

NCP: Neighborhood-Preserving Non-Uniform Circle Packing for Visualization

Duan Li¹, Jun Yuan¹, Xinyuan Guo¹, Xiting Wang², Yang Liu³, Weikai Yang⁴(✉), and Shixia Liu¹

© The Author(s)

Abstract Circle packing is widely used in visualization due to its aesthetic appeal and simplicity, particularly in tasks where the spatial arrangement and relationships between data are of interest, such as understanding proximity relationships (e.g., images with categories) or analyzing quantitative data (e.g., housing prices). Many applications require preserving neighborhood relationships while encoding a quantitative attribute using radii for data analysis. To meet these two requirements simultaneously, we present a *neighborhood-preserving non-uniform circle packing* method, NCP. This method preserves neighborhood relationships between the data represented by non-uniform circles to comprehensively analyze similar data and an attribute of interest. We formulate neighborhood-preserving non-uniform circle packing as a planar graph embedding problem based on the circle packing theorem. This formulation leads to a non-convex optimization problem, which can be solved by the continuation method. We conduct a quantitative evaluation and present two use cases to demonstrate that our NCP method can effectively generate non-uniform circle packing results.

Keywords Circle packing, neighborhood preservation, graph embedding, power diagram, force-directed method

1 Introduction

Circle packing, which represents data items using tightly packed circles, has been widely used in many visualiza-

tion applications due to its visually appealing and intuitive representation [1–6]. Recent research has shed light on the usefulness of neighborhood-preserving circle packing in data analysis. For example, ArchExplorer [7] represents neural network architectures as uniform circles of equal radii and places similar architectures adjacent to each other. This helps users to compare the performance of architectures with similar structures and to gain insights into designing architectures. However, ArchExplorer does not support non-uniform circles, which limits the analysis of quantitative attributes. Consider the task of analyzing label noise in a clothing image dataset. It is difficult to identify images mislabeled as similar classes, such as ‘knitwear’ mislabeled as ‘sweaters’, because they have consistent predictions with their neighboring images. However, since these images usually have higher prediction uncertainty scores due to their ambiguous content, they can be easily identified by simultaneously encoding the uncertainty scores using radii. For example, larger circles along the cluster boundaries between the ‘knitwear’ and ‘sweater’ clusters (see Fig. 1D) clearly highlight these ambiguous images. This example highlights the need for a circle packing method that both preserves neighborhood relationships and encodes a quantitative attribute using radii. In addition, previous studies concerning circle packing and cluster-aware methods highlight the importance of compactness of the layout and the convexity of clusters in data analysis [8–11].

Based on the above analysis, we aim to simultaneously preserve neighborhood relationships and improve compactness and convexity, under the constraints that circles do not overlap and their radii encode the quantitative attribute of interest. However, these objectives and constraints can conflict. For example, improving compactness may compromise the ideal positions for neighborhood preservation. As a result, the key challenge of neighborhood-preserving non-uniform circle packing lies in effectively balancing multiple optimization objectives while satisfying constraints.

To tackle this challenge, we formulate the neighborhood-

- 1 School of Software, BNRist, Tsinghua University, 100084, Beijing, China. E-mail: {ld23, yuanj19, guoxy24}@mails.tsinghua.edu.cn, shixia@tsinghua.edu.cn.
- 2 Gaoling School of Artificial Intelligence, Renmin University, Beijing, China. E-mail: xitingwang@ruc.edu.cn.
- 3 Internet Graphics Group, Microsoft Research Asia, Beijing, China. E-mail: yangliu@microsoft.com.
- 4 Hong Kong University of Science and Technology (Guangzhou), 511453, Guangzhou, China. E-mail: weikaiyang@hkust-gz.edu.cn.

Manuscript received: 2022-01-01; accepted: 2022-01-01

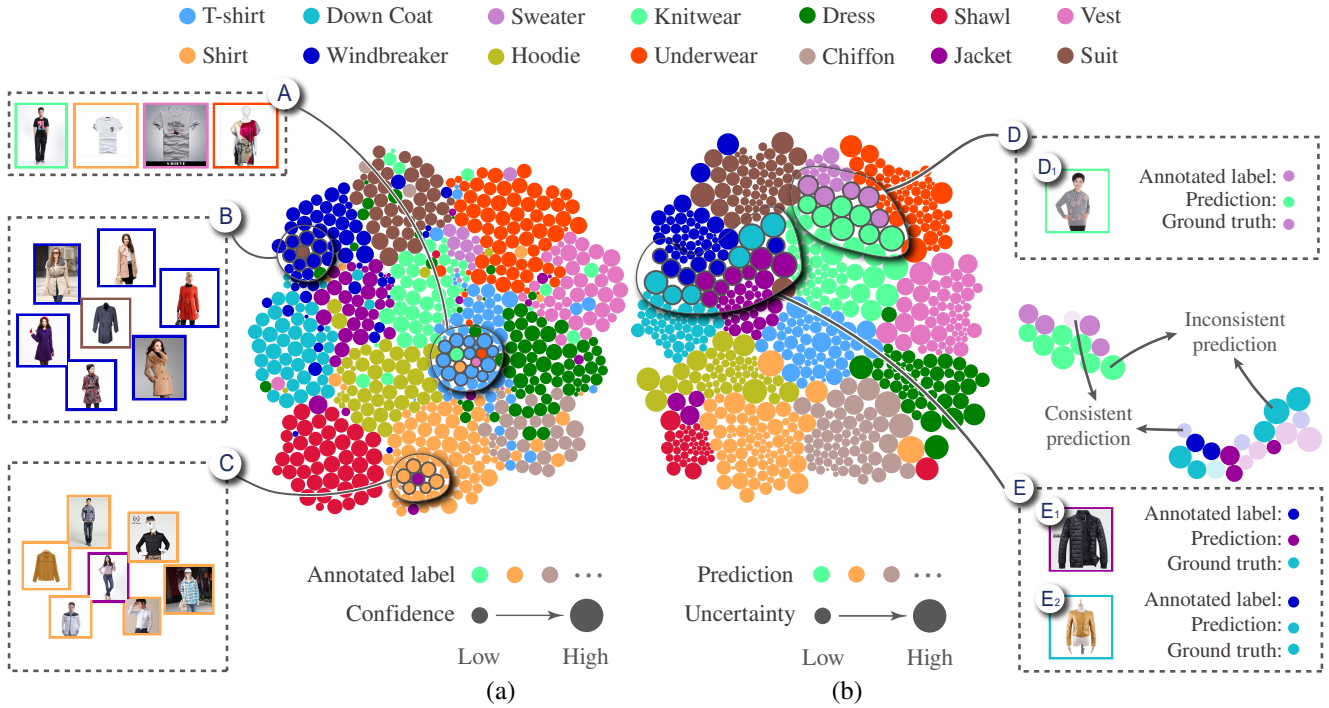


Fig. 1 Using neighborhood-preserving non-uniform circle packing to identify noisy labels in the sampled Clothing dataset: (a) analyzing images with random noise; (b) analyzing images with content-ambiguity-related noise.

preserving non-uniform circle packing problem as a maximal planar graph embedding problem. In this graph, nodes represent circles, edges represent neighborhood relationships between circles, and the embedding provides the 2D coordinates of each node. This formulation is based on the circle packing theorem [12], which establishes a one-to-one correspondence between a non-uniform circle packing and a maximal planar graph. Using this formulation, we develop a *neighborhood-preserving non-uniform circle packing* method which we call NCP. It is designed to effectively balance the objectives of neighborhood preservation, compactness, and convexity, while satisfying size and non-overlap constraints. To achieve this, we solve the associated multi-objective optimization problem with the continuation method [13] that progressively incorporates different optimization objectives. Initially, we project data items onto a 2D plane and generate a maximal planar graph and its initial embedding using the Delaunay triangulation. This planar graph represents the maximal set of neighborhood relationships that we can preserve after projection. Then, we refine this embedding by progressively introducing additional optimization objectives for compactness and convexity while preserving neighborhood relationships. This refinement integrates a power-diagram-based method for compactness and a force-directed method for convexity. The final output is the planar graph embedding and the corresponding circle packing that simultaneously

preserves neighborhood relationships and enhances compactness and convexity. Our quantitative evaluation shows that NCP can better preserve neighborhood relationships between data items than baseline methods. Furthermore, NCP achieves better convexity while obtaining comparable compactness. The usability of our method is demonstrated in two use cases and a user study. Our code is available at <https://github.com/NCP-2024/NCP>. To sum up, the main contributions of this work are:

- formulation of neighborhood-preserving non-uniform circle packing as a maximal planar graph embedding problem,
- an optimization method that simultaneously preserves neighborhood relationships, and improves compactness and convexity, and
- an open-source library for generating neighborhood-preserving non-uniform circle packings.

2 Related Work

Depending on whether spatial efficiency is optimized, circle-related layout methods can be divided into two groups: non-compact layout methods and compact layout (circle packing) methods. Non-compact layout methods, such as the Dorling cartogram [14], place the circles on a 2D plane without maximizing compactness. In contrast, circle packing methods aim to optimize spatial efficiency by tightly packing circles,

and are more closely related to our work. As this is an NP-hard problem, many stochastic optimization methods have been developed [15], based on genetic algorithms [16], simulated annealing [17], and adaptive beam search [18]. However, these methods rely on a great deal of trial-and-error, and require a considerable time to converge. To speed up the process, faster heuristic methods have been proposed, which can be categorized into three classes: front-chain-based methods, power-diagram-based methods, and force-directed methods [19].

Front-chain-based methods place each circle externally tangent to those already on the layout boundary [3]. These boundary circles form the front chain of the layout. Görtler *et al.* [2] utilized this method to generate a bubble treemap. Researchers have also extended this method to visualize time series data, where preserving the temporal sequence of data items is crucial [20, 21]. Front-chain-based methods incrementally place each circle without jointly considering their relative positions. Thus, attempts to preserve neighborhoods may result in sub-optimal results. This is demonstrated by our experiments in Section 5.1 (*SimiFC*). To tackle this issue, power-diagram-based and force-directed methods simultaneously optimize the positions of all circles to preserve neighborhood relationships.

Power-diagram-based methods employ a power diagram [22] to partition the layout region into non-overlapping cells and place circles within the corresponding cells [23]. A power diagram is a weighted Voronoi diagram that partitions the layout region based on a set of weighted distances. By iteratively adjusting circle centers to agree with the maximum inscribed circles of their respective cells and increasing circle radii, the compactness of the resulting circle packing is improved. Yu *et al.* [24] utilized this method to create photo collages, where each photo was represented as a circle with a radius proportional to its importance. A similar idea was also adopted by Liang *et al.* [25] and Rodrigues *et al.* [26]. *Force-directed methods* treat circles as physical objects and employ simulated forces to pack them tightly. Huron *et al.* [27] utilized gravitational forces and collision detection to generate a circle packing. This method has been used in different applications such as enterprise analysis [28] and social media monitoring [29].

Both power-diagram-based and force-directed methods have unique advantages in circle packing. Power-diagram-based methods achieve high compactness efficiently. On the other hand, force-directed methods provide a flexible framework for arranging circles based on simulated physical interactions, allowing the integration of various forces to meet

different optimization objectives. Despite their advantages, each method also has limitations. Power-diagram-based methods primarily focus on compactness and may not adequately preserve neighborhood relationships between circles during the optimization process. However, without good initialization, force-directed methods can face difficulties in balancing multiple competing objectives. For example, attempting to optimize both compactness and neighborhood preservation can introduce forces with conflicting directions, which may lead to slower convergence and sub-optimal local minima. This is demonstrated by our experiments in Section 5.1 (*FD*). To address these issues, our method generates a neighborhood-preserving planar graph and then tries to preserve this planar graph during subsequent optimization. We first employ an improved power-diagram-based method that simultaneously considers compactness and neighborhood preservation. The result is further refined using a force-directed method to obtain better convexity. This hybrid method leverages the strengths of both power-diagram and force-directed methods, producing a better circle packing result.

3 Optimization Objectives and Constraints

To generate a neighborhood-preserving non-uniform circle packing, we distill the design criteria from existing circle packing and cluster-aware layout methods. A detailed list is available in the supplemental materials. We organize these criteria on two levels: global and local.

The global criteria aim to generate a compact circle packing and clearly convey information from data. Previous research has identified three criteria to achieve this goal [3, 23]. First, achieving high compactness (*G1*, Fig. 2A) maximizes the utilization of the available display space [23]. Second, circle radii typically encode quantitative attributes associated with data items (*C1*), which enables efficient comparison of quantitative attributes [20, 21, 24, 30]. Third, ensuring non-overlapping circles (*C2*) reduces visual clutter, enhances the readability of individual circles, and simplifies data analysis [3, 7].

The local criteria focus on two key aspects to facilitate data analysis: enhancing the perceptual clarity of clusters and preserving neighborhood relationships within each cluster. Recent studies have demonstrated that optimizing the convexity of cluster shapes (*G2*, Fig. 2B) can improve the perceptual clarity of clusters and thus facilitate more efficient cluster analysis [8, 10]. This also aligns with the Gestalt law of perceptual grouping [31]. Additionally, preserving neighborhood relationships (*G3*, Fig. 2C) enhances the user's ability to understand inherent structures in data and accurately identify outliers [32].




Fig. 2 Optimization objectives for neighborhood-preserving non-uniform circle packing.




Fig. 3 Correspondence between the circle packing result and the planar graph: (a) example circle packing result, (b) corresponding planar graph. Red circle: referenced circle, green circles: its 1-hop neighbors, blue circles: its 2-hop neighbors.

Accordingly, our circle packing method employs three optimization objectives ($G1$ – $G3$) and two hard constraints ($C1$ – $C2$):

- $G1$: *Compactness*. Create a tightly packed arrangement of circles to maximize spatial efficiency.
- $G2$: *Convexity*. Maintain a convex shape for each cluster to enhance perceptual clarity.
- $G3$: *Neighborhood preservation*. Place circles that represent similar data items in close proximity within each cluster to preserve neighborhood contexts.
- $C1$: *Size constraint*. Ensure that the radius of each circle is proportional to its quantitative attribute for comparative analysis.
- $C2$: *Non-overlap constraint*. Arrange the circles without overlaps to maintain the distinctiveness of each circle.

4 NCP Method

4.1 Planar-Graph-Based Problem Formulation

According to the circle packing theorem [12], a one-to-one correspondence exists between a non-uniform circle packing and a maximal planar graph. In this graph, nodes represent cir-




Fig. 4 The continuation method's basic idea is to transform the original non-convex problem into a sequence of smoother ones and progressively solves them. This strategy guides optimization toward better regions in the solution space and leads to a better solution.

cles, and edges represent neighborhood relationships between circles (Fig. 3). Based on this correspondence, achieving an optimal circle packing is equivalent to identifying an appropriate planar graph and its embedding for placing circles. Given n data items and their associated quantitative attributes $\{w_i\}_{i=1}^n$ as input, our method generates a planar graph embedding and the corresponding circle packing that balances neighborhood preservation (F_p), compactness (F_c), and convexity (F_v), while satisfying the non-overlap (C_o) and size (C_z) constraints. Following common practice, we adopt a weighted sum [33] to scalarize the multi-objective optimization problem as a single-objective problem:

$$\begin{aligned} & \arg \min_{\{\mathbf{p}_i\}, s > 0} F_p + \alpha F_c + \beta F_v \\ & \text{such that } r_i = s w_i, \forall i, & (C_z), \\ & r_i + r_j \leq \|\mathbf{p}_i - \mathbf{p}_j\|, \forall i, & (C_o), \\ & 1 \leq i \leq n. & (1) \end{aligned}$$

Here, \mathbf{p}_i is the center of the i -th circle, which is the embedding coordinate of the corresponding node in the planar graph. The size constraint ($C1$) is inherently satisfied by optimizing the scaling factor s that scales radii r_i proportionally to w_i , while the non-overlap constraint ($C2$) can be achieved by reducing s to prevent circle overlaps. The parameters α and β , which balance the impact of the three terms, are determined by a grid search and set to 0.2 and 1.0 in our implementation. In this setting, our method delivers near-optimal performance across all three goals. A detailed sensitivity analysis is provided in the supplemental material.

