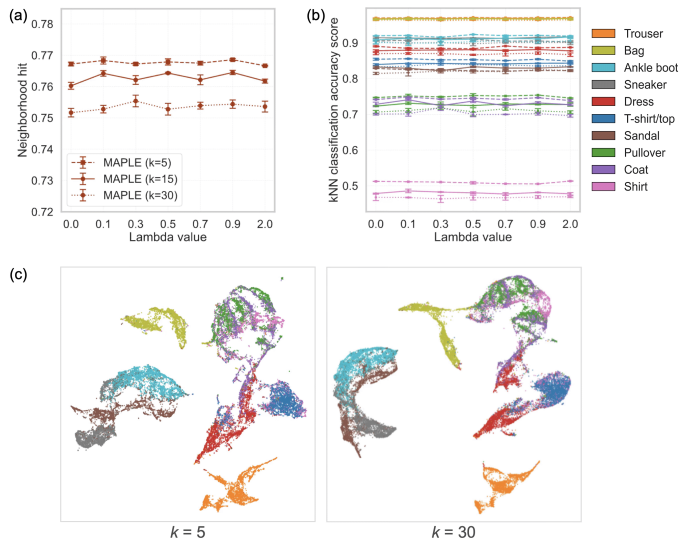


Fig. 6. Parameter sensitivity analysis on the Fashion-MNIST dataset [86] with  $N = 40,000$ . For (a) and (b), the x-axis denotes the parameter  $\lambda$  and different line styles (e.g., dashed) denote neighborhood sizes  $k \in \{5, 15, 30\}$ , averaged across five independent runs. Vertical error bars for each data point represent the standard deviation across five runs. (a) Average neighborhood hit score [44] (underlying data provided in the supplementary materials). (b) Per-class  $k$ NN classification accuracy, where colors represent different class categories consistent with Fig. 3. (c) MAPLE layouts for Fashion-MNIST with neighborhood sizes of  $k = 5$  and  $k = 30$ .

and neighborhood size when applying MAPLE, particularly in high-dimensional settings. Our results show that MAPLE can process datasets with up to 100,000 samples on a standard laptop within minutes, without requiring high-performance hardware.

#### D. Parameter Sensitivity Analysis

To assess the robustness of MAPLE with respect to its two primary hyperparameters, the neighborhood size ( $k$ ) and the trade-off parameter ( $\lambda$ ), we conducted a grid search analysis on the Fashion-MNIST dataset ( $N = 40,000$ ). The search space included  $k \in \{5, 15, 30\}$  and  $\lambda \in \{0, 0.1, 0.3, 0.5, 0.7, 0.9, 2.0\}$ . We evaluated the resulting layouts using our two core quantitative measures, neighborhood hit and  $k$ NN classification accuracy, to examine whether MAPLE requires careful hyperparameter tuning or offers a broad, stable behavior. The results are summarized in Fig. 6, where different line styles (e.g., dashed) denote neighborhood sizes  $k \in \{5, 15, 30\}$ .

**Observation.** The results indicate that MAPLE remains stable across a wide range of hyperparameter settings (see Fig. 6). With a training duration of 20 epochs, all MAPLE configurations outperform the UMAP baselines across the evaluated measures (based on the full results underlying Fig. 6, provided in the supplementary materials). We observe that the trade-off parameter  $\lambda$  has a relatively limited impact on the layout quality within the tested range (Fig. 6(a) and (b)). This behavior is consistent with findings in the original MMCR study [87], where maximizing the global nuclear norm of centroids implicitly encourages compact local manifolds. In our formulation, the global term provides strong separation, while the local term yields more modest refinements.

In contrast, the neighborhood size  $k$  plays a more influential role. MAPLE achieves its highest quantitative scores at smaller neighborhood sizes ( $k = 5$ ), as neighborhood hit is a highly local metric. This behavior is also reflected qualitatively in Fig. 6(c), where structurally similar classes such as ‘Ankle Boot’, ‘Sandal’, and ‘Sneaker’ (the blue–brown–grey cluster on the left) appear more clearly separated than at  $k = 30$ . We hypothesize that smaller  $k$  values yield smaller and more homogeneous neighborhood “views”, which are better aligned with the locality assumptions in MMCR-style self-supervised learning [42], [87]. To contextualize these results, we further compare MAPLE against UMAP using matched neighborhood sizes; detailed neighborhood hit scores for each method are reported in the supplementary materials. MAPLE provides neighborhood improvements at smaller  $k$ , while larger  $k$  values yield stronger relative gains. During development, we also explored soft probabilistic  $k$ NN assignments, but observed only marginal improvements over the standard formulation. This suggests that binary selection of high-confidence neighbors is sufficient for MAPLE to learn meaningful manifold structure, offering a favorable balance between simplicity and accuracy.

