
Gene Expression

Automatic identification of relevant genes from low-dimensional embeddings of single cell RNAseq data

Philipp Angerer^{1,2}, David S. Fischer^{1,2}, Fabian J. Theis¹, Antonio Scialdone^{1,3,4,*} and Carsten Marr^{1,*}

¹Institute of Computational Biology, Helmholtz Zentrum München – German Research Center for Environmental Health, Neuherberg, 85764, Germany,

²TUM School of Life Sciences Weihenstephan, Technical University of Munich, Freising, 85354, Germany,

³Institute of Epigenetics and Stem Cells, Helmholtz Zentrum München – German Research Center for Environmental Health, Neuherberg, 85764, Germany, and

⁴Institute of Functional Epigenetics, Helmholtz Zentrum München – German Research Center for Environmental Health, München, 81377, Germany.

* To whom correspondence should be addressed.

Associate Editor: XXXXXXXX

Received on XXXXX; revised on XXXXX; accepted on XXXXX

Abstract

Dimensionality reduction is a key step in the analysis of single-cell RNA sequencing data. It produces a low-dimensional embedding for visualization and as a calculation base for downstream analysis. Nonlinear techniques are most suitable to handle the intrinsic complexity of large, heterogeneous single cell data. However, with no linear relation between gene and embedding coordinate, there is no way to extract the identity of genes driving any cell's position in the low-dimensional embedding, making it more difficult to characterize the underlying biological processes.

In this paper, we introduce the concepts of local and global gene relevance to compute an equivalent of principal component analysis loadings for non-linear low-dimensional embeddings. *Global gene relevance* identifies drivers of the overall embedding, while *local gene relevance* identifies those of a defined sub-region. We apply our method to single-cell RNAseq datasets from different experimental protocols and to different low dimensional embedding techniques. This shows our method's versatility to identify key genes for a variety of biological processes.

To ensure reproducibility and ease of use, our method is released as part of destiny 3.0, a popular R package for building diffusion maps from single-cell transcriptomic data. It is readily available through Bioconductor.

1 Introduction

Single cell RNA sequencing (scRNAseq) has massively improved the resolution developmental trajectories Baron *et al.* (2016) and allowed unprecedented insights into the heterogeneity of complex tissues Vento-Tormo *et al.* (2018); Tritschler *et al.* (2017). On the flip side, new challenges have arisen due to the amount of data that needs to be processed Angerer *et al.* (2017), higher levels of technical and biological noise Yuan *et al.*

(2017), and identification and interpretation of known and novel cell types Pliner *et al.* (2019). To exploit the new opportunities and deal with the new challenges, a large number of algorithms and tools have been developed Zappia *et al.* (2018).

Dimension reduction methods create a low dimensional embedding of the high dimensional gene expression space. Those embeddings are widely used for their two main applications: For one, they serve as a visual overview of the data on which gene expression profiles and per-cell or per-cluster statistics can be compared. They can also serve as inputs for further downstream computational analysis. E.g., principal component

