

SHORT REPORT

Open Access



A generalization of t-SNE and UMAP to single-cell multimodal omics

Van Hoan Do and Stefan Canzar*

*Correspondence:
canzar@genzentrum.lmu.de
Gene Center,
Ludwig-Maximilians-Universität
München, Feodor-Lynen-Str. 25,
Munich, Germany

Abstract

Emerging single-cell technologies profile multiple types of molecules within individual cells. A fundamental step in the analysis of the produced high-dimensional data is their visualization using dimensionality reduction techniques such as t-SNE and UMAP. We introduce j-SNE and j-UMAP as their natural generalizations to the joint visualization of multimodal omics data. Our approach automatically learns the relative contribution of each modality to a concise representation of cellular identity that promotes discriminative features but suppresses noise. On eight datasets, j-SNE and j-UMAP produce unified embeddings that better agree with known cell types and that harmonize RNA and protein velocity landscapes.

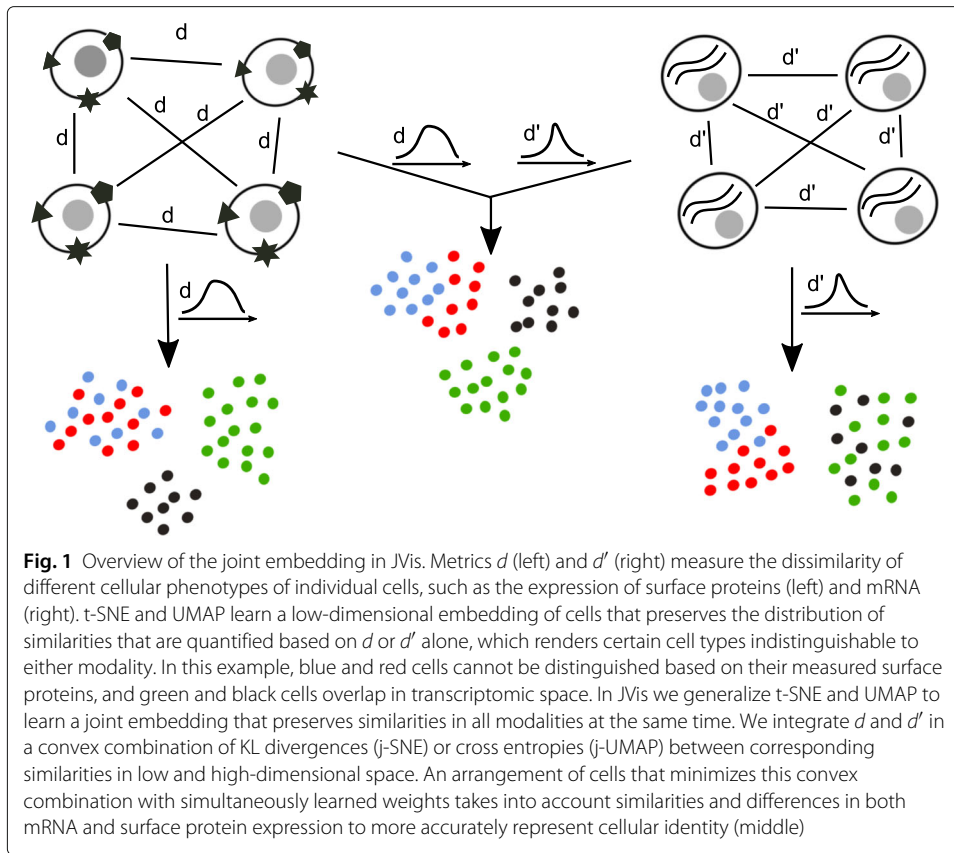
Keywords: Data visualization, Single-cell sequencing, Multimodal omics, t-SNE, UMAP, RNA velocity, Protein velocity