## V. APPLYING MAPLE TO SINGLE-CELL DATASETS

As DR methods play a central role in single-cell bioinformatics [5], [15], [49], [58], [59], we demonstrate how MAPLE supports expert-driven exploratory analysis in this domain. We present an expert-guided evaluation, in which a biologist inspected MAPLE and UMAP layouts side by side using our interactive visual tool. The tool provides standard exploration functionalities, such as zooming and panning, brushing and linking, and details-on-demand. Through this process, the expert confirmed that MAPLE supports more reliable bioinformatics analyses by validating existing annotations and aligning with established single-cell knowledge.

**Data exploration in *C. elegans* embryogenesis.** We selected the *C. elegans* dataset because every individual worm develops through the same sequence of cell divisions, and its cell lineage has been largely mapped. These characteristics make it ideal for closely examining the DR layouts to see how they can help describe gene expression patterns and trace transcriptome dynamics during differentiation and development. The dataset we use is from Packer et al. [60], which contains 86,204 cells with 20,222 measured gene expression dimensions, spanning a wide diversity of cell types. It is fully annotated, and many subsequent studies [30], [68] have validated the accuracy of these annotations. For each cell, annotations are provided at both broad levels (e.g., neurons and glia) and finer subtypes (e.g., amphid sensory neurons such as ADF, AWB), together with lineage information such as “parent” or “neuroblast” states, and rich metadata for analysis. These annotations facilitate comparison of proximity and expression differences across multiple scales of analysis.

In our study, we focused on two major subsets: neurons and glia. The expert interactively explored the layouts produced by MAPLE and UMAP using Euclidean metrics. Fig. 7(a) shows the MAPLE layout of the raw neuron and glia subsets with full metadata (displayed upon mouse hover or box

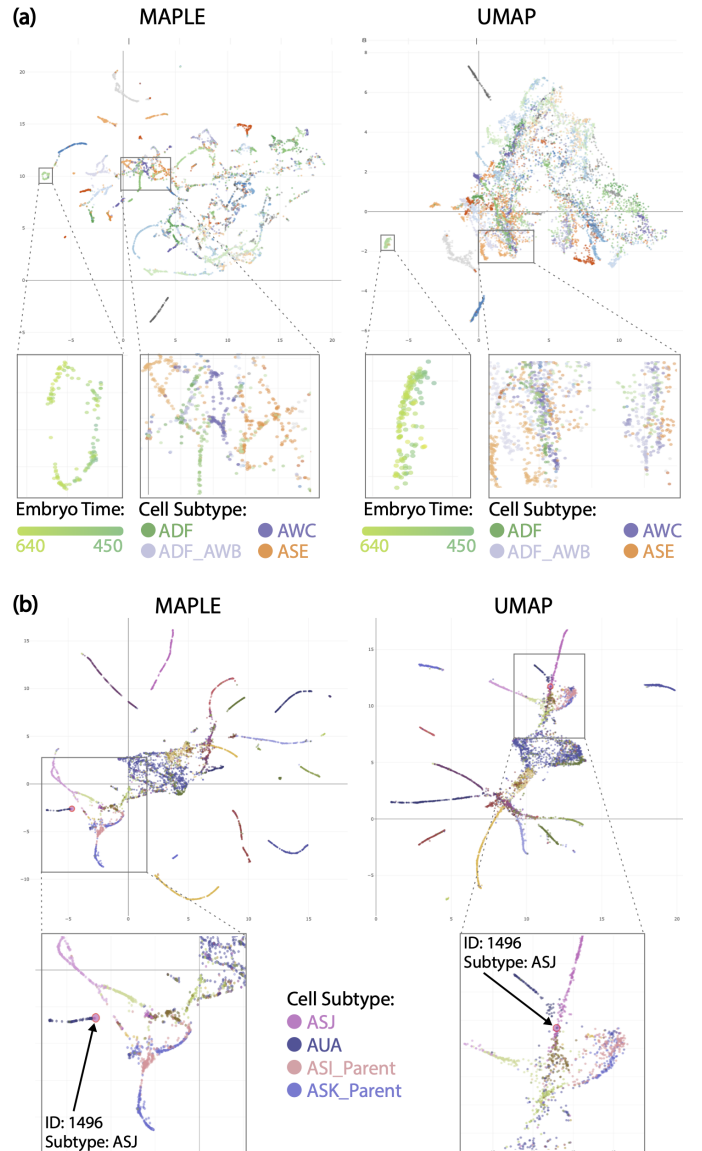


Fig. 7. Comparison of MAPLE and UMAP layouts of *C. elegans* neuron data. (a) Raw neuron subset [60] (20,222 gene dimensions; corresponds to Fig. 3, column 5). MAPLE shows clearer lineage-aligned separability than UMAP, consistent with embryo time and cell subtypes. (b) Monocle-preprocessed neuron subset (50 PCA dimensions; corresponds to Fig. 3, column 7). MAPLE provides cell subtype disambiguation and reveals transitional “bridge” cells that are obscured in UMAP.

selection), where the neuron population is further divided into more than 20 fine-grained cell subtypes, and the top panels are colored accordingly. Fig. 7(b) shows another set of *C. elegans* data after standard preprocessing with Monocle [74] (applying PCA [54], batch correction [34], and regression of the continuous background contamination factors), which reduced the dimensionality from 20,222 to 50 and the granularity of annotations into a smaller set of developmental categories. The side-by-side comparison allowed us to examine how MAPLE preserves detailed lineage information in both raw and preprocessed data while still producing meaningful substructures.

