



Predictive modeling of single-cell DNA methylome data enhances integration with transcriptome data

Yasin Uzun, Hao Wu and Kai Tan

Genome Res. published online November 20, 2020
Access the most recent version at doi:[10.1101/gr.267047.120](https://doi.org/10.1101/gr.267047.120)

P<P	Published online November 20, 2020 in advance of the print journal.
Accepted Manuscript	Peer-reviewed and accepted for publication but not copyedited or typeset; accepted manuscript is likely to differ from the final, published version.
Open Access	Freely available online through the <i>Genome Research</i> Open Access option.
Creative Commons License	This manuscript is Open Access. This article, published in <i>Genome Research</i> , is available under a Creative Commons License (Attribution 4.0 International license), as described at http://creativecommons.org/licenses/by/4.0/ .
Email Alerting Service	Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or click here .



To subscribe to *Genome Research* go to:
<https://genome.cshlp.org/subscriptions>

Published by Cold Spring Harbor Laboratory Press

Predictive modeling of single-cell DNA methylome data enhances integration with transcriptome data

Yasin Uzun^{1,2}, Hao Wu^{3,5} and Kai Tan^{1,2,3,4,5}

¹Center for Childhood Cancer Research, The Children's Hospital of Philadelphia, Philadelphia, PA 19104, USA

²Department of Biomedical and Health Informatics, The Children's Hospital of Philadelphia, Philadelphia, PA 19104, USA

³Department of Genetics, University of Pennsylvania, Philadelphia, PA 19104, USA

⁴Department of Pediatrics, University of Pennsylvania, Philadelphia, PA 19104, USA

⁵ Penn Epigenetics Institute, University of Pennsylvania, Philadelphia, PA 19104, USA

Correspondence: Kai Tan (tank1@email.chop.edu)

Keywords: single cell, bisulfite sequencing, DNA methylation, integration, ensemble learning

Abstract

Single-cell DNA methylation data has become increasingly abundant and has uncovered many genes with a positive correlation between expression and promoter methylation, challenging the common dogma based on bulk data. However, computational tools for analyzing single-cell methylome data are lagging far behind. A number of tasks, including cell type calling and integration with transcriptome data, requires the construction of a robust gene activity matrix as the prerequisite but challenging task. The advent of multi-omics data enables measurement of both DNA methylation and gene expression for the same single cells. Although such data is rather sparse, they are sufficient to train supervised models that capture the complex relationship between DNA methylation and gene expression and predict gene activities at single-cell level. Here, we present MAPLE (Methylome Association by Predictive Linkage to Expression), a computational framework that learns the association between DNA methylation and expression using both gene- and cell-dependent statistical features. Using multiple datasets generated with different experimental protocols, we show that using predicted gene activity values significantly improves several analysis tasks, including clustering, cell type identification and integration with transcriptome data. Application of MAPLE revealed several interesting biological insights into the relationship between methylation and gene expression, including asymmetric importance of methylation signals around transcription start site for predicting gene expression; and increased predictive power of methylation signals in promoters located outside CpG islands and shores. With the rapid accumulation of single-cell epigenomics

data, MAPLE provides a general framework for integrating such data with transcriptome data.

Introduction

Recent advances in single-cell genomics provide the opportunity to capture a variety of epigenomic signatures at single-cell resolution including histone modification (Grosselin et al. 2019), DNA methylation (Smallwood et al. 2014; Urich et al. 2015; Luo et al. 2017, 2018; Mulqueen et al. 2018), chromatin accessibility (Buenrostro et al. 2015; Chen et al. 2018) and long-range chromatin interaction (Nagano et al. 2013; Stevens et al. 2017). In particular, single-cell DNA methylome analysis can provide quantitative and high-resolution measurement of cell-type-specific epigenomic landscape in both development and disease, as the mammalian embryonic development is associated with dynamic changes in DNA methylation at *cis*-regulatory elements and genome-wide deregulation of DNA methylation is associated with many types of cancer (Greenberg and Bourc'his 2019).

Identifying genome-wide methylation signatures at single-cell resolution with an unbiased technique such as bisulfite sequencing (BS-seq) comes with unique challenges. Technically, the data is sparse and genomic coverage is rather limited (~5% of the genome per cell on average), even for deeply sequenced samples with more than 5 million reads per cell (Luo et al. 2017). Biologically, the interpretation of methylome signal is context dependent. Whole genome DNA methylome data suggest that gene body methylation is positively correlated with gene expression in embryonic stem cells

(Lister et al. 2009; Greenberg and Bourc'his 2019), whereas in other cell types such as post-mitotic neurons, genic methylation is negatively correlated with gene expression (Lister et al. 2013; Lee et al. 2019). These observations indicate that the regulatory roles of DNA methylation are both genomic-feature and cell-type specific. Single-cell multi-omics studies suggest significant correlation (both positive and negative) between expression and gene body methylation for only a limited number of genes (Hu et al. 2016; Angermueller et al. 2016a). Additionally, mean promoter methylation is significantly negatively correlated with gene expression only for a fraction of promoters in individual cells (Angermueller et al. 2016a; Clark et al. 2018; Argelaguet et al. 2019).

