

LETTER

A case study on the detailed reproducibility of a Human Cell Atlas project

Kui Hua^{1,2}, Xuegong Zhang^{1,2,3,*}

¹ MOE Key Laboratory of Bioinformatics Division and Center for Synthetic & System Biology, BNRIST, Beijing 100084, China

² Department of Automation, Tsinghua University, Beijing 100084, China

³ School of Life Sciences, Tsinghua University, Beijing 100084, China

* Correspondence: zhangxg@tsinghua.edu.cn

Received November 22, 2018; Accepted December 9, 2018

Background: Reproducibility is a defining feature of a scientific discovery. Reproducibility can be at different levels for different types of study. The purpose of the Human Cell Atlas (HCA) project is to build maps of molecular signatures of all human cell types and states to serve as references for future discoveries. Constructing such a complex reference atlas must involve the assembly and aggregation of data from multiple labs, probably generated with different technologies. It has much higher requirements on reproducibility than individual research projects. To add another layer of complexity, the bioinformatics procedures involved for single-cell data have high flexibility and diversity. There are many factors in the processing and analysis of single-cell RNA-seq data that can shape the final results in different ways.

Methods: To study what levels of reproducibility can be reached in current practices, we conducted a detailed reproduction study for a well-documented recent publication on the atlas of human blood dendritic cells as an example to break down the bioinformatics steps and factors that are crucial for the reproducibility at different levels.

Results: We found that the major scientific discovery can be well reproduced after some efforts, but there are also some differences in some details that may cause uncertainty in the future reference. This study provides a detailed case observation on the on-going discussions of the type of standards the HCA community should take when releasing data and publications to guarantee the reproducibility and reliability of the future atlas.

Conclusion: Current practices of releasing data and publications may not be adequate to guarantee the reproducibility of HCA. We propose building more stringent guidelines and standards on the information that needs to be provided along with publications for projects that evolved in the HCA program.

Keywords: Human Cell Atlas; reproducibility; single cell; bioinformatics

Author summary: The Human Cell Atlas (HCA) project aims to build a comprehensive set of “Google Maps” for all cell types of a healthy human body. As a reference for future studies, it is important to guarantee that all maps are reproducible by third-party labs at high fidelity. Building the cell atlas or its parts involves complex bioinformatics procedures besides the bench work. Subtle differences in the bioinformatics processing may cause big differences in the resulted maps, but many current publications paid less attention to reporting the details of bioinformatics processing than the bench protocols. To study how the reproducibility can be reached under current practices, we conducted a detailed case study of a recent cell atlas work. The experiment provides observations helpful for safeguarding the reproducibility of the future HCA projects.

INTRODUCTION

The Human Genome Project (HGP) has provided a complete list of virtually all nucleic acid sequences of the

human genome [1,2]. Such a list, together with the annotations completed by HGP as well as follow-up projects like ENCODE, provided a fundamental reference for current biological and medical studies on human [3–

8]. The reference is a blueprint program for all cells in a human body. Not all cells in the same human body implement the program in the same way, and it is crucial to understand the genomic, transcriptomic, epigenomic, proteomic and metabolomics characteristics of all different types of cells at different time and organs to build a completing understanding of the human body in health and diseases. The Human Cell Atlas (HCA) program initiated from 2016 is the worldwide collaborative effort toward this goal [9,10]. The program was largely enabled by the recent development of single-cell genomics, especially single-cell RNA-sequencing (scRNA-seq) technology. It provides efficient ways to measure the expression of thousands of genes in each of thousands to millions of single cells [11]. Other single-cell technologies like single-cell ATAT-seq, single-cell Hi-C, single-cell metabolomics and technologies for single-cell resolution spatial transcriptomics are also under rapid development [12–16]. These single-cell technologies provide detailed observations on cells from multiple angles. They convert a biological cell to a data point in the high-dimensional space of multiple omics vectors. Cells that belong the same cell type or share similar molecular signatures may occupy a closed or open region in the high-dimensional space as clusters or distributions, and cell developments and state transitions can be reflected as linear or nonlinear trajectories in the space or in a certain subspace.