**MAPLE brings better cell subtype disambiguation.** During the evaluation, the expert confirmed that the sepa-

rations produced by MAPLE align mostly with the known *C. elegans* cell lineage. For example, as shown in Fig. 7(a), MAPLE provides clearer separation of both embryo time and known cell subtypes compared to UMAP. Amphid sensory neurons such as ADF and AWC, which often appear clustered together in UMAP, are also resolved into distinct trajectories in MAPLE. At the same time, MAPLE preserves alignments between closely related categories such as ADF and their progenitor ADF\_AWB, which are consistent with developmental biology findings [60]. This combination of proper separation and alignment with biological events gives domain experts confidence that MAPLE’s higher-resolution results more faithfully capture lineage structure than UMAP.

**MAPLE facilitates hypothesis generation.** Beyond validation, MAPLE also supports the discovery of potential new insights. As illustrated in Fig. 7(b), MAPLE not only disambiguates progenitor states such as ASI\_parent and ASK\_parent, but also highlights individual cells that act as “bridges” between trajectories [60]. For instance, a cell that UMAP merges into the AUA cluster is instead positioned by MAPLE between ASJ and AUA lineages. Such bridge-like cells may represent transitional populations where progenitors begin differentiating, which may be captured through the expression of marker genes of both descendant types [82]. Identifying these rare states is valuable because it enables researchers to generate and test hypotheses about the molecular signatures of ancestor cells and the order in which lineages branch. Those insights would remain hidden in other DR layouts.

MAPLE can be used as a stand-alone method, or as a supplementary panel to mitigate the “hallucination” of connectivity that has been reported for UMAP [69]. While UMAP is widely used in single-cell studies, it has also been criticized for generating overly confident continuous trajectories that may not exist [15], [69], [85]. This can mislead analysts into forming incorrect hypotheses, especially in datasets where annotations are lacking and uncertainty is high. MAPLE’s local compression and diversification strategy produces clearer separations, which may raise new questions about connectivity while also providing additional perspective to help validate or challenge the results of other DR methods.

## VI. DISCUSSION AND FUTURE WORK

While UMAP remains highly competitive in terms of efficiency and overall robustness, MAPLE offers a new perspective by explicitly learning fuzzy neighborhood relations through self-supervised learning.

Conceptually, MAPLE belongs to a growing line of work that focuses on the often-overlooked yet critical aspect of input-space modeling in DR methods [9], [84]. By reinterpreting neighborhood graphs as implicit views, MAPLE establishes a connection between MVSSL and DR, while introducing geometry-aware objectives to regularize local structure. At the same time, MAPLE represents a hybrid approach, uniquely combining parametric graph construction with non-parametric projection. We believe this opens up broader discussions on the multifaceted relationships between DR and SSL, as discussed in prior theoretical and empirical studies [21], [40].

Furthermore, MAPLE highlights the need to reconsider how DR methods are evaluated. For methods like ours, quantitative evaluation should move beyond statistical consistency checks between input and output space and instead develop new measures—potentially inspired by computer vision and human perception [3], [70], [89]—that can better capture downstream, task-oriented goals. In addition, approaches developed in medical image segmentation [70], where evaluation is designed to quantitatively reflect perceptual quality and task performance, could potentially be adapted to expand the current set of evaluation measures in MAPLE.

### Strengths and limitations across data characteristics.

A key characteristic that differentiates MAPLE from existing DR methods [24] is its explicit focus on resolving intra-cluster variation. By revealing higher-resolution structures within clusters, MAPLE makes internal organization more visible in the resulting layouts, which often also increases overall class separability, as shown in our quantitative evaluation results.

If the primary objective is class separability, beyond the experiments performed on real-world datasets presented in Tab. I, we observe that MAPLE’s effectiveness depends strongly on the structure of the input data. We define “simple” datasets as those with already high manifold capacity [87] and clean geometry (e.g., synthetic manifolds like the Helix or Swiss roll). On such simple datasets, MAPLE tends to provide limited additional benefit, as a standard  $k$ NN graph already captures the local manifold structure reasonably well. In contrast, MAPLE consistently outperforms standard methods on more “complex” datasets (e.g., MNIST, Fashion-MNIST), where neighborhoods exhibit noise and significant class overlap. These results suggest that MAPLE is especially effective when the graph construction benefits from additional geometric regularization.