1

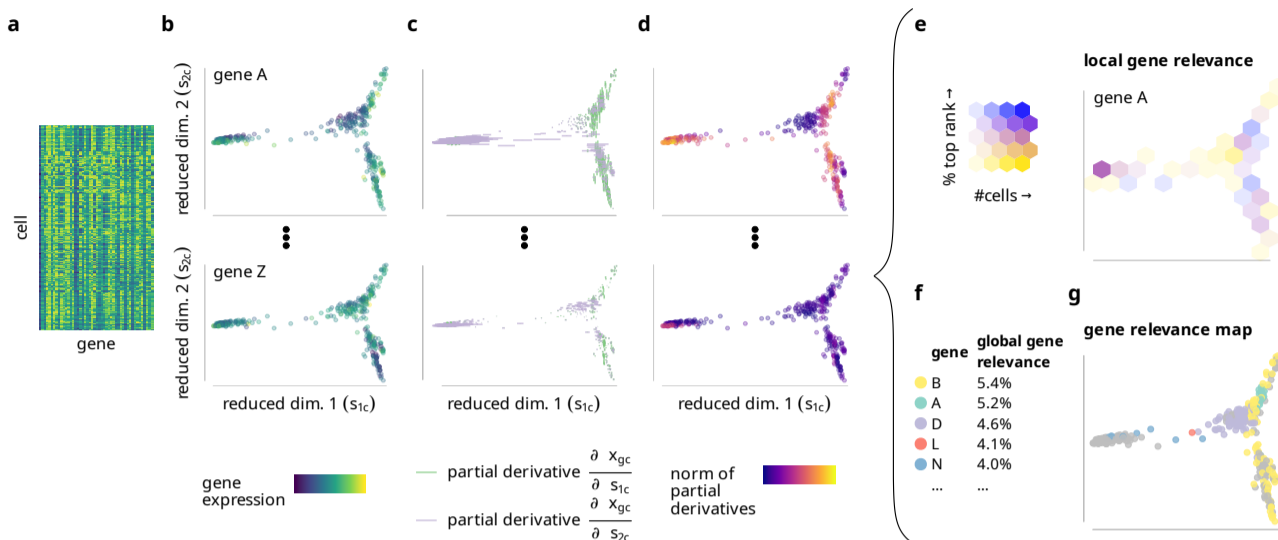


Fig. 1. The gene relevance concept. (a) A gene expression matrix from a single cell RNA sequencing experiment is (b) reduced to a low-dimensional embedding s_{pc} , with each dot representing a cell, and the color representing the expression x_{gc} of gene $g \in \{A, B, \dots, Z\}$ in cell c . (c) Expression changes are calculated from estimates of partial derivatives with respect to the embedding, which results in one value per cell \times gene \times dimension combination. (d) We score the relevance of each gene in each cell according to the partial derivatives' euclidean norm. This score indicates how relevant each gene is within its neighborhood. (e) For local gene relevance scores, we subdivide the embedding into bins and determine the fraction of cells per bin for which a given gene is among the (e.g. 10) most relevant genes (indicated by “%top rank” in the figure legend). We color bins of the embedded cells according to their local gene relevance score, and fade them according to the number of cells they contain (indicated by “#cells” in the legend). (f) For the global gene relevance score, we determine local relevance for all cells instead of a bin. In our illustrative example, gene B has been ranked among the top 10 genes in 5.4% of all cells. (g) A gene relevance map indicates all cells where a given gene has the largest norm of partial derivatives (with or without a smoothing step – see Methods). Such cells mark the areas where that gene has high local relevance.

analysis (PCA) is a popular technique to identify orthogonal linear combinations of genes that explain variance in the data. PCA loadings quantify the contribution of genes to each principal component and help understand the genetic drivers of the underlying molecular processes. However, linear methods are often not able to capture the complexity of high-dimensional datasets Haghverdi *et al.* (2015), which is why nonlinear dimension reduction methods have become the standard for scRNAseq data analysis (see e.g. t-SNE Husnain *et al.* (2019), diffusion maps Husnain *et al.* (2019); Coifman *et al.* (2005); Haghverdi *et al.* (2015), UMAP McInnes *et al.* (2018); Becht *et al.* (2018), and graph-based methods Islam *et al.* (2011)). However, no intrinsic measure of individual genes' contribution to each embedding dimension exists for non-linear embeddings. Without such a measure, the identification of genes that drive the variability in the data requires tedious manual inspection and prior knowledge about possible target genes.

Here, we introduce *gene relevance*, a measure for a gene's contribution to variance in low dimensional embeddings, and present a method to infer a local as well as a global gene relevance score from any kind of low-dimensional embedding. To demonstrate the utility of the method, we apply gene relevance to several datasets prepared with different droplet-based and plate-based protocols (see Suppl. Table ??). Gene relevance is available as part of the R package *destiny* Angerer *et al.* (2016).

2 Methods

We define gene relevance as a measure of how much a gene contributes to the cell-to-cell variability in a low dimensional embedding of a scRNAseq dataset (see Fig. 1). It can be interpreted as a generalization of PCA loadings to non-linear dimensionality reduction techniques. Note that PCA loadings are constant with respect to the PC space while feature importance in a non-linear embedding is naturally a non-constant function of the embedding coordinates. A ranking of genes based on their relevance is built for every cell of the embedding. These rankings, can be combined to

obtain a measurement of the *local* or *global* relevance of each gene (see Fig. 1e-f), which highlight genes relevant in defined sub-regions of the embedding and all cells, respectively. To explore and visualize the results further, the method also provides a *gene relevance map*, where the locally most relevant genes are displayed along with their corresponding neighborhoods in the embedding (see Fig. 1g). Below, we describe in details every step in the estimation of global and local gene relevance.