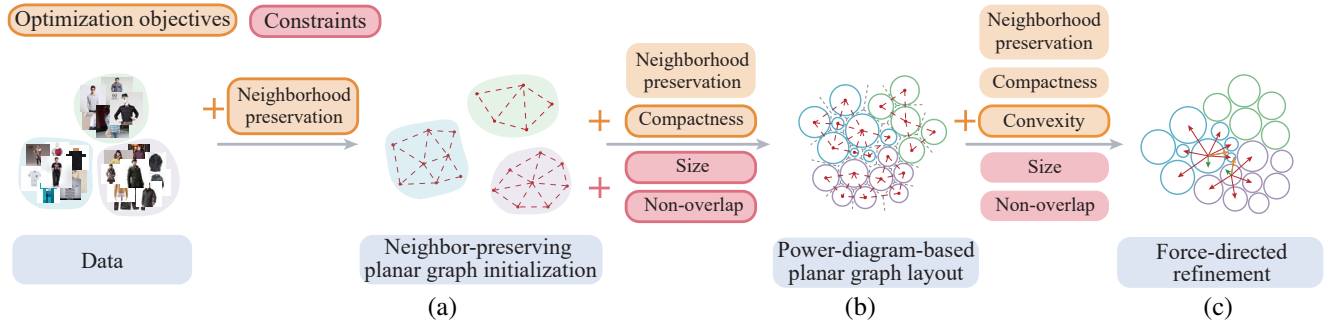


Fig. 5 Pipeline of the *NCP* method: (a) *neighborhood-preserving planar graph initialization* generates an initial planar graph that connects similar data items, (b) *power-diagram-based planar graph layout* refines the planar graph and produces an intermediate result preserving both compactness and neighborhood, (c) *force-directed refinement* compacts the circles and improves convexity.

4.2 Method Overview

Despite having a well-defined optimization objective, the non-convex nature of the problem and the large number of local minima present a challenge [34]. An effective way to overcome this issue is the continuation method [13], which has been widely used in many learning tasks such as curriculum learning [35] and semi-supervised learning [36]. As Fig. 4 illustrates, the basic idea of this method is to transform the original non-convex problem into a sequence of smoother problems that are easier to solve. Then, it sequentially solves them from the easiest one to the original one (Fig. 4A–C); each solution serves as a starting point for solving the next problem [37]. This strategy guides the optimization toward more favorable regions in the solution space and thus speeds convergence [35, 38, 39]. Its effectiveness is demonstrated by a quantitative evaluation in Section 5.1.

Accordingly, we design a three-step optimization method that progressively incorporates our three objectives. The pipeline is shown in Fig. 5. In the first step, *neighborhood-preserving planar graph initialization* (Fig. 5(a)), we create a maximal planar graph that preserves neighborhood relationships as well as possible by maximizing connections between similar data items. This initial graph serves as the starting point for subsequent optimization. The two subsequent steps involve refining the embedding of the graph to optimize compactness and convexity while preserving neighborhood relationships. Compactness is prioritized over convexity because it is affected by all circles, whereas convexity is only affected by circles on cluster boundaries [10]. Thus, in the second step, *power-diagram-based planar graph layout* (Fig. 5(b)), we generate an intermediate result to balance compactness and neighborhood preservation. In the third step, *force-directed refinement* (Fig. 5(c)), we improve convexity while also preserving neighborhood relationships and compactness by using simulated forces to adjust the circle positions.

4.3 Neighborhood-Preserving Planar Graph Initialization

To generate a neighborhood-preserving planar graph, we begin by using a projection method to place similar data items in close proximity. We then generate the edges using Delaunay triangulation, which is effective in connecting similar data items without intersections [40]. In this way, we generate a maximal planar graph in which edges represent neighborhood relationships between data items, to be preserved during the following steps.

The projection method is crucial because it decides which neighborhood relationships will be preserved. To select the most suitable method, we conducted an experiment on eight high-dimensional datasets with cluster structures, as used in Xia *et al.*'s work [41]. We identified five candidate projection methods: t-SNE [42], UMAP [43], PCA [44], MDS [45], and NMF [46], based on previous studies [41, 47]. To evaluate the ability of these projection methods to preserve neighborhoods, we used the neighborhood preservation degree [48, 49], which calculates how many neighborhoods in the high-dimensional space are preserved as neighbors in the generated planar graph.

$$F_p = -\frac{1}{n} \sum_{i=1}^n \frac{|\Gamma_G(i) \cap \Gamma_D(i, k_i)|}{|\Gamma_G(i) \cup \Gamma_D(i, k_i)|}, \quad k_i = |\Gamma_G(i)|, \quad (2)$$

where G is the generated planar graph, and $\Gamma_G(i) = \{j \mid d_G(i, j) = 1, l_i = l_j, \forall 1 \leq j \leq m, j \neq i\}$ denotes the 1-hop neighbors of the i -th data item with the same cluster label. $d_G(i, j)$ is the graph distance between the i -th and j -th data items, and l_i and l_j are their cluster labels. $\Gamma_D(i, k_i)$ denotes the k_i nearest neighbors of the i -th data item in the high-dimensional space based on data similarity.

The experimental results show that t-SNE best preserves neighborhoods, achieving an average neighborhood preservation degree of 0.392. Notably, t-SNE yielded the best result in seven out of eight cases, and ranked second in the remaining




Fig. 6 Enhancing neighborhood preservation in the power-diagram-based algorithm: (a) initial edge p_1p_2 (green); (b) node movements in Step 1 cause an edge change: the original edge p_1p_2 (green) is removed, and a new edge p_3p_4 (red) is added; (c) after moving p_1 in the Step 2, the change in edges is mitigated.

case. The circle packing results for various initialization methods are also available in the supplemental material, further illustrating t-SNE's qualitative advantages. Therefore, we use t-SNE in our implementation. Users could instead employ other methods based on their analysis needs.

4.4 Power-Diagram-Based Planar Graph Layout

Following the optimization of neighborhood preservation, the second step aims to improve compactness with the size and non-overlap constraints. To meet this additional objective and the constraints, we refine the embeddings of the previously generated neighborhood-preserving planar graph:

$$\begin{aligned} \arg \min_{\{p_i\}_{i=1}^n, s>0} \quad & F_p + \alpha F_c \\ \text{such that} \quad & C_z, C_o, \end{aligned} \quad (3)$$

The first term (F_p) denotes neighborhood preservation, and the second term (F_c) denotes compactness. For the first term, instead of directly optimizing the discrete term F_p as defined in Eq. (2), we use a continuous form F'_p to facilitate optimization. Here, $F'_p = \sum_{i=1}^m \sum_{j \in \Gamma_G(i)} \|p_i - p_j\|/s$, which encourages circle pairs connected by edges to remain close. For the second term, we follow the method in MosaicSets [9] and use $F_c = \sum_{i=1}^m \|p_i - O\|/s$, which encourages circles to be placed near the center O to enhance compactness.

To solve the optimization problem defined in Eq. (3), we leverage the power-diagram-based method developed by Zhao *et al.* [23] to improve compactness while ensuring the constraints. This method moves nodes toward the centers of the maximum inscribed circles within their respective Voronoi cells and enlarges the scaling factor. However, this may change neighborhood relationships between circles, which leads to the removal or addition of corresponding edges in the planar graph. Fig. 6 provides an illustrative example where the original edge p_1p_2 (in green) is removed, and a new edge p_3p_4 (in red) is added after moving p_1 . This is because p_1 now lies outside the weighted circumcircle of $p_2p_3p_4$ (the orange dotted circle) [50]. To address this issue, we move p_1 towards




Fig. 7 Parallel computation of power diagrams: (a) cluster nodes based on their initial positions and partition the layout region; (b) redistribute the nodes in each cluster with 2D transformations; (c) construct the power diagram for each cluster.

the center of their corresponding weighted circumcircles $p_2p_3p_4$ (Figs. 6(b)–(c)). For each adjustment, we start from a predefined movement distance and then iteratively reduce it by half until the objective function ($F'_p + \alpha F_c$) is improved. This ensures fast convergence when alternatively optimizing compactness and neighborhood preservation.

In addition, we propose two practical improvements to accelerate this method. First, given the time-consuming nature of calculating the maximum inscribed circle center within a cell, we approximate it by calculating its centroid, which is less computationally demanding [51]. Second, since *NCP* focuses on preserving neighborhood relationships within each cluster, we parallelize the computation of power diagrams across different clusters (Fig. 7). We start by clustering the nodes based on their positions using *k*-means [52], with the optimal number of clusters selected via the silhouette score [53]. The layout region is divided into several sub-regions based on the clustering result (Fig. 7(a)). To improve compactness, we restrict the size of each sub-region proportionally to the total area of circles within each cluster using the capacity-constraint power diagram [54]. To more evenly distribute the corresponding nodes in each sub-region, we apply 2D transformations, including translation, rotation, and scaling, to the sub-graph within each cluster (Fig. 7(b)), aiming to enlarge it and fill the sub-region as much as possible. These transformations preserve the edges in the initial planar graph and minimize the need for additional node movements in subsequent optimization, which contributes to better neighborhood preservation. Finally, the power diagram of each cluster is generated concurrently (Fig. 7(c)).

4.5 Force-Directed Refinement

The final step incorporates convexity, which brings us to the original problem of jointly optimizing all objectives while satisfying the size and non-overlap constraints:

$$\begin{aligned} \arg \min_{\{p_i\}_{i=1}^n} \quad & F'_p + \alpha F_c + \beta F_v, \\ \text{such that} \quad & C_z, C_o. \end{aligned} \quad (4)$$

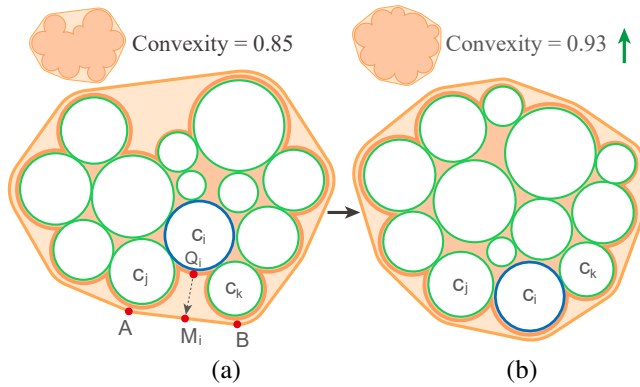


Fig. 8 Example of improving the convexity of cluster shapes: (a) move the blue circle on the cluster boundary toward the convex hull of the cluster; (b) calculate the movement direction of the blue circle to form a more convex boundary.

The first two terms encourage neighborhood preservation (F'_p) and compactness (F_c), as defined in the second step. The third term (F_v) is introduced to promote the formation of convex cluster shapes by adjusting the circles on cluster boundaries. Inspired by EulerSmooth [55], the basic idea is to move the circles on the cluster boundary toward the convex hull, such as the blue circle c_i in Fig. 8(a). The key is to determine the movement directions for these circles. For each circle on the cluster boundary, we identify two neighboring circles on the left and right that are also on the boundary and intersect the convex hull. The goal is to improve the convexity in this local region without compromising neighborhood preservation. To achieve this, we move the circle to close this gap while minimizing movements of other circles to preserve the neighborhood structure. For example, in Fig. 8(a), we move c_i to close the gap between adjacent circles c_j and c_k that touch the convex hull. These two circles require only minimal movement to accommodate c_i between them. Specifically, suppose the convex hull touches c_j and c_k at points A and B , with M_i representing the midpoint of AB . We identify the point Q_i on c_i that is closest to M_i . We then move c_i from Q_i towards M_i . Accordingly, $F_v = \sum_{i \in \delta} \|Q_i - M_i\|$, where δ represents the set of circles on the cluster boundaries (Fig. 8(b)).

To solve the optimization problem defined in Eq. (4), we employ a force-directed method to refine the circle positions. Here, the forces are set to the negative gradient of the objective function. To avoid overlaps between circles, we utilize a 2D physics engine, Box2D [56], to simulate forces. The simulation velocity is adjusted using a linear annealing decay for better convergence.

5 Evaluation

In this section, we first present a quantitative evaluation to show the effectiveness of the proposed *NCP* method. We then present two use cases to demonstrate how *NCP* can be used to facilitate data analysis.

5.1 Quantitative Evaluation

5.1.1 Experimental Setting

Datasets The quantitative evaluation was conducted on eight high-dimensional datasets with cluster structures, previously used in Xia *et al.*'s work [41]. These datasets vary in size from 155 to 1083 data points with 7 to 192 dimensions. In practical applications, the number of samples simultaneously displayed in non-uniform circle packing typically does not exceed one thousand [1]. Displaying more samples reduces the circle radii, making it difficult to discern radius differences on standard monitors. Similarity between data items was computed based on the features provided in the datasets. The radii of the circles were generated from two distinct uniform distributions, one spanning from 0.1 to 1 and the other from 0.5 to 1. The two distributions were employed to simulate cases where the circle radii had large or small variance.

Baseline methods As no existing circle packing methods simultaneously consider non-uniform circles and neighborhood preservation, we extended two representative circle packing methods to provide baselines. In addition, we considered a force-directed method, to demonstrate the effectiveness of the continuation method. We did not include a power-diagram-based method because it fails to produce a compact circle packing result.

SimiFC (Similarity-aware Front-Chain). We extended the front-chain-based method [3] to improve neighborhood preservation by improving its strategy for placing circles. In each iteration, it first identifies the position outside the front chain that is closest to the layout center and then places the circle that maximizes neighborhood preservation.

AEF (ArchExplorer + Force-directed). The neighborhood-preserving circle packing method in ArchExplorer places uniform circles on a hexagonal grid and then greedily swaps them to maximize neighborhood preservation. We extended this method to support the packing of non-uniform circles by combining it with the force-directed method. The extension is referred to as *AEF*. *AEF* initially determines the positions of circles using the original method without considering their radii. The positions are then refined using the force-directed method described in Section 4.5.

Table 1 Performance comparison in terms of neighborhood preservation (NP_1 , NP_2), compactness, and convexity.

Dataset	NP_1				NP_2				Compactness				Convexity			
	Baseline		Ours		Baseline		Ours		Baseline		Ours		Baseline		Ours	
	SimiFC	AEF	FD	NCP	SimiFC	AEF	FD	NCP	SimiFC	AEF	FD	NCP	SimiFC	AEF	FD	NCP
Boston	0.255	0.235	0.285	0.341	0.270	0.287	0.394	0.423	0.865	0.882	0.889	0.889	0.507	0.665	0.837	0.809
Dermatology	0.220	0.256	0.290	0.311	0.238	0.278	0.370	0.385	0.858	0.868	0.876	0.874	0.582	0.683	0.796	0.790
Ecoli	0.248	0.233	0.298	0.344	0.242	0.271	0.378	0.411	0.845	0.876	0.880	0.874	0.633	0.821	0.802	0.831
ExtYaleB	0.371	0.321	0.406	0.422	0.364	0.355	0.476	0.487	0.850	0.876	0.872	0.868	0.590	0.782	0.788	0.806
MNIST64	0.203	0.203	0.276	0.306	0.207	0.224	0.349	0.372	0.833	0.866	0.867	0.863	0.455	0.713	0.821	0.827
Olive	0.210	0.199	0.271	0.328	0.213	0.235	0.382	0.414	0.844	0.866	0.875	0.862	0.554	0.713	0.784	0.817
Weather	0.369	0.320	0.352	0.448	0.355	0.366	0.527	0.567	0.850	0.872	0.877	0.868	0.389	0.683	0.843	0.795
World12D	0.194	0.253	0.368	0.379	0.252	0.305	0.462	0.463	0.866	0.888	0.900	0.896	0.596	0.630	0.782	0.809
Average	0.259	0.253	0.318	0.360	0.268	0.290	0.417	0.440	0.851	0.874	0.880	0.874	0.538	0.711	0.807	0.811

(a) Performance when circle radii have a large variance.

Dataset	NP_1				NP_2				Compactness				Convexity			
	Baseline		Ours		Baseline		Ours		Baseline		Ours		Baseline		Ours	
	SimiFC	AEF	FD	NCP	SimiFC	AEF	FD	NCP	SimiFC	AEF	FD	NCP	SimiFC	AEF	FD	NCP
Boston	0.257	0.232	0.315	0.365	0.251	0.260	0.396	0.437	0.851	0.870	0.881	0.888	0.557	0.657	0.842	0.833
Dermatology	0.233	0.284	0.319	0.364	0.259	0.278	0.378	0.399	0.845	0.868	0.872	0.875	0.589	0.691	0.842	0.845
Ecoli	0.222	0.257	0.360	0.350	0.236	0.286	0.415	0.416	0.841	0.862	0.870	0.866	0.516	0.767	0.820	0.821
ExtYaleB	0.393	0.375	0.442	0.476	0.349	0.392	0.486	0.505	0.842	0.864	0.870	0.862	0.500	0.757	0.801	0.808
MNIST64	0.215	0.235	0.274	0.327	0.203	0.231	0.335	0.375	0.828	0.856	0.874	0.854	0.460	0.720	0.828	0.808
Olive	0.219	0.235	0.243	0.342	0.215	0.245	0.322	0.422	0.838	0.861	0.870	0.858	0.489	0.615	0.809	0.839
Weather	0.408	0.355	0.177	0.483	0.363	0.365	0.232	0.578	0.837	0.865	0.891	0.860	0.366	0.658	0.743	0.816
World12D	0.306	0.305	0.399	0.380	0.310	0.320	0.476	0.477	0.857	0.870	0.896	0.885	0.480	0.632	0.800	0.796
Average	0.281	0.285	0.316	0.386	0.273	0.297	0.380	0.451	0.842	0.864	0.878	0.868	0.495	0.687	0.810	0.825

(b) Performance when circle radii have a small variance.