Furthermore, we observe that preprocessing and the resulting dimensionality influenced MAPLE’s performance. Theoretical and empirical results on manifold capacity [67], together with our MAPLE experiments, suggest that performance degrades when data are heavily whitened and the embedding dimension is large but contains little orthogonal variance. For instance, after applying PCA and retaining more than 200 principal components on a simple dataset, if neighborhoods are fuzzy, MAPLE may degenerate, meaning that the global nuclear norm has no additional variance to distribute, so centroids in the embedding fail to spread out, and the low-dimensional layout contracts into compact clusters.

In summary, MAPLE is more effective for data where local neighborhoods exhibit rich, high-dimensional variation within clusters, and less effective when data are heavily whitened or already well modeled by linear similarities. Exploring these boundaries systematically and connecting them to manifold capacity theory remains an interesting direction. Meanwhile, multimodal datasets (e.g., image–text), where heterogeneous manifolds are common, present a promising application area for future work with MAPLE.

**Broader connections.** From an information-theoretic perspective, MAPLE conceptually achieves lower information loss about neighborhood connectivity compared to UMAP. Although one of MAPLE’s objectives is to compress neigh-

borhoods, what is compressed are noisy or curved directions that do not contribute to the intrinsic manifold geometry. At the same time, MAPLE enhances global separability by suppressing irrelevant local variance and reallocating it toward informative between-neighborhood differences. In this way, MAPLE preserves more task-relevant information, and the resulting low-dimensional layouts reveal a finer intra-cluster structure. This perspective also suggests opportunities for future work on developing information-theoretic measures that evaluate not only the projection process, but also the role of input-space modeling and graph refinement in DR. Developing entropy- or mutual-information-based measures could help quantify how much structure is maintained or lost at each stage, and potentially offer a principled way [65] to compare methods that differ in their treatment of neighborhood graphs or general input-space modeling.

Moreover, since MAPLE provides a parametric geometry-aware approach to construct graphs and is fully compatible with UMAP’s CE optimization, the same pipeline could in principle be adapted to other graph-based DR methods. Since methods such as *t*-SNE [76] and Laplacian Eigenmaps [6] begin with a neighborhood graph or similarity matrix before applying their respective optimization objectives, MAPLE could serve as a modular graph refinement stage that improves the reliability of the neighborhood structure while leaving the downstream projection objectives unchanged. This opens the door to integrating geometry-aware SSL objectives into a broader range of graph-based DR methods.

Finally, through our use cases, we see potential for further research on using geometry-aware SSL to improve DR in practice. In particular, such approaches can produce layouts with finer internal resolution, supporting closer inspection of subtle structure in complex, high-dimensional data. In biological applications, most existing workflows rely on linear preprocessing steps before applying DR [34], [69]. Our findings suggest that MAPLE brings the most quantitative improvements when applied directly to raw single-cell data, indicating that SSL-based DR methods could be further developed and adapted to domain-specific settings. More broadly, MAPLE illustrates how geometry- or manifold-aware SSL can enhance the expressiveness of DR and contribute to uncovering new insights in scientific data.

## VII. CONCLUSION

In this work, we introduced MAPLE, a new nonlinear dimensionality reduction method that extends UMAP with self-supervised learning objectives for neighborhood graph modeling. Through the reinterpretation of maximum manifold capacity representations, MAPLE effectively improves neighborhood construction in the embedding space. Our quantitative and qualitative evaluations show that MAPLE leads to clearer visual cluster separations and finer resolution of subclusters compared to UMAP. These results suggest that MAPLE provides a promising direction for geometry-aware SSL in DR for visual analysis. Future work would include addressing scalability through parallelization and further optimization strategies, providing formal theoretical explanations,

developing new task-oriented DR measures, and extending the framework to broader graph-based DR methods as well as multimodal contexts.

## ACKNOWLEDGMENTS

The authors wish to thank the anonymous reviewers for their valuable and insightful feedback. This work was supported through the ELLIIT environment for strategic research in Sweden.

