BMC Bioinformatics

RESEARCH

Open Access

scBubbletree: computational approach for visualization of single cell RNA-seq data

Simo Kitanovski[1*], Yingying Cao[1], Dimitris Ttoouli[1], Farnoush Farahpour[1,2], Jun Wang[1,3] and Daniel Hoffmann[1*]

*Correspondence:
simo.kitanovski@uni-due.de;
daniel.hoffmann@uni-due.de

[1] Bioinformatics and Computational Biophysics, Faculty of Biology and Centre for Medical Biotechnology (ZMB), University of Duisburg-Essen, 45141 Essen, Germany
[2] Institute of Cell Biology (Cancer Research), University Hospital Essen, University of Duisburg-Essen, 45147 Essen, Germany
[3] National Clinical Research Centre for Infectious Diseases, The Third People's Hospital of Shenzhen and The Second Affiliated Hospital of Southern University of Science and Technology, Shenzhen 518112, Guangdong Province, China

Abstract

Background: Visualization approaches transform high-dimensional data from single cell RNA sequencing (scRNA-seq) experiments into two-dimensional plots that are used for analysis of cell relationships, and as a means of reporting biological insights. Yet, many standard approaches generate visuals that suffer from overplotting, lack of quantitative information, and distort global and local properties of biological patterns relative to the original high-dimensional space.

Results: We present scBubbletree, a new, scalable method for visualization of scRNA-seq data. The method identifies clusters of cells of similar transcriptomes and visualizes such clusters as "bubbles" at the tips of dendrograms (bubble trees), corresponding to quantitative summaries of cluster properties and relationships. scBubbletree stacks bubble trees with further cluster-associated information in a visually easily accessible way, thus facilitating quantitative assessment and biological interpretation of scRNA-seq data. We demonstrate this with large scRNA-seq data sets, including one with over 1.2 million cells.

Conclusions: To facilitate coherent quantification and visualization of scRNA-seq data we developed the R-package scBubbletree, which is freely available as part of the Bioconductor repository at: https://bioconductor.org/packages/scBubbletree/

Keywords: scRNA-seq, Visualization, Transcriptomics

Introduction

Single cell RNA sequencing (scRNA-seq) data can convey an unprecedented richness of biological information, which has led to an explosion of scRNA-seq experiments [1]. However, the complexity of scRNA-seq data makes analysis also notoriously difficult: the transcriptome of each cell is characterized by a high-dimensional vector of gene expressions, and we have many cells and hence many vectors. How can we visually present such complex data so that essential biological information becomes easily accessible? Standard workflows perform reduction of the high-dimensional transcription data into printable two dimensional scatterplots [2, 3] with each cell corresponding to one dot and the geometric arrangement of the dots reflecting aspects of transcriptional similarity of cells. While this approach works reasonably well for small samples, visualization of large and complex samples runs into the well-known problem of overplotting [4], i.e. massive

© The Author(s) 2024. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

piling up of dots on top of each other severely hampers user's ability to identify patterns in the plotted data. This issue is compounded by distortions in the sample's distance structure that arise as we squeeze high-dimensional transcription data into two dimensions [5–7]. Thus, visuals produced with the standard workflow are in general inept for presenting quantitative information, including also transcriptional similarity.

Here we present scBubbletree, a method for quantitative visual exploration of scRNA-seq data. When designing scBubbletree our guiding question was: Which properties of scRNA-seq data should be visualized? We identified three classes of properties: (i) local and global transcriptional structure, (ii) cell density distribution and (iii) cell attributes (marker gene expressions, biological condition, cell type labels, etc.). Furthermore, scalability was an important factor in the design of scBubbletree, i.e. our method was developed to avoid overplotting. These goals were achieved by integrating in scBubbletree state-of-the-art methods for clustering of scRNA-seq data and for visualization. scBubbletree is available as an R-package that is open-source, easy-to-use and simple to integrate with popular approaches for scRNA-seq data analysis.

We demonstrate the added value of visual exploration with scBubbletree compared to popular methods, such as UMAP and t-SNE, by analyzing two public datasets.

Results

scBubbletree is implemented as an R package that provides an easy-to-use workflow for visual exploration of single cell RNA-seq data. A typical application, as presented here with two scRNA-seq samples (studies A and B), runs in less than one hour on a standard personal computer. Furthermore, scBubbletree can be integrated seamlessly with the R package Seurat [8], a popular tool for scRNA-seq analysis. In the next we describe the implementation of scBubbletree in detail.

Input

The pipeline takes as first input a matrix $A^{n\times f}$ that represents a low-dimensional projection of the high-dimensional scRNA-seq data, with n rows as cells and f columns as low-dimensional features. scBubbletree works directly with $A^{n\times f}$ and is agnostic about the initial data processing protocol. For the generation of the low-dimensional projection, we recommend the use of linear techniques that approximately conserve pairwise distances between cellular transcriptome vectors. Example techniques are principal component analysis (PCA) [9], non-negative matrix factorization [10], or multi-dimensional scaling [11]. *Non-linear* dimensionality reduction techniques, such as t-SNE [3] and UMAP [2], should be avoided as these approaches tend to distort long-range distances [12]. In the present work we used PCA for the low-dimensional projection.

Algorithm

In short, the algorithm of scBubbletree identifies clusters of transcriptionally similar cells, and then visualizes these clusters as leaf-nodes (bubbles) of a hierarchical dendrogram (bubbletree). The workflow comprises four steps (Fig. 1A): 1. determining the clustering resolution, 2. clustering, 3. hierarchical cluster grouping, and 4. visualization. We explain each step in the following.