Neighborhoods. If a k nearest neighbor (kNN) search has been performed as part of the embedding, it can be efficiently used for estimating the gene relevance. To perform the kNN search, *destiny* offers the choice between euclidean distance, cosine distance, and spearman rank correlation distance. The latter was used in all analyses performed for this paper.

Gene expression changes. The differential d_{gc} of gene g in cell c describes the change in gene expression x_{gc} along a change in embedding coordinates s_{pc} , where $p \in \{1, \dots, P\}$ is the embedding dimension and d_{gc} corresponds to the partial derivatives of the gene expression with respect to each embedding coordinate:

$$d_{gc} = \left(\frac{\partial x_{gc}}{\partial s_{1c}}, \dots, \frac{\partial x_{gc}}{\partial s_{pc}} \right) \quad (1)$$

We estimated d_{gc} from the cells' neighborhood $NN_k(c)$ in gene expression space using finite differences. To address the high dropout rate present in scRNAseq data, we do not define d_{gc} for $x_{gc} = 0$.

$$\widehat{(d_{gc})_p} = \begin{cases} NA, & \text{if } x_{gc} = 0 \\ \text{median}_{n \in NN_k(c) \wedge n \neq c} \frac{x_{gc} - x_{gn}}{s_{pn} - s_{pc}}, & \text{otherwise} \end{cases} \quad (2)$$

Local gene relevance. The basis of local and global gene relevance is the score of gene $g \in \{1, \dots, G\}$ in cell $c \in \{1, \dots, C\}$, defined as the euclidean norm $\|d_{gc}\|_2$ of the differential d_{gc} :

$$\|d_{gc}\|_2 = \sqrt{\sum_{p=1}^P (d_{gc})_p^2} \quad (3)$$

In each cell c , genes can be ranked according to their score $\|d_{gc}\|_2$, from most to least relevant. Given the ranks $\text{rg}_{\|d_{gc}\|_2}$ of gene g and a rank cutoff rg_{\max} , we define the local gene relevance $\text{LR}_{\text{rg}_{\max}}(g, \Psi)$ of a gene for a set of cells $\Psi \subseteq \{1, \dots, C\}$ as:

$$\text{LR}_{\text{rg}_{\max}}(g, \Psi) = \frac{\sum_{c \in \Psi} [\text{rg}_{\|d_{gc}\|_2} < \text{rg}_{\max}]}{|\Psi|} \quad (4)$$

with the iverson bracket notation

$$[P] = \begin{cases} 1, & \text{if } P \text{ is true} \\ 0, & \text{otherwise} \end{cases}, \text{ for any predicate } P \quad (5)$$

In our analyses, we used the default $\text{rg}_{\max} = 10$. The method is robust against the cutoff used, with the most relevant genes stable even for higher cutoffs > 100 (see Suppl. Fig. ??).

Global gene relevance. The global relevance $\text{GR}_{\text{rg}_{\max}}(g)$ can simply be defined as the local gene relevance for the set of all cells $\{1, \dots, C\}$.

$$\text{GR}_{\text{rg}_{\max}}(g) = \text{LR}_{\text{rg}_{\max}}(g, \{1, \dots, C\}) \quad (6)$$

Local gene relevance plots. The local gene relevance of genes can be visualized by evenly dividing the embedding space into bins $b \in \{1, \dots, B\}$ and calculating $\text{LR}_{\text{rg}_{\max}}(g, \Psi_b)$ for the set of cells falling into each bin Ψ_b . This visualization is shown in Fig. 1e and Fig. 2d.

Gene relevance maps. For a set of genes of interest $\Omega \subseteq \{1, \dots, G\}$ (which can be chosen, e.g., among those with highest global relevance) and each cell c , we define the locally most relevant gene l_c^m after a number of smoothing steps m :

$$l_c^m = \arg \max_{g \in \Omega} \begin{cases} \|d_{gc}\|_2, & \text{if } m = 0 \\ \frac{1}{k} \sum_{n \in \text{NN}_k(c)} [l_n^{m-1} = g], & \text{otherwise.} \end{cases} \quad (7)$$

During a smoothing iteration, we replace the local gene relevance score of cell c and gene g with the fraction of neighbours that have g as the most relevant gene.

