OXFORD

Gene expression

ACTINN: automated identification of cell types in single cell RNA sequencing

Feiyang Ma [iD] [1] **and Matteo Pellegrini**[1,2,]*

[1]Molecular Biology Institute and [2]Institute of Genomics and Proteomics, University of California, Los Angeles, CA 90095, USA

*To whom correspondence should be addressed.

Associate Editor: Mark Robinson

Abstract

Motivation: Cell type identification is one of the major goals in single cell RNA sequencing (scRNA-seq). Current methods for assigning cell types typically involve the use of unsupervised clustering, the identification of signature genes in each cluster, followed by a manual lookup of these genes in the literature and databases to assign cell types. However, there are several limitations associated with these approaches, such as unwanted sources of variation that influence clustering and a lack of canonical markers for certain cell types. Here, we present ACTINN (Automated Cell Type Identification using Neural Networks), which employs a neural network with three hidden layers, trains on datasets with predefined cell types and predicts cell types for other datasets based on the trained parameters.

Results: We trained the neural network on a mouse cell type atlas (Tabula Muris Atlas) and a human immune cell dataset, and used it to predict cell types for mouse leukocytes, human PBMCs and human T cell sub types. The results showed that our neural network is fast and accurate, and should therefore be a useful tool to complement existing scRNA-seq pipelines.

Availability and implementation: The codes and datasets are available at https://figshare.com/articles/ACTINN/8967116. Tutorial is available at https://github.com/mafeiyang/ACTINN. All codes are implemented in python.

Contact: matteop@mcdb.ucla.edu

Supplementary information: Supplementary data are available at *Bioinformatics* online.

1 Introduction

Single cell RNA sequencing (scRNA-seq) enables the profiling of the transcriptomes of individual cells, thus characterizing the heterogeneity of samples in manner that was not possible using traditional bulk RNA-Seq (Hwang *et al.*, 2018). However, scRNA-seq experiments typically yield high volumes of data, especially when the number of cells is large (often many thousands). Thus, fast and efficient computational methods are essential for scRNA-seq analyses.

One common goal of scRNA-seq analyses is to identify the cell type of each individual cell that has been profiled. To accomplish this, typically cells are first grouped into different clusters in an unsupervised way, and the number of clusters allows us to approximately determine how many distinct cell types are present in the sample. Each cluster should contain cells with similar expression profiles, and so the aggregated profile of a cluster increases the signal to noise of the expression estimates. To attempt to interpret the identity of each cluster, marker genes are found as those that are uniquely highly expressed in a cluster, compared to all the other clusters. These canonical markers are then used to assign the cell types for the clusters, by cross referencing the markers with lists of previously characterized cell type specific markers. While this process is able to identify cell types, there are some limitations: (i) Since the clustering method is unsupervised, all sources of variation influence the formation clusters, including effects that are not directly related to cell types such as differential expression induced by cell cycles. (ii) It is often difficult to find an optimal match between the

marker genes associated with each cluster and the canonical markers for specific cell types. Moreover, depending on the clustering parameters used, one cluster might contain multiple cell types, or one cell type could be split into multiple clusters. (iii) Using canonical markers to assign cell types requires background knowledge of cell type specific markers, and sometimes these are not well characterized or difficult to find in the literature. Moreover, some canonical markers may be expressed by more than one cell type, and some cell types may have no known markers. (iv) The same types of cells processed by two distinct scRNA-seq techniques tend to cluster separately due to technical batch effects, which complicates cell type identification in composite datasets. (v) Cell subtypes are often very similar to each other, which limits efforts to separate them accurately into different clusters. To overcome many of the limitations of existing approaches, new methods need to be developed.

Neural networks provide a popular framework for machine learning algorithms which can be used to interpret complex datasets. As a result, neural networks have been widely used in many fields, including for the analysis of scRNA-seq data (Cho *et al.*, 2018; Lin *et al.*, 2017; Lopez *et al.*, 2018; Shaham *et al.*, 2017). Since the output data from scRNA-seq is feature-enriched and well-structured, it is well suited as an input for neural networks. Here, we present ACTINN (Automated Cell Type Identification using Neural Networks) for scRNA-seq cell type identification. To overcome many of the limitations of traditional cell type identification approaches described above, we used a neural network with three hidden layers, trained it on scRNA-seq datasets with predefined cell types, and predicted cell types in other datasets based on the trained parameters. We tested our neural network with several published datasets and show that it is fast, efficient and accurate.