As a large-scale fundamental scientific program, HCA shares many features with HGP. But HCA has many unique and more challenging characteristics. Unlike the HGP which has a very specific and well-defined goal from the very beginning, for HCA, specifying what should be the essential defining components of the cell atlas and how they should be represented is one of the key questions to be answered by the program. There is a quick accumulation of publications on single-cell genomics studies in different human tissues [17–21]. They brought many advancements in the study of our immune systems, neuron systems, etc., and also in early embryonic development. Some of these data and discoveries are becoming the earliest parts of the future Human Cell Atlas. The most typical form of these discoveries are novel cell types or subtypes that have physiological or pathological importance, and gene expression markers of the discovered cell types [20,22,23]. Usually the discoveries are described in the published papers and supplementary materials, and some authors also release their raw or processed gene expression matrix data. There are major on-going efforts in the HCA community for building Data Coordination Platforms to serve as future banks for HCA data [24]. Multiple steps of bioinformatics processing and high-level analysis are involved in all single-cell studies. The bioinformatics methods are

usually described in the supplementary materials of published papers as they do for wet-lab protocols, mostly for the readers to understand what they have done. Most readers read the method parts to find hints to help their work on in-house data, without strong interests in reproducing published results. Many factors in the bioinformatics can have dramatic influences in the analysis of single-cell data, but a comprehensive guideline for how such factors should be documented and controlled for HCA has not been available yet.

After several attempts to work on the data of some recent work in this field, we realized that reproducing the final results from the original data can often be a challenging task. This may not imply any flaw in the published work, but usually the information provided is not sufficient for reproduction. To study the major factors in bioinformatics pipelines that can influence various results, and the level of disclosure that is necessary to guarantee full reproducibility, we chose a well-documented recent publication on the atlas of human blood dendritic cells [23] as an example to conduct a breakup reanalysis on the data. The level of details provided with this paper is among the most sufficient ones in recent publications in this field.

We found that with this level of details disclosed, the major scientific conclusion can be well reproduced, but still with noticeable differences in some potentially important details. This detailed case study provides experimental observations to support the design of guidelines and standards of the HCA community on the documentation and quality control of bioinformatics pipelines.

RESULTS

The original paper updated the taxonomy of human blood dendritic cells (DCs) and monocytes by finding new cell types in a single-cell study of ~2400 cells. In brief, they identified a new subtype of DC (named DC5), a new subdivision of a known subtype of DC (named DC2 and DC3), the existence of a conventional dendritic cell (cDC) progenitor, and two additional subtypes of monocytes by collecting and analyzing scRNA-seq data of 768 DCs and 372 monocytes derived from a healthy individual. The existence of these cell types is further confirmed in a subsequent study of additionally cells. They also studied the function of these newly discovered cell types and revealed the relationship between some of them. We mainly focused on the new discoveries on the taxonomy of the DCs in this reproduction experiment, *i.e.*, the discovery of DC5 and the separation of previously known CD1C⁺ to the new subtypes DC2 and DC3 (Figure 1A).

The original paper had provided not only the expression data matrix, but also detailed information on the

bioinformatics pipeline they used, including descriptions of the analyzing procedures and the computer codes for the key analyzing steps. With all these materials, the reproduction procedure is quite straightforward, except for a few minor bugs that need to be fixed (see Supplementary Materials) [25]. Figure 1B shows a diagram of the bioinformatics pipeline we adopted in this experiment according to the description in the original paper and supplementary file.