The amount of smoothing is controlled by the smoothing parameter m . Decreasing m will result in more locality and more genes with a high relevance in a small region will appear on the map, while genes with a medium relevance in more or larger regions will vanish. When determining the locally most relevant gene using the globally most relevant genes as Ω , the aforementioned parameter rg_{\max} can be used as a sensitivity parameter, with higher values resulting in more genes being selected as relevant. Finding the globally most relevant genes is robust to both parameters (see Suppl. Fig. ??), while locally relevant genes are affected to a larger extent. Importantly, due to the short computation time of gene relevance maps, one can explore several combinations of parameters. For cell sizes in the range of few hundreds, the running time is less than one to a few seconds with 40000 cell \times gene combinations per second processed. For cell sizes in the thousands 20000 cell \times gene combinations per second are processed, resulting in run times of a few minutes. (Run times have been measured on a single 3.6GHz CPU core.)

Single cell RNA sequencing data. We demonstrate our method on four datasets (see **Results** and Suppl. Table ?? for details).

In the mouse gastrulation data from Scialdone *et al.* (2016), we used count data from 271 cells mostly of the neural plate (embryonic day 7.5)

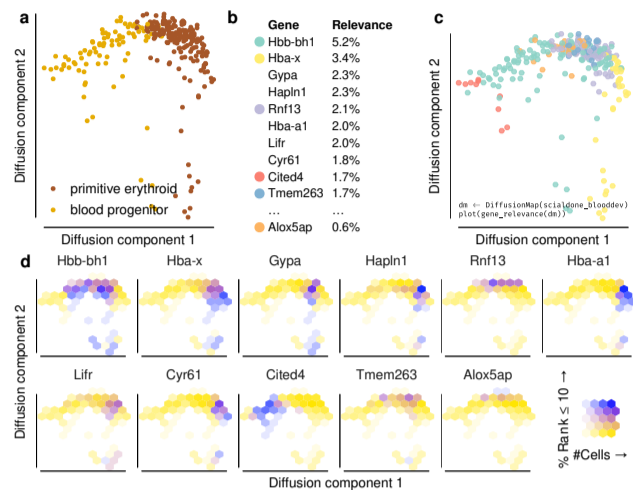


Fig. 2. Gene relevance automatically detects drivers of embryonic blood development. (a) Diffusion map of 271 single hematopoietic progenitor cells from mostly day 7.5 and 7.75 mouse embryos, profiled in Scialdone *et al.* (2016). (b) Global gene relevance identifies Hbb-bh1 and Hba-x as genes that change most dramatically during hematopoietic development. (c) A gene relevance map identifies the contribution of relevant genes in specific regions of the process and the corresponding code to create it. The genes corresponding to each color are shown in panel b. (d) A local gene relevance plot details the areas where the contribution of genes is highest. Alox5ap shows a high local relevance in the top region of the diffusion map and has been implicated with early blood development Ibarra-Soria *et al.* (2018).

and head fold (embryonic day 7.75) development stages of mouse embryos. There, the libraries were constructed using the Smart-seq2 protocol, read counts were obtained via HTseq-count. The 271 cells we used correspond to the clusters annotated as “blood progenitor” and “primitive erythroid” in the original publication. We selected highly variable genes using the method of Brennecke *et al.* (2013) because of its stable performance Yip *et al.* (2018), and embedded the log-transformed data using the diffusion map implementation destiny Angerer *et al.* (2016).

For the two human cell datasets, we applied the same analysis steps, starting from the highly variable gene selection (see Suppl. Fig. ??). The human endocrine cell data from Veres *et al.* (2019) was sequenced using the inDrops platform, while the human brain organoid data from Gray Camp *et al.* (2015) used the SMARTer Ultra Low RNA Kit in combination with an Illumina sequencer. The data from Kolodziejczyk *et al.* (2015) used Nextera kits together with an Illumina sequencer. For further details about the preprocessing of these data, please refer to the individual publications.