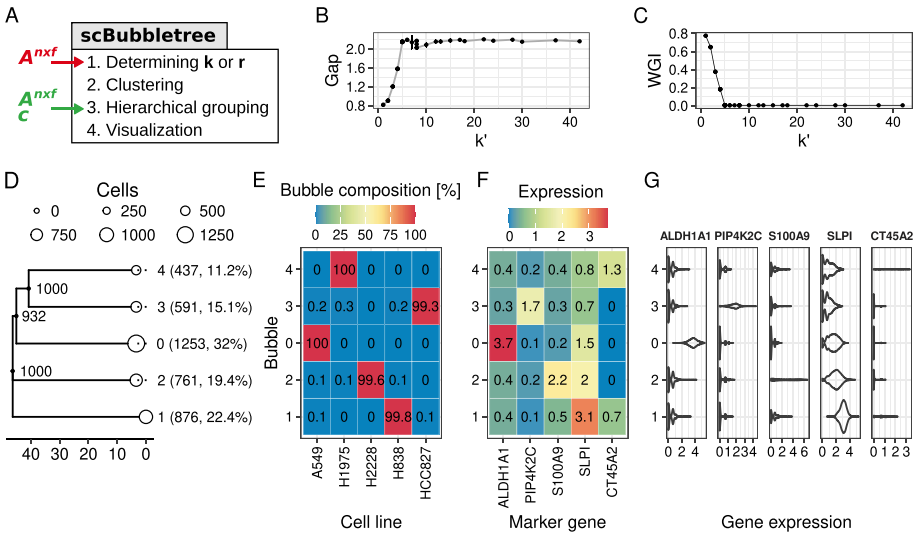

Fig. 1 scBubbletree analysis of dataset A. **A** scBubbletree workflow with main and alternative inputs shown with red and green labels, respectively. **B** Gap statistic for clustering solutions generated by the Louvain method with varying clustering resolutions. Vertical error bars are 95% confidence intervals of the gap statistic. **C** WGI for clustering solutions in panel B and the predicted cell lines of dataset A. **D** bubbletree (tree structure) annotated with a (**E**) heatmap displaying as rows the within-bubble relative frequencies of different cell lines. The bubbletree has five bubbles (white points) shown as leaves. Bubble radii scale linearly with the number of cells in the bubbles. Bubble identities, absolute and relative cell frequencies are shown as labels between bubbletree and heatmap. Gray branch labels are branch bootstrap support values. Rows in the heatmap integrate to 100%. **F** Mean normalized expression of five marker genes (x-axis) in each bubble (y-axis). **G** Density distribution of five marker gene (column panels) with normalized expressions of the cells in each bubble shown as violins. Numbers in heatmap tiles are rounded to the nearest tenth

Determining the clustering resolution with gap statistic

scBubbletree can cluster scRNA-seq data in two ways, namely by graph-based community detection (GCD) algorithms such as Louvain [13] or Leiden [14], and by k-means [15]. Benchmarking studies have demonstrated that approaches for GCD, such as Louvain, perform well with scRNA-seq data and have relatively short runtimes [16]. These findings were independently verified by our benchmarking study of the impact of clustering algorithm on biological cluster homogeneity (Supplementary Section 1). We recommend using the original Louvain algorithm for clustering of large scRNA-seq datasets. For smaller and simpler scRNA-seq datasets k-means generates comparable results.

Each clustering approach requires a resolution as input. In k-means, the clustering resolution is specified by the parameter k ($k \in \mathbb{N}$) that determines the number of clusters in the data, whereas Louvain (or Leiden) uses a resolution parameter r ($r \in \mathbb{R}_{>0}$), where higher r lead to more clusters and lower r lead to fewer clusters. Hence, r can be mapped onto a number k' of communities, with a range of r values of mapping onto the same value of k'. The choice of k and r should be guided primarily by the sample heterogeneity: the higher the cellular diversity in the sample, the higher the clustering resolution needed to resolve that diversity. The choice of k and r is also impacted by the research objective, for instance, high clustering resolution might be necessary to identify e.g. rare T cell subsets in a large sample of immune cells, whereas lower clustering resolution might be sufficient to characterize canonical immune subsets of that sample.

To determine k and r, scBubbletree relies on the *gap statistic* method [17] as implemented in the R-package cluster (function *clusGap*, version 2.1.2). This functionality is implemented in scBubbletree's functions *get_k* and *get_r*. The gap statistic compares the within cluster sum of squares (WCSS) to its expectation under an appropriate null reference distribution for different values of k and r. To estimate k and r we need to examine the gap curve for a) clear maximum or b) an 'elbow' (a bend in the curve from high to low slope). The functions *get_k* and *get_r* report the average gap and its 95% confidence interval for each clustering resolution computed based on B_{gap} simulations.

Many complementary approaches for determining the clustering resolution are available in the literature [18], and these can be used alongside the gap statistic. To avoid under- or over-clustering, which leads to the failure to detect rare populations or the false discovery of novel populations, respectively, we recommend that the estimated k' (associated with r) and k be compared against prior biological knowledge about the cellular composition based on data from e.g. the human protein atlas (https://www.proteinatlas.org) or from relevant publications. The estimated k' and k can also be compared against the clustering resolution inferred by the clustering approach sc-SHC [19], which has built-in hypothesis testing to automatically identify hierarchically organized clusters representing distinct populations, preventing both under-clustering and over-clustering.

Determining the clustering resolution with Gini impurity

For some scRNA-seq datasets it is possible to assess the subtypes of individual cells by mapping their transcriptional profiles onto reference atlases [20]. scBubbletree quantifies the homogeneity (or 'purity') of individual clusters i of n_i cells in terms of the their composition of subtype labels with the Gini impurity (GI) index:

$$\text{GI}_i = \sum_{j=1}^{L} \pi_{ij}(1 - \pi_{ij}), \tag{1}$$

where L is the total number of subtype labels, and π_{ij} the relative frequency of label j in cluster i. In homogeneous clusters, GI takes on a small value close to zero, and in mixed clusters GI can approach $1 - (1/L)$. Given a clustering result with k clusters, *get_gini* also computes the weighted Gini impurity (WGI) index by calculating the weighted average of the GIs:

$$\text{WGI} = \sum_{i=1}^{k} \text{GI}_i \frac{n_i}{n}, \tag{2}$$

with $n = \sum_i n_i$. We expect low WGIs for clustering results at optimal clustering resolutions, and higher WGIs for clustering results at insufficient resolutions. Hence, the WGI can be used like the gap statistic to identify appropriate resolutions.

Clustering

For clustering, scBubbletree uses as main inputs matrix $A^{n \times f}$, a clustering algorithm, and the clustering resolution (k or r) estimated from the previous step of this workflow. Four algorithms for GCD are available via the function *get_bubbletree_graph*, including, the original Louvain algorithm [13], the Louvain algorithm with multilevel

refinement (LMR) [21], the smart local moving (SLM) algorithm [22], and the Leiden algorithm [14]. In addition to this, k-means [15] clustering can be performed with the function *get_bubbletree_kmeans*. Auxiliary input parameters can be used to employ a specific variant of the clustering algorithms and to fine-tune their performance. To perform GCD, scBubbletree first constructs a shared nearest neighbor (SNN) graph using the function *FindNeighbors* (R-package Seurat, version 4.1.0). GCD is applied with Seurat's function *FindClusters* and k-means clustering by the function *kmeans* (R-package stats, version 4.2).

The output of the clustering step is a vector $c = (c_1, c_2, \ldots, c_n)$ of cluster assignments for the n cells, with $c_i \in \{1, 2,..., k\}$. The clustering result allows us to partition $A^{n\times f}$ into k data subsets $C_1^{n_1\times f}, C_2^{n_2\times f}, \ldots, C_k^{n_k\times f}$, with n_i as the number of cells in cluster i and $n = \sum_i n_i$. These matrices are used to compute inter-cluster distances in the next step of this workflow.

To facilitate the use of novel clustering approaches [23], scBubbletree provides the function *get_bubbletree_dummy*, which allows the user to provide as input matrix $A^{n\times f}$ and vector c of cluster assignments (Fig. 1A, green input) and proceed with hierarchical grouping of the clusters. This enables seamless integration of scBubbletree with computational pipelines that employ other clustering approaches, such as PhenoGraph [24], TooManyCells [25], sc-SHC [19] or Dirichlet process mixture models [26].