FD (Force-Directed). This basic method skips the power-diagram-based planar graph layout step. Specifically, we first placed circles based on t-SNE projection and then directly optimized their positions using the force-directed method. The weights in the optimization objective Eq. (4) are the same as those used in the *NCP* method.

Parameters The perplexity of t-SNE was set to 15 to give the best neighborhood preservation, as determined by a grid search. The force-directed method ran 1,250 iterations for *NCP* and *AEF*, and 10,000 iterations for *FD* to guarantee convergence on all datasets.

5.1.2 Measures

We evaluated the circle packing results based on neighborhood preservation, compactness, and convexity. The scores of these measures range from 0 to 1; higher scores indicate better results.

Neighborhood preservation This measure quantifies neighbor overlap between the high-dimensional data and their 2D embeddings. It is defined based on the neighborhood




Fig. 9 Envelope (thick orange curve) and convex hull (thin orange curve) of a set of tightly packed circles.

preservation degree from Eq. (2):

$$NP_1 = \frac{1}{m} \sum_{i=1}^m \frac{|\Gamma_G(i) \cap \Gamma_D(i, k_i)|}{|\Gamma_G(i) \cup \Gamma_D(i, k_i)|}, \quad k_i = |\Gamma_G(i)|.$$

Moreover, to assess the capability to preserve a larger range of neighborhoods, we introduce NP_2 to consider 2-hop neighbors of each node in the planar graph:

$$NP_2 = \frac{1}{m} \sum_{i=1}^m \frac{|\Gamma'_G(i) \cap \Gamma'_D(i, k'_i)|}{|\Gamma'_G(i) \cup \Gamma'_D(i, k'_i)|}, \quad k'_i = |\Gamma'_G(i)|.$$

Here $\Gamma'_G(i) = \{j \mid d_G(i, j) \leq 2, i \neq j, \forall 1 \leq j \leq m, j \neq i\}$ is the set of 2-hop neighbors of the i -th node.

Compactness We evaluated compactness following Liang *et al.* [25]. Compactness is defined as the ratio of the total area covered by all circles to the area of their envelope:

$$\text{Compactness} = \frac{\text{Area}(\cup_i c_i)}{\text{Area}(\text{Env}(\cup_i c_i))},$$

where c_i denotes the i -th circle. For any region Ω , $\text{Area}(\Omega)$ denotes its area. $\text{Env}(\Omega)$ denotes its envelope, which is defined as the shape formed by its outer boundary [2] (see the thick orange curve in Fig. 9).

Convexity We evaluated convexity based on a popular measure that is defined as the ratio of the area of a shape's envelope to that of its convex hull [57]. Since circles within a cluster may be divided into multiple connected components in the results, the envelope of a cluster is defined as the union of the envelopes of such connected components. The convex hull of a cluster is defined in the same manner. Consequently, the convexity of the j -th cluster is expressed as $\text{Area}(\text{Env}(\cup_{i=j} c_i)) / \text{Area}(\text{CH}(\cup_{i=j} c_i))$, where $\text{CH}(\Omega)$ denotes the convex hull of Ω (the light orange curve in Fig. 9). The convexity of a circle packing is defined as the average convexity across all clusters:

$$\text{Convexity} = \frac{1}{L} \sum_{j=1}^L \frac{\text{Area}(\text{Env}(\cup_{i=j} c_i))}{\text{Area}(\text{CH}(\cup_{i=j} c_i))},$$

where L is the number of clusters.

5.1.3 Results

Our experimental results are presented in Table 1. Overall, our method (*NCP*) achieved performed better than the two baseline methods (*SimiFC* and *AEF*) in terms of neighborhood preservation and convexity while showing comparable performance on compactness. The reason for the lower neighborhood preservation and convexity of the two baseline methods is that they adopt greedy strategies instead of projection methods to place circles. Such strategies result in suboptimal circle arrangements for preserving neighborhood relationships and do not guarantee the formation of convex cluster shapes. Consequently, we mainly focus on the comparison between *NCP* and the basic method *FD*.

The results show that *NCP* achieves better convexity and neighborhood preservation without sacrificing compactness: *NCP* obtains an improved solution via the continuation method. In contrast, *FD* directly optimizes the final objective function, making it prone to becoming trapped in unfavorable local optima, a common pitfall in circle packing optimization. Specifically, for neighborhood preservation, *NCP* achieved

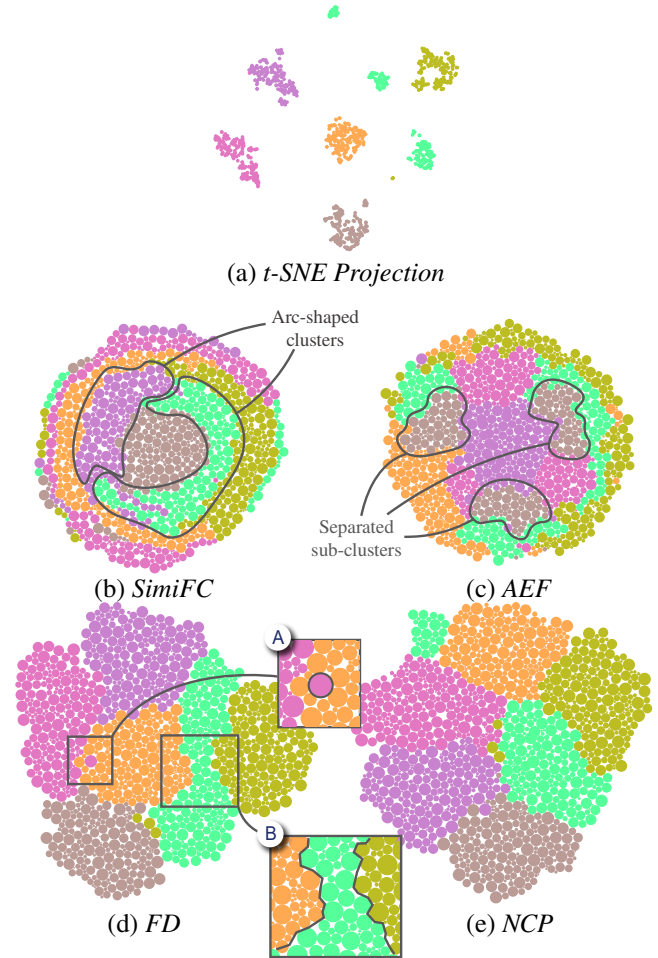


Fig. 10 Comparison of circle packing results generated by different methods on the MNIST64 dataset.

average improvements of 0.042 and 0.070 for NP_1 , and 0.023 and 0.071 for NP_2 under the two cases of radii variances. Since *NCP* and *FD* use the same planar graph initialization, this gap was caused by the subsequent optimization process. Unlike *FD*, which optimizes these competing objectives simultaneously, *NCP* utilizes the continuation method to add new objectives progressively, reducing the risk of becoming trapped in local minima in the complex original problem. *NCP* achieves a similar average compactness to *FD*: both methods adopt the force-directed method as the final step in their optimization processes, which leads to similar compactness. This conclusion is also supported by the result of *AEF*, which also employs the force-directed method and achieves similar compactness. *NCP* showed average convexity improvements of 0.004 and 0.015 in the two cases of radii variances. The main reason is the parallelized computation strategy in our enhanced power-diagram-based planar graph layout which creates the power diagram for each cluster in a convex sub-region, providing an initialization in which the clusters have

more convex shapes for subsequent force-directed refinement. Therefore, it is easier for *NCP* to achieve better convexity.

We also visually compared the circle packing results of these methods to provide a more intuitive explanation of their differences. Full results for all datasets are provided in the supplemental material. Fig. 10 shows the results for MNIST64 where the circle radii have a large variance. *SimiFC* generates arc-shaped clusters (Fig. 10(b)), and *AEF* may separate a cluster into several sub-clusters (the three sub-clusters of brown circles in Fig. 10(c)). Compared to the projection result (Fig. 10(a)), these two layout results distort the cluster structures and hinder cluster analysis. In contrast, *FD* (Fig. 10(d)) and *NCP* (Fig. 10(e)) well preserve the cluster structures and their relative positions, making clusters easy to recognize. However, *FD* displaces circles more frequently. In Fig. 10A, one pink circle with thick borders is incorrectly placed inside the cluster of orange circles. This displacement could lead to incorrect identification of outliers. Instead, *NCP* correctly places these circles in their respective clusters. Furthermore, *FD* creates irregular and concave boundaries between clusters, such as the boundaries of the light green cluster in Fig. 10B. In contrast, *NCP* forms smooth and convex boundaries between the clusters.

5.1.4 Running Time

We evaluated the running time of different circle packing methods on a desktop PC with a 3.00 GHz Intel i9-13900K CPU. We report results averaged over five trials and two cases of radii variances per dataset to reduce randomness. Table 2 shows that *SimiFC* was the fastest since it employed a simple greedy strategy to place circles. However, it does not produce satisfactory packing results, as reported in Table 1. The other three methods take longer time because they incorporate a force-directed method. Among them, *NCP* is the fastest and generates circle packing results for a dataset with about 1,000 data items in 2 seconds. Its high efficiency is due to the integration of the power-diagram-based method, which provides a good starting point for the force-directed method and thus leads to faster convergence (1,250 iterations vs. 10,000 iterations for *FD*).

5.2 Use Cases

We showcase the application of *NCP* to data analysis using two real-world datasets: Clothing [58] and Boston Housing [59].

5.2.1 Clothing

In this use case, we illustrate how *NCP* can be applied to identify and analyze label noise in an image classification dataset. We used the Clothing dataset, which contains 14

Table 2 Comparison of running time (in seconds).

Dataset	Size	SimiFC	AEF	Ours-FD	Ours-NCP
Boston	155	0.004	0.717	4.073	0.212
Dermatology	259	0.010	1.509	6.376	0.480
Ecoli	336	0.018	2.161	9.598	0.534
ExtYaleB	320	0.018	1.745	9.399	0.760
MNIST64	1,083	0.219	12.108	17.587	2.002
Olive	572	0.059	3.352	9.744	0.965
Weather	366	0.023	2.404	8.906	0.706
World12D	151	0.004	0.782	2.761	0.349




Fig. 11 Visualizing the Clothing dataset with a scatterplot.

categories of clothing images, with 38.5% of them reported as mislabeled. Alice is a graduate student. She aimed to identify and analyze the label noise in this dataset. To achieve this, Alice trained a ResNet-50 model (accuracy: 76.4%) on this dataset for feature extraction and similarity calculation. Subsequently, she randomly sampled 700 images for further analysis. Based on her prior knowledge, Alice focused on two types of label noise. The first type is *random noise*, which is often introduced by spammer annotators and is, therefore, irrelevant to the image content. For example, an image of ‘sweater’ is mislabeled as ‘shawl’, even though there are clear visual differences between the two categories. Despite such mislabeling, these images can usually be correctly predicted as their ground-truth labels with high confidence scores. The second type is *content-ambiguity-related noise*. They mainly come from images with ambiguous content that is hard to categorize, such as the images of ‘knitwear’ and “sweater.” These images are usually predicted with low confidence scores, which also means high uncertainty scores.

Analyzing random noise. Initially, Alice visualized the projection results of the sampled images using a scatterplot (Fig. 11). She used colors to encode their annotated labels and radii to encode the confidence scores, a larger radius indicating a higher score. This facilitated the identification of images with annotated labels that differ from their neighbors and have high prediction confidence scores. Alice observed a cluster mainly consisting of light blue circles, while it also included a few circles of different colors. However, the circles

overlapped each other, causing visual clutter which hindered the analysis. To address this issue, she then employed the *NCP* method to generate a neighborhood-preserving circle packing result with non-overlapping circles that facilitated sample-level analysis (Fig. 1(a)). The result preserved the cluster structures well, as circles of the same color are mostly grouped together. Alice turned her attention to the previously identified cluster and easily identified those large circles of different colors placed inside this cluster. She suspected that they represented images with random noise because their confidence scores were high. After examining the associated images, Alice confirmed that several images of ‘T-shirt’ were mislabeled as other labels, such as ‘knitwear’, ‘shirt’, ‘vest’, and ‘underwear’ (Fig. 1(a)). She also checked other regions with circles that showed similar visual patterns and identified more images with random noise. For example, in region B, a brown circle is placed inside the cluster of other blue circles. It represents a ‘windbreaker’ image mislabeled as a ‘suit’. Similarly, in region C, an image of a ‘shirt’ is mislabeled as a ‘jacket’ (purple circle), placed inside the ‘shirt’ cluster (orange circles).

Analyzing content-ambiguity-related noise In Fig. 1(a), circles representing images with low confidence scores are hard to recognize due to their small radii. Therefore, it is hard to identify images with content-ambiguity-related noise since their confidence scores are usually low. Alice then changed the encoding and used radii to encode the uncertainty scores, which effectively highlights those ambiguous samples. Since these images are typically located on the decision boundaries between different predictions, Alice used colors to encode their predicted labels to facilitate the identification of these samples (Fig. 1(b)).

First, Alice analyzed region D in Fig. 1, where there were large circles on the boundary between the ‘knitwear’ and ‘sweater’ clusters. This indicates the model’s difficulty in predicting the correct labels for these images. Upon inspection, Alice identified several cases where the model’s predictions were inconsistent with the annotated labels. Amongst them, some images were correctly labeled but mispredicted by the model. For example, D_1 is labeled correctly as a ‘sweater’, but the model mispredicts it as ‘knitwear’.

Second, Alice moved to a complex case, region E, characterized by many large circles from three clusters. The corresponding images belonged to three categories, ‘windbreaker’, ‘down coat’, and ‘jacket’. Some were predicted inconsistently with their annotated labels due to similarity of appearance. For example, both E_1 and E_2 are a ‘down

coat’ but are mislabeled as a ‘windbreaker’. E_1 is incorrectly predicted as a ‘jacket’, while E_2 is predicted as ground-truth label ‘windbreaker’. Alice noticed that the content-ambiguity-related noise within these categories led to many inconsistent predictions with the annotated labels and thus hindered the model from correctly classifying such images.

5.2.2 Boston Housing

In the second use case, we illustrate how *NCP* can be applied to analyze tabular datasets. We used the Boston Housing dataset, containing 506 samples. Each sample represents a town associated with its housing price along with 13 quantitative attributes derived from multiple factors that affect the housing price, such as educational resources and the environment. Following previous works [60, 61], these quantitative attributes were normalized to derive the features of each sample and calculate their similarity. Bob is a real estate agent. He wanted to analyze how different factors affect housing prices and compare housing prices in similar towns. To this end, he employed *NCP* to generate a circle packing that placed similar towns together, simultaneously encoding the housing price and a selected quantitative attribute using different channels: color and radius (Fig. 12). Here, a larger radius corresponds to a higher housing price, and a darker color indicates a higher value for the selected attribute.

Initially, Bob wanted to explore how educational resources affected housing prices. Therefore, he used color to encode ‘pupil-teacher ratio’. A lighter color represents a lower ratio, indicating more educational resources. He quickly identified a region where several large, light-colored circles were gathered (region A in Fig. 12(a)). This finding suggested that abundant educational resources contributed to high housing prices. Adjacent to this region, Bob noted another group of large circles with darker colors (region B in Fig. 12(a)). He was curious why these towns had fewer educational resources but still maintained high housing prices compared to those in region A. To explore this, he encoded other quantitative attributes using colors. He then analyzed the color distributions of these attributes in regions A and B. This analysis revealed a clear difference in the distribution of ‘nitrogen oxide concentrations’. As shown in Fig. 12(b), the towns in region B had lighter colors than those in region A, which indicated lower pollution levels. This suggested that the towns in region B had a better natural environment, which played a key role in their high housing prices. Such findings are difficult to obtain through statistical correlations alone, as they involve only a small subset of samples. In contrast, by visually placing similar houses together and encoding housing prices as circle radii, these findings are more readily identifiable.