Background

Single-cell RNA sequencing has enabled gene expression profiling at single-cell resolution and provided novel opportunities to study cellular heterogeneity, cellular differentiation and development. Emerging single-cell technologies assay multiple modalities such as transcriptome, genome, epigenome, and proteome at the same time [1–3]. The joint analysis of multiple modalities has allowed to resolve subpopulations of cells at higher resolution [4, 5], has helped to infer the “acceleration” of RNA dynamics [6] and to extend time periods over which cell states can be predicted [7], and has linked dynamic changes in chromatin accessibility to transcription during cell-fate determination [8]. A fundamental step in the analysis of high-dimensional single-cell data is their visualization in two dimensions. Arguably the most widely used nonlinear dimensionality reduction techniques are t-distributed stochastic neighbor embedding (t-SNE) [9] and uniform manifold approximation and projection (UMAP) [10]. Currently, these techniques are applied to each modality one at a time [1, 8, 11], and separate views of the data need to be reconciled by manual inspection. Here, we generalize t-SNE and UMAP to the joint visualization of multimodal single-cell measurements. While t-SNE and UMAP seek a low-dimensional



© The Author(s). 2021 **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/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.



embedding of cells that preserves similarities in the original (e.g., gene expression) space as well as possible, we propose j-SNE and j-UMAP that simultaneously preserve similarities across all modalities (Fig. 1). Through Python package *JVis*, they will allow to combine different views of the data into a unified embedding that can help to uncover previously hidden relationships among them. At the same time, our joint embedding schemes learn the relative importance of each modality from the data to reveal a concise representation of cellular identity.

Results and discussion

In j-SNE, we want to learn a joint embedding \mathcal{E} of cells for each of which we have measured multiple modalities. Analog to t-SNE [9], we want to arrange cells in low-dimensional space such that similarities observed between points in high-dimensional space are preserved, but in all modalities at the same time. Generalizing the objective of t-SNE, we aim to minimize the convex combination of KL divergences of similarities in the original high-dimensional (distribution P) and similarities in the embedding low-dimensional space (distribution Q) for each modality k :

$$C(\mathcal{E}) = \sum_k \alpha_k KL(P^{(k)} || Q) + \lambda \sum_k \alpha_k \log \alpha_k, \quad (1)$$

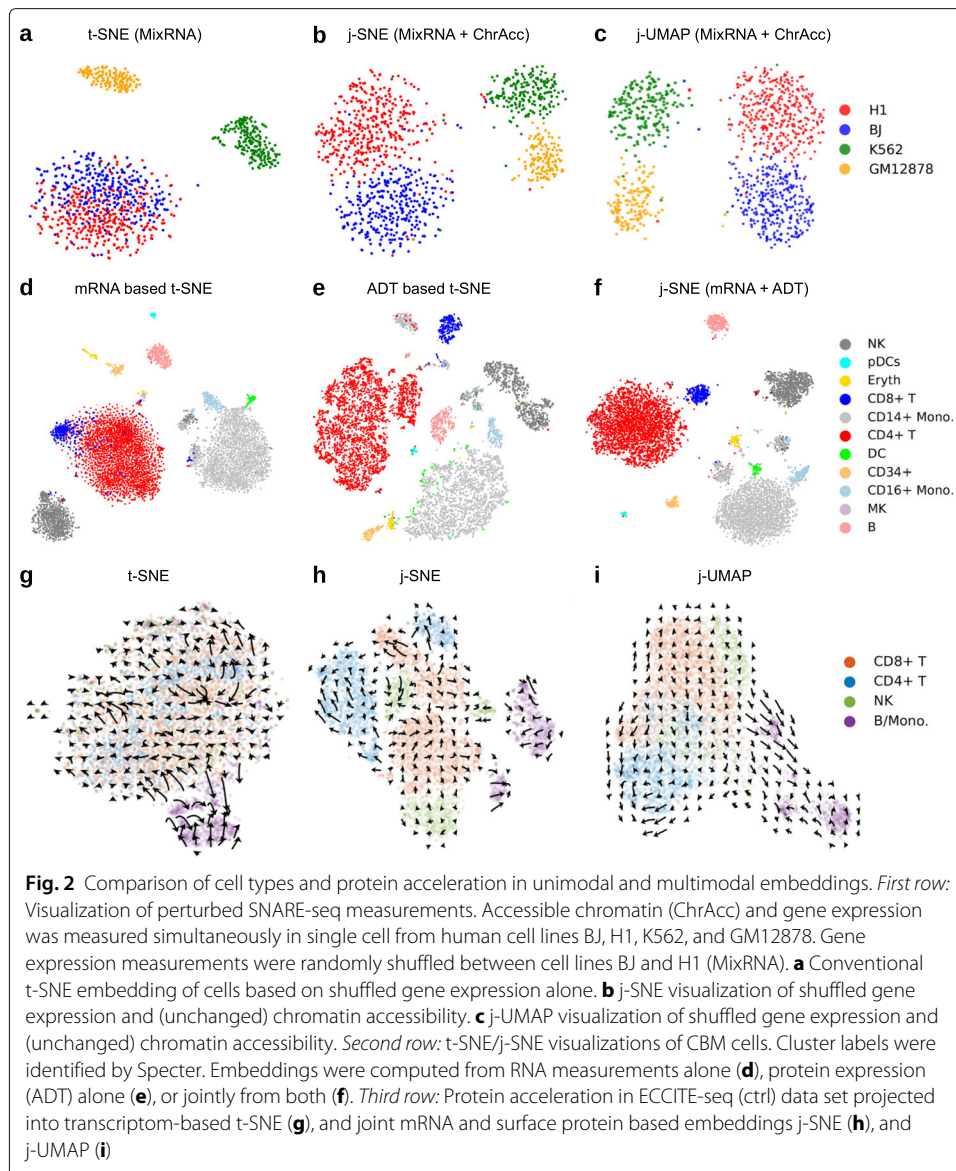
where coefficients α of the convex combination represent the importance of individual modalities towards the final location of points in the embedding. We add a regularization term (with regularization parameter λ) that prevents the joint embedding from