Hierarchical grouping of clusters

Clusters of cells are arranged in a dendrogram by performing agglomerative hierarchical clustering with average linkage [27] (function *hclust*, R-package stats, version 4.2). Starting at the lowest level with k clusters (or communities) of cells, the clustering procedure selects pairs of clusters with the smallest inter-cluster Euclidean distance and groups them at the next higher level of the hierarchy. This is repeated until a complete dendrogram is built with only one cluster at the highest level that contains the full data. Other commonly used linkage functions and the Manhattan distance metric are also implemented in scBubbletree.

The distance between clusters i and j can be computed by estimating the average Euclidean distance between all pairs of cells in $C_i^{n_i\times f}$ and $C_j^{n_j\times f}$. Hence, the complete hierarchical clustering procedure has time complexity $O(n^2)$ and requires approximately $\Omega(n^2)$ memory to store the resulting distance matrix. Assuming that each matrix cell takes up 8 bytes of memory, for a dataset of 10^6 cells this operation requires 7,450 gigabytes of memory, and considerable computational cost.

To avoid such high memory demands scBubbletree employs a bootstrapping approach with B iterations (Supplementary Algorithm 1). In iteration b, the algorithm draws with replacement a random subset with $n_i' = \min(n_i, N_{eff})$ and $n_j' = \min(n_j, N_{eff})$ number of cells from cluster i and j, respectively, where N_{eff} is a user-defined parameter that controls the maximum number of cells to be drawn from each cluster. The algorithm computes the average distance between cluster i and j based on $n_i' n_j'$ inter-cell Euclidean distances, and stores the result in distance matrix $D_b^{k\times k}$:

$$D_b^{ij} = \frac{1}{n_i' n_j'} \sum_{\substack{\forall \mathbf{x} \in X \\ \forall \mathbf{y} \in Y}} \sqrt{(\mathbf{x} - \mathbf{y})^2}, \tag{3}$$

where $X^{n_i' \times f}$ and $Y^{n_j' \times f}$ are subsets drawn from $C_i^{n_i \times f}$ and $C_j^{n_j \times f}$, respectively, and $\mathbf{x}$ and $\mathbf{y}$ are cell transcriptome features (row vectors) from each matrix. At the end of bootstrap iteration b, the distance matrix $D_b^{k \times k}$ is provided as input for hierarchical clustering with average linkage to generate the dendrogram H_b.

From the collection of B distance matrices we compute an average distance matrix ($\hat{D}^{k \times k}$), and use $\hat{D}^{k \times k}$ to generate a consensus hierarchical dendrogram (bubbletree; $\hat{H}$) by hierarchical clustering with average linkage. The collection of bootstrap dendrograms are used to assess the robustness of the bubbletree topology by quantifying the number of times each branch in $\hat{H}$ was found among the topologies of the bootstrap dendrograms (function *prop.clades*, R-package ape, version 5.6.2). Branches can have variable degrees of support ranging between values close to 0 (no support) and B (complete support). Distances between individual bubbles are visualized quantitatively as sums of branch lengths in root-to-leaf direction between these bubbles. In a horizontally arranged bubbletree, the vertical dimension has no significance and is only used to visually lay out the tree with the bubbles evenly spaced.

Visualization of the bubbletree

The main visual output of scBubbletree is a bubbletree (Figs. 1D and 2C). The bubbles represent clusters of cells of similar transcriptomes, and the tree describes hierarchical relationships and distances between the bubbles. Bubbletrees are visualized with ggtree (R-package, version 3.2.1), which offers a rich syntax for visualization of dendrograms [28]. ggtree extends the flexible ggplot2 (R-package, version 3.3.5) visualization framework [29], and thereby allows multiple layers of annotations to be attached to the bubbletree.

scBubbletree scales the bubble radii linearly as the number of cells in the corresponding bubbles. As linear increase in bubble radii leads to quadratic increase in bubble areas, this scaling visually emphasizes bubbles with higher numbers of cells. Each bubble is labeled with an identifier and its absolute and relative numbers of cells.

In summary, a bubbletree gives a quantitative visual summary of scRNA-seq data in terms of transcriptome cluster sizes and relationships.

Visualization of cell attributes

To make bubbles biologically interpretable, bubbletrees facilitate addition of further information, including *numeric* cell attributes, such as expression levels of marker genes, number of unique molecular identifiers (UMIs), or prevalence of UMIs from mitochondrial genes; or *categorical* cell attributes, such as predicted subtype labels, treatment groups (cancer vs. control cell), sample name, or cell cycle phase. For visualization of numeric and categorical cell attributes scBubbletree relies on the R-package ggplot2. Assembly of plots from multiple cell attributes is done with the R-package patchwork (version 1.1.1), which enables attaching of user-generated ggplot2 figures to the bubbletree.