2 Materials and methods

2.1 Data normalization
We used several publicly available datasets in our analyses. The mouse cell atlas datasets were collected from https://tabula-muris.ds.czbiohub.org. The CD45 sorted leukocyte datasets were published in Winkels *et al*. The T cell subtypes and PBMC datasets were collected from https://support.10xgenomics.com/single-cell-gene-expression/datasets. To filter and normalize the data, we first identified genes that were detected in both training set and test set. The training set and the test set were then merged into one matrix based on the common genes. Next, each cell's expression value was normalized to its total expression value and multiplied by a scale factor of 10 000. The counts were increased by 1, and the log2 value was calculated. To filter out outlier genes, the genes with the highest 1% and lowest 1% expression were removed. The gene with the highest 1% and the lowest 1% standard deviation were also removed. Finally, the matrix was split into the training set and the test set.

2.2 Neural network configuration
We used a neural network that contains an input layer, three hidden layers and an output layer. The input layer had a number of nodes equal to the number of genes in the training set. The 3 hidden layers had 100, 50 and 25 nodes, respectively. The output layer had a number of nodes equal to the number of cell types in the training set. Forward propagation was implemented as:

$$x^{[i]} = g\big(W^{[i]}x^{[i-1]} + b^{[i-1]}\big)$$

where $x^{[i]}$ represents the output of the ith layer ($x^{[0]}$ represents the input layer), $b^{[i]}$ represents the intercept of the ith layer, $W^{[i]}$

represents the weight matrix of the ith layer and g represents the activation function used in the neural network. Specifically, for the activation function, the rectified linear unit (ReLU) function was used for the input and hidden layers, which is defined as:

$$ReLU(x) = max(0, x)$$

For the output layer, the softmax function was used, which is defined as:

$$softmax(x_{[j]}) = \frac{\exp(x_{[j]})}{\sum_{j=1}^{k} \exp(x_{[j]})}$$

where $x_{[j]}$ represents the jth element of the input vector for the output layer, which has k elements, representing a total of k cell types in the training set. For the loss function, we used the cross-entropy function, which is defined as:

$$H(y', y) = \sum_{j=1}^{k} \big(y_{[j]}\log(y'_{[j]}) + (1 - y_{[j]})\log(1 - y'_{[j]})\big)$$

where vector y represents the true label for the cell, $y_{[j]}$ is defined to be 1 if the cell is the jth cell type, and the other elements in y are defined to be 0. y' represents the output of the output layer, and $y'_{[j]}$ represents the posterior probability that the cell is the jth cell type. L2 regularization was added to the loss function.

2.3 Parameters used in the neural network
The neural network model was implemented using TensorFlow (https://www.tensorflow.org), and the code was written in python. The parameters were initialized with Xavier initializer (Xavier *et al.*, 2010). The starting learning rate was set to 0.0001 with staircase exponential decay, the decay rate was set to 0.95 and the decay step was set to 1000. This means that after every 1000 global steps, the learning rate would be the original learning rate multiplied by 0.95. 50 epochs were used to train the neural network with a mini batch size of 128, which is the number of samples used in training at every global step. The L2 regularization rate was set to 0.005.

2.4 Unsupervised single cell analysis
To identify different cell types and find signature genes for each cell type, Seurat (Butler *et al.*, 2018) was used to analyze the digital expression matrix generated by scRNA-seq. Specifically, in Seurat, cells with less than 1000 unique molecular identifiers (UMIs) and genes detected in less than 10 cells were first filtered out. Second, highly variable genes were detected and used for further analysis. Third, the data was scaled for sequencing depth of each cell. Fourth, principle component analysis (PCA) and t-distributed stochastic neighbor embedding (tSNE) were used to reduce the dimension and plot the data on a two-dimensional graph. Lastly, a graph-based clustering approach was used to cluster the cells, then signature genes were found and used to define cell type for each cluster.