being biased towards individual modalities. In j-UMAP, we generalize UMAP to multimodal data analogously, minimizing a convex combination of cross entropies instead of KL divergences. We jointly optimize the location of points in the embedding and the importance coefficients α of modalities through an alternating optimization scheme: We fix coefficients α and find the best point locations by gradient descent, and in turn find optimal coefficients α for fixed locations by solving a convex optimization problem. Our approach is described in detail in Additional file 1: Supplementary Methods.

As proof of concept, we first demonstrate the ability of JVis to integrate modalities with different signal strengths. scRNA-seq, for example, often allows a finer mapping of cell states than single-cell ATAC-seq [12]. We used JVis to compute a joint embedding of accessible chromatin and gene expression measured simultaneously by SNARE-seq [11] in 1047 single cells from cultured human cell lines BJ, H1, K562, and GM12878. Similar to the conventional t-SNE and UMAP embeddings of transcriptomes or chromatin state alone, our joint j-SNE and j-UMAP embeddings clearly separate cells into four distinct clusters (Additional file 1: Fig. S1). Even when randomly shuffling gene expression measurements between cell lines BJ and H1 in a toy experiment, JVis employs chromatin accessibility to disentangle mixed mRNA measurements and separate all four cell lines (Fig. 2a–c and Additional file 1: Fig. S2).

To examine the effectiveness of the joint optimization scheme underlying JVis, we devise a simulation study following a similar strategy as [13]. We used Splatter [14] to simulate joint gene and ADT counts based on model parameters estimated from a real CITE-seq data set [15] in which mRNA and surface protein (ADT) expression were measured in human peripheral blood mononuclear cells (PBMC). We added a third modality that is obtained from gene expression measurements that are randomly shuffled between a random subset of cells. We generated eight synthetic multimodal data sets that vary in the relative abundance of (five) cell types, number of cells, and in the number of genes (Additional file 1: Table S1). In contrast to its conventional counterparts, j-SNE and j-UMAP learn weights for each modality from the data that reflect their relevance to the final embedding. Additional file 1: Figs. S3 and S4 show that these weights distinguish informative from noisy modalities. With an increasing amount of perturbation of the third modality, i.e., an increasing number of cells with shuffled gene expression, JVis assigns a lower weight to the corresponding modality. The rate of weight decrease (and the simultaneous increase mostly in ADT weight) is higher for data sets with a larger number of cells and, as expected, depends on the regularization coefficient λ . For λ close to 0, weights essentially include a single most informative modality (here ADT, see Additional file 1: Table S2) (Additional file 1: Fig. S5 and Supplementary Methods). Higher penalties associated with non-uniform weights result in a weaker adjustment of weights by the joint optimization scheme. The absolute adjustment of weights associated with cross entropy terms in j-UMAP is less pronounced than the adjustment of weights associated with KL divergences in j-SNE. Finally, in Additional file 1: Fig. S6 we show that weights computed by JVis are robust with respect to the precise subset of cells sampled in an experiment.