## REFERENCES

- [1] 10x Genomics. Datasets. <https://www.10xgenomics.com/datasets>, 2025. Accessed: 2025-9-7.
- [2] D. Atzberger, T. Cech, W. Scheibel, J. Döllner, M. Behrisch, and T. Schreck. A large-scale sensitivity analysis on latent embeddings and dimensionality reductions for text spatializations. *IEEE Trans Vis Comput Graph*, 31(1):305–315, 2025. doi: 10.1109/TVCG.2024.3456308
- [3] S. S. Bae, T. Fujiwara, C. Tseng, and D. A. Szafir. Uncovering how scatterplot features skew visual class separation. In *Proc. CHI*, article no. 681, 21 pages, 2025. doi: 10.1145/3706598.3713976
- [4] D. Bear, C. Fan, D. Mrowca, Y. Li, S. Alter, et al. Learning physical graph representations from visual scenes. *Adv Neural Inf Process Syst*, 33:6027–6039, 2020.
- [5] 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. *Nat Biotechnol*, 37(1):38–44, 2019. doi: doi.org/10.1038/nbt.4314
- [6] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput*, 15(6):1373–1396, 2003. doi: 10.1162/089976603321780317
- [7] L. Bergman, N. Cohen, and Y. Hoshen. Deep nearest neighbor anomaly detection. *arXiv:2002.10445*, 2020. doi: 10.48550/arXiv.2002.10445
- [8] K. Beyer, J. Goldstein, R. Ramakrishnan, and U. Shaft. When is “nearest neighbor” meaningful? In *Proc ICDT*, pp. 217–235. Springer, 1999. doi: doi.org/10.1007/3-540-49257-7\_15
- [9] J. N. Böhm, P. Berens, and D. Kobak. Unsupervised visualization of image datasets using contrastive learning. In *Proc ICLR*, 2023. doi: 10.48550/arXiv.2210.09879
- [10] A. P. Bradley. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognit*, 30(7):1145–1159, 1997. doi: 10.1016/S0031-3203(96)00142-2
- [11] J. N. Böhm, P. Berens, and D. Kobak. Attraction-repulsion spectrum in neighbor embeddings. *J Mach Learn Res*, 23(95):1–32, 2022.
- [12] T. Caliński and J. Harabasz. A dendrite method for cluster analysis. *Commun Stat Theory Methods*, 3(1):1–27, 1974. doi: 10.1080/03610927408827101
- [13] M. Caron, P. Bojanowski, A. Joulin, and M. Douze. Deep clustering for unsupervised learning of visual features. In *Proc ECCV*, 2018.
- [14] Cellarium. Cellarium cell annotation service (CAS). <https://cellarium.ai/tool/cellarium-cell-annotation-service-cas/>, 2025. Accessed: 2025-9-7.
- [15] T. Chari and L. Pachter. The specious art of single-cell genomics. *PLOS Comput Biol*, 19(8):e1011288, 2023.
- [16] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *Proc ICML*, vol. 119, pp. 1597–1607, 2020.
- [17] S. Chung, D. D. Lee, and H. Sompolinsky. Classification and geometry of general perceptual manifolds. *Phys Rev X*, 8:031003, 26 pages, Jul 2018. doi: 10.1103/PhysRevX.8.031003
- [18] A. Coates, A. Ng, and H. Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proc AISTATS*, vol. 15, pp. 215–223, 2011.
- [19] C. Conwell, J. S. Prince, K. N. Kay, G. A. Alvarez, and T. Konkle. A large-scale examination of inductive biases shaping high-level visual representation in brains and machines. *Nat Commun*, 15(1):9383, 2024. doi: doi.org/10.1038/s41467-024-53147-y
- [20] J. P. Cunningham and Z. Ghahramani. Linear dimensionality reduction: Survey, insights, and generalizations. *J Mach Learn Res*, 16(89):2859–2900, 2015.
- [21] S. Damrich, N. Böhm, F. A. Hamprecht, and D. Kobak. From *t*-SNE to UMAP with contrastive learning. In *Proc ICLR*, 2023.