3 Results

3.1 Overview of the neural network
We used a neural network with 3 hidden layers, each containing 100, 50 and 25 nodes, respectively (Supplementary Fig. S1). For the activation functions, we used the softmax function for the ouput layer and the rectified linear unit (ReLU) function for the other layers. We used the cross-entropy function as the loss function. The neural network model was implemented using TensorFlow, and the code was written in python. We trained the neural network on 6

Intel(R) Xeon(R) CPU E5-2660 v3 nodes, and the training process took 0.5 min to complete with 1000 cells, 11 min with 32 000 cells and 21 min with 56 000 cells. The maximum memory used in training with 56 000 cells was 18 GB. The code and datasets used in this study are available at https://github.com/mafeiyang/ACTINN.

3.2 ACTINN model for murine cell types

We used 2 datasets from the Tabula Muris Consortium (2018) to train and test our neural network. The datasets contain 100 605 cells from 20 mouse organs, and were sequenced by two distinct techniques, 10× Genomics (10×) and Smart-seq2 (SS2). To ensure we are using cells with high quality, we filtered out cells with less than 300 detected genes, clustered the cells and identified marker genes for each cluster using Seurat (Butler *et al.*, 2018). The details of the Seurat analysis can be found in the methods section. We manually assigned cell types for each cluster based on canonical markers (Fig. 1A). To make the analysis easier to interpret, we merged similar cell types into one single cell type. For example, we merged B cells, naïve B cells, immature B cells, pro-B cells and late pro-B cells from the TMA datasets into B cells. We focused on 12 cell types and selected cells that have the same labels between our analyses and the Tabula Muris Consortium's. This process resulted in 56 112 cells (Fig. 1B). Cells processed by 10× have a median of 4787 unique molecular identifiers (UMIs) and 1558 genes detected, and cells processed by SS2 have a median of 623 799 UMIs and 2448 genes detected.
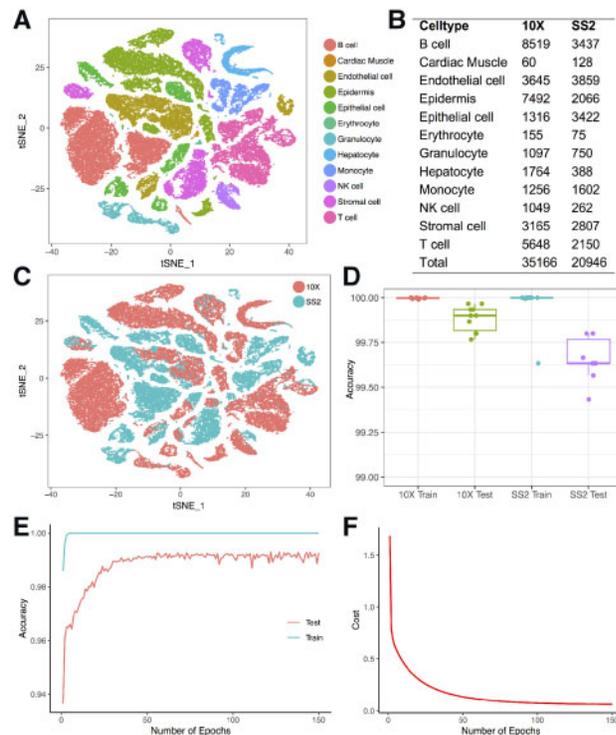
Fig. 1. Training and testing of the neural network on the Tabula Muris Atlas. (**A**) Cell types obtained from the TMA. (**B**) Number of cells obtained for each cell type from each technique. (**C**) The same cell type tends to cluster separately by techniques. (**D**) Training and testing accuracy of the neural network when trained and tested using cells processed by the same technique. (**E**) Training and testing accuracy after each epoch when trained with 5000 10× cells and tested with 5000 SS2 cells. (**F**) Cost after each epoch when trained with 5000 10× cells and tested with 5000 SS2 cells

To test the robustness of our neural network's performance, we first trained and tested it on cells processed by each scRNA-seq platform separately. To this end, we randomly sampled 3000 cells for testing, and used the remainder of cells for training. We repeated this process 10 times, and the average training accuracies for the 10× dataset and the SS2 dataset were 99.997 and 99.963%, respectively, and the average testing accuracies were 99.883 and 99.660%, respectively (Fig. 1D). These results show that our neural network can achieve very high accuracy when training and testing on datasets generated by the same technique.