We measure the positive effect the learned weights have on the accuracy of the final embedding using the Silhouette score [16] that measures how well separated cell types are in the embedding, and metric KNI that we introduce as the fraction of k -nearest neighbors in the embedding that are of the same type, averaged over all cells. A high KNI value indicates homogeneous neighborhoods of cell types, while a random mixing of cells



would cause low KNI values. We compared the performance of JVis to conventional t-SNE and UMAP applied to the concatenation of modalities that were normalized by dividing them by the Frobenius norm of the count matrix and to the embedding obtained when assigning (fixed) uniform weights to each modality ($\alpha_i = 1/3$ in (1)). Additional file 1: Figs. S7-S10 demonstrate the benefit of borrowing information across modalities by the joint optimization scheme implemented in JVis. Compared to the normalized concatenation and the (uniform) averaging approach, the distinction between meaningful and noisy modalities in j-SNE and j-UMAP yields more accurate embeddings, across various noise levels. For data sets containing 5000 cells the separation of cell types in the embeddings obtained with (fixed) uniform weights continuously decreased with increasing noise levels. In contrast, the joint optimization scheme in j-SNE was able to retain a high accuracy on these data sets (Additional file 1: Figs. S7 and S9), especially for smaller penalties assigned to non-uniform weights, i.e. small values of λ . This is consistent with the sharper

drop in the weight associated with the noise modality observed for data sets containing 5000 cells and small values of λ (Additional file 1: Fig. S3). For j-UMAP, on the other hand, computed weights were close to uniform (for $\lambda \geq 0.5$) on data sets GpbmcN5k and GpbmcN5kD1k (Additional file 1: Fig. S4) and thus yielded embeddings with similar accuracy as the uniform weighting scheme on these data sets (Additional file 1: Figs. S8 and S10). The normalized concatenation approach works reasonably well on data sets with 1000 cells and large number of genes (N1k), but its performance varies substantially between different data sets and is even less accurate than its unnormalized version on data sets with 5000 cells. On most data sets, concatenation-based approaches show a sharp initial drop in accuracy for small levels of noise.

t-SNE and UMAP often produce embeddings that are in good agreement with known cell types or cell types computed by unsupervised clustering [17, 18] of high-dimensional molecular measurements such as mRNA expression. The simultaneous measurement of multiple types of molecules such as RNA and protein can refine cell types and JVis seeks to capture this refinement in their low-dimensional embedding. We compared unimodal and multimodal embeddings of mRNA and surface protein (ADT) expression measured in 4292 healthy human PBMCs [15] and in 8617 cord blood mononuclear cells (CBMC) [2] using CITE-seq [2]. Cell type labels were inferred by methods Specter [4] or CiteFuse [5], which have recently been introduced for the joint clustering of CITE-seq data.

Consistent with observations in [4, 5], t-SNE and UMAP visualizations of transcriptomic data alone does not show a clear distinction of CD4+ T cells and CD8+ T cells in the CBMC data set, while the embedding of protein expression mixes dendritic cells with CD14+ cells (Fig. 2d–f, Additional file 1: Figs. S11, S12). In contrast, JVis makes use of both modalities to compute a joint embedding that accurately separates CD4+ and CD8+ T cells as well as dendritic and CD14+ cells. Again, we confirm the visual interpretation quantitatively using the same metrics as above (Additional file 1: Table S2). The joint embedding of mRNA and ADT by JVis yields substantially larger Silhouette scores than the two unimodal t-SNE and UMAP embeddings.