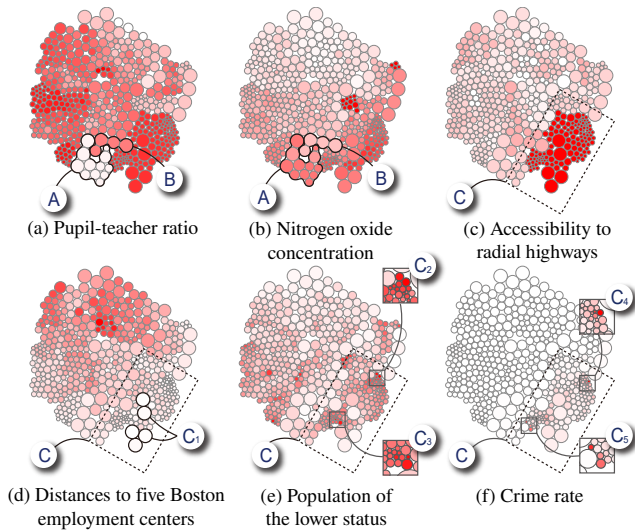


Fig. 12 Circle packing results for the Boston Housing dataset. Larger circles represent higher housing prices; darker colors indicate higher values for the selected quantitative attribute.

Subsequently, Bob wanted to explore the influence of transportation on housing prices. To this end, he encoded ‘accessibility to radial highways’ using colors. He observed that the circles in region C had the highest values (Fig. 12(c)). However, the large variance in their radii suggested the presence of other factors that affected their housing prices. To investigate these factors, Bob encoded other quantitative attributes using colors and examined their distributions within this region. The large circles were associated with shorter ‘distances to five Boston employment centers’ (C_1 in Fig. 12(d)). This factor also contributes to their high housing prices. On the contrary, the small circles are associated with larger ‘populations with lower status’ (C_2, C_3 in Fig. 12(e)) or higher ‘crime rates’ (C_4, C_5 in Fig. 12(f)), which account for their low housing prices.

5.3 User Study

In addition to the quantitative evaluation and the use cases, we also conducted a user study to demonstrate the effectiveness and usefulness of NCP in data analysis. The study used 9 real-world datasets: 8 used in Sec. 5.1 and 1 used in Sec. 5.2.1.

5.3.1 Study Setup

Participants We recruited 16 participants (P1–P16), aged 24 to 32 years, comprising graduate students and professors with extensive experience (≥ 2 years) in visual analytics and information visualization. Upon completion, each participant received \$15 compensation, regardless of performance.

Baseline methods We used 5 baseline methods: 3 circle packing methods (*SimiFC*, *AEF*, and *FD*) used in Sec. 5.1, and 2 scatterplot methods (*Scatter-S* and *Scatter-L*) that allow

overlap between circles. In the scatterplot methods, circle positions were determined by t-SNE, consistent with the initialization results of our method. Scatter-S uses the *size* of circles to encode quantitative attributes, while Scatter-L uses color *lightness*. These two methods are widely used in data analysis [62].

Task Design Our study consisted of three tasks commonly used in visual analytics [63–65].

Task 1—Cluster Identification: Participants were required to identify the cluster to which a highlighted circle belonged. They could use a lasso or click to select the circles in the identified cluster and submit their answer.

Task 2—Outlier Identification: Participants were required to judge whether a highlighted circle was an outlier. They could analyze the neighborhood relationships around the highlighted circle before submitting their answer.

Task 3—Quantitative Attribute Comparison: Participants were required to compare three highlighted circles and identify the one corresponding to the largest quantitative value. For *Scatter-L*, this was achieved by comparing the lightness of their colors, while for the other methods, it was done by comparing their sizes.

Study Protocol Participants started by signing a consent form and watching a tutorial video about the study procedure, system interactions, and tasks. Following a within-subjects design, each participant was required to evaluate six different methods and finish all three tasks sequentially. Each task consisted of both a practice and a test session. In the practice session, participants were required to finish six trials, one for each method. After completing the practice session and confirming that they fully understood the tasks and methods, participants proceeded to the test session. Participants were allowed to take short breaks whenever they requested one. Upon completing each task, we assessed the participants’ workloads and fatigue levels using NASA’s Task Load Index [66] with brief descriptions of the six methods, and collected their feedback on the methods. For each trial, we recorded participants’ answers and completion times.

To control the experiment duration and reduce the learning effect, the nine datasets were evenly distributed across three tasks, with each task assigned three datasets. Thus, each participant completed 54 trials (3 tasks \times 3 datasets \times 6 methods). The entire study lasted 45–60 minutes. The method order was also counterbalanced to reduce the learning effect.

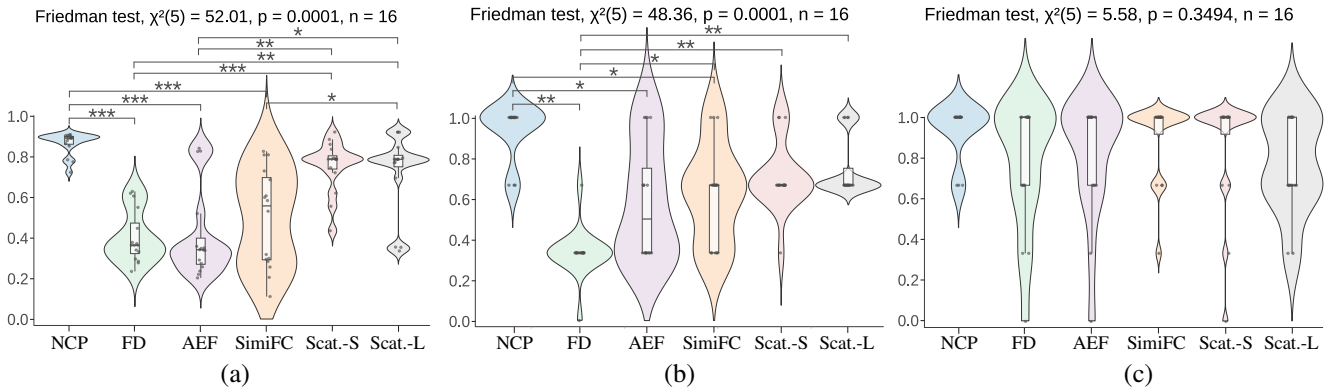


Fig. 13 User study results for three tasks: (a) *Task 1—Cluster Identification*; (b) *Task 2—Outlier Identification*; (c) *Task 3—Quantitative Attribute Comparison*. Here, * indicates $p < 0.05$, ** indicates $p < 0.01$, *** indicates $p < 0.001$. Scat.-S stands for Scatter-S, and Scat.-L stands for Scatter-L.

5.3.2 Result Analysis

We analyzed both task accuracy and participants' subjective ratings for workload and fatigue.

Accuracy For each task, we computed participants' average accuracy across different methods. We conducted Friedman tests and pair-wise Wilcoxon signed-rank tests with Bonferroni correction for multiple comparisons. Statistical test results are shown in Fig. 13. The Friedman test results indicate significant differences between methods in the first two tasks: identify the cluster ($\chi^2(5) = 52.01$, $p < 0.0001$), identify the outlier ($\chi^2(5) = 48.36$, $p < 0.0001$), and no significant difference in compare quantitative values ($\chi^2(5) = 5.58$, $p = 0.3494$). In the subsequent analysis, we focus on pairwise comparisons.

Task 1—Cluster Identification: *NCP* delivers superior performance to *SimiFC*, *AEF*, and *FD* while performing comparably to *Scatter-S* and *Scatter-L*. Participants consistently praised *NCP* for its ability to produce “clear cluster boundaries.” This is because *NCP* explicitly optimizes cluster convexity during force-directed refinement. Comparing performance between *NCP* and scatterplot methods (*Scatter-S* and *Scatter-L*) indicates that reducing the space between clusters does not hamper the identification of clusters but enhances layout compactness. This optimized use of space improves visual clarity and thus facilitates deeper and more effective data analysis.

Task 2—Outlier Identification: *NCP* significantly outperforms *SimiFC*, *AEF*, and *FD*, while performing comparably to *Scatter-S* and *Scatter-L*. Participants described outliers in *NCP* as “clearly noticeable.” Compared to other circle packing methods, *NCP* achieves better cluster convexity and clearer cluster boundaries, making outliers that deviate from clusters with the same class label more noticeable. Compared

to *Scatter-S* and *Scatter-L*, the lack of overlap and compactness of *NCP* is a double-edged sword. On the one hand, it ensures that every circle is clearly visible, preventing outliers from being too small to notice or occluded by others. On the other hand, the compactness reduces the space between circles, potentially drawing outliers that originally deviated from clusters back toward the cluster boundaries, making some of them harder to identify as outliers.

Task 3—Quantitative Attribute Comparison: Analysis results show no significant difference between methods. However, *NCP* achieved the highest mean accuracy (0.933) with the lowest standard deviation (0.138), highlighting its effectiveness in helping users to compare attribute values.

Workload Fig. 14 shows participant workload and fatigue levels measured using NASA's Task Load Index, including mental demand, physical demand, temporal demand, effort, performance, and frustration. For all three tasks, *NCP* performed better than other circle packing methods across all six measures, with lower mental and physical demands, reduced effort, less temporal pressure, higher performance, and lower frustration. P3 commented that *NCP* could “balance the cluster overview and local details well”, enhancing analysis efficiency and reducing workload. When compared to scatterplot methods, *Scatter-S* and *Scatter-L* resulted in greater physical demands and higher frustration, particularly in *Task 3—Quantitative Attribute Comparison*. This is because scatterplot methods with lower compactness have reduced space efficiency, leading to poorer readability. P5 commented: “At the beginning, I could hardly compare the highlighted circles since they were too small.” Even with the provided zooming in and out interactions, P7 noted: “I have to zoom in several times”, which caused a loss of broader context. In contrast, the lack of overlap and compactness of *NCP* significantly enhance

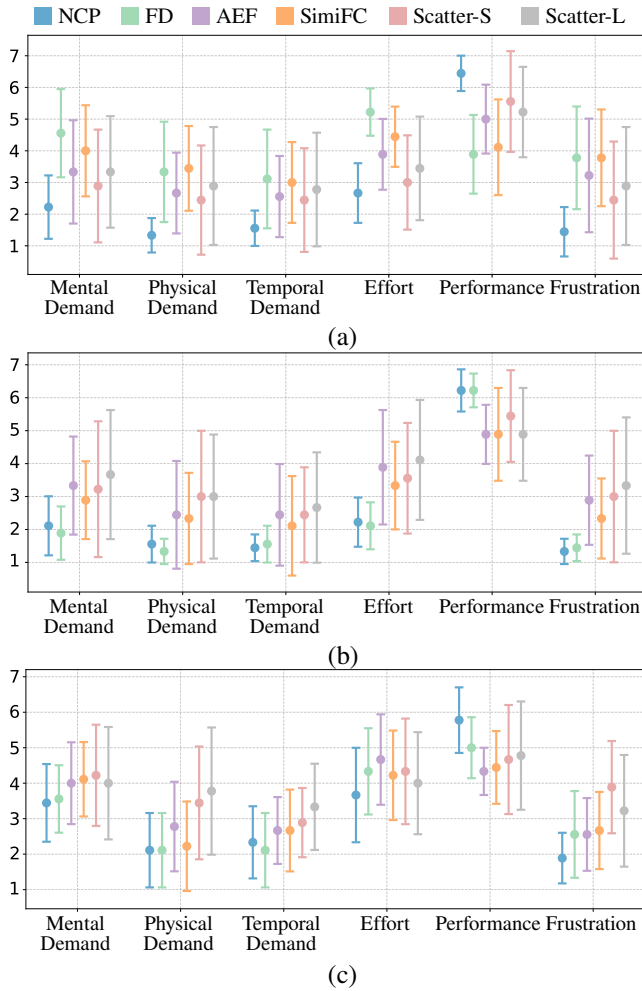


Fig. 14 Participants' workload and fatigue levels according to NASA's Task Load Index: (a) *Task 1—Cluster Identification*; (b) *Task 2—Outlier Identification*; (c) *Task 3—Quantitative Attribute Comparison*. Here, error bars show the 95% confidence intervals.

readability, enabling participants to locate their targets and perform comparisons with greater ease and efficiency.

6 Expert Feedback and Discussion

We interviewed three experts (E_1 – E_3) specializing in visual analytics and machine learning. Each interview took approximately an hour, including a 10-minute introduction to *NCP*, a 30-minute session where we presented our use cases and collected expert feedback, and a 20-minute discussion. Overall, the experts highlighted the utility of *NCP* in data analysis and its convenient integration with existing algorithms. We also identified several potential directions for future research based on the interviews.

6.1 Usability

Enhancing data analysis Overall, the experts acknowledged the capability of *NCP* to facilitate the analysis and




Fig. 15 Two alternative layouts for the Clothing dataset.

comparison of similar data items of interest. They commented that they could easily understand the cluster structures and identify outliers with *NCP*. For example, in Fig. 1(a), the cluster structures are clearly shown since circles of the same colors are mostly grouped together. Circles representing mislabeled images stand out due to the color differences from their neighbors. This prompts users to examine whether they are mislabeled. The experts also noted that the convexity of the cluster shapes was beneficial for understanding clusters, and consequently for identifying outliers efficiently.

The experts were also asked to compare the circle packing results generated by *NCP* and the baseline methods, *SimiFC* and *AEF*. They agreed that all these methods could generate compact circle packing results, but the visual patterns differed a lot in each case. Generally, they favored *NCP* and noted that its results preserved cluster structures well. In contrast, the experts identified distortions of clusters in the results of the baseline methods. Using the clothing dataset from the first use case as an example, *SimiFC* generated arc-shaped clusters, as shown in Fig. 15(a). E_3 pointed out that these visual patterns might lead to misunderstandings of the data relationships. For example, users may perceive a certain sequential order in the placement of circles from the inside to the outside. Both *SimiFC* and *AEF* might separate the clusters into different sub-clusters. For example, the cluster of light blue circles is disjoint in *SimiFC* (Figs. 15A₁ and A₂) and *AEF* (Figs. 15B₁ and B₂). In addition, the green, orange, pink, and red circles representing the samples with random noise are placed near the center of Fig. 1A. However, in Figs. 15A₁, B₁, and B₂, they are on the boundary of the cluster consisting of light blue circles. This hinders the identification of these outliers.

Pipeline flexibility While our *NCP* method is primarily designed for visualization, its flexible optimization pipeline enables potential applications to other domains by supporting diverse algorithms and objective functions. First, our pipeline supports the integration of different projection methods, power-diagram-based methods, and utilized forces. This

flexibility allows users to tailor their choices in each phase according to their specific needs. For example, in the case of textual data analysis, users can employ least-squares projection in the planar graph initialization, which has been shown to be effective in preserving neighborhood relationships in textual data [67]. Second, the continuation method can be extended to accommodate more objectives, such as ensuring efficient connectivity while minimizing interference in wireless sensor network design, and enhancing the aesthetic arrangement of circles in design and art applications. By carefully determining the sequence of objectives and constraints to be incorporated into the optimization process, it achieves a well-balanced integration of multiple criteria. This adaptability improves the effectiveness of our method in a variety of optimization contexts, making it applicable to other domains.

6.2 Limitations and Future Work

Dynamic parameter tuning Our *NCP* method achieves a balance between three optimization objectives using the parameters α and β . In our implementation, we determined their values using a grid search, which is time-consuming. We would like to explore automatic parameter tuning methods to ease this process, such as multi-task learning [68]. In addition, E_2 pointed out that user preferences for optimization objectives could vary in different real-world applications. For example, E_2 said, “When the model prediction is reliable, I would prioritize optimizing convexity to enhance the perceptual clarity of clusters. Otherwise, I would focus more on neighborhood preservation, which helps me identify prediction errors more easily.” Therefore, integrating user feedback to dynamically adjust these parameters is essential, to allow users to tailor the circle packing to their specific requirements.

Integration with other visualizations techniques Our experts have identified several opportunities to enhance *NCP* through integration with other visualization techniques. First, they proposed that the layout could be enhanced to provide more guidance for data exploration. A possible method would be to design informative glyphs displayed within circles, to offer more details of data items. For example, in the second use case, pie charts can be used to show how different factors contribute to housing prices instead of encoding each quantitative attribute individually. Another method would be to select representative data items and show them in the empty space, to help users better understand the data. Second, E_2 and E_3 suggested employing *NCP* for hierarchical exploration

of large data. By building a hierarchy for the data items, a subset of data items can be sampled and visualized as in the first use case. Alternatively, circles can represent groups of similar data items instead of individual ones. Users can navigate the whole dataset with the zooming function.