- [22] S. Damrich and F. A. Hamprecht. On UMAP's true loss function. In *Adv Neural Inf Process Syst*, vol. 34, pp. 5798–5809, 2021.
- [23] L. Deng. The MNIST database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Process Mag*, 29(6):141–142, 2012. doi: 10.1109/MSP.2012.2211477
- [24] M. Espadoto, R. M. Martini, A. Kerren, N. S. T. Hirata, and A. C. Telea. Toward a quantitative survey of dimension reduction techniques. *IEEE Trans Vis Comput Graph*, 27(3):2153–2173, 2021. doi: 10.1109/TVCG.2019.2944182
- [25] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proc KDD*, 6 pages, p. 226–231, 1996.
- [26] R. Etemadpour, B. Olk, and L. Linsen. Eye-tracking investigation during visual analysis of projected multidimensional data with 2D scatterplots. In *Proc IVAPP*, pp. 233–246, 2014.
- [27] L. Feng, C. Wang, P. Liu, K. Ge, and J. Zhang. NCLDR: Nearest-neighbor contrastive learning with dual correlation loss for dimensionality reduction. *Neurocomputing*, 594:127848, 2024. doi: 10.1016/j.neucom.2024.127848
- [28] D. Francois, V. Wertz, and M. Verleysen. The concentration of fractional distances. *IEEE Trans Knowl Data Eng*, 19(7):873–886, 2007. doi: 10.1109/TKDE.2007.1037
- [29] T. Fujiwara, X. Wei, J. Zhao, and K.-L. Ma. Interactive dimensionality reduction for comparative analysis. *IEEE Trans Vis Comput Graph*, 28(1):758–768, 2022. doi: 10.1109/TVCG.2021.3114807
- [30] W. Fung, L. Wexler, and M. G. Heiman. Cell-type-specific promoters for *C. elegans* glia. *J Neurogenet*, 34(3-4):335–346, 2020. doi: 10.1080/01677063.2020.1781851
- [31] G. D. Garson. *Factor Analysis and Dimension Reduction in R: A Social Scientist's Toolkit*. Routledge, 2022.
- [32] A. Gisbrecht, A. Schulz, and B. Hammer. Parametric nonlinear dimensionality reduction using kernel t-SNE. *Neurocomputing*, 147:71–82, 2015. doi: 10.1016/j.neucom.2013.11.045
- [33] J.-B. Grill, F. Strub, F. Alché, C. Tallec, P. Richemond, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Adv Neural Inf Process Syst*, 33:21271–21284, 2020.
- [34] L. Haghverdi, A. T. Lun, M. D. Morgan, and J. C. Marioni. Batch effects in single-cell RNA-sequencing data are corrected by matching mutual nearest neighbors. *Nat Biotechnol*, 36(5):421–427, 2018. doi: 10.1038/nbt.4091
- [35] G. Hamerly and C. Elkan. Learning the k in k-means. *Adv Neural Inf Process Syst*, 16, 2003.
- [36] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc CVPR*, June 2016.
- [37] J. Healy and L. McInnes. Uniform manifold approximation and projection. *Nat Rev Methods Primers*, 4(1):82, 2024. doi: 10.1038/s43586-024-00363-x
- [38] G. E. Hinton and S. Roweis. Stochastic neighbor embedding. *Adv Neural Inf Process Syst*, 15, 2002.
- [39] H. Hotelling. Analysis of a complex of statistical variables into principal components. *J Educ Psychol*, 24(6):417, 1933. doi: 10.1037/h0071325
- [40] T. Hu, Z. LIU, F. Zhou, W. Wang, and W. Huang. Your contrastive learning is secretly doing stochastic neighbor embedding. In *Proc ICLR*, 2023.
- [41] D. Huttenlocher, G. Klanderman, and W. Rucklidge. Comparing images using the Hausdorff distance. *IEEE Trans Pattern Anal Mach Intell*, 15(9):850–863, 1993. doi: 10.1109/34.232073
- [42] B. Isik, V. Lecomte, R. Schaeffer, Y. LeCun, M. Khona, et al. An information-theoretic understanding of maximum manifold capacity representations. In *Pror. UniReps*, 2023.
- [43] M. T. Islam and J. W. Fleischer. The shape of attraction in UMAP: Exploring the embedding forces in dimensionality reduction. *arXiv:2503.09101*, 2025. doi: doi.org/10.48550/arXiv.2503.09101
- [44] H. Jeon, A. Cho, J. Jang, S. Lee, J. Hyun, et al. ZADU: A python library for evaluating the reliability of dimensionality reduction embeddings. In *Proc VIS*, pp. 196–200, 2023. doi: 10.1109/VIS54172.2023.00048
- [45] H. Jeon, Y.-H. Kuo, M. Aupetit, K.-L. Ma, and J. Seo. Classes are not clusters: Improving label-based evaluation of dimensionality reduction. *IEEE Trans Vis Comput Graph*, 30(1):781–791, 2024. doi: 10.1109/TVCG.2023.3327187
- [46] H. Jeon, H. Lee, Y.-H. Kuo, T. Yang, D. Archambault, et al. Unveiling high-dimensional backstage: A survey for reliable visual analytics with dimensionality reduction. In *Proc CHI*, article no. 394, 24 pages, 2025. doi: 10.1145/3706598.3713551
- [47] M. Jung, T. Fujiwara, and J. Jo. GhostUMAP2: Measuring and analyzing  $(r, d)$ -stability of UMAP. *IEEE Trans Vis Comput Graph*, 2025. doi: 10.48550/arXiv.2507.17174
- [48] Y. Kalantidis, C. E. R. K. Lassance, J. Almazán, and D. Larlus. TLDR: Twin learning for dimensionality reduction. *Trans Mach Learn Res*, 2022.
- [49] D. Kobak and G. C. Linderman. Initialization is critical for preserving global data structure in both t-SNE and UMAP. *Nat Biotechnol*, 39(2):156–157, Feb. 2021. doi: 10.1038/s41587-020-00809-z
- [50] M. Kocabas, S. Karagoz, and E. Akbas. Self-supervised learning of 3D human pose using multi-view geometry. In *Proc CVPR*, June 2019.
- [51] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. *Adv Neural Inf Process Syst*, 25, 2012.
- [52] X. Liu, F. Zhang, Z. Hou, L. Mian, Z. Wang, et al. Self-supervised learning: Generative or contrastive. *IEEE Trans Knowl Data Eng*, 35(1):857–876, 2023. doi: 10.1109/TKDE.2021.3090866
- [53] A. Machado, M. Behrisch, and A. Telea. Necessary but not sufficient: Limitations of projection quality metrics. *Comput Graph Forum*, 44(3):e70101, 2025. doi: 10.1111/cgf.70101
- [54] A. Maćkiewicz and W. Ratajczak. Principal components analysis (PCA). *Comput Geosci*, 19(3):303–342, 1993.
- [55] L. McInnes, J. Healy, and J. Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv:1802.03426*, 2018. doi: 10.48550/arXiv.1802.03426
- [56] M. Meilä and J. Shi. A random walks view of spectral segmentation. In *Proc AISTATS*, vol. R3, pp. 203–208, 04–07 Jan 2001.
- [57] S. Mika, B. Schölkopf, A. Smola, K.-R. Müller, M. Scholz, and G. Rätsch. Kernel PCA and de-noising in feature spaces. *Adv Neural Inf Process Syst*, 11, 1998.
- [58] K. R. Moon, D. van Dijk, Z. Wang, S. Gigante, D. B. Burkhardt, et al. Visualizing structure and transitions in high-dimensional biological data. *Nat Biotechnol*, 37(12):1482–1492, 2019. doi: 10.1038/s41587-019-0336-3
- [59] A. Narayan, B. Berger, and H. Cho. Assessing single-cell transcriptomic variability through density-preserving data visualization. *Nat Biotechnol*, 39(6):765–774, June 2021. doi: 10.1038/s41587-020-00801-7
- [60] J. S. Packer, Q. Zhu, C. Huynh, P. Sivaramakrishnan, E. Preston, et al. A lineage-resolved molecular atlas of *C. elegans* embryogenesis at single-cell resolution. *Science*, 365(6459):eaax1971, 2019. doi: 10.1126/science.aax1971
- [61] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, et al. PyTorch: An imperative style, high-performance deep learning library. *Adv Neural Inf Process Syst*, 32, 2019.
- [62] 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 Trans Vis Comput Graph*, 14(3):564–575, 2008. doi: 10.1109/TVCG.2007.70443
- [63] F. Pedregosa, G. Varoquaux, A. Gramfort, et al. Scikit-learn: Machine learning in Python. *J Mach Learn Res*, 12:2825–2830, 2011.
- [64] X. Peng, C. Lu, Z. Yi, and H. Tang. Connections between nuclear-norm and Frobenius-norm-based representations. *IEEE Trans Neural Netw Learn Syst*, 29(1):218–224, 2016. doi: 10.1109/TNNLS.2016.2608834
- [65] B. Rodríguez Gálvez, A. Blaas, P. Rodríguez, A. Golinski, X. Suau, et al. The role of entropy and reconstruction in multi-view self-supervised learning. In *Proc ICML*, vol. 202, pp. 29143–29160, 2023.
- [66] T. Sainburg, L. McInnes, and T. Q. Gentner. Parametric UMAP embeddings for representation and semisupervised learning. *Neural Comput*, 33(11):2881–2907, 10 2021. doi: 10.1162/neco\_a\_01434
- [67] R. Schaeffer, V. Lecomte, D. B. Pai, A. Carranza, B. Isik, et al. Towards an improved understanding and utilization of maximum manifold capacity representations. In *Proc ME-FoMo*, 2024. doi: 10.48550/arXiv.2406.09366
- [68] P. Sivaramakrishnan, C. Watkins, and J. I. Murray. Transcript accumulation rates in the early *Caenorhabditis elegans* embryo. *Sci Adv*, 9(34):ead1270, 2023. doi: 10.1126/sciadv.adi1270
- [69] S. Sun, J. Zhu, Y. Ma, and X. Zhou. Accuracy, robustness and scalability of dimensionality reduction methods for single-cell RNA-seq analysis. *Genome Biol*, 20(1):269, 2019. doi: doi.org/10.1186/s13059-019-1898-6
- [70] A. A. Taha and A. Hanbury. Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool. *BMC Med Imaging*, 15(1):29, 2015. doi: doi.org/10.1186/s12880-015-0068-x
- [71] J. Tang, J. Liu, M. Zhang, and Q. Mei. Visualizing large-scale and high-dimensional data. In *Proc WWW*, 11 pages, p. 287–297, 2016. doi: 10.1145/2872427.2883041
- [72] J. B. Tenenbaum, V. d. Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. doi: 10.1126/science.290.5500.2319
- [73] C. Trapnell. Constructing single-cell trajectories. <https://cole-trapnell-lab.github.io/monocle3/docs/trajectories/>, 2022. Accessed: 2025-9-7.