Similarly, the joint embeddings of cells in the PBMC data set by JVis separate naïve and memory CD4+ T cell that are mixed in the ADT based t-SNE and UMAP embeddings as well as CD4+ and CD8+ T cells that are mixed in the mRNA based embeddings (Additional file 1: Figs. S13, S14). Again, joint embeddings are more accurate in terms of Silhouette scores than unimodal embeddings (Additional file 1: Table S2), even though overall the additional information provided by RNA measurements is limited relative to ADT counts on this data set.

RNA velocity [19] describes the rate of change of mRNA abundance estimated from the ratio of mature and pre-mRNA. While RNA velocity points to the future state of a cell, the recently introduced protein velocity [6] extends this concept and utilizes the joint measurement of RNA and protein abundance to infer the past, present, and future state of a cell. In [6], the authors used PCA and t-SNE to visualize RNA and protein velocity as well as the resulting *protein acceleration* in six PBMC data sets that were generated using four different technologies: CITE-seq, REAP-seq [20], ECCITE-seq [15] (data sets “CTCL”, a cutaneous T cell lymphoma patient, and “ctrl”, a healthy control), and 10X Genomics (data sets 1k and 10k). The authors observed strong velocity signals offered by the CITE-seq and 10x Genomics technologies, while REAP-seq and ECCITE-seq yielded noisier acceleration landscapes. Both RNA and protein velocity, however, were projected into the same

t-SNE embedding of transcriptomic measurements alone, rendering their interpretation difficult. We therefore repeated the analysis of the six different data sets but projected velocities into the joint embedding of both modalities computed by JVis. The noisy acceleration landscapes observed in [6] in the ECCITE-seq and REAP-seq data sets become aligned across cell types in their joint embeddings by JVis (Fig. 2g–i and Additional file 1: Figs. S15, S16). Consistent with the improved distinction of transcriptionally similar CD4 and CD8 T cells in the joint embeddings above, acceleration landscapes in all six data sets are projected onto an embedding that more clearly separates CD4 and CD8 T cells compared to the original ones proposed in [6] (Fig. 2g–i and Additional file 1: Figs. S15–S19). RNA and protein velocities (without Bézier curve fitting for acceleration) for all six data sets are shown in Additional file 1: Figs. S20–S25.

The noisy acceleration landscapes reported in [6] for the REAP-seq and ECCITE-seq data sets might be a result of the larger number of measured surface proteins (44 and 49 antibodies versus 13 and 17 antibodies in CITE-seq and 10X, respectively) that provide a finer distinction of subpopulations of cells. In fact, we observed lower agreement between RNA and protein based clusterings for the ECCITE-seq data set shown in Fig. 2 (ARI 0.21), compared to the clusterings obtained from the two modalities in the CITE-seq data set that agree well (ARI 0.82). Since protein acceleration is computed from both RNA and protein abundances, their joint embedding can help to reduce visualization artifacts that arise when protein velocities are projected into a purely transcriptome based t-SNE embedding as in [6].

The complexity of Barnes-Hut based t-SNE is $\mathcal{O}(n \log n)$, where n is the number of input cells [21]. Although no theoretical complexity bounds have been established for UMAP, its empirical complexity is $\mathcal{O}(n^{1.14})$ [10]. Since in addition the alternating minimization in j-SNE and j-UMAP requires only a few iterations of (conventional) t-SNE and UMAP calculations to converge to its final estimation of modality weights (Additional file 1: Fig. S26), JVis is expected to scale well to large data sets. For example, it took JVis less than 5 minutes to compute an embedding of the 10,000 cells contained in the largest data set used in this study (10x 10k). Memory usage and running time of j-SNE and j-UMAP shown in Additional file 1: Figs. S27 and S28 as a function of number of cells with 2 and 4 simulated modalities demonstrate practicability of both approaches in the analysis of larger and more complex multimodal data sets. The scalability of our approach to large data sets can be further improved by combining it with the recently proposed FFT-accelerated Interpolation-based t-SNE method [22], that scales linearly with the number of cells.