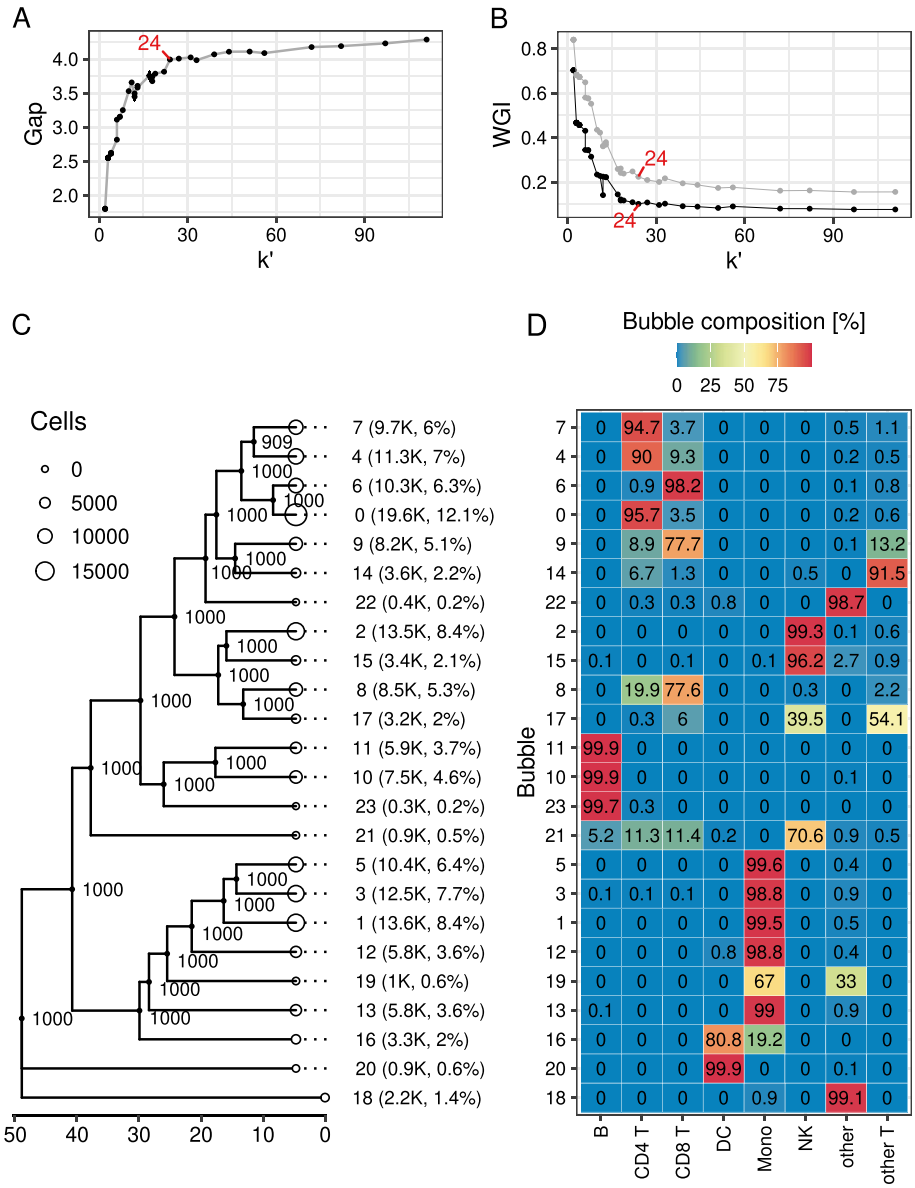

Fig. 2 scBubbletree analysis of dataset B. **A** Gap statistic for clustering solutions generated by Louvain clustering algorithm with varying clustering resolutions. Vertical error bars are 95% confidence intervals of the gap statistic. **B** WGI index for clustering solutions by Louvain algorithm and the cell type annotation sets *l1* (gray points and lines) and *l2* (black points and lines). **C** bubbletree with 24 bubbles (white points) shown as leaves. Bubble radii scale linearly with the number of cells in the bubbles. Bubble identities, absolute (in thousands) and relative cell frequencies are shown as labels. Gray branch labels are branch bootstrap support values. Clades of lymphocytes and monocytes are labeled with symbols 'a' and 'b'. **D** Within-bubble relative frequencies (%) of cell types (*l1*; x-axis) rounded to the nearest tenth. Rows integrate to 100% (up to rounding error)

scBubbletree provides three functions for visualization of numeric cell attributes: the function *get_num_tiles* computes statistical summaries (mean, median, sum, number of zero/nonzero values, etc.) of numeric cell attributes in each bubble and visualizes them with tile plots (heatmaps); the function *get_num_violins* visualizes the distributions of numeric cell attributes in each bubble with violin plots; the function *get_num_cell_tiles*

visualizes numeric attributes of individual cells in each bubble using heatmaps, with the cells ordered by the magnitude of their numeric attributes.

Categorical cell attributes are visualized using a matrix of tiles in which columns represent specific attribute categories. Tile colors and annotations can be configured for two types of interpretation using the logical parameter *integrate_vertical*. For *integrate_vertical=T* the tiles will show the relative frequencies of each attribute label across the different bubbles, which answers questions such as: are cells carrying a specific label (e.g. 'T cell' or sample label 'S1') enriched in a particular set of bubbles? Alternatively, for *integrate_vertical=F* the tiles will show the within-bubble relative frequencies of different labels. With this information we can answer questions such as: what is the label composition in a particular bubble?

Evaluation of scBubbletree using publicly available data

scBubbletree was evaluated with two publicly available scRNA-seq datasets, which we call A and B; for detailed descriptions of datasets and workflow see Supplementary Section 2.

Dataset A contains 3,918 cell transcriptomes from five human lung adenocarcinoma cell lines (HCC827, H1975, A549, H838 and H2228) [30]. The predicted cell line labels for each cell are available as part of the meta data. Dataset B contains transcriptomes of 161,764 peripheral blood mononuclear cells (PBMCs) from eight healthy volunteers enrolled in an HIV vaccine trial [8]. PBMC subtype predictions are available for each cell at two levels of resolution: annotation set *l1* and *l2* with 8 and 31 subtype categories, respectively.

In both datasets we saw that the first 15 principal components capture most of the variance in the data, and the proportion of variance explained by each subsequent principal component was negligible (Supplementary Fig. S1A-B). Thus, we used the single cell projections in 15-dimensional feature space, $A^{3,918\times15}$ and $A^{161,764\times15}$, as input of scBubbletree.

To assess scBubbletree's versatility in analyzing scRNA-seq data from solid tumors, we studied scRNA-seq samples from glioblastoma tissues. A detailed description of the datasets and workflow is available in Supplementary Section 3.

Results

Study A: exploring sample of five cancer cell lines

Generating the bubbletree

We know that dataset A contains a mixture of five cancer cell lines. Hence, we expect $k = 5$ cell clusters or communities. To verify this we computed the gap statistic for the Louvain method with resolution parameter $r = 10^R$, where R was initialized using a sequence of values from -4 to 1 in increments of 0.1. At each resolution r we recorded the number k' of identified communities (Supplementary Fig. S2A). WGIs were computed based on the cluster assignments and the vector of predicted cell line labels.

The gap curve had a distinct 'elbow' at $k' = 5$ ($r \in [0.003, 0.1]$) (Fig. 1B). In line with these observations, we saw a steep drop in the WGIs to a value close to 0 at $k' = 5$ ($r = 0.003$) (Fig. 1C) These results suggested that there is little benefit in using clustering resolutions that yield more than five communities.

In the next step of the scBubbletree workflow we applied clustering with the original version of Louvain [13] (function *get_bubbletree_graph*) and resolution parameter $r = 0.003$ ($k' = 5$). For input, we constructed a SNN graph based on $A^{3,918\times15}$, where each vertex (cell) was connected to its 50 nearest neighbors. Clustering was performed with 20 random starts and up to 100 iterations in each run. The clusters were organized in a bubbletree by hierarchical clustering with average linkage (tree in Fig. 1D). We ran $B = 1000$ bootstrap iterations and drew samples with up to $N_{eff} = 200$ cells from each cluster to estimate inter-cluster distances and their robustness.

Examining the bubbles

Bubble sizes range from the smallest (bubble 4, 437 cells, 11.2% of the sample cells) to the largest (bubble 0, 1,253 cells, 32% of the sample cells). By visualizing the cell lines as categorical attributes, we saw clear mapping between different cell lines and bubbles (heatmap in Fig. 1E), i.e. cells from a specific cell lines were enriched in specific bubbles. For instance, 100% of the cells in bubble 0 and 4 belonged to cell line A549 and H1975, respectively, and between 99.5% and 99.8% of the cells in bubbles 1, 2 and 3 belong to cell lines H838, H2228 and HC827, respectively.