3 Results

We demonstrate our method on a scRNAseq dataset of blood progenitors and blood cells from mouse embryos Scialdone *et al.* (2016) (see Fig. 2a). In the original publication, this data was used to reconstruct a trajectory representing primitive erythropoiesis, along which blood marker expression increases and other markers (such as endothelial cells) decrease. There, an ad-hoc method was devised to find important genes in the 2D diffusion map embedding of the data. Here, we show how our method can be used “out of the box” to rank genes based on their local and global relevance.

First, we ranked all highly variable genes according to their global gene relevance (see Fig. 2b). As expected, the high-ranking genes are mostly associated with blood development, including the hemoglobin genes Hba-a1, Hba-x, Hbb-bh1, and the erythrocyte membrane genes Gypa and Cited4

Yahata *et al.* (2002). The genes *Cyr61* and *Hapln1* are involved in extracellular matrix and important for development of the cardiovascular system Latinkić *et al.* (2001). The top of the list has a good overlap with the ad-hoc method in Scialdone *et al.* (2016): 4 genes are shared between the top ten of both lists, and we find a Rank-Biased Overlap of $RBOp = 0.48$, where we used $p = 0.9$, which assigns ~86% of the weight to the first 10 genes Webber *et al.* (2010).

Second, we created a gene relevance map (Fig. 2c). Five out of the six locally most relevant genes (see Fig. 2c) are among the ten most globally relevant ones. Interestingly, *Alox5ap* is included only in the gene relevance map, because its contribution is confined to a small region of the diffusion space (bottom right panel in Fig. 2d) and hard to detect at the level of gene expression (see Suppl. Fig. ??). This gene was not discovered by the ad-hoc method of Scialdone *et al.* (2016), but it has been recently found to be important in early blood development Ibarra-Soria *et al.* (2018). Locally and globally relevant genes can also be inferred in other embeddings such as t-SNE Maaten and Hinton (2008) and UMAP Becht *et al.* (2018), with a stable recovery of the most relevant genes between comparable embeddings (See Suppl. Fig. ??).

Applied to other scRNAseq data sets, we showcase versatility and ease of application of our method. In droplet-sequenced data of human endocrine cells Veres *et al.* (2019), gene relevance maps detect genes driving the separation of sub-regions of the embedding (see Suppl. Fig. ??a), in accordance to the markers identified in the original paper. In human brain organoid cells Gray Camp *et al.* (2015), we detect relevant genes different from the markers specified in the paper because of a low density region between mesenchymal cells and neurons/neural progenitors (see Suppl. Fig. ??b). The genes therefore seem to be selected for driving the difference between progenitors and neurons: *TXNRD1* plays a vital role for neuron progenitor cells Soerensen *et al.* (2008), the selenoprotein *SELT* protects neurons against oxidative stress in mouse models Boukhzar *et al.* (2016), and *CRABP1* modulates the neuronal cell cycle in mice Lin *et al.* (2017).

Finally, we applied gene relevance to mouse embryonic stem cells grown in three different pluripotency retaining media Kolodziejczyk *et al.* (2015). As expected for cells in a relatively homogenous pluripotent steady state, the relevant genes for diffusion map embedding of all three media were enriched for housekeeping, metabolic and proliferation pathways (see Suppl. Fig. ??).

4 Discussion

We presented a method that is able to reliably detect relevant genes from low dimensional embeddings of scRNAseq data. More specifically, our method computes both a local and a global gene relevance score: local gene relevance identifies the main drivers of the cell-to-cell variability in defined sub-regions of the embedding, while global gene relevance identifies those of the whole embedding. In addition to a gene ranking based on global relevance, the method also provides graphic tools to visualize the local gene relevance (see Fig. 1e) and the changes in gene expression levels within the embedding (see Fig. 1c and Suppl. Fig. ??). It can be used for any single cell data set and any dimensionality reduction technique.

We applied our method to three datasets, including one from mouse embryonic blood progenitors, where we show that it performs comparably to a technique custom-made for the dataset. Interestingly, our method identifies *Alox5ap* (Fig. 2), a gene that was recently shown to be important for blood development in a later publication Ibarra-Soria *et al.* (2018). In two other examples, we used human cells, endocrine Veres *et al.* (2019) and from brain organoids Gray Camp *et al.* (2015), showing that the method works robustly in varied conditions.