Conclusions

t-SNE and UMAP are routinely used to explore high-dimensional measurements of single cells in low-dimensional space. We have introduced method JVis that generalizes t-SNE and UMAP to the joint visualization of single-cell multimodal omics data. We have demonstrated that JVis combines multiple omics measurements of single cells into a unified embedding that exploits relationships among them that are not visible when applying conventional t-SNE or UMAP to each modality separately. Higher expected levels of noise in the measurements can be counteracted by smaller regularization coefficients λ that allow to downweight noisy modalities. Not surprisingly, projecting RNA and protein velocities into the joint embedding of both modalities yielded less noisy acceleration

landscapes compared to embeddings of mRNA measurements alone. We therefore anticipate that JVis will aid in the meaningful visual interpretation of data generated by emerging multimodal omics technologies such as CITE-seq [2] and SHARE-seq [23], the latter allowing to combine RNA velocity with *chromatin potential*.

Methods

A formal description of our generalizations j-SNE and j-UMAP as well as the algorithm to solve the underlying optimization problem can be found in Additional file 1: Supplementary Methods. The maximal number of iterations in our alternating optimization approach was set to 10 in all experiments (*maxIter*=10). Guided by the results of our simulation study and by visual inspection of known cell types, we set the regularization parameter λ to 3 for j-SNE and to 1 for j-UMAP in all experiments.

To simulate mRNA and ADT counts for each cell in data sets listed in Additional file 1: Table S1, we followed the strategy proposed in [13] and ran Splatter [14] with model parameters estimated from a real PBMC CITE-seq data set [15]. In particular, the same number of genes and antibodies were used, and ADT counts were simulated based on estimated dropout rate, library size, expression outlier, and dispersion across features. We added a third modality by duplicating gene expression measurements and randomly permuting expression vectors between a variable size random subset of cell. The larger the subset of cells, the larger the artificially introduced level of noise in this third modality. For the runtime and memory experiments, we generated a fourth modality by applying the same strategy to ADT counts, shuffling measurements between 40% of cells.

We measured the accuracy of an embedding using two different metrics. We introduce the k -nearest neighbor index (KNI), which denotes the fraction of k -nearest neighbors in the embedding that are of the same type. We used $k = 10$ if not specified otherwise and computed the average across all points. Different values of k yielded consistent results (Additional file 1: Figs. S29-S32). In addition, we used the Silhouette score [16] that ranges between -1 and 1 to measure how much cell types overlap (score 0) or how well separated (score 1) they are. We used the Adjusted Rand Index (ARI) [24] to measure the agreement between RNA and protein based clusterings. Clusterings of cells were computed using the Louvain algorithm [17] where the resolution parameter is tuned to match the number of annotated cell types. Following best practice [25], we used standard preprocessing of the input data including log-transformation of the expression matrix followed by principal component analysis (PCA) and applied j-SNE and j-UMAP as well as their conventional counterparts to 20 or 50 principle components. In all protein velocity experiments, preprocessed data was taken from the original study [6], no further preprocessing was performed. We computed protein acceleration using the protaccel Python package introduced in [6].

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s13059-021-02356-5>.

Additional file 1: Supplementary Methods, Tables S1, S2 and Fig. S1-Fig. S32.

Additional file 2: Review history

Peer review information

Barbara Cheifet was the primary editor of this article and managed its editorial process and peer review in collaboration with the rest of the editorial team.

Review history

The review history is Available as Additional file 2.

Authors' contributions

V.H.D. and S.C. conceived the algorithm. V.H.D. implemented the software and performed the computational experiments. S.C. guided the research. All authors wrote the manuscript.

Authors' information

Twitter handles: @vanhoan310 (Van Hoan Do); @StefanCanzar (Stefan Canzar).

Funding

V.H.D. was supported by a Deutsche Forschungsgemeinschaft fellowship through the Graduate School of Quantitative Biosciences Munich.

Availability of data and materials

The SNARE-seq and CBMC CITE-seq data sets were downloaded from Gene Expression Omnibus with accession codes GSE126074 and GSE126310, respectively. The six data sets used in the protein acceleration experiments were from [6]. The implementations of j-SNE and j-UMAP are based on the *scikit-learn* v0.23.1 library [26] and the UMAP v0.4.5 Python package [10], respectively. The *JVis* Python package can be installed through PyPi [27] and its open-source code is maintained at <https://github.com/canzarlab/JVis-learn> [28] under the 3-clause BSD license. The source code has been deposited in archived format at <https://doi.org/10.5281/zenodo.4682805> [29]. Python scripts to reproduce all results in this paper are available at https://github.com/canzarlab/JVis_paper [30].