Figure 1C–1E show the comparison of the t-SNE plot obtained in the reproduction experiment (Figure 1D) and the original one (Figure 1C). The major discovery that the DCs are of 6 clusters is well re-discovered. The two sets of clusters can be well aligned except some confusion between clusters DC2 and DC3, and a few tiny confusions between several other clusters. The relative relations of the 6 clusters are also mostly recovered, except that the new cluster DC5 is shown to have closer relations to the cluster DC2, which is not shown in the original t-SNE

map. (Note that the t-SNE map does not preserve global distance information so differences in the overall layout of the plot do not indicate inconsistency [26–29].)

The major noticeable difference between the reproduced result and the original one is the separation of DC2 and DC3. Figure 2A and 2C show the comparison of the two heatmaps, drawn with the selected marker genes of each cluster with discriminating power of $AUC \geq 0.85$. The reproduction experiment assigned more cells to DC2 and fewer cells to DC3. From the enlarged heatmap (Figure 2B), we can see that the boundary between DC2 and DC3 on the marker genes are quite blurred. The slight differences in the clustering boundaries affected the identification of marker genes (There can also be other factors that caused the difference, see Supplementary Materials). To see how differences in cluster assignments would affect downstream analyses, we conducted standard differential expression genes (DE gene) analysis. Identifying DE genes is an old but still challenging




Figure 1. Major reproduction and comparison of the human dendritic cell atlas. (A) Overview of the original work from its paper. The red frame shows the focus of this reproduction study. (B) Workflow of the computational part in the original work (extracted from the code given in the original supplementary materials). (C) T-SNE plot of DCs from the original paper. (D) Reproduced t-SNE plot of DCs. (E) Confusion matrix between the reproduced label of DCs and the label provided by the original paper.




Figure 2. Comparison of the cell-type-specific marker genes. (A) Heatmap of cluster-specific markers from the original paper. (B) Heatmap of some markers of DC2 and DC3 (markers are selected according to original Figure 2A). Side color bars illustrate clustering results using different gene lists. (C) Reproduced heatmap of cluster-specific markers. (D) Comparison of cluster-specific markers between the original results and reproduced results.

question in single cell studies. Different methods can give quite different results [29,30]. Here we adopted the “one vs the rest” strategy and selected DE genes based the AUC score as in the original paper. The table in Figure 2D summarizes the numbers of marker genes discovered for each cluster in the original paper and in the reproduction experiment. We can see although most of the marker genes ranked on the top are shared in the two experiments, the two gene lists do have some differences. The list of marker genes and their expression distributions in the two studies are shown in the Supplementary Materials.

We conducted a step-by-step comparison of the reproduction experiment with the original experiment to pinpoint the factors that can cause the observed differences. The detailed description and analysis of these steps are provided in the Supplementary Materials.

The original paper did several downstream experiments with the discovered 6 DC clusters. They can be seen as examples for future applications of the 6 DC clusters as the reference of human DC atlas. In Figure 3A (the original Figure 3B), they showed the relations of the 4 monocyte clusters with the 6 DC clusters. Figure 3B shows the same analysis in the reproduction study. The relationship among mono1, mono2, mono3 and mono4 are basically preserved, but it is interesting to notice that with the presence of the monocyte cells, the DC5 cluster in the original study became connected with DC2, but it

becomes separated from DC2 in the reproduction.

The original paper recruited 10 independent donors to confirm the existence of AXL + SIGLEC6 + cells (“AS DCs”), and used Figure 3C (original Figure 4C) to show their relation with the 6 DC clusters. Figure 3D shows the reproduction result. We can see the although the two t-SNE maps seem not identical, but the relationships revealed between the clusters are consistent, including the observed closeness of the AS DCs to the DC5 cluster.