Interactive circle packing User interactions could also be introduced to improve the packing result of *NCP*. Specifically, E_3 expressed the need to author functions to manipulate the relative positions of circles. “If I notice that some circles should form a cluster but are not placed adjacently, I want to adjust their positions to make this cluster more evident.” We should consider allowing users to move a few representative circles using drag-and-drop. During this process, the proximity between circles changes. As a result, further investigation is required to ensure the stability of circle packing and maintain users’ mental maps. We also should consider supporting users directly specifying must-link constraints between data items and incorporating these constraints into the optimization process, to ensure that the generated circle packing places circles of interest adjacent to one another.

7 Conclusions

In this paper, we have developed a new layout method, *NCP*, for generating a neighborhood-preserving non-uniform circle packing. We formulate circle packing as a planar graph embedding problem and solve it using the continuation method. By progressively incorporating multiple optimization objectives and constraints, this method steers the optimization towards a more favorable solution. Our quantitative comparison to baselines shows that *NCP* performs better in terms of neighborhood preservation and convexity while achieving comparable compactness. Two use cases further demonstrate its application in data analysis.

1 Appendix A: Projection Method and Associated Parameters

We conducted quantitative experiments to select the best projection method and associated parameters for the neighborhood-preserving planar graph initialization.

Experimental Setting

Datasets We employed eight real-world datasets used in Xia *et al.* [41], including Boston [69], Dermatology [70], Ecoli [70], ExtYaleB [71], MNIST64 [70], Olive [72], Weather [73], and World12D [69].

Measure We used the neighborhood preservation degree [48, 49] to measure how well the neighborhood re-

Table 3 Neighborhood preservation comparison for different projection methods.

Dataset	t-SNE	UMAP	PCA	MDS	NMF
Boston	0.374	0.346	0.184	0.230	0.167
Dermatology	0.366	0.328	0.115	0.151	0.087
Ecoli	0.380	0.338	0.159	0.186	0.146
ExtYaleB	0.474	0.373	0.109	0.200	0.101
MNIST64	0.346	0.288	0.074	0.082	0.039
Olive	0.375	0.335	0.165	0.177	0.143
Weather	0.473	0.400	0.208	0.263	0.199
World12D	0.350	0.400	0.280	0.313	0.240
Average	0.392	0.351	0.162	0.200	0.140

relationships in the high-dimensional data space are preserved in the projection results.

Methods We identified five projection methods: t-SNE [42], UMAP [43], PCA [44], MDS [45], and NMF [46], based on previous studies [41, 47].

Parameters We explored *perplexity* of t-SNE and *neighbors* of UMAP, which are key factors that affect the neighborhood preservation of these methods. The candidate parameters were set the same as in Xia *et al.* [41]. Specifically, the *perplexity* of t-SNE was chosen in {5, 15, 30, 40, 50}, and the *neighbors* of UMAP was chosen in {4, 7, 10, 13, 16}. PCA, MDS, and NMF applied default parameter settings.

Results

For t-SNE and UMAP, we first identified the best parameter settings based on the average neighborhood preservation across the eight datasets. The perplexity of t-SNE was set to 15, and the Neighbors of UMAP was set as 4. We then compared the neighborhood preservation for different datasets. As shown in Table 3, t-SNE and UMAP significantly outperformed the other methods, with t-SNE outperforming UMAP on seven datasets and on average (0.392 v.s. 0.351). Fig. 16 provides a visual comparison between these methods. In datasets like ExtYaleB and MNIST64, t-SNE and UMAP clearly separated different samples with different labels, but the other three methods tended to confuse them. This explains the clear gap between t-SNE/UMAP and the other methods. When comparing the results of t-SNE and UMAP, the points in t-SNE are more evenly distributed compared to UMAP, which better reflects the neighborhood relationships. Therefore, we selected t-SNE with a *perplexity* of 15 in our layout method.

Appendix B: Comparison of Different Initialization Methods

We also conducted a study to compare circle packing results generated using different initialization methods. We used the

same datasets and metrics as those in Sec. 5.1.

Results The experimental results are presented in Table 4. Overall, t-SNE achieved the best average performance in terms of neighborhood preservation. This finding aligns with results in Appendix A, supporting the importance of initialization methods in neighborhood preservation. All methods performed comparably in terms of compactness and convexity. This is because these objectives are incorporated into optimization in subsequent steps, which are identical for all methods reported in this section.

We also visually compared the circle packing results of the reported methods in Figs. 17 and 18 to provide a more intuitive explanation of their differences. Overall, all methods were able to form reasonable clusters, as indicated by their comparable convexity. However, there were differences in neighborhood preservation within clusters. Taking the MNIST64 dataset in Fig. 17 as an example, two highly similar digits ‘1’ were correctly placed together in the results of UMAP and t-SNE, whereas in the results of other methods, they were separated. This separation complicates the accurate interpretation of proximity within clusters.

Appendix C: Grid Search and Parameter Analysis

Our optimization problem is formulated in Eq. 1. Since the parameters α and β , which balance the impact of the three terms, affect the optimization result, we performed a grid search to determine the optimal parameters.

Our measures involve neighborhood preservation (NP_1 and NP_2), compactness, and convexity, consistent with those in our quantitative evaluation. We conducted a grid search to investigate the relationships between different measures and choices of weights α and β . Here, we considered $\alpha \in [0.10, 0.20, 0.50, 1.00, 2.00, 5.00, 10.00]$ and $\beta \in [0.10, 0.20, 0.50, 1.00, 2.00, 5.00, 10.00]$. The results are shown in Fig. 19. For neighborhood preservation, NP_1 and NP_2 decrease as α and β increase: larger α and β give less weight to neighborhood preservation F_n . For compactness, the results are quite stable (between 0.867 and 0.874) when α and β change due to use of sufficient iterations in the force-directed layout to guarantee convergence. Convexity decreases as α increases or β decreases, because larger α or smaller β give less weight to F_v . The observed trends in how the measures vary with changes in α and β also demonstrate the effectiveness of adjusting optimization preferences by setting different values for α and β .

To comprehensively compare the performance of our methods across different parameter settings, we designed a balanced index, which combines NP_1 , NP_2 , compactness, and

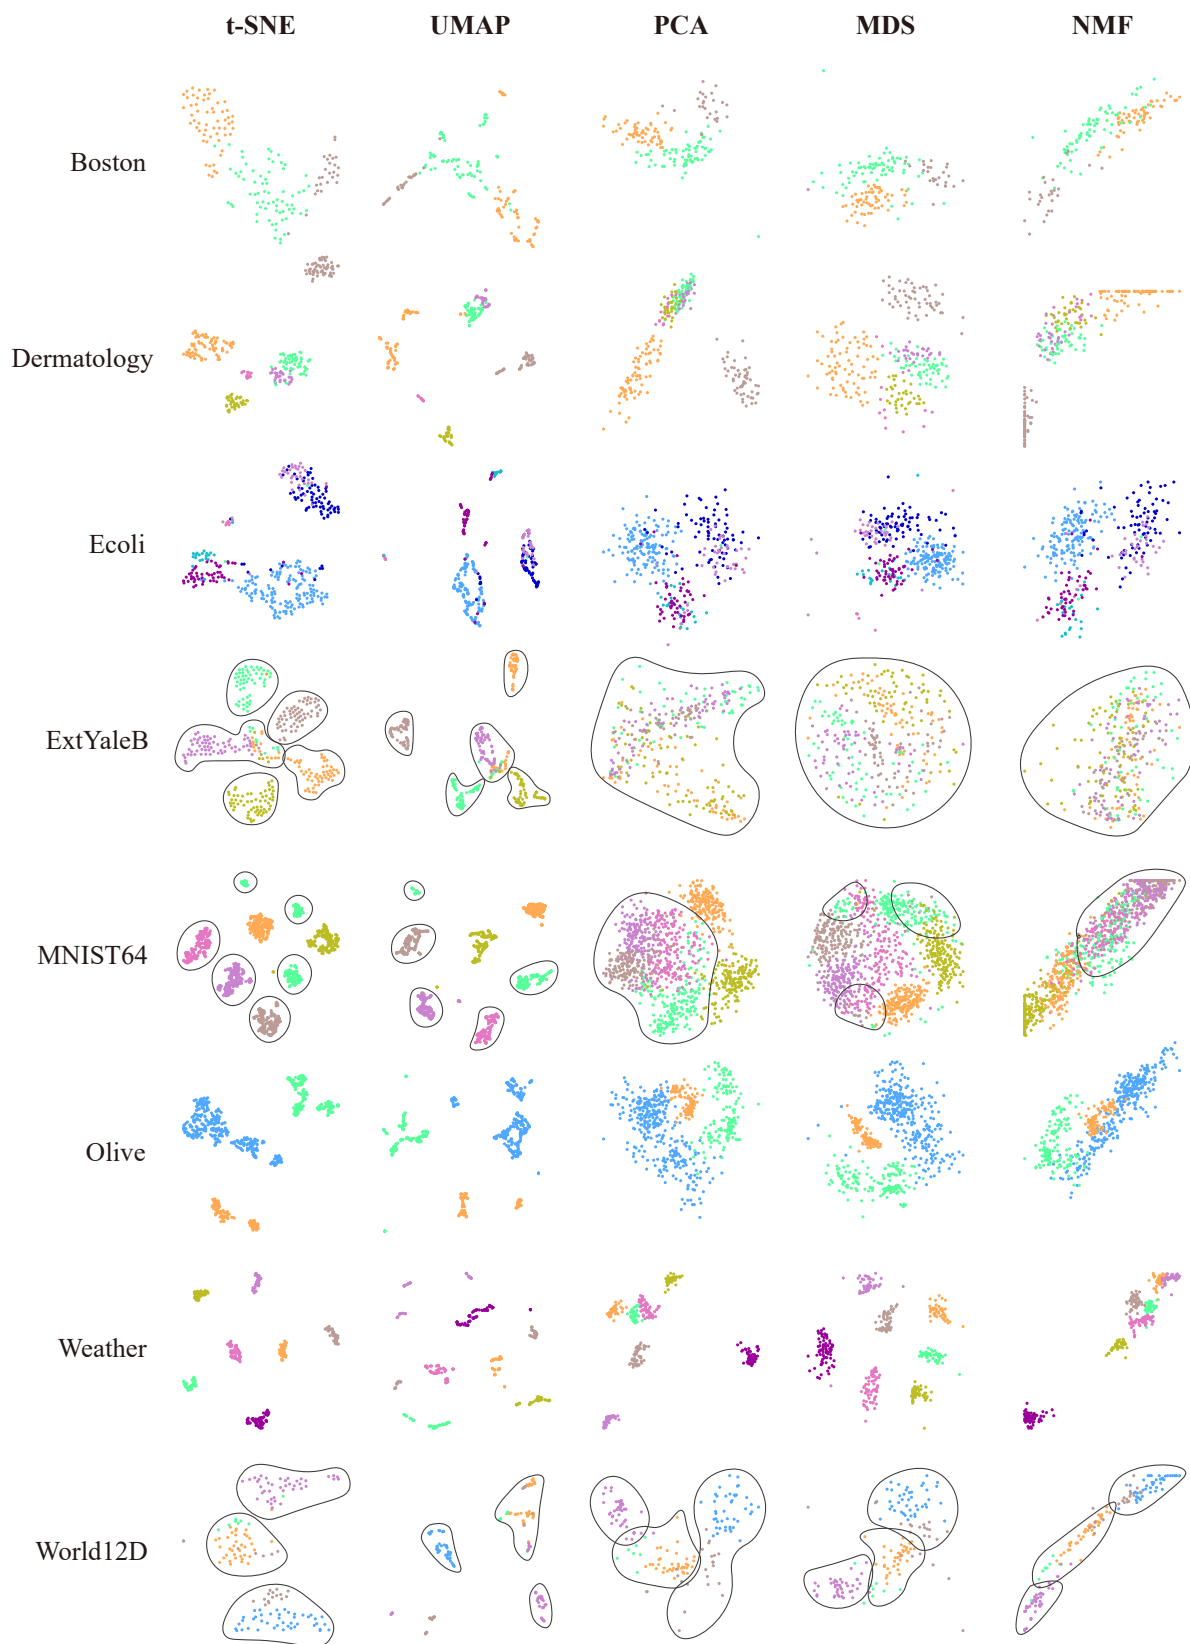


Fig. 16 Projection results for the eight datasets.

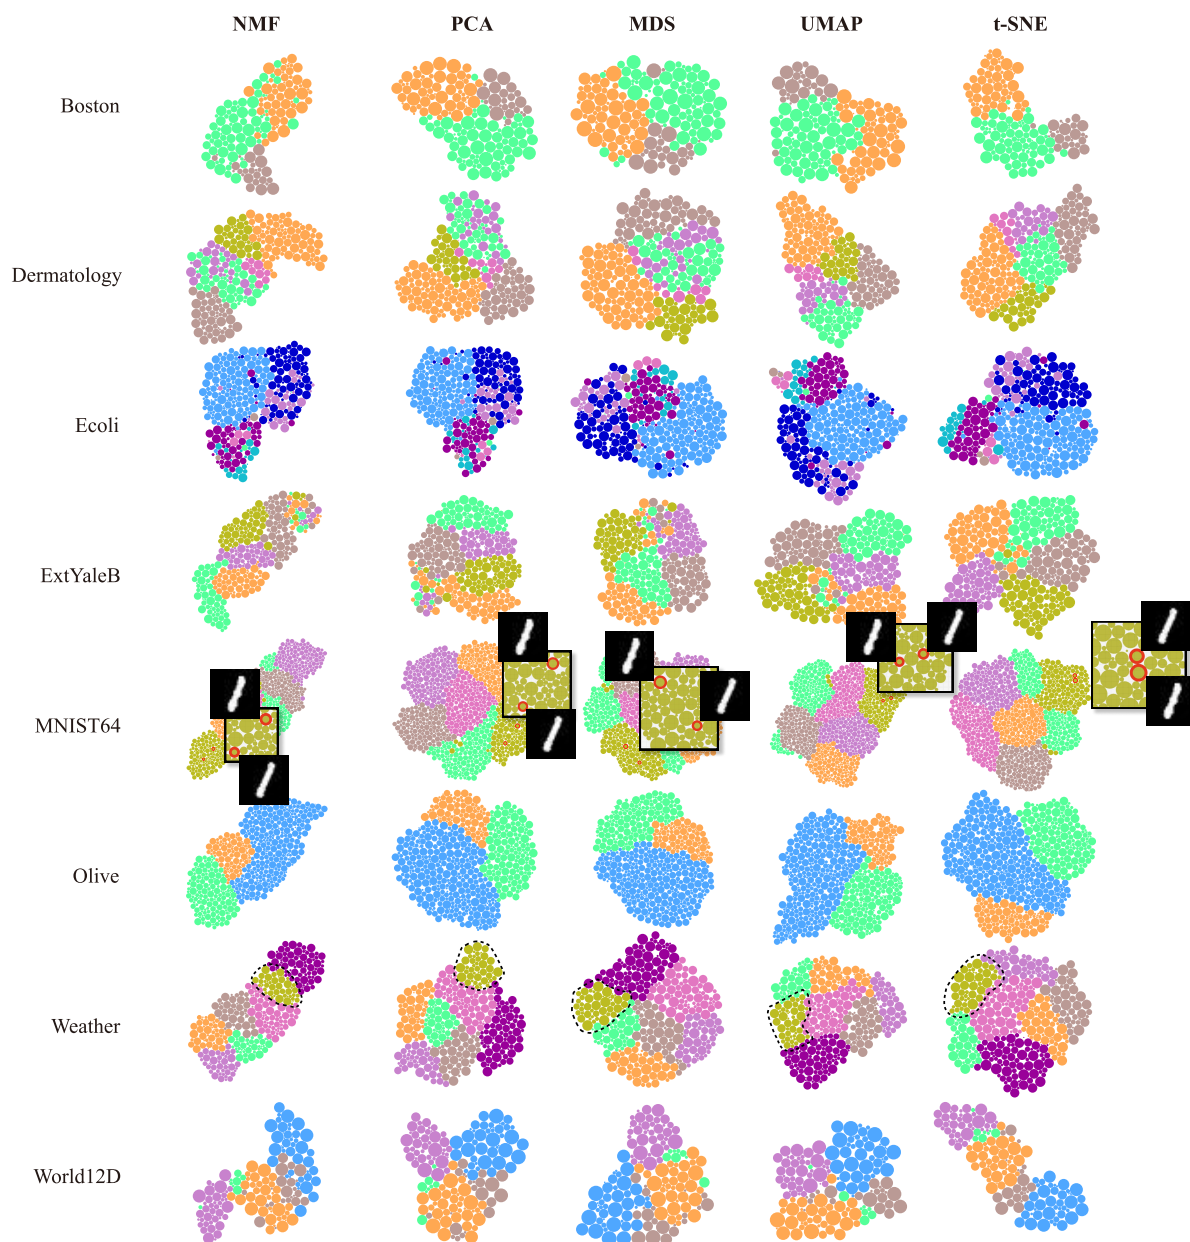


Fig. 17 Circle packing results for the eight datasets generated by alternative initialization methods where the circle radii have a small variance.