Declarations**Ethics approval and consent to participate**

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Received: 6 January 2021 Accepted: 19 April 2021

Published online: 03 May 2021

References

- Cao J, Cusanovich DA, Ramani V, Aghamirzaie D, Pliner HA, Hill AJ, Daza RM, McFaline-Figueroa JL, Packer JS, Christiansen L, Steemers FJ, Adey AC, Trapnell C, Shendure J. Joint profiling of chromatin accessibility and gene expression in thousands of single cells. *Science*. 2018;361(6409):1380–5. <https://doi.org/10.1126/science.aau0730>.
- Stoeckius M, Hafemeister C, Stephenson W, Houck-Loomis B, Chattopadhyay PK, Swerdlow H, Satija R, Smibert P. Simultaneous epitope and transcriptome measurement in single cells. *Nat Methods*. 2017;14(9):865–8. <https://doi.org/10.1038/nmeth.4380>.
- Zhu C, Preissl S, Ren B. Single-cell multimodal omics: the power of many. *Nat Methods*. 2020;17(01):11–4. <https://doi.org/10.1038/s41592-019-0691-5>.
- Do VH, Ringeling FR, Canzar S. Linear-time cluster ensembles of large-scale single-cell rna-seq and multimodal data. *Genome Res*. 2021;31(4):677–688. <https://doi.org/10.1101/gr.267906.120>.
- Kim HJ, Lin Y, Geddes TA, Yang JYH, Yang P. CiteFuse enables multi-modal analysis of CITE-seq data. *Bioinformatics*. 2020;36(14):4137–43. <https://doi.org/10.1093/bioinformatics/btaa282>.
- Gorin G, Svensson V, Pachter L. Protein velocity and acceleration from single-cell multiomics experiments. *Genome Biol*. 2020;21:39. <https://doi.org/10.1186/s13059-020-1945-3>.
- Qiu X, Zhang Y, Hosseinzadeh S, Yang D, Pogson AN, Wang L, Shurtleff M, Yuan R, Xu S, Ma Y, Replogle JM, Darmanis S, Bahar I, Xing J, Weissman JS. Mapping transcriptomic vector fields of single cells. *bioRxiv* 696724. 2019. <https://doi.org/10.1101/696724>.
- Argelaguet R, Clark SJ, Mohammed H, Stapel LC, Krueger C, Kapourani C-A, Imaz-Rosshandler I, Lohoff T, Xiang Y, Hanna CW, Smallwood S, Ibarra-Soria X, Buettner F, Sanguinetti G, Xie W, Krueger F, Göttgens B, Rugg-Gunn PJ, Kelsey G, Dean W, Nichols J, Stegle O, Marioni JC, Reik W. Multi-omics profiling of mouse gastrulation at single-cell resolution. *Nature*. 2019;576(7787):487–91. <https://doi.org/10.1038/s41586-019-1825-8>.
- van der Maaten L, Hinton G. Visualizing data using t-SNE. *J Mach Learn Res*. 2008;9:2579–605.
- McInnes L, Healy J, Melville J. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv*. 2018;1802:arXiv Prepr arXiv:180203426 at <https://arxiv.org/abs/1802.03426>. 2018.
- Chen S, Lake BB, Zhang K. High-throughput sequencing of the transcriptome and chromatin accessibility in the same cell. *Nat Biotechnol*. 2019;37(12):1452–7. <https://doi.org/10.1038/s41587-019-0290-0>.
- Stuart T, Butler A, Hoffman P, Hafemeister C, Papalexi E, Mauck III WM, Hao Y, Stoeckius M, Smibert P, Satija R. Comprehensive Integration of Single-Cell Data. *Cell*. 2019;177(7):1888–902.
- Wang X, Sun Z, Zhang Y, Xu Z, Xin H, Huang H, Duerr RH, Chen K, Ding Y, Chen W. BREM-SC: a bayesian random effects mixture model for joint clustering single cell multi-omics data. *Nucleic Acids Res*. 2020;48(11):5814–24. <https://doi.org/10.1093/nar/gkaa314>.