Compared to the classical method of looking at PCA loadings, gene relevance provides a generalization not tied to this dimension reduction method, as it is also suited for non-linear dimension reduction methods. It can therefore be used to explore the differences in the embeddings obtained by different dimension reduction methods (see Suppl. Fig. ??). Applied to a PCA embedding, gene relevance recovers a similar list of genes to those with the highest PC loadings (see Suppl. Fig. ??). Other methods to identify important genes specifically from scRNAseq data exist, but most of them aim to find marker genes that can best distinguish different cell types Delaney *et al.* (2019). Conversely, the method we presented is unsupervised and does not rely on cell type annotation.

Recently, two computational methods have been developed to identify variable genes in spatial RNAseq datasets, trendsceek and SpatialDE Edsgård *et al.* (2018); Svensson *et al.* (2018). While these methods were designed to find patterns in spatial transcriptomic datasets, they can also be used to identify relevant genes in low-dimensional embeddings of scRNAseq datasets (see Suppl. Fig. ?? in Edsgård *et al.* (2018)). We compared our approach to trendsceek and found similar genes (see Suppl. Fig. ??) in a considerably shorter running time: our method took 6.5 seconds, while trendsceek needed 1080 seconds. (Run times measured on a single 2.0GHz CPU core.) SpatialDE returned a perfect score for too many genes, making gene ranking impossible. This is probably related to both methods being optimized towards identifying spatial patterns. Moreover, neither method allows estimation of local gene relevance.

To summarize, our gene relevance method is a fast and versatile exploratory tool that can help identify the biological processes and reveal the presence and driving genes of potentially rare cell sub-populations. It is available online, easily applicable, and faster than model fitting approaches. While we focussed our discussion on scRNAseq datasets, our method can be applied to virtually any kind of dataset where low-dimensional embeddings are obtained, including, for instance, single-cell epigenomic Shema *et al.* (2018) and mass cytometry data Spitzer and Nolan (2016).

5 Availability of data and materials

Gene relevance has been developed as part of the Bioconductor package *destiny*: bioconductor.org/packages/destiny. API documentation for analysis and plotting and are available at <https://theislab.github.io/destiny/>

The datasets analysed within this publication are available from their original publications as described in Suppl. Table ??.

6 Competing interests

The authors declare that they have no competing interests.

7 Authors' contributions

PA designed the analysis, implemented the method. AS interpreted the results and wrote the paper with PA and CM. DF contributed to the mathematical description of the gene relevance concept. FT contributed the initial idea of gene relevance. CM supervised the study.

Acknowledgements and Funding

D.S.F. acknowledges financial support by a German research foundation (DFG) fellowship through the Graduate School of Quantitative Biosciences Munich (QBM) (GSC 1006) and by the Joachim Herz Stiftung. FT and CM acknowledge support from the DFG funded collaborative research

center SFB 1243. We thank Moritz Thomas and Ali Boushehri for helpful feedback on the manuscript.