The original paper mapped pathogenic cells from blastic plasmacytoid dendritic cell neoplasm (BPDCN) patients to the healthy DC atlas. Figure 4A (original Figure 6G) showed that the BPDCN cells are mixed with DC6 cells in the PCA map. This result was generated using 174 out of 269 BPDCN cells. Since neither the final list of the used cells nor the filtering details were provided, we first used all 269 BPDCN cells to do the reproduction. In Figure 4B, we can see that cells from the 4 BPDCN patients are scattered aside the DCs in the map. They are close to DC6 but not mixed with it. We observed that the overall layout of the BPDCN cells are quite scattered, which also causes the 6 DC clusters squeezed, but some of the BPDCN cells seem to form a denser cluster which tends to be closer to DC6 than the more scattered ones. We therefore removed those scattered BPDCN cells as Suspected Materials low-quality cells and redid the PCA (see Supplementary Materials for details). After this

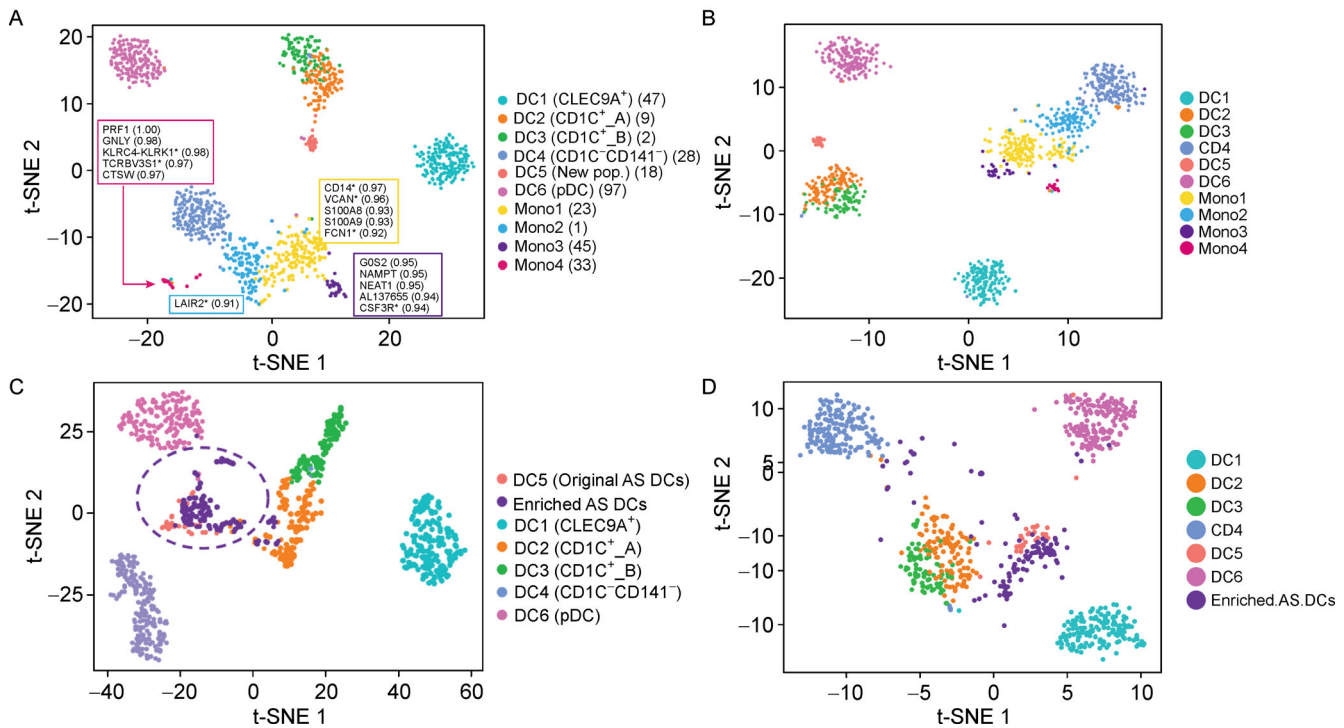


Figure 3. Reproduction of the human blood monocyte heterogeneity and the identification of AS DCs. (A) Original figure 3B. (B) Reproduction of the original figure 3B. (C) Original figure 4C. (D) Reproduction of the original figure 4C.

removing step, we got a similar PCA map where BPDCN cells are mixed with DC6 as in the original result. We noted that the description on the bioinformatics processing was less detailed for this part of the experiment than for the data used in the major discovery part in the original paper. Although we managed to get a similar result, this reproduction was achieved mostly by guessing and retries.