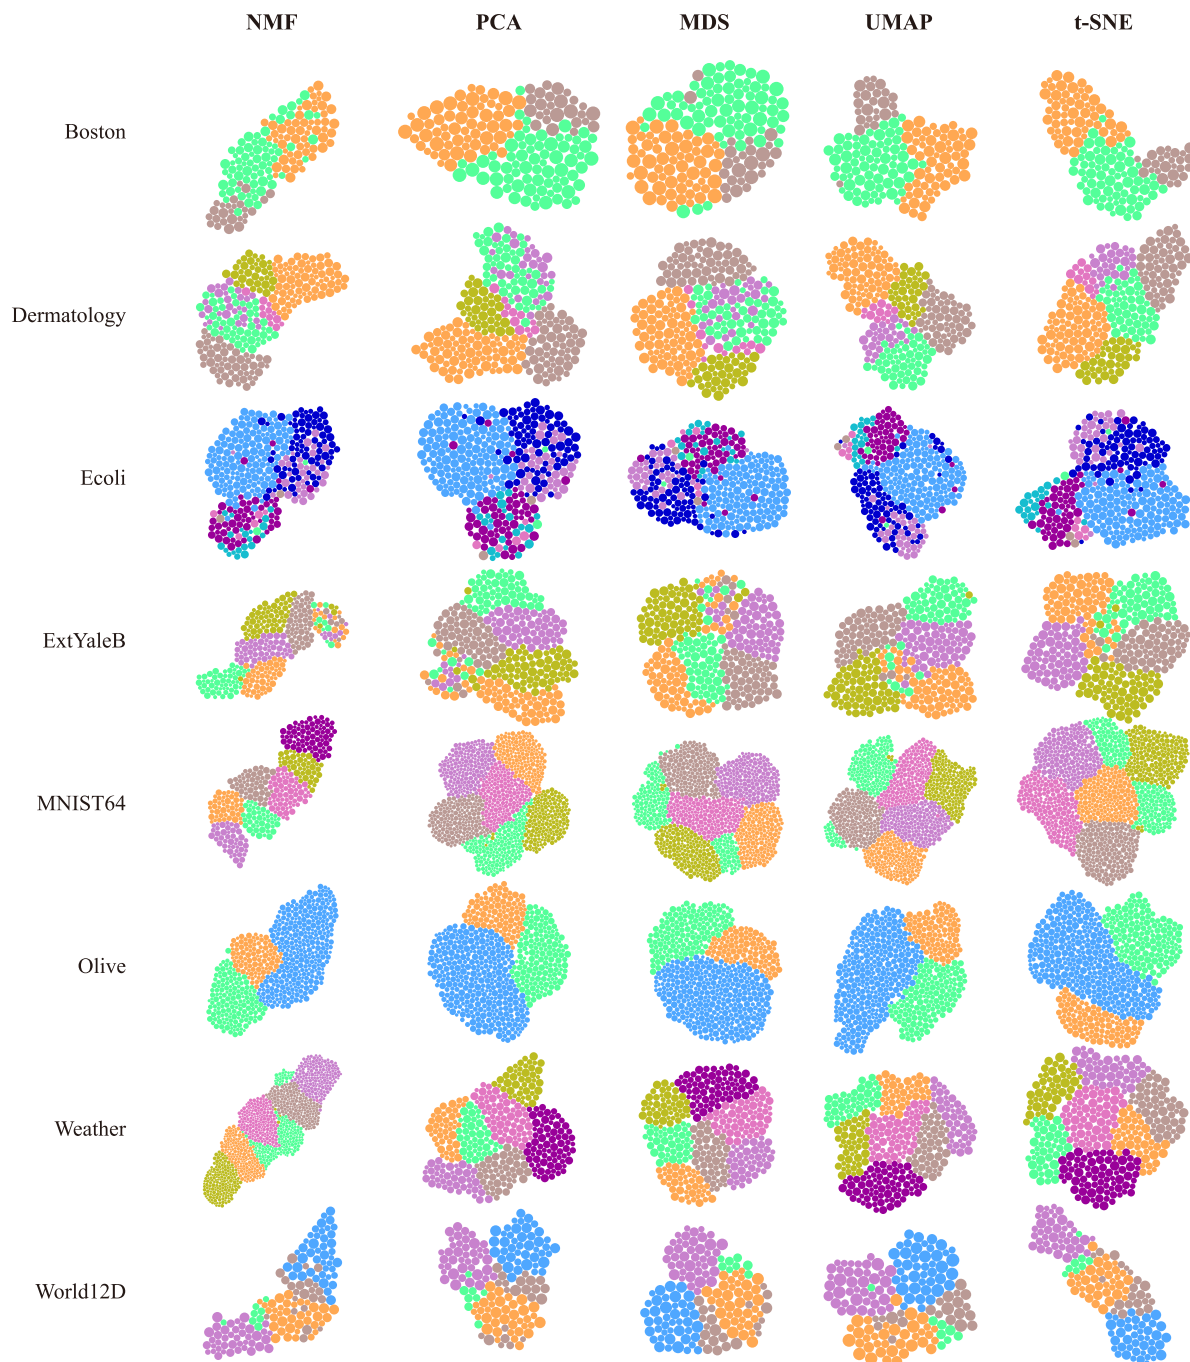


Fig. 18 Circle packing results for the eight datasets generated by alternative initialization methods where the circle radii have a large variance.

Table 4 Performance comparison in terms of neighborhood preservation (NP_1 , NP_2), compactness, and convexity.

Dataset	NP_1					NP_2					Compactness					Convexity				
	Baseline				Ours	Baseline				Ours	Baseline				Ours	Baseline				Ours
	NMF	PCA	MDS	UMAP	NCP	NMF	PCA	MDS	UMAP	NCP	NMF	PCA	MDS	UMAP	NCP	NMF	PCA	MDS	UMAP	NCP
Boston	0.150	0.170	0.219	0.301	0.341	0.305	0.336	0.389	0.422	0.423	0.890	0.881	0.879	0.876	0.889	0.810	0.851	0.797	0.807	0.809
Dermatology	0.119	0.125	0.160	0.280	0.311	0.256	0.259	0.281	0.399	0.385	0.886	0.881	0.868	0.872	0.874	0.795	0.771	0.802	0.803	0.790
Ecoli	0.164	0.163	0.191	0.316	0.344	0.276	0.292	0.308	0.393	0.411	0.870	0.875	0.862	0.871	0.874	0.784	0.801	0.794	0.812	0.831
ExtYaleB	0.201	0.225	0.241	0.387	0.422	0.296	0.322	0.357	0.469	0.487	0.881	0.868	0.865	0.870	0.868	0.781	0.778	0.763	0.777	0.806
MNIST64	0.073	0.090	0.092	0.213	0.306	0.143	0.172	0.186	0.325	0.372	0.871	0.860	0.861	0.869	0.863	0.804	0.800	0.834	0.787	0.827
Olive	0.147	0.157	0.159	0.269	0.328	0.277	0.280	0.285	0.371	0.414	0.868	0.861	0.866	0.865	0.862	0.825	0.824	0.829	0.789	0.817
Weather	0.179	0.188	0.223	0.351	0.448	0.332	0.350	0.359	0.527	0.567	0.868	0.864	0.868	0.867	0.868	0.811	0.797	0.811	0.811	0.795
World12D	0.235	0.312	0.289	0.319	0.379	0.375	0.442	0.447	0.483	0.463	0.901	0.883	0.881	0.876	0.896	0.734	0.819	0.807	0.828	0.809
Average	0.158	0.179	0.197	0.304	0.360	0.282	0.307	0.326	0.423	0.440	0.879	0.872	0.869	0.871	0.874	0.793	0.805	0.805	0.802	0.811

(a) Performance in the case where circle radii have a large variance.

Dataset	NP_1					NP_2					Compactness					Convexity				
	Baseline				Ours	Baseline				Ours	Baseline				Ours	Baseline				Ours
	NMF	PCA	MDS	UMAP	NCP	NMF	PCA	MDS	UMAP	NCP	NMF	PCA	MDS	UMAP	NCP	NMF	PCA	MDS	UMAP	NCP
Boston	0.175	0.181	0.206	0.332	0.365	0.309	0.345	0.366	0.442	0.437	0.895	0.876	0.870	0.879	0.888	0.819	0.824	0.844	0.826	0.833
Dermatology	0.115	0.126	0.153	0.320	0.364	0.240	0.246	0.272	0.398	0.399	0.875	0.877	0.862	0.875	0.875	0.834	0.797	0.860	0.828	0.845
Ecoli	0.172	0.170	0.171	0.311	0.350	0.273	0.278	0.297	0.404	0.416	0.869	0.866	0.862	0.862	0.866	0.806	0.841	0.830	0.816	0.821
ExtYaleB	0.194	0.220	0.216	0.385	0.476	0.302	0.308	0.354	0.475	0.505	0.870	0.862	0.855	0.869	0.862	0.775	0.760	0.791	0.765	0.808
MNIST64	0.070	0.091	0.100	0.225	0.327	0.138	0.165	0.185	0.326	0.375	0.853	0.852	0.850	0.854	0.854	0.829	0.834	0.835	0.800	0.808
Olive	0.161	0.162	0.158	0.273	0.342	0.274	0.281	0.273	0.371	0.422	0.861	0.857	0.855	0.859	0.858	0.819	0.848	0.851	0.789	0.839
Weather	0.192	0.212	0.241	0.379	0.483	0.325	0.353	0.352	0.539	0.578	0.877	0.867	0.862	0.859	0.860	0.850	0.841	0.847	0.805	0.816
World12D	0.264	0.268	0.323	0.363	0.380	0.373	0.431	0.450	0.473	0.477	0.881	0.885	0.879	0.876	0.885	0.768	0.800	0.808	0.845	0.796
Average	0.168	0.179	0.196	0.323	0.386	0.280	0.301	0.319	0.429	0.451	0.875	0.867	0.862	0.866	0.868	0.813	0.818	0.833	0.809	0.825

(b) Performance in the case where circle radii have a small variance.

convexity. The balanced index is defined as the sum of all four measures. We calculated the balanced index across all datasets, varying the parameters α and β and taking the average. As illustrated in Fig. 20, the analysis revealed a local maximum region with relatively high and stable balanced indexes ($\alpha \in [0.10, 0.20, 0.50]$ and $\beta \in [0.50, 1.00, 2.00]$), with the peak occurring at $\alpha = 0.20$ and $\beta = 1.00$. With this optimal parameter setting, the method delivers near-optimal performance across all evaluated dimensions: 99.5% of the best NP_1 , 99.8% of the best NP_2 , 99.6% of the best compactness, and 98.7% of the best convexity. Therefore, we chose it as the default parameter setting.

Appendix D: The Layout Results

Figs. 21 and 22 show the original t-SNE projection results (the first column) and circle packing generated by SimiFC, AEF, FD, and NCP (the remaining columns) for two settings: large and small radius variance, respectively. Observing the shapes of clusters, it is easy to identify that SimiFC tends to generate arc-shaped clusters, and AEF tends to separate a cluster into multiple components. Both failed to provide a clear cluster structure, hindering analysis. In contrast, both FD and NCP preserve the cluster structure well. However, compared to the t-SNE projection result, FD often fails to preserve the relative positions of clusters. Take the Boston dataset in Fig. 21 as

an example. In the t-SNE projection result (Fig. 21A), the green cluster was placed in between the orange cluster and the brown cluster, but FD dragged these three clusters together (Fig. 21B). This would mislead users in understanding the similarity relationships between the three clusters. Instead, NCP preserved the relative positions of clusters in the t-SNE projection well (Fig. 21C). In addition, NCP usually forms smoother and more convex boundaries between the clusters (e.g., Fig. 21D versus Fig. 21E, and Fig. 21F versus Fig. 21G). This further enhances the perception of cluster structures and facilitates data analysis.

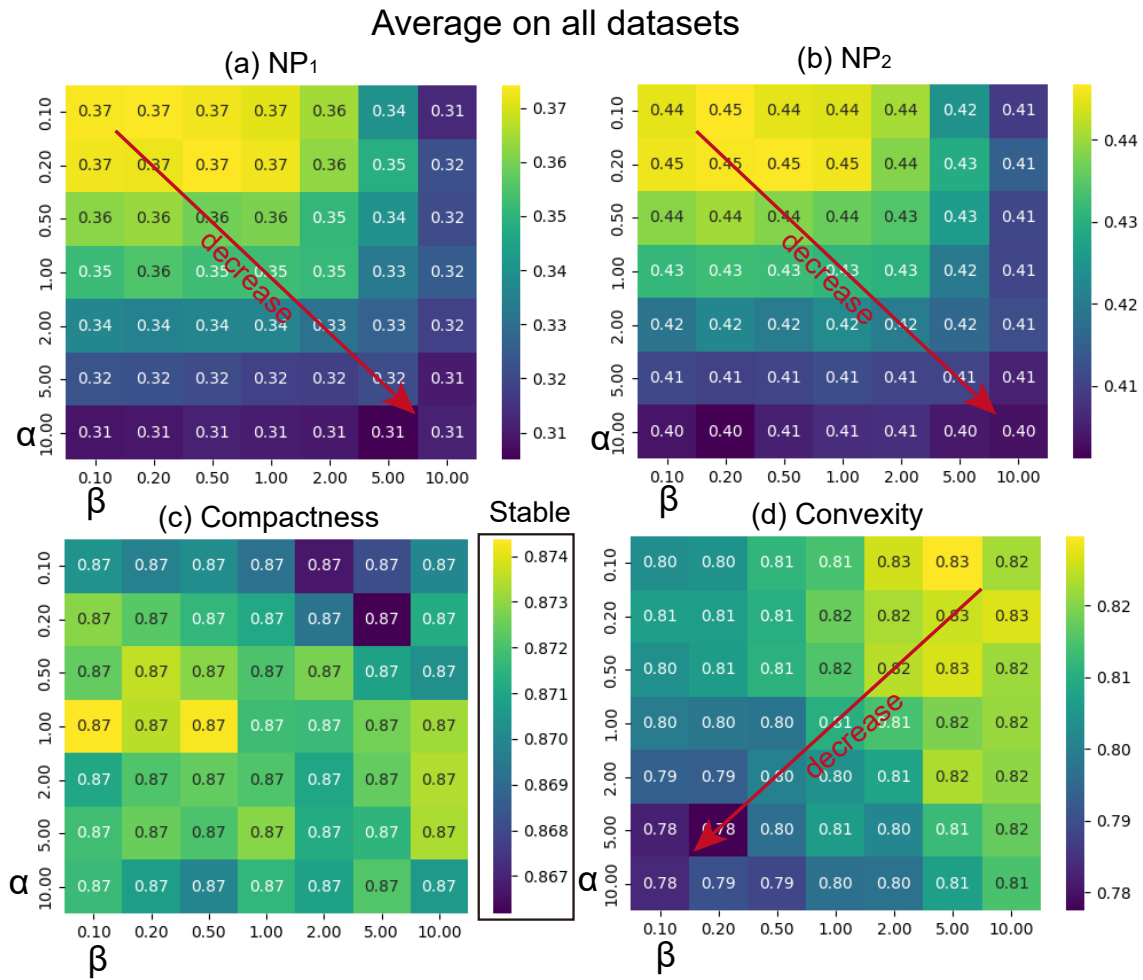
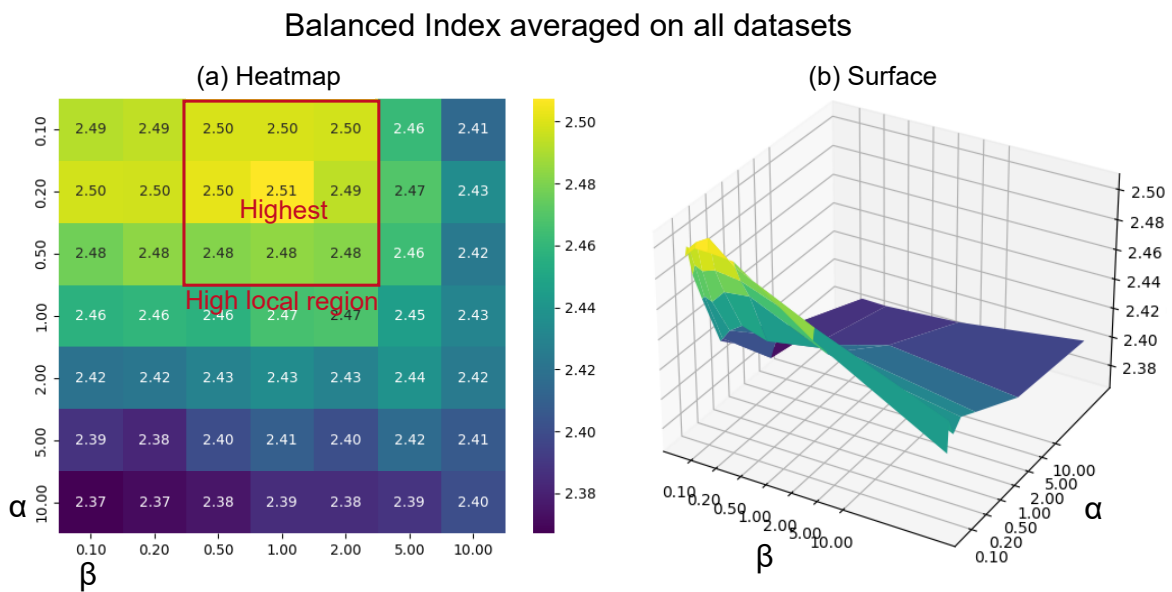
Appendix E: Papers Used to Summarize the Design Criteria for NCP

Circle packing methods: [2–4, 7, 20, 21, 23, 24, 26–29, 74, 75].