References

- Angerer, P., Haghverdi, L., Büttner, M., Theis, F. J., Marr, C., and Buettner, F. (2016). destiny: diffusion maps for large-scale single-cell data in R. *Bioinformatics*, **32**(8), 1241–1243.
- Angerer, P., Simon, L., Tritschler, S., Alexander Wolf, F., Fischer, D., and Theis, F. J. (2017). Single cells make big data: New challenges and opportunities in transcriptomics.
- Baron, M., Veres, A., Wolock, S. L., Faust, A. L., Gaujoux, R., Vetere, A., Ryu, J. H., Wagner, B. K., Shen-Orr, S. S., Klein, A. M., Melton, D. A., and Yanai, I. (2016). A Single-Cell transcriptomic map of the human and mouse pancreas reveals inter- and intra-cell population structure. *Cell Syst*, **3**(4), 346–360.e4.
- Becht, E., McInnes, L., Healy, J., Dutertre, C.-A., Kwok, I. W. H., Ng, L. G., Ginhoux, F., and Newell, E. W. (2018). Dimensionality reduction for visualizing single-cell data using UMAP. *Nat. Biotechnol.*
- Boukhar, L., Hamieh, A., Cartier, D., Tanguy, Y., Alsharif, I., Castex, M., Arabo, A., El Hajji, S., Bonnet, J.-J., Errami, M., Falluel-Morel, A., Chagraoui, A., Lihmann, I., and Anouar, Y. (2016). Selenoprotein T exerts an essential oxidoreductase activity that protects dopaminergic neurons in mouse models of parkinson's disease.
- Brennecke, P., Anders, S., Kim, J. K., Kolodziejczyk, A. A., Zhang, X., Proserpio, V., Baying, B., Benes, V., Teichmann, S. A., Marioni, J. C., and Heisler, M. G. (2013). Accounting for technical noise in single-cell RNA-seq experiments. *Nat. Methods*, **10**(11), 1093–1095.
- Coifman, R. R., Lafon, S., Lee, A. B., Maggioni, M., Nadler, B., Warner, F., and Zucker, S. W. (2005). Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps.
- Delaney, C., Schnell, A., Cammarata, L. V., Yao-Smith, A., Regev, A., Kuchroo, V. K., and Singer, M. (2019). Combinatorial prediction of gene-marker panels from single-cell transcriptomic data.
- Edsgård, D., Johnsson, P., and Sandberg, R. (2018). Identification of spatial expression trends in single-cell gene expression data. *Nat. Methods*, **15**(5), 339–342.
- Gray Camp, J., Badsha, F., Florio, M., Kanton, S., Gerber, T., Wilsch-Bräuninger, M., Lewitus, E., Sykes, A., Hevers, W., Lancaster, M., Knoblich, J. A., Lachmann, R., Pääbo, S., Huttner, W. B., and Treutlein, B. (2015). Human cerebral organoids recapitulate gene expression programs of fetal neocortex development. *Proc. Natl. Acad. Sci. U. S. A.*, **112**(51), 15672–15677.
- Haghverdi, L., Buettner, F., and Theis, F. J. (2015). Diffusion maps for high-dimensional single-cell analysis of differentiation data. *Bioinformatics*, **31**(18), 2989–2998.
- Husnain, M., Missen, M., Mumtaz, S., Luqman, M., Coustaty, M., and Ogier, J.-M. (2019). Visualization of High-Dimensional data by pairwise fusion matrices using t-SNE.
- Ibarra-Soria, X., Jawaid, W., Pijuan-Sala, B., Ladopoulos, V., Scialdone, A., Jörg, D. J., Tyser, R. C. V., Calero-Nieto, F. J., Mulas, C., Nichols, J., Vallier, L., Srinivas, S., Simons, B. D., Göttgens, B., and Marioni, J. C. (2018). Defining murine organogenesis at single-cell resolution reveals a role for the leukotriene pathway in regulating blood progenitor formation. *Nat. Cell Biol.*, **20**(2), 127–134.
- Islam, S., Kjällquist, U., Moliner, A., Zajac, P., Fan, J.-B., Lönnerberg, P., and Linnarsson, S. (2011). Characterization of the single-cell transcriptional landscape by highly multiplex RNA-seq. *Genome Res.*, **21**(7), 1160–1167.
- Kolodziejczyk, A. A., Kim, J. K., Tsang, J. C. H., Ilicic, T., Henriksson, J., Natarajan, K. N., Tuck, A. C., Gao, X., Bühler, M., Liu, P., Marioni, J. C., and Teichmann, S. A. (2015). Single cell RNA-Sequencing of pluripotent states unlocks modular transcriptional variation. *Cell Stem Cell*, **17**(4), 471–485.