By visualizing in each bubble the mean (normalized) expression (Fig. 1F) and the expression distributions (Fig. 1G) of five marker genes (one per cell line), we were able to confirm that bubble 0 was not only enriched with cells from cell line A549, but also had high ALDH1A1 expression which is a known A549 marker [31]. Similarly, bubble 4 had high expression of CT45A2 which is a marker of H1975 [32]. Bubbles 1, 2, and 3 were associated with high expression of SLPI, S100A9, and PIP4K2C, respectively, and based on data from the Cancer Dependency Map Portal (https://depmap.org; accessed 23.06.2022) we were able to confirm that these genes are typically over-expressed in cell lines H838, H2228, and HCC827, respectively.

Examining the tree

The distance between bubble 3 and 4 was smaller than that of any other pair of bubbles in the bubbletree (Fig. 1D). This hinted at relatively higher transcriptional similarity between the corresponding cell lines HCC827 and H1975. Furthermore, bubble 3 and 4 were in a robust sub-tree that was found in all 1,000 bootstrap dendrograms. The sub-tree of bubbles 0, 3, and 4 was less robust (932 of 1,000 bootstrap dendrograms).

The robust sub-tree of bubbles 3 (cell line HCC827) and 4 (cell line H1975) with the relatively short distance between these bubbles indicates transcriptional similarity of these cell lines. We checked this aspect of the tree by comparison with a set of independently measured transcriptional profiles of 69 adenocarcinoma cell lines from the Cancer Cell Line Encyclopedia [33], including HCC827 and H1975 (Supplementary Section 4). Euclidean distances were computed between the vectors of normalized gene expressions of all pairs of cell lines. The distance between HCC827 and H1975 was the 6th lowest among 2,346 pairwise comparisons (Supplementary Fig. S3), whereas the distances between the remaining cell lines of dataset A fell well within the general distribution of distances between the different adenocarcinoma cell lines. This result is consistent with the structure of the bubbletree, especially the robust closeness of bubbles 3 and 4.

Dataset A had been measured with an artificially mixture of five distinct cell line populations, which explains the simple structure of the resulting bubbletree. The recovery of the ground truth of five cell lines (Fig. 1), was a sanity check for scBubbletree. In this simple case, even UMAP and t-SNE were able recover aspects of the overall structure (Supplementary Fig. S4A,B). The advantage of bubbletrees over these standard methods was the quantitative visualization of tree structure and cluster sizes. However, scBubbletree has been developed for the analysis of more complex scRNA-seq data that are typically observed with real tissue samples, as in the following study B.

Study B: exploring a sample of human PBMCs

Clustering resolution

To determine the clustering resolution of dataset B we computed the gap statistic for the Louvain method with resolution parameter $r = 10^R$, where R was initialized using a sequence of values from -4 to 1 in steps of 0.1 (Fig. 2A). At each resolution we recorded the number k' of identified communities (Supplementary Fig. S2B).

The gap increased rapidly between $k' = 1$ and $k' = 11$, followed by a dip at $k' = 12$ and a further ascend to about 4 at $k' = 24$ ($r \approx 0.79$), and a much slower, approximately linear increase at higher k' (Fig. 2A). The increase in the gap between $k' = 22$ and $k' = 24$ was about two-fold larger than that in the interval between $k' = 22$ and $k' = 57$, even though only two communities were added in the first interval and 23 communities (about 10-fold more) were added in the second interval. This suggested that there is little benefit in using values of r larger than 0.794 ($k' = 24$).

The value of $k' = 24$ is close to the number of 18 canonical cell populations identified in PBMCs in the Human Protein Atlas (HPA) database (https://www.proteinatlas.org/humanproteome/immune+cell).

Clustering resolution analysis based on WGI

To compute WGI curves we used as inputs the PBMC subtype labels from annotation set *l1* and *l2* and the Louvain clustering assignments generated for the gap analysis. We saw that annotation set *l1* (Fig. 2B) was associated with lower WGIs compared to *l2*. The relatively lower WGIs of *l1* are likely due to the fact that *l1* contains labels of 8 major PBMC subtypes that have distinct transcriptional profiles. Hence, in each bubble of dataset B we expect a low degree of mixing in terms of different *l1* labels (Fig. 2D). In contrast to this, annotation set *l2* contains labels of 31 PBMC subtypes including, for instance, 13 T cell subtypes with comparable transcriptional profiles. Hence, we expect relatively higher degree of mixing of different *l2* labels in common clusters in comparison to the *l1* labels.

The WGI curve converged to low WGIs at $k' \approx 20$ for both annotation sets. This is similar to the clustering resolution of $k' = 24$ identified with the gap curve.

Clustering and generating the bubbletree

For clustering of dataset B we chose the original Louvain method and a resolution parameter $r = 0.794$ ($k' = 24$). First, we constructed a SNN graph based on $A^{161,764\times15}$, where each vertex (cell) in the SNN graph was connected to its 50 nearest neighbors. Clustering was performed with 20 random starts and up to 100 iterations in each run.

The clusters were organized in a bubbletree using hierarchical clustering with average linkage (Fig. 2C). For this we used $B = 1000$ bootstrap iterations and drew samples with up to $N_{eff} = 200$ cells from each cluster to compute inter-cluster distances.

The resulting bubbletree (Fig. 2C) had 24 bubbles with sizes ranging from 19,600 cells (12%, cluster 0) down to 300 cells (0.2%, cluster 23). The bubbletree had two major clades and two small outgroup bubbles (Fig. 2C). Clade 'a' contained 15 bubbles that accounted for about 65% of the cells in the sample, whereas clade 'b' contained seven bubbles that accounted for about 32% of the cells in the sample. Nearly all branches in the bubbletree were completely robust. The branch between bubble 4 and 7 had lower branch support as it was found in only 909 out of 1000 bootstrap dendrograms. A biological interpretation of the different clades and their within-clade branching patterns are provided in the following.

Bubbletree evaluation with annotation set l1 with 8 canonical PBMC subtypes

To determine which PBMC subtypes are enriched in each bubble we visualized the within-bubble relative frequencies of 8 major PBMC subtypes from annotation set *l1* (heatmap in Fig. 2D). Note that Figs. 2C, D demonstrate that bubbletrees (Fig. 2C) can be easily and coherently stacked with additional quantitative information (Fig. 2D), thus facilitating biological interpretation (Supplementary Fig. S5 for more extensive stacking).

Figure 2D shows that most bubbles are enriched with cells from either one PBMC subtype or a combination of related subtypes (e.g. $CD4^+$ and $CD8^+$ T cells). Furthermore, bubbles enriched with cells from related PBMC subtypes are close to each other in the bubbletree; for instance, bubbles 10, 11, and 23 are enriched with B cells and constitute a robust subclade of the bubbletree. Similarly, bubbles 5, 3, 1, 12, 19, 13 and 16 form clade 'b' and are enriched with monocytes and dendritic cells (DCs).