3.3 ACTINN overcomes batch effects introduced by different techniques

Different scRNA-seq techniques can introduce significant batch effects (Haghverdi *et al.*, 2018) with the same cell types clustering separately due to technical artifacts (Fig. 1C). To test our neural network's performance accounting for the batch effects introduced by different techniques, we trained it on cells processed by one platform and tested it on cells processed by the other. We first trained the neural network on all the 10× cells and tested in on all the SS2 cells. The training accuracy was 99.997% and the testing accuracy was 98.625%. Among the 288 incorrectly predicted cells, 118 monocytes were predicted as B cells, 64 monocytes were predicted as epithelial cells, 47 natural killer (NK) cells were predicted T cells (Supplementary Table S1). We then trained the neural network on the SS2 dataset and tested it on the 10× dataset. The training accuracy was 100% and the testing accuracy was 99.195%. Among the 283 incorrectly predicted cells, 150 endothelial cells were predicted as epidermis, 46 T cells were predicted as NK cells, and there were several other mispredictions (Supplementary Table S2).

3.4 Early stopping prevents overfitting of the training set

To prevent overfitting the parameters on the training set, we randomly sampled 5000 cells from the 10× dataset and 5000 cells from the SS2 dataset. We trained the neural network on the 10× cells and tested it on the SS2 cells. During the training process, we recorded the accuracy and the cost after each epoch. The accuracy was defined as the percentage of cells whose cell type was correctly predicted, and the cost was the output of the cost function after each epoch. We found that the training accuracy saturated early (5 epochs), and the testing accuracy saturated at around 50 epochs (Fig. 1E), and the cost decreased very slowly after 50 epochs (Fig. 1F). These results indicate that early stopping can be used to reduce training time and prevent overfitting.

3.5 Cell type prediction using the mouse cell atlas

Since the cell types from the two mouse cell atlas datasets can be accurately predicted, we combined the two datasets and used the combined dataset as the reference to predict cell types for other datasets. We first tried to predict cell types for a dataset that contains flow cytometry sorted leukocytes from mouse aorta (Winkels *et al.*, 2018). All cells were predicted as leukocytes except for one erythrocyte, which we think is a doublet of an erythrocyte and B cell as high expression of hemoglobin genes was detected (Fig. 2A). We also carried out unsupervised analysis on the dataset and clustered the cells using Seurat. Then we used the canonical markers to assign the cell types for each cluster (Fig. 2B). Most cells had the same cell type assignment by the two methods. However, our neural network detected some NK cells, which were in the same cluster with the T cells, and were assigned as T cells in the unsupervised clustering.

We checked the expression of CD3D, CD8A and GZMA (Fig. 2C), and found no expression of CD3D and CD8A, but high expression of GZMA in the NK cells, which suggests that these are likely NK cells. To test if ACTINN produces consistent results from run to run, we trained the neural network on the combined TMA dataset, tested it on the mouse leukocytes dataset, and repeated this process 10 times. We found that most of the cells were assigned the same label across all 10 runs (Supplementary Fig. S2A), and the frequency for each cell type was also consistent between different runs (Supplementary Fig. S2B).

It is generally thought that human and mouse share similar cell types, and the same cell type from human and mouse share similar expression profiles. To test this, we trained our neural network on the mouse cell atlas datasets and used the parameters to predict the cell types for a human peripheral blood mononuclear cell (PBMC) dataset. We found 4 main populations in the PBMC dataset, namely, B cells, monocytes, NK cells and T cells (Fig. 2D). We plotted the canonical markers for these 4 populations (Fig. 2E) and found that the predicted cell types matched the expected marker expression. These results suggest that the mouse cell atlas datasets can be used as a reference to identify cell types for both human and mouse cells.