- Latinkić, B. V., Mo, F.-E., Greenspan, J. A., Copeland, N. G., Gilbert, D. J., Jenkins, N. A., Ross, S. R., and Lau, L. F. (2001). Promoter function of the angiogenic inducer Cyr61 Gene in transgenic mice: Tissue specificity, inducibility during wound healing, and role of the serum response element. *Endocrinology*, **142**(6), 2549–2557.
- Lin, Y.-L., Persaud, S. D., Nhieu, J., and Wei, L.-N. (2017). Cellular retinoic Acid-Binding protein 1 modulates stem cell proliferation to affect learning and memory in male mice. *Endocrinology*, **158**(9), 3004–3014.
- Maaten, L. v. d. and Hinton, G. (2008). Visualizing data using t-SNE. *J. Mach. Learn. Res.*, **9**(Nov), 2579–2605.
- McInnes, L., Healy, J., Saul, N., and Großberger, L. (2018). UMAP: Uniform manifold approximation and projection.
- Pliner, H. A., Shendure, J., and Trapnell, C. (2019). Supervised classification enables rapid annotation of cell atlases.
- Scialdone, A., Tanaka, Y., Jawaid, W., Moignard, V., Wilson, N. K., Macaulay, I. C., Marioni, J. C., and Göttgens, B. (2016). Resolving early mesoderm diversification through single-cell expression profiling. *Nature*, **535**(7611), 289–293.
- Shema, E., Bernstein, B. E., and Buenrostro, J. D. (2018). Single-cell and single-molecule epigenomics to uncover genome regulation at unprecedented resolution. *Nat. Genet.*, **51**(1), 19.
- Soerensen, J., Jakupoglu, C., Beck, H., Förster, H., Schmidt, J., Schmahl, W., Schweizer, U., Conrad, M., and Brielmeier, M. (2008). The role of thioredoxin reductases in brain development. *PLoS One*, **3**(3).
- Spitzer, M. H. and Nolan, G. P. (2016). Mass cytometry: Single cells, many features. *Cell*, **165**(4), 780–791.
- Svensson, V., Teichmann, S. A., and Stegle, O. (2018). SpatialDE: identification of spatially variable genes. *Nat. Methods*, **15**(5), 343–346.
- Tritschler, S., Theis, F. J., Lickert, H., and Böttcher, A. (2017). Systematic single-cell analysis provides new insights into heterogeneity and plasticity of the pancreas. *Mol Metab*, **6**(9), 974–990.
- Vento-Tormo, R., Efremova, M., Botting, R. A., Turco, M. Y., Vento-Tormo, M., Meyer, K. B., Park, J.-E., Stephenson, E., Polański, K., Goncalves, A., Gardner, L., Holmqvist, S., Henriksson, J., Zou, A., Sharkey, A. M., Millar, B., Innes, B., Wood, L., Wilbrey-Clark, A., Payne, R. P., Ivarsson, M. A., Lisgo, S., Filby, A., Rowitch, D. H., Bulmer, J. N., Wright, G. J., Stubbington, M. J. T., Haniffa, M., Moffett, A., and Teichmann, S. A. (2018). Single-cell reconstruction of the early maternal-fetal interface in humans. *Nature*, **563**(7731), 347–353.
- Veres, A., Faust, A. L., Bushnell, H. L., Engquist, E. N., Kenty, J. H.-R., Harb, G., Poh, Y.-C., Sintov, E., Gürtler, M., Pagliuca, F. W., Peterson, Q. P., and Melton, D. A. (2019). Charting cellular identity during human in vitro β -cell differentiation. *Nature*, **569**(7756), 368–373.
- Webber, W., Moffat, A., and Zobel, J. (2010). A similarity measure for indefinite rankings.
- Yahata, T., Takedatsu, H., Dunwoodie, S. L., Bragança, J., Swingler, T., Withington, S. L., Hur, J., Coser, K. R., Isselbacher, K. J., Bhattacharya, S., and Shioda, T. (2002). Cloning of mouse *cited4*, a member of the CITED family p300/CBP-binding transcriptional coactivators: induced expression in mammary epithelial cells. *Genomics*, **80**(6), 601–613.
- Yip, S. H., Sham, P. C., and Wang, J. (2018). Evaluation of tools for highly variable gene discovery from single-cell RNA-seq data. *Brief. Bioinform.*
- Yuan, G.-C., Cai, L., Elowitz, M., Enver, T., Fan, G., Guo, G., Irizarry, R., Kharchenko, P., Kim, J., Orkin, S., Quackenbush, J., Saadatpour, A., Schroeder, T., Shivdasani, R., and Tirosh, I. (2017). Challenges and emerging directions in single-cell analysis. *Genome Biol.*, **18**(1), 84.
- Zappia, L., Phipson, B., and Oshlack, A. (2018). Exploring the single-cell RNA-seq analysis landscape with the scRNA-tools database. *PLoS*