The bubbles of clade 'a' are mostly enriched with lymphocytes, and robust subclades within clade 'a' are formed by sets of bubbles that are enriched with specific lymphocyte subtypes, including T cells, B cells, and NK cells. The small outgroup bubble 21 contains a mixture of different lymphocytes. Clade 'b' has seven bubbles that are enriched in monocytes and DCs. The branches in that clade form a ladder-like topology in contrast to the more complex nested topology of clade 'a'. The simpler topology in 'a' may be due to the lower heterogeneity of circulating monocytes compared to lymphocytes, with classical monocytes constituting about 85% of all circulating monocyte pool in humans [34].

Clearly separated from clades 'a' an 'b' is the small outgroup of bubbles 20 of DCs and 18 consisting of a mixture of cells ('other').

Supplementary Section 5 contains a bubbletree analysis based on the more detailed annotation set *l2*, revealing enrichment of specific PMBC subtypes in the different bubbles.

Comparison with known approaches

We compared scBubbletree with UMAP, t-SNE and other popular methods for scRNA-seq data visualization with respect to visualization clarity, conservation of local and global data structure.

Conservation of local structure: qualitative analysis

In the following we compare the ability of the Louvain method to conserve local structure with those of UMAP and t-SNE. For this analysis we generated conventional UMAP and t-SNE figures for datasets A and B (Supplementary Fig. S4A–D).

In each figure, the different bubbles in dataset A and B were represented by heavily overplotted groups of dots in 2D UMAP and t-SNE space so that group size cannot be discerned from the figure. To our surprise, some cell lines from dataset A were split by UMAP and t-SNE into multiple subclusters; for instance, the cell line H1975 (bubble 4) had three distinct subclusters. The apparent distances between these subclusters in the drawing plane was more pronounced in t-SNE space (Supplementary Fig. S4B). Similarly, cell lines A549 (bubble 0) and H838 (bubble 1) were split into at least two subclusters of cells (Supplementary Fig. S4A–B). We realized that the cells from the smaller subclusters of A549 and H838 had systematically lower numbers of detected RNA molecules compared to the cells from the larger subclusters (Supplementary Fig. S6C–F). In contrast to this, the three subclusters in H1975 had comparable numbers of detected RNAs (Supplementary Fig. S6A–B), i.e. the subcluster structure of cell line H1975 could not be explained by differences in the number of detected RNA molecules. By using higher clustering resolution as input of Louvain we recovered these subclusters (data not shown). However, this was associated with negligible improvement in the gap and WGI (Fig. 1B–C).

Several conclusions can be drawn from this qualitative analysis. First, visualization of the local structure of scRNA-seq data with UMAP and t-SNE can be challenging even for relatively small samples. Second, UMAP and t-SNE are prone to generating clearly separated visual clusters despite few underlying biological differences [7]. Furthermore, UMAP and t-SNE use a number of hyperparameters, which can have severe impact on the 2D projection [2, 7] (Supplementary Fig. S7). Taken together, it is difficult to ascertain whether the subclusters in dataset A represent biologically distinct subpopulations of the main cell lines or technical artifacts resulting from biological/technical noise.

Conservation of distance structure

Users of conventional methods such as UMAP or t-SNE are tempted to interpret distances between dots or apparent clusters as indicators of transcriptional similarity of the corresponding cells. Yet, several studies [2, 35] have warned that while both, UMAP and t-SNE, consistently conserve distances on a small-scale, accurate conservation of large-scale distances is not always guaranteed.

To highlight the challenge in conserving large-scale distances with conventional methods, we analyzed UMAP and t-SNE plots of a dataset with over 1.3 million mouse brain cells (Supplementary Fig. S8A–B). The UMAP and t-SNE plots suffer from excessive overplotting, which hinders visual evaluation of the distances between cells and clusters. In addition to this, we see large spread (indicated by the wide and overlapping error bars in Supplementary Fig. S8C–D) in the distances between cells in 20D PCA space over the entire range of distances in 2D UMAP and t-SNE space, which indicates ambiguous mapping of distances in 2D UMAP and t-SNE space relative to distances in 20D PCA space. We saw similar patterns in the 2D UMAP and t-SNE plots of dataset B (Supplementary

Fig. S4G–H), even though dataset B is about one order of magnitude smaller than the dataset of mouse brain cells. As currently available workflows for scRNA-seq are able to generate data at similar scale and complexity as in these examples, we conclude that inference of transcriptional relationships between cells and clusters from UMAP and t-SNE will in general not be reliable.

Bubbletrees by construction should better reflect distances in PCA space in the dendrogram distances. In fact, Supplementary Figs. S4I–J and S9B demonstrate that bubbletree distances are consistent with the average distances in PCA space between cells from the corresponding bubbles. This was true for distances on all scales except for intra-bubble distances, which correspond to bubbletree distance = 0 (distances between bubbles to themselves). For bubbletree distances equal to 0 we saw consistently small average distances between cells in PCA space. It is noteworthy that the spread of distances between dots in PCA space was not uniform over all pairs of bubbles, presumably because the underlying clusters deviate more or less from homogeneous spherical shapes in PCA space, leading to the dispersion of inter-bubble cell distances.

Visualization clarity

We assessed the visualization clarity of scBubbletree with respect to two aspects: (1) *scalability*: the ability to handle large RNA-seq datasets; (2) *flexibility*: the ease of incorporating diverse cell attributes as part of the main visuals.

Scalability We investigated scalability in the context of the overplotting problem. A high degree of overplotting was observed in the 2D UMAP and t-SNE maps generated based on dataset A and B (Supplementary Fig. S4A–D). The overplotting was particularly excessive in the 2D UMAP and t-SNE maps of dataset B, where individual dots (cells) were barely visible and clusters of cells appeared as colored blobs. There are several consequences of such massive overplotting: 1. failure to grasp data's local and global structure; 2. inadequate description of cluster homogeneity in terms of different cell attributes; 3. loss of information about the density of cells across the different clusters; and 4. visual biases, i.e. users can deliberately or inadvertently emphasize or de-emphasize specific cell attributes by assigning the corresponding cells to the top of the cell piles [36].

Various techniques have been proposed to mitigate the overplotting problem, such as, reducing the size of the dots, adding an alpha channel for point transparency, data subsetting [37], hexagonal cell binning [36], applying density-conserving dimensionality reduction [38] or exploring the data with interactive graphical tools [39]. However, as the scale of scRNA-seq datasets continues to increase these countermeasures are of limited help, and most of them even have a negative impact on the clarity of the resulting figures.

Instead of visualizing individual cells, scBubbletree uses a high-level abstraction (bubbletree) to summarize scRNA-seq datasets and thereby avoids the overplotting problem. To demonstrate the high degree of scalability of scBubbletree in comparison to UMAP and t-SNE, we visualized with each approach a dataset of over 1.3 million mouse brain cells (Supplementary Fig. S8A–B and S9A). While both visuals are complex, the bubbletree approach at least offers an opportunity to understand the dataset, whereas the UMAP and t-SNE visuals lump all cell clusters into densely packed areas, making it

impossible to understand the structure of the data. scBubbletree is not the only approach that employs trees to visualize scRNA-seq data, i.e. approaches such as TooManyCells [25] and clustree [40] use similar strategies to mitigate the overplotting problem.