3.6 ACTINN accurately identifies cell types not in the reference

An scRNA-seq experiment may be performed on tissues where not all the cell types in the data of interest are included in the reference dataset. If a cell cannot be classified as a known cell type in the

training data, we would label it 'uncertain'. To test if ACTINN can identify cell types that are not in the reference, we trained the neural network on the TMA datasets and tested it on the mouse leukocytes plus 109 mouse nerve cells (the nerve cells are not in the training data). We output the probabilities for each cell being one of the cell types in the training data, and labeled the cell 'uncertain' if its highest probability is smaller than 0.95. We found that most of the B cells, T cells, NK cells, monocytes and granulocytes were assigned correctly (Supplementary Fig. S3A). By contrast, 105 out of 107 nerve cells were assigned 'uncertain' (Supplementary Fig. S3B). These results show that ACTINN is able to identify cell types that are not in the training dataset.

3.7 ACTINN accurately predicts cell subtypes

Although it is relatively easy to distinguish different cell types in scRNA-seq using the unsupervised clustering methods, it is more difficult to further divide one cell type into cell subtypes. Here, we collected five publicly available datasets (Zheng *et al.*, 2017), each containing one flow cytometry sorted T cell subtype. We merged these datasets and selected the cells that have the same labels between our analyses and the flow cytometry sorting, and then used these cells as a reference for the neural network. We then clustered the selected cells and identified markers (Fig. 3A and B) for each sub cell type using Seurat. For the test set, we used the T cells from the human PBMC datasets mentioned above.
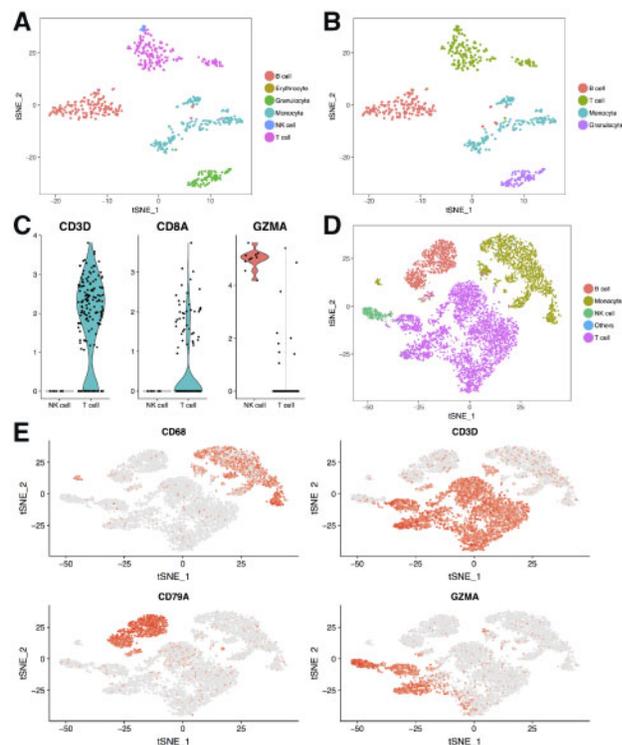
Fig. 2. Neural network predicts cell types for human and mouse datasets. (**A**) Cell types predicted by the neural network for the mouse leukocyte dataset. (**B**) Cell types identified by unsupervised clustering and canonical markers for the mouse leukocyte dataset. (**C**) Violin plots showing three genes' expression level in the NK and T cells from the mouse leukocytes. (**D**) Cell types predicted by the neural network for the human PBMC dataset. (**E**) TSNE plots showing four marker genes' expression for the human PBMC dataset
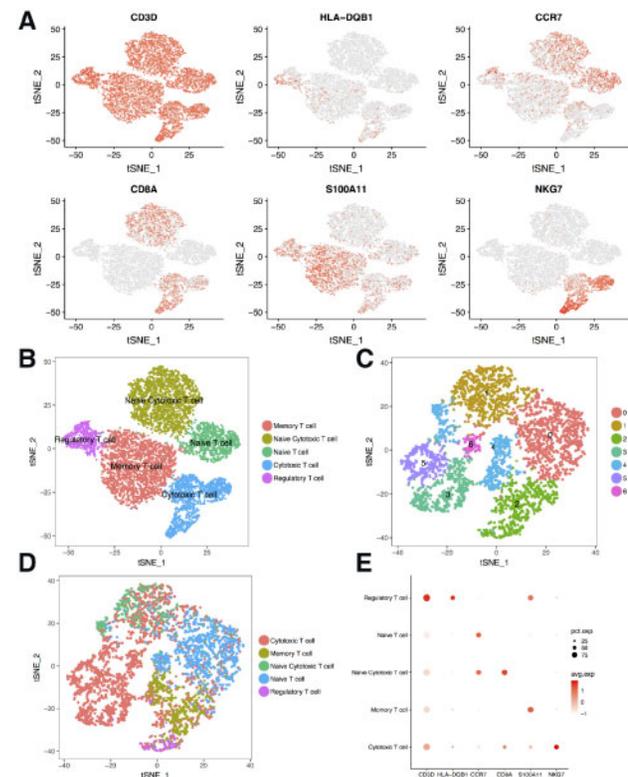
Fig. 3. Neural network predicts sub cell types. (**A**) TSNE plots showing six maker genes' expression for the reference T cell subtypes. (**B**) T cell subtypes obtained to train the neural network. (**C**) T cells from the human PBMC were grouped into seven clusters by the unsupervised method. (**D**) Subtypes predicted for the T cells from the human PBMC. (**E**) Dot plot showing the expression of six genes for the predicted subtypes, dot size represents the percentage of cells expressing the gene, color scale represents the expression level of the gene