- [74] C. Trapnell, D. Cacchiarelli, J. Grimsby, P. Pokharel, S. Li, et al. The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. *Nat Biotechnol*, 32(4):381–386, 2014. doi: doi.org/10.1038/nbt.2859
- [75] Y.-H. H. Tsai, Y. Wu, R. Salakhutdinov, and L.-P. Morency. Self-supervised learning from a multi-view perspective. In *Proc ICLR*, 2021.
- [76] L. van der Maaten and G. Hinton. Visualizing data using t-SNE. *J Mach Learn Res*, 9(Nov):2579–2605, 2008.
- [77] L. van der Maaten, E. Postma, and J. van den Herik. Dimensionality reduction: A comparative review. Technical Report TiCC-TR 2009-005, Tilburg University Technical Report, 2009.
- [78] J. Venna and S. Kaski. Local multidimensional scaling. *Neural Netw*, 19(6-7):889–899, 2006. doi: 10.1016/j.neunet.2006.05.014
- [79] A. A. Vergani and E. Binaghi. A soft Davies-Bouldin separation measure. In *Proc FUZZ-IEEE*, pp. 1–8, 2018. doi: 10.1109/FUZZ-IEEE.2018.8491581
- [80] N. X. Vinh, J. Epps, and J. Bailey. Information theoretic measures for clusterings comparison: is a correction for chance necessary? In *Proc ICML*, 8 pages, p. 1073–1080, 2009. doi: 10.1145/1553374.1553511
- [81] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, et al. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nat Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2
- [82] D. E. Wagner and A. M. Klein. Lineage tracing meets single-cell omics: opportunities and challenges. *Nat Rev Genet*, 21(7):410–427, 2020. doi: 10.1038/s41576-020-0223-2
- [83] Y. Wang, H. Huang, C. Rudin, and Y. Shaposhnik. Understanding how dimension reduction tools work: An empirical approach to deciphering t-SNE, UMAP, TriMAP, and PaCMAP for data visualization. *J Mach Learn Res*, 22(201):1–73, 2021.
- [84] Y. Wang, Y. Sun, H. Huang, and C. Rudin. Dimension reduction with locally adjusted graphs. In *Proc AAAI*, vol. 39, pp. 21357–21365, 2025. doi: 10.1609/aaai.v39i20.35436
- [85] L. Xia, C. Lee, and J. J. Li. Statistical method scDEED for detecting dubious 2D single-cell embeddings and optimizing t-SNE and UMAP hyperparameters. *Nat Commun*, 15(1):1753, 2024. doi: 10.1038/s41467-024-45891-y
- [86] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv:1708.07747*, 2017. doi: 10.48550/arXiv.1708.07747
- [87] T. Yerxa, Y. Kuang, E. Simoncelli, and S. Chung. Learning efficient coding of natural images with maximum manifold capacity representations. In *Adv Neural Inf Process Syst*, vol. 36, pp. 24103–24128, 2023.
- [88] J. Zbontar, L. Jing, I. Misra, Y. LeCun, and S. Deny. Barlow Twins: Self-supervised learning via redundancy reduction. In *Proc ICML*, vol. 139, pp. 12310–12320, 2021.
- [89] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc CVPR*, 2018.