Flexibility The flexibility of scBubbletree was compared with that of UMAP, t-SNE, TooManyCells and clustree. Specifically, we assessed the ease with which each tool is able to integrate and combine multiple cell attributes. It is important to note that while these tools are functionally related to scBubbletree, there are inherent differences in the tasks they aim to solve. Therefore, a direct comparison of their results was not possible.

Cell attributes are visualized in UMAP and t-SNE plots by color-coding each point according the value of a particular attribute (e.g. expression level of a marker gene). For simultaneous visualization of two or more cell attributes, it is current practice to generate multiple and often small copies of the UMAP or t-SNE plots in which the individual attributes are visualized [41]. Although this method of integrating and presenting scRNA-seq data is widely used, it suffers from the same and even more shortcomings than individual UMAP and t-SNE plots.

TooManyCells is more flexible as it uses multiple facets to visualize cell attributes alongside its dendrogram e.g. by adjusting the color-code or the widths of its branches to for visual summaries (e.g. averages) of continuous cell attributes. TooManyCells can also attach pie charts to each leaf-node to visualize the within-leaf composition of a categorical attribute. Plots that encode information in such a diverse set of visual elements make it difficult for the viewer to decode and discover relationships.

The clustree algorithm visualizes cell attributes as part of its main dendrogram by adjusting the different visual elements such as colors, sizes, shapes and labels of the leaf-nodes. However, clustree is less flexible than TooManyCells as only few cell attributes can be integrated simultaneously alongside the main dendrogram. One of the key advantages of clustree is that it allows users to explore relationships between clusterings at multiple resolutions, which can be used determining the optimal number of clusters. Similar functionality is implemented in scBubbletree's function *compare_bubbletrees*, which enables comparison of pairs of bubbletrees generated from the same scRNA-seq data but with different input parameters such as e.g. clustering resolutions.

With scBubbletree we have therefore limited the number of types of visual elements (dendrogram, bubbles, and annotations) but made it easy to attach to the tree simple, matrix-like plots (heatmaps, violins) with one matrix row for each bubble (e.g. Figs. 1D or 2D). So the tree can be virtually read from left to right or top to bottom in a consistent way. Several of these matrix-like elements can be combined without hampering analysis by visual overloading (Supplementary Fig. S5).

Discussion

To avoid some of the common problems associated scRNA-seq data visualization, we designed scBubbletree to convey key properties of cells in ways that are visually easily accessible and scalable. These properties include quantitative information about local and global similarity of transcription patterns and cell density. The properties can be flexibly linked to stacks of visually easily accessible additional information, e.g. on marker gene expression, cell type, or information from multiomics experiments such as immune cell receptors, surface-protein expression or methylation levels. The quantitative nature

of its output lends scBubbletree to applications not only to research but also to quality control (Supplementary Section 2.4).

Yet, scBubbletree has clearly limitations. For instance, as a part of computing the dendrogram, the user has to choose a resolution. Although scBubbletree provides tools that support this choice, some familiarity with the underlying clustering concepts is required which precludes use of scBubbletree in a black-box manner. Furthermore, estimating the clustering resolution with *get_k* and *get_r* can be challenging in terms of computing times and memory for datasets of millions of cells, as these functions perform repetitive clustering simulations for a set of clustering resolutions. While the computational burden can be mitigated by parallelization, scBubbletree does not yet implement data representations that minimize memory requirements. In terms of its visualization capabilities, the current version of scBubbletree provides only the most essential functions for visualization of cell attributes. Future developments of the code will enrich its visualization repertoire in accordance with new developments in single-cell multiomics.

Conclusions

scBubbletree improves quantitative exploration of scRNA-seq data in several ways. First, in contrast to UMAP or t-SNE, the expressiveness and accessibility of scBubbletree output is robust against increasing number of cells, a feature that is becoming more and more important. Second, scBubbletree directly conveys quantitative information. Third, scBubbletree allows stacking of various types of information in easily accessible form. For instance, scBubbletree allows seamless integration of scRNA-seq data with other omics-modalities (e.g. genomic, transcriptomic, or epigenomic data). This feature is flexibly customizable and makes the analysis of the rich data available in current experiments straightforward.

Supplementary Information

The online version contains supplementary material available at https://doi.org/10.1186/s12859-024-05927-y.

Additional file 1:

Acknowledgements
Not applicable

Author contributions
SK, YC and DH conceived the study. SK developed the method and analyzed the data. SK and DH wrote the manuscript. YC, DT, JW and FF contributed to data interpretation. YC contributed to data acquisition. DT contributed code and evaluated the software. FF, JW, DT and DH critically revised the manuscript. All authors discussed the results and their implications and were involved in manuscript editing. All authors have read and approved the final manuscript.

Funding
Open Access funding enabled and organized by Projekt DEAL. This work was supported by DFG grants HO 1582/12-1 and GRK 2762, subprojects M1 und M3.

Availability of data and materials
The code and results are available at: https://github.com/snaketron/scBubbletree_paper_data. Dataset A is available (as object sce_sc_10x_5cl_qc) at https://github.com/LuyiTian/sc_mixology/raw/master/data/sincell_with_class_5cl.RData; and dataset B at https://atlas.fredhutch.org/data/nygc/multimodal/pbmc_multimodal.h5seurat.

Declarations

Ethics approval and consent to participate
Not applicable.

Consent for publication
Not applicable.

Competing interests
Not applicable.

Received: 26 June 2024 Accepted: 9 September 2024
Published online: 13 September 2024