DISCUSSION

Single-cell genomics studies involve multiple steps of bioinformatics processing, including filtering of cells and selection of genes for quality control, batch-effect correction, normalization and dimension reductions for extracting major information, clustering, visualization and marker selection, etc. Some of the steps may need to be run more than once and need to be adapted according to the data and the specific question. Due to the nature of variability between many biological problems and between single-cell genomics technologies, it may not be feasible to fix a “standard” bioinformatics pipeline. Therefore, providing standard descriptions of the pipeline with sufficient details is crucial for the reproducibility. Through this detailed reproduction case study, we observed that the extraordinarily detailed information of

the bioinformatics procedures provided in the selected work has guaranteed the reproduction of the major scientific discovery, but still there are some unneglectable differences in the reproduced results. Some of the differences can be due to version changes of some third-party packages, and some may still be due to insufficient record or description in some minor details. Some seemingly subtle inaccuracies in the early and intermediate steps of processing may cause noticeable differences in the downstream results.

The HCA program is just at its beginning phase. All the current exciting discoveries based on single-cell genomics are just the beginning of the footstone construction of the future skyscraper. This distinguishes HCA projects with other single-cell studies that aim for only answering specific scientific questions. The construction of the full atlas will be done by joint efforts of many labs throughout the world, and it may span years or even decades. New discoveries are the focus of each individual publication, but the data used to support the discoveries are what accumulate in HCA, which will become part of the reference atlas for all future studies. The level of differences between results we observed in this case study can be subtle in terms of the major scientific discovery, but such tiny differences will be enlarged when the data act as part of the earliest footstones in the




Figure 4. Reproduction of the incorporated result of monocyte, DCs and the four BPDCN patient samples. (A) Original figure 6G. (B) Reproduction of the original figure 6G using all BPDCN cells ($n = 269$). (C) Reproduction of the original figure 6G using filtered BPDCN cells ($n = 175$).

construction. There have been many recent discussions on conventions that are needed for journal publications to guarantee the reproducibility of scientific discoveries [30–32]. From the observations of this case study, it is evident that an even higher standard of reproducibility is needed for HCA as a reference for future sciences.

The importance of HCA is comparable with HGP but the scale is much larger. At the beginning of HGP, scientists had estimations of the final goal and had estimations on how much is still remaining during the progress of the program. For HCA, the final goal is more open and is of multifaceted nature. And the technology is evolving and expanding in revolutionary paces. For many studies, it is too early to judge which part of the data or intermediate results are unimportant details. At this phase, keeping a reasonable level of diversity of bioinformatics pipelines and methods is healthy for the community. But this should be accompanied by sufficient record and sharing of the details to guarantee full reproduction. In this reproduction experiment, to guarantee all results we

got can be reproduced again by a third-party, we have created a notebook recording all details of the bioinformatics processing together with the corresponding live codes. Such efforts for guaranteeing reproducibility can serve as examples for all HCA projects. We propose building more stringent guidelines and standards on the information that needs to be provided along with publications for projects that will be taken as part of the HCA program. And a mechanism of systematic third-party reproduction validation should be introduced.

METHODS

Data availability

All results and the major bioinformatics processing are described in the Supplementary File. Data used in this reproduction and a jupyter notebook with all details of the reproduction are available at <https://github.com/Xuegon-Lab/HCA-reproducibility>.

SUPPLEMENTARY MATERIALS

The supplementary materials can be found online with this article at <https://doi.org/10.1007/s40484-018-0164-3>.