**Angelos Chatzimparmpas** is a Tenured Assistant Professor in the Department of Information and Computing Sciences at Utrecht University, NL. His research interests are Information Visualization and Visual Analytics. His main research interests include visual exploration of the inner parts and the quality of machine learning (ML) models with a specific focus on making complex ML models better understandable and explainable, as well as providing reliable trust in the ML models and their results.



**Wandrille Duchemin** is a bioinformatician at the Center for Scientific Computing (sciCORE), University of Basel, Switzerland, and a trainer at the Swiss Institute of Bioinformatics. His current research interests include the relationship between gene and species phylogenies, the analysis of DNA/RNA sequencing data, and the application of machine learning methods to biological data. Duchemin received his Ph.D. degree in evolutionary biology from Lyon University, France.



**Andreas Kerren** is a full professor with the Department of Science and Technology, Linköping University, and the Department of Computer Science and Media Technology, Linnaeus University, Sweden. He holds the Chair of Information Visualization at LiU and is head of the ISOVIS group at LNU. His research interests include several areas of information visualization and visual analytics, especially visual network analytics, text visualization, and the use of visual analytics for explainable AI.



**Zeyang Huang** received the MS degree in computer science from Northwestern University, US. She is currently working toward the PhD degree at Linköping University, Sweden, focusing on Computer Science and Visualization under the guidance of Takanori Fujiwara. Her research centers around dimensionality reduction and visual analytics, with an emphasis on interpretable representation learning.



**Takanori Fujiwara** is an Assistant Professor in the Department of Computer Science at the University of Arizona, where he co-directs the Data Representation Laboratory (D-REP Lab). His research spans network science, machine learning, and visual analytics, with a particular focus on interactive dimensionality reduction and the use of 3D spaces and interfaces for high-dimensional data analysis. He is a recipient of the 2026 IEEE PacificVis Early Career Award for his foundational contributions to the field.