14. Zappia L, Phipson B, Oshlack A. Splatter: simulation of single-cell RNA sequencing data. *Genome Biol.* 2017;18(1):174. <https://doi.org/10.1186/s13059-017-1305-0>.
15. Mimitou EP, Cheng A, Montalbano A, Hao S, Stoeckius M, Legut M, Roush T, Herrera A, Papalexi E, Ouyang Z, Satija R, Sanjana NE, Koralov SB, Smibert P. Multiplexed detection of proteins, transcriptomes, clonotypes and CRISPR perturbations in single cells. *Nat Methods.* 2019;16(5):409–12. <https://doi.org/10.1038/s41592-019-0392-0>.
16. Rousseeuw PJ. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J Comput Appl Math.* 1987;20:53–65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7).
17. Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. *J Stat Mech: Theor Exp.* 2008;2008(10):10008.
18. Kiselev VY, Kirschner K, Schaub MT, Andrews T, Yiu A, Chandra T, Natarajan KN, Reik W, Barahona M, Green AR, Hemberg M. SC3: consensus clustering of single-cell RNA-seq data. *Nat Methods.* 2017;14(5):483–6. <https://doi.org/10.1038/nmeth.4236>.
19. La Manno G, Soldatov R, Zeisel A, Braun E, Hochgerner H, Petukhov V, Lidschreiber K, Kastri ME, Lönnberg P, Furlan A, Fan J, Borm LE, Liu Z, van Bruggen D, Guo J, He X, Barker R, Sundström E, Castelo-Branco G, Cramer P, Adameyko I, Linnarsson S, Kharchenko PV. RNA velocity of single cells. *Nature.* 2018;560(7719):494–8. <https://doi.org/10.1038/s41586-018-0414-6>.
20. Peterson VM, Zhang KX, Kumar N, Wong J, Li L, Wilson DC, Moore R, McClanahan TK, Sadekova S, Klappenbach JA. Multiplexed quantification of proteins and transcripts in single cells. *Nat Biotechnol.* 2017;35(10):936–9. <https://doi.org/10.1038/nbt.3973>.
21. van der Maaten L. Accelerating t-SNE using tree-based algorithms. *J Mach Learn Res.* 2014;15(93):3221–45.
22. Linderman GC, Rachh M, Hoskins JG, Steinerberger S, Kluger Y. Fast interpolation-based t-SNE for improved visualization of single-cell RNA-seq data. *Nat Methods.* 2019;16(3):243–5. <https://doi.org/10.1038/s41592-018-0308-4>.
23. Ma S, Zhang B, LaFave L, Chiang Z, Hu Y, Ding J, Brack A, Kartha VK, Law T, Lareau C, Hsu Y-C, Regev A, Buenrostro JD. Chromatin potential identified by shared single-cell profiling of RNA and chromatin. *Cell.* 2020;183(4):1103–16. <https://doi.org/10.1016/j.cell.2020.09.056>.
24. Hubert L, Arabie P. Comparing partitions. *J Classif.* 1985;2(1):193–218.
25. Luecken MD, Theis FJ. Current best practices in single-cell RNA-seq analysis: a tutorial. *Mol Syst Biol.* 2019;15(6):8746. <https://doi.org/10.15252/msb.20188746>.
26. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E. Scikit-learn: Machine learning in Python. *J Mach Learn Res.* 2011;12:2825–30.
27. Do VH, Canzar S. A generalization of t-SNE and UMAP to single-cell multimodal omics. PyPi; software. 2021. <https://pypi.org/project/Jvis-learn/>.
28. Do VH, Canzar S. A generalization of t-SNE and UMAP to single-cell multimodal omics. Github. 2021. <https://github.com/canzarlab/Jvis-learn>.
29. Do VH, Canzar S. A generalization of t-SNE and UMAP to single-cell multimodal omics. Zenodo. 2021. <https://doi.org/10.5281/zenodo.4682805>.
30. Do VH, Canzar S. Reproducibility of Jvis results. Github. 2021. https://github.com/canzarlab/Jvis_paper.

Publisher's Note

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

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