AUTHORIZATION STATEMENT

Figure 1A, Figure 1C, Figure 2A, Figure 3A, Figure 3C and Figure 4A in this paper are reprinted from Ref. [23]. The use of these figures is authorized with the license number 4502220162373.

ACKNOWLEDGEMENTS

We thank Nir Hacohen, Alexandra-Chloé Villani and Orit Rozenblatt-Rosen (authors of the original paper) for their helpful discussion. This work is supported by CZI Human Cell Atlas Pilot Project and the National Natural Science Foundation of China (Nos. 61673231 and 61721003).

COMPLIANCE WITH ETHICS GUIDELINES

The authors Kui Hua and Xuegong Zhang declare that they have no conflict of interests.

All procedures performed in studies were in accordance with the ethical standards of the institution or practice at which the studies were conducted, and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

REFERENCES

1. Watson, J. D. (1990) The human genome project: past, present, and future. *Science*, 248, 44–49
2. Collins, F. S., Morgan, M. and Patrinos, A. (2003) The Human Genome Project: lessons from large-scale biology. *Science*, 300, 286–290
3. Gibbs, R. A., Belmont, J. W., Hardenbol, P., Willis, T. D., Yu, F., Zhang, H., Zeng, C., Matsuda, I., Fukushima, Y., Macer, D. R., *et al.* (2003) The International HapMap Project. *Nature*, 426, 789–796
4. Feingold, E. A., Good, P. J., Guyer, M. S., Kamholz, S., Liefer, L., Wetterstrand, K., Collins, F. S., Gingeras, T. R., Kampa, D., Sekinger, E. A. *et al.* (2004) The ENCODE (ENCyclopedia of DNA Elements) project. *Science*, 306, 636–640
5. Haines, J. L., Hauser, M. A., Schmidt, S., Scott, W. K., Olson, L. M., Gallins, P., Spencer, K. L., Kwan, S. Y., Noureddine, M., Gilbert, J. R., *et al.* (2005) Complement factor H variant increases the risk of age-related macular degeneration. *Science*, 308, 419–421
6. The 1000 Genomes Project Consortium. (2010) A map of human genome variation from population-scale sequencing. *Nature*, 467, 1061–1073
7. The 1000 Genomes Project Consortium. (2012) An integrated map of genetic variation from 1092 human genomes. *Nature*, 491, 56–65
8. Kellis, M., Wold, B., Snyder, M. P., Bernstein, B. E., Kundaje, A., Marinov, G. K., Ward, L. D., Birney, E., Crawford, G. E., Dekker, J., *et al.* (2014) Defining functional DNA elements in the human genome. *Proc. Natl. Acad. Sci. USA*, 111, 6131–6138
9. The Human Cell Atlas Meeting Participants. (2017) The Human Cell Atlas. *eLife*, 6, e27041
10. Rozenblatt-Rosen, O., Stubbington, M. J. T., Regev, A. and Teichmann, S. A. (2017) The Human Cell Atlas: from vision to reality. *Nature*, 550, 451–453
11. Svensson, V., Vento-Tormo, R. and Teichmann, S. A. (2018) Exponential scaling of single-cell RNA-seq in the past decade. *Nat. Protoc.*, 13, 599–604
12. Cusanovich, D. A., Daza, R., Adey, A., Pliner, H. A., Christiansen, L., Gunderson, K. L., Steemers, F. J., Trapnell, C. and Shendure, J. (2015) Multiplex single cell profiling of chromatin accessibility by combinatorial cellular indexing. *Science*, 348, 910–914
13. Nagano, T., Lubling, Y., Stevens, T. J., Schoenfelder, S., Yaffe, E., Dean, W., Laue, E. D., Tanay, A. and Fraser, P. (2013) Single-cell Hi-C reveals cell-to-cell variability in chromosome structure. *Nature*, 502, 59–64
14. Zenobi, R. (2013) Single-cell metabolomics: analytical and biological perspectives. *Science*, 342, 1243259
15. Crosetto, N., Bienko, M. and van Oudenaarden, A. (2015) Spatially resolved transcriptomics and beyond. *Nat. Rev. Genet.*, 16, 57–66
16. Lein, E., Borm, L. E. and Linnarsson, S. (2017) The promise of spatial transcriptomics for neuroscience in the era of molecular cell typing. *Science*, 358, 64–69
17. Zhong, S., Zhang, S., Fan, X., Wu, Q., Yan, L., Dong, J., Zhang, H., Li, L., Sun, L., Pan, N., *et al.* (2018) A single-cell RNA-seq survey of the developmental landscape of the human prefrontal cortex. *Nature*, 555, 524–528
18. Muraro, M. J., Dharmadhikari G., Grün, D., Groen, N., Dielen, T., Jansen, E., van Gurp, L., Engelse, M. A., Carlotti, F., de Koning, E. J. P. *et al.* (2016) A single-cell transcriptome atlas of the human pancreas. *Cell Syst.*, 3, 385–394 e3
19. Darmanis, S., Sloan, S. A., Zhang, Y., Enge, M., Caneda, C., Shuer, L. M., Hayden Gephart, M. G., Barres, B. A. and Quake, S. R. (2015) A survey of human brain transcriptome diversity at the single cell level. *Proc. Natl. Acad. Sci. USA*, 112, 7285–7290
20. Macosko, E. Z., Basu, A., Satija, R., Nemesh, J., Shekhar, K., Goldman, M., Tirosh, I., Bialas, A. R., Kamitaki, N., Martersteck, E. M., *et al.* (2015) Highly parallel genome-wide expression profiling of individual cells using nanoliter droplets. *Cell*, 161, 1202–1214
21. Stubbington, M. J. T., Rozenblatt-Rosen, O., Regev, A. and Teichmann, S. A. (2017) Single-cell transcriptomics to explore the immune system in health and disease. *Science*, 358, 58–63
22. Jaitin, D. A., Kenigsberg, E., Keren-Shaul, H., Elefant, N., Paul, F., Zaretsky, I., Mildner, A., Cohen, N., Jung, S., Tanay, A., *et al.* (2014) Massively parallel single-cell RNA-seq for marker-free decomposition of tissues into cell types. *Science*, 343, 776–779
23. Villani, A. C., Satija, R., Reynolds, G., Sarkizova, S., Shekhar, K., Fletcher, J., Griesbeck, M., Butler, A., Zheng, S., Lazo, S., *et al.* (2017) Single-cell RNA-seq reveals new types of human blood dendritic cells, monocytes, and progenitors. *Science*, 356, eaah4573