Cluster-aware layout methods: [8–10, 76, 77].

Availability of data and materials

All data and materials are available on Github at <https://github.com/NCP-2024/NCP>. In particular, they include datasets, source code, a video, experimental results concerning the selection of projection methods and parameters, grid search analysis for algorithm parameters, detailed layout results, and the paper list used to summarize the design criteria for NCP.

Fig. 19 Measure distributions with varying α and β Fig. 20 Distribution of balanced index with varying α and β

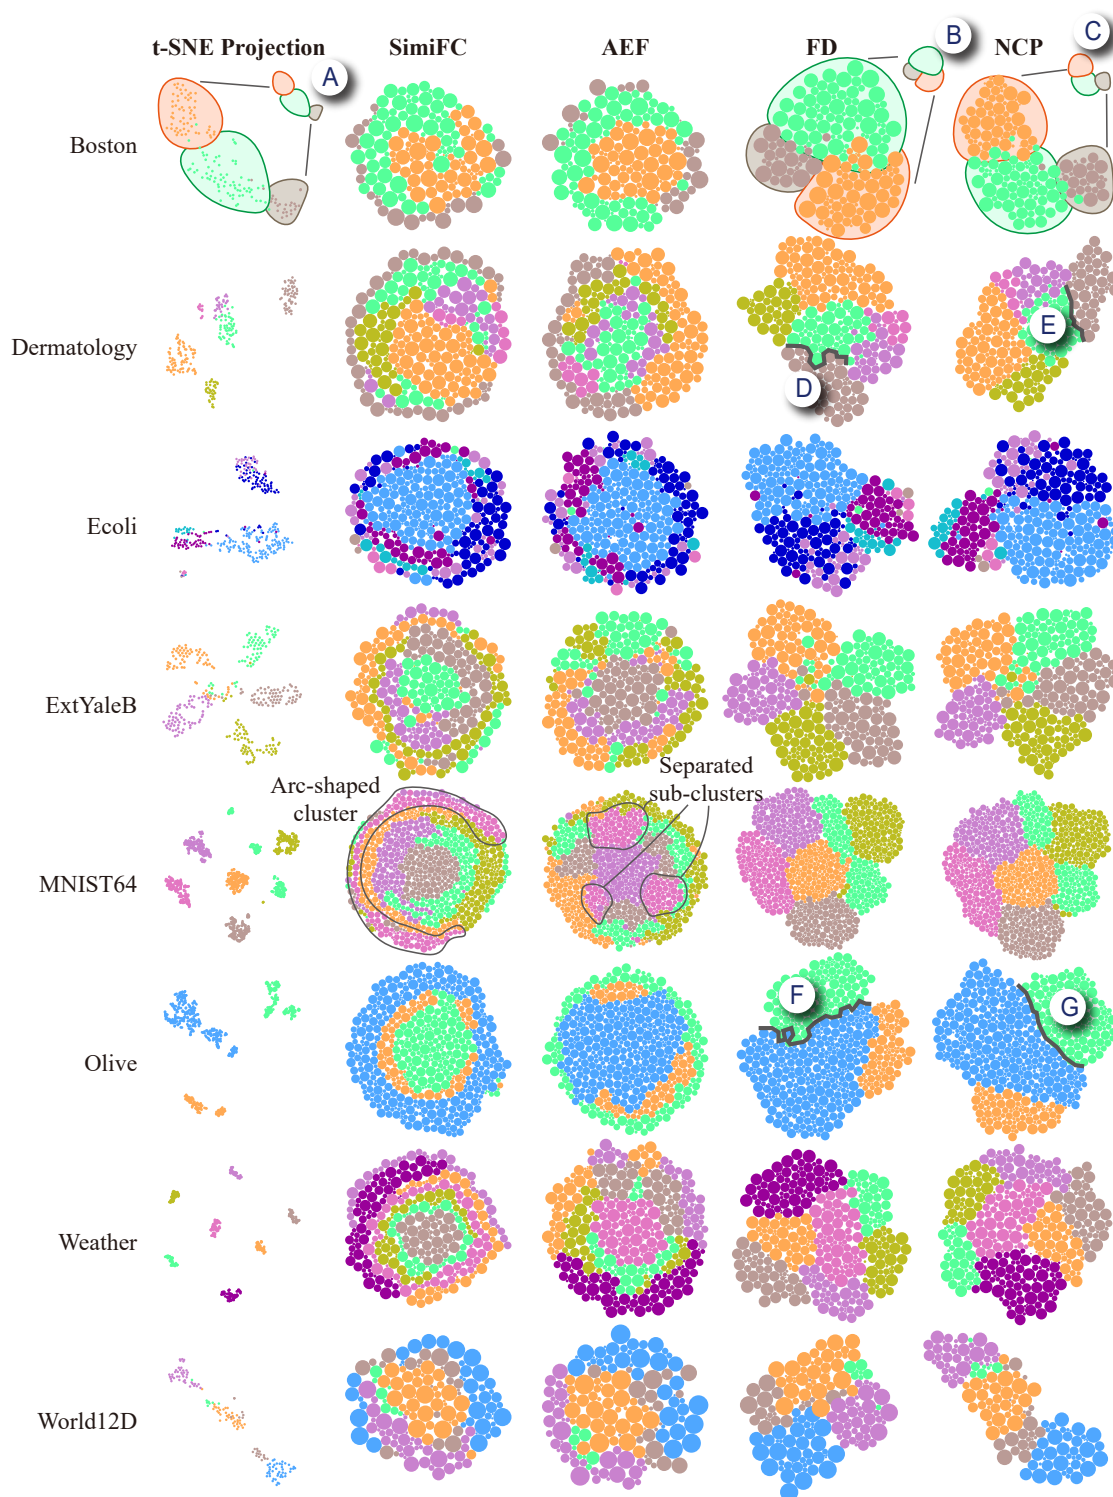


Fig. 21 Circle packing results for the eight datasets generated by SimiFC, AEF, FD, and NCP for circle radii with large variance.

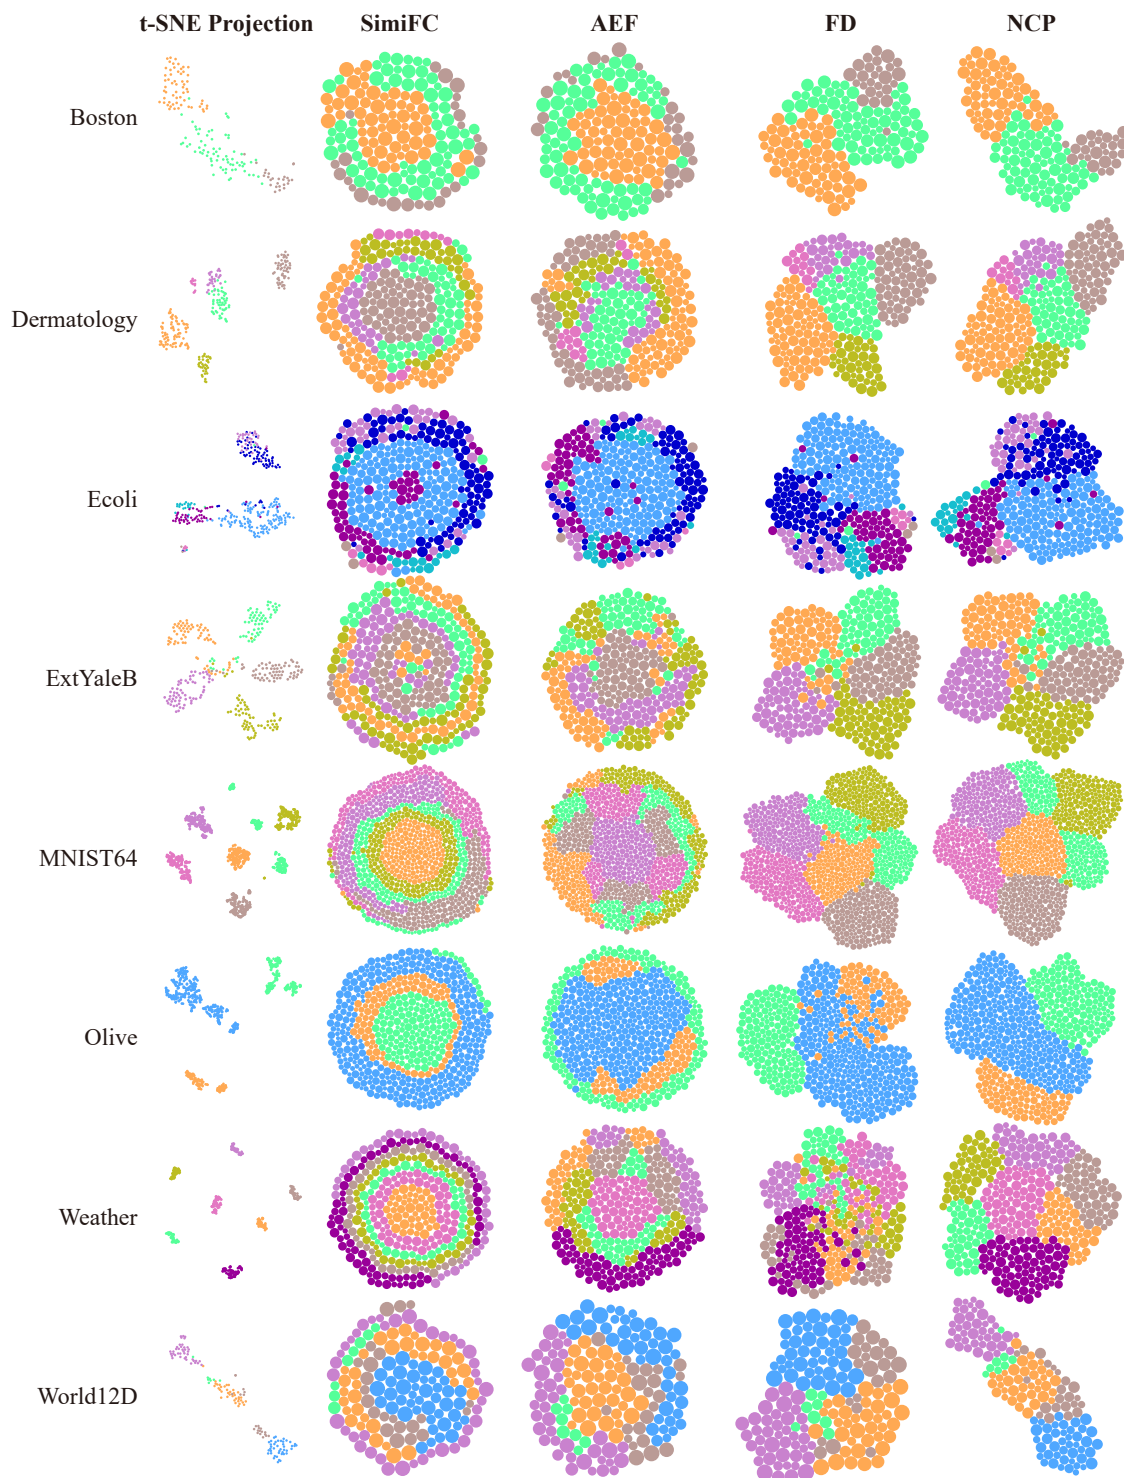


Fig. 22 Circle packing results of the eight datasets generated by SimiFC, AEF, FD, and NCP for circle radii with small variance.

Declaration of competing interest

The authors have no competing interests to declare relevant to the content of this article.

Funding

This work was supported by the National Natural Science Foundation of China (U21A20469, 61936002), and in part by Tsinghua University-China Telecom Wanwei Joint Research Center.

Authors' contributions

Duan Li: implementation, algorithm design, writing—original draft, writing—review editing.

Jun Yuan: implementation, algorithm design, writing—original draft, writing—review editing.

Xinyuan Guo: implementation, algorithm design.

Xiting Wang: writing—review editing

Yang Liu: algorithm design, writing—review editing

Weikai Yang: writing—original draft, writing—review editing.

Shixia Liu: algorithm design, supervision, writing—original draft, writing—review editing, funding acquisition.

Acknowledgements

The authors thank Dr. Changjian Chen, Dr. Qianwen Wang, Yukai Guo, Haoze Wang, and Jiangning Zhu for their valuable comments.