To test our neural network's ability to predict cell subtypes, we trained it on the T cell subtype reference, and predicted the subtypes for the T cells from the PBMC dataset (Fig. 3D). We then identified marker genes for each predicted subtype. As expected, the marker genes matched the ones from the reference (Fig. 3E). These results show that our neural network can be used to accurately identify cell subtypes. We found that the subtypes predicted by the neural network did not perfectly match the cell types associated with the Seurat clusters (Fig. 3C). Some clusters contained different subtypes and some subtypes were composed of several clusters. We think the difference was influenced by two factors: (i) Unsupervised clustering considers all variance in the data, while the neural network is trained to find the difference between the subtypes; (ii) It is difficult to set the parameters optimally for the unsupervised analysis, which can result in multiple cell types in one cluster or multiple clusters for one cell type.

3.8 Comparison to other cell type identification tools

As the field of scRNA-seq is evolving rapidly, new ideas and methods are being published frequently. Several supervised scRNA-Seq cell type identification methods were proposed recently. SuperCT (Xie *et al.*, 2019) uses a neural network, CaSTLe (Lieberman *et al.*, 2018) uses XGboost and SingleCellNet (Tan *et al.*, 2018) uses a random forest to annotate cell types in scRNA-Seq experiment. We found that these three methods convert the expression values to binary signals (SuperCT and XGboost) or four categories (CaSTLe) before training the data. This conversion may significantly decrease the complexity of the expression data, which makes it difficult to distinguish between small changes in expression. We compared the performance of the three methods to ACTINN in sub cell type identification. We trained CaSTLe and SingleCellNet using the T cell subtype reference, and trained SuperCT on its human cell reference as it does not allow user defined reference. Then we predicted the subtypes for the T cells from the PBMC dataset. CaSTLe and SingleCellNet failed to define most of the naïve T cells and regulatory T cells, and SuperCT failed to distinguish T cell subtypes (Supplementary Fig. S4A–D). Based on the predictions and marker gene expression, we manually set the labels for the T cell subtypes (Supplementary Fig. S4E). Then we calculated the prediction accuracy for ACTINN (73%), CaSTLe (59%) and SingleCellNet (57%) (Supplementary Fig. S4F). These results show that ACTINN outperforms the 3 tools in finding small changes between subtypes.

4 Discussion

scRNA-seq provides high resolution profiling of the transcriptomes of single cells. Typically, the first step in scRNA-seq analysis is to assign each cell a cell type based on our prior knowledge of marker genes. Current methods for cell type assignment first cluster the cells in an unsupervised manner and rely on the canonical markers to identify the cell types for each cluster. However, this approach has several limitations, including the fact that the clusters may not optimally segregate single cell types, and certain cell types may not have previously characterized markers. Moreover, these methods are computationally intensive, especially when the number of cells becomes large. To render cell type identification in scRNA-seq more efficient, we employed a neural network, trained it on cells with predefined cell types, and used it to predict cell types for new datasets.