24. Data Coordination – Human Cell Atlas (<https://www.humancellatlas.org/data-sharing>)
25. Butler, A., Hoffman, P., Smibert, P., Papalexi, E. and Satija, R. (2018) Integrating single-cell transcriptomic data across different conditions, technologies, and species. *Nat. Biotechnol.*, 36, 411–420
26. van der Maaten, L. and Hinton, G. (2008) Visualizing data using t-SNE. *J. Mach. Learn. Res.*, 9, 2579–2605.
27. Wattenberg, F. V. M. and Johnson, I. (2016) How to use t-SNE effectively. *Distill*<http://doi.org/10.23915/distill.00002>
28. Yuansheng Zhou, T. O. S. (2018) Using global t-SNE to preserve inter-cluster data structure. *bioRxiv*, Doi: <https://doi.org/10.1101/331611>
29. Kobak, D. and Berens, P. (2018) The art of using t-SNE for single-cell transcriptomics. *bioRxiv*, Doi: <https://doi.org/10.1101/453449>
30. Baker, M. (2016) Is there a reproducibility crisis? *Nature*. 533, 452–454
31. Berg, J. (2018) Progress on reproducibility. *Science*, 359, 9
32. Stark, P. B. (2018) Before reproducibility must come preproducibility. *Nature*, 557, 613