References

- [1] Abdullah NAS, Zulkepely N, Idrus Z. Malaysian Budget Visualization Using Circle Packing. In *Proceedings of the Soft Computing in Data Science*, 2019, 81–90, doi:10.1007/978-981-15-0399-3_7.
- [2] Görtler J, Schulz C, Weiskopf D, Deussen O. Bubble Treemaps for Uncertainty Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 2018, 24(1): 719–728, doi:10.1109/TVCG.2017.2743959.
- [3] Wang W, Wang H, Dai G, Wang H. Visualization of Large Hierarchical Data by Circle Packing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2006, 517–520, doi:10.1145/1124772.1124851.
- [4] Wang Q, Xu Z, Chen Z, Wang Y, Liu S, Qu H. Visual analysis of discrimination in machine learning. *IEEE Transactions on Visualization and Computer Graphics*, 2021, 27(2): 1470–1480, doi:10.1109/TVCG.2020.3030471.
- [5] Li D, Guo X, Shu X, Xiao L, Yu L, Liu S. RouteFlow: Trajectory-Aware Animated Transitions. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, 2025, doi:10.1145/3706598.3714300.
- [6] Liu S, Yang W, Wang J, Yuan J. *Visualization for Artificial Intelligence*. Springer Cham, 2025.
- [7] Yuan J, Liu M, Tian F, Liu S. Visual Analysis of Neural Architecture Spaces for Summarizing Design Principles. *IEEE Transactions on Visualization and Computer Graphics*, 2023, 29(1): 288–298, doi:10.1109/TVCG.2022.3209404.
- [8] Chen C, Wang Z, Wu J, Wang X, Guo LZ, Li YF, Liu S. Interactive Graph Construction for Graph-Based Semi-Supervised Learning. *IEEE Transactions on Visualization and Computer Graphics*, 2021, 27(9): 3701–3716, doi:10.1109/TVCG.2021.3084694.
- [9] Rottmann P, Wallinger M, Bonerath A, Gedicke S, Nöllenburg M, Haunert JH. MosaicSets: Embedding Set Systems into Grid Graphs. *IEEE Transactions on Visualization and Computer Graphics*, 2023, 29(1): 875–885, doi:10.1109/TVCG.2022.3209485.
- [10] Zhou Y, Yang W, Chen J, Chen C, Shen Z, Luo X, Yu L, Liu S. Cluster-Aware Grid Layout. *IEEE Transactions on Visualization and Computer Graphics*, 2024, 30(1): 240–250, doi:10.1109/TVCG.2023.3326934.
- [11] Zhou Y, Chen C, Shen Z, Zhu J, Chen J, Yang W, Liu S. Hierarchical Fuzzy-Cluster-Aware Grid Layout for Large-Scale Data. *IEEE Transactions on Visualization and Computer Graphics*, 2025: 1–14, doi:10.1109/TVCG.2025.3566558.
- [12] Beardon AF, Stephenson K. The uniformization theorem for circle packings. *Indiana University Mathematics Journal*, 1990, 39(4): 1383–1425, doi:10.1512/iumj.1990.39.39062.
- [13] Allgower EL, Georg K. *Numerical continuation methods: an introduction*. Springer, 2012.
- [14] Sun H, Li Z. Effectiveness of cartogram for the representation of spatial data. *The Cartographic Journal*, 2010, 47(1): 12–21, doi:10.1179/000870409X12525737905169.
- [15] Hifi M, M'Hallah R. A Literature Review on Circle and Sphere Packing Problems: Models and Methodologies. *Advances in Operations Research*, 2009, 2009(1): 150624, doi:10.1155/2009/150624.
- [16] George JA, George JM, Lamar BW. Packing different-sized circles into a rectangular container. *European Journal of Operational Research*, 1995, 84(3): 693–712, doi:10.1016/0377-2217(95)00032-L.
- [17] Zhang D, Huang W. A simulated annealing algorithm for the circles packing problem. In *International Conference on Computational Science*, 2004, 206–214, doi:10.1007/978-3-540-24685-5_26.
- [18] Akeb H, Hifi M, M'Hallah R. A beam search algorithm for the circular packing problem. *Computers & Operations Research*, 2009, 36(5): 1513–1528, doi:10.1016/j.cor.2008.02.003.
- [19] Scheibel W, Limberger D, Döllner J. Survey of treemap layout algorithms. In *Proceedings of the International Symposium on Visual Information Communication and Interaction*, 2020, 1–9, doi:10.1145/3430036.3430041.
- [20] Liu S, Yin J, Wang X, Cui W, Cao K, Pei J. Online Visual Analytics of Text Streams. *IEEE Transactions on Visualization*

- and *Computer Graphics*, 2016, 22(11): 2451–2466, doi:10.1109/TVCG.2015.2509990.
- [21] Zhao J, Cao N, Wen Z, Song Y, Lin YR, Collins C. #FluxFlow: Visual Analysis of Anomalous Information Spreading on Social Media. *IEEE Transactions on Visualization and Computer Graphics*, 2014, 20(12): 1773–1782, doi:10.1109/TVCG.2014.2346922.
- [22] Aurenhammer F. Power Diagrams: Properties, Algorithms and Applications. *SIAM Journal on Computing*, 1987, 16(1): 78–96, doi:10.1137/0216006.
- [23] Zhao H, Lu L. Variational circular treemaps for interactive visualization of hierarchical data. In *Proceedings of the IEEE Pacific Visualization Symposium*, 2015, 81–85, doi:10.1109/PACIFICVIS.2015.7156360.
- [24] Yu Z, Lu L, Guo Y, Fan R, Liu M, Wang W. Content-Aware Photo Collage Using Circle Packing. *IEEE Transactions on Visualization and Computer Graphics*, 2014, 20(2): 182–195, doi:10.1109/TVCG.2013.106.
- [25] Liang Y, Wang X, Zhang SH, Hu SM, Liu S. PhotoRecomposer: Interactive Photo Recomposition by Cropping. *IEEE Transactions on Visualization and Computer Graphics*, 2018, 24(10): 2728–2742, doi:10.1109/TVCG.2017.2764895.
- [26] Rodrigues N, Schulz C, Döring S, Baumgartner D, Krake T, Weiskopf D. Relaxed Dot Plots: Faithful Visualization of Samples and Their Distribution. *IEEE Transactions on Visualization and Computer Graphics*, 2023, 29(1): 278–287, doi:10.1109/TVCG.2022.3209429.
- [27] Huron S, Vuillemot R, Fekete JD. Visual Sedimentation. *IEEE Transactions on Visualization and Computer Graphics*, 2013, 19(12): 2446–2455, doi:10.1109/TVCG.2013.227.
- [28] Liu X, Xu A, Gou L, Liu H, Akkiraju R, Shen HW. SocialBrands: Visual analysis of public perceptions of brands on social media. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, 2016, 71–80, doi:10.1109/VAST.2016.7883513.
- [29] Blumenstein K, Leitner B, Thür N, Kirchknopf A, Seidl M, Aigner W. LiveVis: Visualizing Results of Second Screen Surveys in Real Time at TV Stages. In *Proceedings of the IEEE Workshop on Visualization in Practice*, 2017.
- [30] Meihoefer HJ. The utility of the circle as an effective cartographic symbol. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 1969, 6(2): 105–117, doi:10.3138/J04Q-1K34-26X1-7244.
- [31] Wagemans J, Elder JH, Kubovy M, Palmer SE, Peterson MA, Singh M, von der Heydt R. A Century of Gestalt Psychology in Visual Perception I. Perceptual Grouping and Figure-Ground Organization. *American Psychological Association*, 2012, 138(6): 1172, doi:10.1037/a0029333.
- [32] Xia J, Huang L, Lin W, Zhao X, Wu J, Chen Y, Zhao Y, Chen W. Interactive visual cluster analysis by contrastive dimensionality reduction. *IEEE Transactions on Visualization and Computer Graphics*, 2023, 29(1): 734–744, doi:10.1109/TVCG.2022.3209423.
- [33] Marler RT, Arora JS. The weighted sum method for multi-objective optimization: new insights. *Structural and multidisciplinary optimization*, 2010, 41: 853–862, doi:10.1007/s00158-009-0460-7.
- [34] Lu L, Choi YK, Sun F, Wang W. Variational circle packing based on power diagram. *Technical report, The University of Hong Kong*, 2011.
- [35] Bengio Y, Louradour J, Collobert R, Weston J. Curriculum learning. In *Proceedings of the International Conference on Machine Learning*, 2009, 41–48, doi:10.1145/1553374.1553380.
- [36] Chapelle O, Chi M, Zien A. A continuation method for semi-supervised SVMs. In *Proceedings of the International Conference on Machine Learning*, 2006, 185–192, doi:10.1145/1143844.1143868.
- [37] Yang W, Springer M, Strohaber J, Kolomenski A, Schuessler H, Kattawar G, Sokolov A. Spectral phase retrieval from interferometric autocorrelation by a combination of graduated optimization and genetic algorithms. *Optics express*, 2010, 18(14): 15028–15038, doi:10.1364/OE.18.015028.
- [38] Wang X, Chen Y, Zhu W. A Survey on Curriculum Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022, 44(9): 4555–4576, doi:10.1109/TPAMI.2021.3069908.
- [39] Chen J, Yang W, Jia Z, Xiao L, Liu S. Dynamic Color Assignment for Hierarchical Data. *IEEE Transactions on Visualization and Computer Graphics*, 2025, 31(1): 338–348, doi:10.1109/TVCG.2024.3456386.
- [40] Misue K, Eades P, Lai W, Sugiyama K. Layout adjustment and the mental map. *Journal of Visual Languages & Computing*, 1995, 6(2): 183–210, doi:10.1006/jvlc.1995.1010.
- [41] Xia J, Zhang Y, Song J, Chen Y, Wang Y, Liu S. Revisiting Dimensionality Reduction Techniques for Visual Cluster Analysis: An Empirical Study. *IEEE Transactions on Visualization and Computer Graphics*, 2022, 28(1): 529–539, doi:10.1109/TVCG.2021.3114694.
- [42] van der Maaten L, Hinton G. Visualizing Data Using t-SNE. *Journal of Machine Learning Research*, 2008, 9(86): 2579–2605.
- [43] McInnes L, Healy J, Melville J. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018, doi:10.48550/arXiv.1802.03426.
- [44] Wold S, Esbensen K, Geladi P. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 1987, 2(1–3): 37–52, doi:10.1016/0169-7439(87)80084-9.
- [45] Kruskal JB, Wish M. *Multidimensional scaling*. Sage, 1978.
- [46] Lee DD, Seung HS. Learning the parts of objects by non-negative matrix factorization. *Nature*, 1999, 401(6755): 788–791, doi:10.1038/44565.
- [47] Vernier EF, Garcia R, Silva IPd, Comba JLD, Telea AC. Quantitative Evaluation of Time-Dependent Multidimensional


- Projection Techniques. *Computer Graphics Forum*, 2020, 39(3): 241–252, doi:10.1111/cgf.13977.
- [48] Kruijer JF, Rauber PE, Martins RM, Kerren A, Kobourov S, Telea AC. Graph Layouts by t-SNE. *Computer Graphics Forum*, 2017, 36(3): 283–294, doi:10.1111/cgf.13187.
- [49] Zhong F, Xue M, Zhang J, Zhang F, Ban R, Deussen O, Wang Y. Force-Directed Graph Layouts Revisited: A New Force Based on the T-Distribution. *IEEE Transactions on Visualization and Computer Graphics*, 2023, doi:10.1109/TVCG.2023.3238821, to be published.
- [50] Goes Fd, Memari P, Mullen P, Desbrun M. Weighted Triangulations for Geometry Processing. *ACM Transactions on Graphics*, 2014, 33(3): 1–13, doi:10.1145/2602143.
- [51] Passenbrunner TE, del Re L. SkyObserver: Decentralized, real-time algorithm for deployment of a swarm of Unmanned Aircraft Systems. In *Proceedings of the IEEE International Conference on Control and Automation*, 2011, 633–638, doi:10.1109/ICCA.2011.6138044.
- [52] Lloyd S. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 1982, 28(2): 129–137, doi:10.1109/TIT.1982.1056489.
- [53] Rousseeuw PJ. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 1987, 20(1): 53–65, doi:10.1016/0377-0427(87)90125-7.
- [54] Aurenhammer F, Hoffmann F, Aronov B. Minkowski-type theorems and least-squares partitioning. *Algorithmica*, 1998, 20(1): 61–76, doi:10.1007/PL00009187.
- [55] Simonetto P, Archambault D, Scheidegger C. A Simple Approach for Boundary Improvement of Euler Diagrams. *IEEE Transactions on Visualization and Computer Graphics*, 2016, 22(1): 678–687, doi:10.1109/TVCG.2015.2467992.
- [56] Catto E. Box2D, 2010, <https://box2d.org>.
- [57] Zunic J, Rosin P. A new convexity measure for polygons. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2004, 26(7): 923–934, doi:10.1109/TPAMI.2004.19.
- [58] Xiao T, Xia T, Yang Y, Huang C, Wang X. Learning from massive noisy labeled data for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, 2691–2699, doi:10.1109/CVPR.2015.7298885.
- [59] Harrison Jr D, Rubinfeld DL. Hedonic housing prices and the demand for clean air. *Journal of environmental economics and management*, 1978, 5(1): 81–102, doi:10.1016/0095-0696(78)90006-2.
- [60] Adetunji AB, Akande ON, Ajala FA, Oyewo O, Akande YF, Oluwadara G. House price prediction using random forest machine learning technique. *Procedia Computer Science*, 2022, 199: 806–813, doi:10.1016/j.procs.2022.01.100.
- [61] Ghatnekar A, Shanbhag AD. Explainable, multi-region price prediction. In *2021 International conference on electrical, computer and energy technologies (ICECET)*, 2021, 1–7, doi:10.1109/ICECET52533.2021.9698641.
- [62] Munzner T. *Visualization analysis and design*. CRC press, 2014.
- [63] Brehmer M, Munzner T. A Multi-Level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics*, 2013, 19(12): 2376–2385, doi:10.1109/TVCG.2013.124.
- [64] Yang W, Liu M, Wang Z, Liu S. Foundation models meet visualizations: Challenges and opportunities. *Computational Visual Media*, 2024, 10(3): 399–424, doi:10.1007/s41095-023-0393-x.
- [65] Yang W, Chen C, Zhu J, Li L, Liu P, Liu S. A Survey of Visual Analytics Research for Improving Training Data Quality. *Journal of Computer-Aided Design & Computer Graphics*, 2023, 35(11): 1629–1642, doi:10.3724/SP.J.1089.2023.2023-00321.
- [66] Sandra G H. Nasa-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 2006, 50(9): 904–908, doi:10.1177/154193120605000909.
- [67] Paulovich FV, Minghim R. Hipp: A novel hierarchical point placement strategy and its application to the exploration of document collections. *IEEE Transactions on Visualization and Computer Graphics*, 2008, 14(6): 1229–1236, doi:10.1109/TVCG.2008.138.
- [68] Liu S, Liang Y, Gitter A. Loss-balanced task weighting to reduce negative transfer in multi-task learning. In *Proceedings of the AAAI conference on artificial intelligence*, 2019, 9977–9978, doi:10.1609/aaai.v33i01.33019977.
- [69] Sedlmair M, Tatu A, Munzner T, Tory M. A Taxonomy of Visual Cluster Separation Factors. *Computer Graphics Forum*, 2012, 31(3pt4): 1335–1344, doi:10.1111/j.1467-8659.2012.03125.x.
- [70] Dua D, Graff C. UCI Machine Learning Repository, 2017.
- [71] Georgiades A, Belhumeur P, Kriegman D. From few to many: illumination cone models for face recognition under variable lighting and pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2001, 23(6): 643–660, doi:10.1109/34.927464.
- [72] Forina M, Armanino C, Lanteri S, Tiscornia E. Classification of olive oils from their fatty acid composition. In *Proceedings of Food research and data analysis*, 1983.
- [73] Ventocilla E, Riveiro M. A comparative user study of visualization techniques for cluster analysis of multidimensional data sets. *Information visualization*, 2020, 19(4): 318–338, doi:10.1177/1473871620922166.
- [74] Viégas F, Wattenberg M, Hebert J, Borggaard G, Cichowlas A, Feinberg J, Orwant J, Wren C. Google+ripples: A native visualization of information flow. In *Proceedings of the 22nd international conference on World Wide Web*, 2013, 1389–1398.
- [75] Wang G, Zhang X, Tang S, Zheng H, Zhao BY. Unsupervised clickstream clustering for user behavior analysis. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, 2016, 225–236.

- [76] Kehlbeck R, Görtler J, Wang Y, Deussen O. SPEULER: Semantics-preserving Euler Diagrams. *IEEE Transactions on Visualization and Computer Graphics*, 2022, 28(1): 433–442, doi:10.1109/TVCG.2021.3114834.
- [77] Efrat A, Hu Y, Kobourov SG, Pupyrev S. MapSets: Visualizing embedded and clustered graphs. In *Proceedings of Graph Drawing*, 2014, 452–463.

Author biographies



Duan Li is a Ph.D. student at Tsinghua University, where he received his B.S. degree. His research interests lie in visual analytics.



Jun Yuan is currently a Ph.D. student at Tsinghua University, where he received his B.S. degree. His research interests lie in explainable artificial intelligence.




Xinyuan Guo is currently a Ph.D. student at Tsinghua University. His research interests include visual analytics and animated transition design.



Xiting Wang is an assistant professor at Renmin University of China. Her research interests include explainable machine learning and visual text analytics. She has academic papers in outlets such as KDD, IEEE Transactions on Knowledge and Data Engineering, AAAI, IJCAI, IEEE Transactions on Visualization and Computer Graphics and VAST.



Yang Liu is a principal researcher at Microsoft Research Asia. He received his Ph.D. degree from The University of Hong Kong, and master's and bachelor's degrees from The University of Science and Technology of China. His recent research focuses on geometric computation and learning-based geometry processing and generation. He is an associate editor of Transactions on Visualization and Computer Graphics.



Weikai Yang is an assistant professor in Hong Kong University of Science and Technology (Guangzhou). His research interests lie in visual analytics, machine learning, and data quality improvement. He received B.S. and Ph.D. degrees from Tsinghua University.



Shixia Liu is a professor at Tsinghua University. Her research interests include visual text analytics, visual social analytics, interactive machine learning, and text mining. She has worked as a research staff member at IBM China Research Lab and was a lead researcher at Microsoft Research Asia. She received B.S. and M.S. degrees from Harbin Institute of Technology, and a Ph.D. from Tsinghua University. She is a fellow of IEEE and an associate editor-in-chief of Transactions on Visualization and Computer Graphics.

Graphical abstract