References

1. Svensson V, Vento-Tormo R, Teichmann SA. Exponential scaling of single-cell RNA-seq in the past decade. Nat Protoc. 2018;13(4):599–604.
2. McInnes L, Healy J, Melville J. Umap: Uniform manifold approximation and projection for dimension reduction. 2018.
3. Maaten L, Hinton G. Visualizing data using t-SNE. J Mach Learn Res. 2008;9:2579–605.
4. Carr DB, Littlefield RJ, Nicholson W, Littlefield J. Scatterplot matrix techniques for large N. J Am Stat Assoc. 1987;82(398):424–36.
5. Marx V. 2024, Seeing data as t-sne and umap do. Nature Methods, 1–4
6. Chari T, Pachter L. The specious art of single-cell genomics. PLoS Comput Biol. 2023;19(8):1011288.
7. Huang H, Wang Y, Rudin C, Browne EP. Towards a comprehensive evaluation of dimension reduction methods for transcriptomic data visualization. Commun Biol. 2022;5(1):719. https://doi.org/10.1038/s42003-022-03628-x.
8. Hao Y, Hao S, Andersen-Nissen E, Mauck WM III, Zheng S, Butler A, Lee MJ, Wilk AJ, Darby C, Zager M, et al. Integrated analysis of multimodal single-cell data. Cell. 2021;184(13):3573–87.
9. Hotelling H. Analysis of a complex of statistical variables into principal components. J Educ Psychol. 1933;24(6):417.
10. Paatero P, Tapper U. Positive matrix factorization: a non-negative factor model with optimal utilization of error estimates of data values. Environmetrics. 1994;5(2):111–26.
11. Kruskal JB. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. Psychometrika. 1964;29(1):1–27.
12. Becht E, McInnes L, Healy J, Dutertre C-A, Kwok IW, Ng LG, Ginhoux F, Newell EW. Dimensionality reduction for visualizing single-cell data using UMAP. Nat Biotechnol. 2019;37(1):38–44.
13. Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. J Stat Mech: Theory Exp. 2008;2008(10):10008.
14. Traag VA, Waltman L, Van Eck NJ. From Louvain to Leiden: guaranteeing well-connected communities. Sci Rep. 2019;9(1):1–12.
15. MacQueen J *et al.* Some methods for classification and analysis of multivariate observations. In: Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, 1967;1:281–297. Oakland, CA, USA
16. Duò A, Robinson MD, Soneson C. A systematic performance evaluation of clustering methods for single-cell RNA-seq data. F1000Research. 2018;7:1141.
17. Tibshirani R, Walther G, Hastie T. Estimating the number of clusters in a data set via the gap statistic. J R Stat Soc: SerB (Statistical Methodology). 2001;63(2):411–23.
18. Yu L, Cao Y, Yang JY, Yang P. Benchmarking clustering algorithms on estimating the number of cell types from single-cell RNA-sequencing data. Genome Biol. 2022;23(1):1–21.
19. Grabski IN, Street K, Irizarry RA. Significance analysis for clustering with single-cell RNA-sequencing data. Nat Methods. 2023;20(8):1196–202.
20. Aran D, Looney AP, Liu L, Wu E, Fong V, Hsu A, Chak S, Naikawadi RP, Wolters PJ, Abate AR, et al. Reference-based analysis of lung single-cell sequencing reveals a transitional profibrotic macrophage. Nat Immunol. 2019;20(2):163–72.
21. Rotta R, Noack A. Multilevel local search algorithms for modularity clustering. J Exp Algorithmics (JEA). 2011;16:2.
22. Waltman L, Van Eck NJ. A smart local moving algorithm for large-scale modularity-based community detection. Eur Phys J B. 2013;86(11):1–14.
23. Kiselev VY, Andrews TS, Hemberg M. Challenges in unsupervised clustering of single-cell RNA-seq data. Nat Rev Genet. 2019;20(5):273–82.
24. Levine JH, Simonds EF, Bendall SC, Davis KL, El-ad DA, Tadmor MD, Litvin O, Fienberg HG, Jager A, Zunder ER, et al. Data-driven phenotypic dissection of AML reveals progenitor-like cells that correlate with prognosis. Cell. 2015;162(1):184–97.
25. Schwartz GW, Zhou Y, Petrovic J, Fasolino M, Xu L, Shaffer SM, Pear WS, Vahedi G, Faryabi RB. TooManyCells identifies and visualizes relationships of single-cell clades. Nat Methods. 2020;17(4):405–13.
26. Adossa NA, Rytkönen KT, Elo LL. Dirichlet process mixture models for single-cell rna-seq clustering. Biol Open. 2022;11(4):059001.
27. Hastie T, Tibshirani R, Friedman JH, Friedman JH. The elements of statistical learning data mining, inference, and prediction. Berlin: Springer; 2009. p. 520–8.
28. Yu G. Using ggtree to visualize data on tree-like structures. Curr Protoc Bioinform. 2020;69(1):96.
29. Wickham H. ggplot2: elegant graphics for data analysis. Berlin: Springer; 2016.
30. Tian L, Dong X, Freytag S, Lê Cao K-A, Su S, JalalAbadi A, Amann-Zalcenstein D, Weber TS, Seidi A, Jabbari JS, et al. Benchmarking single cell RNA-sequencing analysis pipelines using mixture control experiments. Nat Methods. 2019;16(6):479–87.
31. Park JW, Jung K-H, Lee JH, Moon SH, Cho YS, Lee K-H. Inhibition of aldehyde dehydrogenase 1 enhances the cytotoxic effect of retinaldehyde on A549 cancer cells. Oncotarget. 2017;8(59):99382.

32. Yang K, Chen Y, Zhou J, Ma L, Shan Y, Cheng X, Wang Y, Zhang Z, Ji X, Chen L, et al. Ursolic acid promotes apoptosis and mediates transcriptional suppression of CT45A2 gene expression in non-small-cell lung carcinoma harbouring EGFR T790M mutations. Br J Pharmacol. 2019;176(24):4609–24.
33. Barretina J, Caponigro G, Stransky N, Venkatesan K, Margolin AA, Kim S, Wilson CJ, Lehár J, Kryukov GV, Sonkin D, et al. The cancer cell line encyclopedia enables predictive modelling of anticancer drug sensitivity. Nature. 2012;483(7391):603–7.
34. Patel AA, Zhang Y, Fullerton JN, Boelen L, Rongvaux A, Maini AA, Bigley V, Flavell RA, Gilroy DW, Asquith B, et al. The fate and lifespan of human monocyte subsets in steady state and systemic inflammation. J Exp Med. 2017;214(7):1913–23.
35. Kobak D, Linderman GC. Initialization is critical for preserving global data structure in both t-SNE and UMAP. Nat Biotechnol. 2021;39(2):156–7.
36. Freytag S, Lister R. SCHEX avoids overplotting for large single-cell RNA-sequencing datasets. Bioinformatics. 2020;36(7):2291–2.
37. Hao Y, Stuart T, Kowalski M, Choudhary S, Hoffman P, Hartman A, Srivastava A, Molla G, Madad S, Fernandez-Granda C et al. Dictionary learning for integrative, multimodal, and scalable single-cell analysis. 2022.
38. Narayan A, Berger B, Cho H. Assessing single-cell transcriptomic variability through density-preserving data visualization. Nat Biotechnol. 2021;39(6):765–74.
39. Hillje R, Pelicci PG, Luzi L. Cerebro: interactive visualization of scRNA-seq data. Bioinformatics. 2020;36(7):2311–3.
40. Zappia L, Oshlack A. Clustering trees: a visualization for evaluating clusterings at multiple resolutions. Gigascience. 2018;7(7):083.
41. Stuart T, Butler A, Hoffman P, Hafemeister C, Papalexi E, Mauck WM III, Hao Y, Stoeckius M, Smibert P, Satija R. Comprehensive integration of single-cell data. Cell. 2019;177(7):1888–902.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

DuEPublico

Duisburg-Essen Publications online

UNIVERSITÄT DUISBURG ESSEN

Offen im Denken

ub universitäts bibliothek

This text is made available via DuEPublico, the institutional repository of the University of Duisburg-Essen. This version may eventually differ from another version distributed by a commercial publisher.

DOI: 10.1186/s12859-024-05927-y
URN: urn:nbn:de:hbz:465-20251212-100530-9

This work may be used under a Creative Commons Attribution 4.0 License (CC BY 4.0).