We first obtained and cleaned two datasets from the Tabula Muris Consortium, then trained and tested our neural network on these datasets with or without batch effect introduced by different scRNA-seq platforms. The training accuracy always approached 100%, and the testing accuracy was around 99.8% within a platform and 99.0% when testing and training are performed across different platforms. As the cell types in the two Tabula muris atlas datasets can be mutually predicted using our neural network, we merged them and used the combined datasets as the reference to predict cell types for other datasets. The predicted cell types were well matched with the cell types assigned using the canonical markers for both the mouse and human datasets. We also trained and tested the neural network on 5 T cell subtypes and found that the predicted subtypes showed the same markers as the reference subtypes, which suggests that our neural network can be used to predict sub cell types as well.

Compared to the traditional unsupervised methods used for cell type identification, our neural network has the following advantages: (i) It uses all the genes to capture the features for each cell type instead of relying on a limited number of canonical markers. (ii) It focuses the analysis on the signal associated with the variance between cell types, while unsupervised clustering tends to be affected by other sources of cell type independent variation (i.e. platform or cell cycle). (iii) It requires no background knowledge of cell type markers, while the unsupervised method requires users to have prior knowledge of canonical markers for each cell type in their data. (iv) It is much more computationally efficient than the traditional approach. Moreover, users can subsample the reference cells to make the computation of the neural network less compute intensive and more memory efficient. We also compared ACTINN to three other cell type prediction tools, and the results showed that ACTINN performs better in finding small changes between subtypes.

There are some aspects of our approach that could be improved in the future. As the neural network is supervised, the quantity and quality of the reference data are critical. We anticipate that with time more cell types from larger atlases should be used to train a more comprehensive neural network. Also, better pairing of reference and test sets will undoubtedly improve performance. For example, the soon to be developed human cell atlas should be used to predict human cell types instead of the mouse cell atlas. Nonetheless, we showed that even with the current reference data our neural network is computationally efficient and accurate, and should improve cell type identification pipelines.

Funding

References

Butler,A. *et al*. (2018) Integrating single-cell transcriptomic data across different conditions, technologies, and species. *Nat. Biotechnol*., **36**, 411–420.

Cho *et al*. (2018) Generalizable and scalable visualization of single cell data using neural networks. *Cell Syst*., **7**, 185–191.e4.

Haghverdi,L. *et al*. (2018) Batch effects in single-cell RNA-sequencing data are corrected by matching mutual nearest neighbors. *Nat Biotechnol*., **36**, 421–427.

Hwang,B. *et al*. (2018) Single-cell RNA sequencing technologies and bioinformatics pipelines. *Exp. Mol. Med*., **50**, 96.

Lieberman,Y. *et al*. (2018) CaSTLe - Classification of single cells by transfer learning: harnessing the power of publicly available single cell RNA

sequencing experiments to annotate new experiments. *PLoS One*, **13**, e0205499.

Lin,C. *et al.* (2017) Using neural networks for reducing the dimensions of single-cell RNA-Seq data. *Nucleic Acids Res.*, **45**, e156.

Lopez,R. *et al.* (2018) Deep generative modeling for single-cell transcriptomics. *Nat. Methods*, **15**, 1053–1058.

Shaham,U. *et al.* (2017) Removal of batch effects using distribution-matching residual networks. *Bioinformatics*, **33**, 2539–2546.

Tan,Y. *et al.* (2018) SingleCellNet: a computational tool to classify single cell RNA-Seq data across platforms and across species. *bioRxiv*, doi: 10.1101/508085.

Winkels,H. *et al.* (2018) Atlas of the immune cell repertoire in mouse atherosclerosis defined by single-cell RNA-sequencing and mass cytometry. *Circ. Res.*, **122**, 1675–1688.

Xavier. *et al.* (2010) Proceedings of the thirteenth international conference on artificial intelligence and statistics. *PMLR*, **9**, 249–256.

Xie,P. *et al.* (2019) SuperCT: a supervised-learning framework for enhanced characterization of single-cell transcriptomic profiles. *Nucleic Acids Res.*, **47**, e48.

Zheng,G.X.Y. *et al.* (2017) Massively parallel digital transcriptional profiling of single cells. *Nat. Commun.*, **8**, 14049.