The lack of well-defined association between DNA methylation and gene expression poses obstacles for the analysis and integration of this type of epigenomic data. Identifying cell types is relatively straightforward in scRNA-seq data using clustering techniques and marker genes. However, the same approach does not work well for single-cell DNA methylation data due to the lack of clear association between methylation and gene expression. Consequently it is difficult to use marker genes to assign cell types in this case. In terms of data integration, several computational methods have been developed for integrating different types of single-cell data (Stuart et al. 2019; Welch et al. 2019; Korsunsky et al. 2019; Tran et al. 2020). However, without accurate input representing gene activities in single cell methylation data, these methods cannot achieve reliable integration, except for certain cell types, such as neurons, which have strong methylation signals in the gene bodies. Therefore, robust estimate of methylation-based gene activity score remains a bottleneck for accurate

integration of DNA methylation data with other types of single-cell omics data. Resolving this bottleneck can pave the way for comparative analysis of gene regulation mechanisms across cell populations in a complex tissue.

Recently, true multi-omics protocols (i.e. joint profiling of transcriptome and DNA methylome of the same cells) have started to emerge. It has already been demonstrated that several hundreds of cells can be sequenced for gene expression and DNA methylation in the same cell in a single experiment (Angermueller et al. 2016a; Clark et al. 2018). Here, we aim to take advantage of such data and develop a supervised learning framework called MAPLE (Methylome Association by Predictive Linkage to Expression), for predicting gene activity score based on single-cell DNA methylation data. To this end, we aim to develop gene-dependent and cell-dependent statistical features as the input to an ensemble learning framework.

Results

Overview of the method

We independently verified the findings related to association between DNA methylation and gene expression by computing the correlation between the expression and promoter and gene body methylation in single cell multi-omics datasets. As a result, we found that for only a small fraction of the promoters, the negative correlation between the methylation and the gene expression was significant. In addition, methylation of many promoters are positively correlated (Supplemental Fig. S1, Supplemental Table S1, S2). Similarly, many more genes have negative correlation between body

methylation and gene expression than those with positive correlation. Taken together, both published results and our own analyses suggest that there is not a straightforward approach for inferring gene activity levels using single-cell DNA methylation datasets. Instead, the correlation between the gene expression and methylation is gene dependent.

We hypothesize that the common patterns of promoter methylation-gene expression association for groups of genes can be modeled by using a supervised learning model. We developed two classes of statistical features as input to the supervised predictor. Promoters overlapping with CpG islands or shores have distinctive response to methylation (Deaton and Bird 2011; Jones 2012; Weber et al. 2007; Greenberg and Bourc'his 2019) and overall promoter CpG frequency has been used as a feature for predicting gene expression using bulk DNA methylation data (Kapourani and Sanguinetti 2016). In MAPLE, we used CpG frequency at higher resolution, as a vector of CpG frequencies of multiple genomic bins tiled across the promoter. This step generates the gene-dependent but cell-independent feature set to be used in the learning model. For the gene- and cell-dependent feature set, we computed the methylation level of each promoter bin for all cells and genes.

The sparsity of single-cell bisulfite sequencing data poses a significant challenge for machine learning approaches. Many promoter regions have a limited number of CpG sites covered in each single cell and dividing the promoter region into multiple bins further exacerbates the sparsity problem, resulting in either bins with no overlapping

calls at all, or insufficient calls to make a reliable estimation of the methylation level of the bins. To alleviate this problem, we resorted to the concept of ‘meta-cell’, essentially borrowing information from neighboring single cells. Such an approach has been used to analyze other types of single cell data (Wagner et al. 2017; van Dijk et al. 2018; Gong et al. 2018; Zhu et al. 2020). Each meta-cell represents a group of individual cells that are in a similar state with a specific cell in the data (Fig. 1A, Methods). Although this approach results in a slight loss of data resolution, it provides a reliable estimation of CpG level for the vast majority of meta-cells, even with small neighborhood sizes (Supplemental Fig. S2). Moreover, each single cell has its own unique neighborhood, thus no two meta-cells are expected to be identical and the single-cell nature of the data is preserved.

Once the meta-cells are identified, the CpG methylation level for each meta-cell is computed for 500bp bins across the promoter region, resulting in 20 features for each gene and each meta-cell. Together with the CpG frequencies of these bins, they constitute the two classes of features for each gene-meta-cell pair (Fig. 1B, Methods). These features are used as the input to the supervised learning model to infer the gene activity level in the meta-cell. For downstream analysis, a DNA methylation-based gene activity matrix is constructed and a desired method can be used to integrate the methylome data with gene expression data, for reliable assignment of cell types in the data and other joint multi-omics analyses (Fig. 1C).

Many supervised learning methods can be used to predict the gene activity level of the cells. We selected one commonly used, representative predictor from three broad categories: artificial neural networks, regression based models and decision tree based models. Convolutional Neural Networks (CNNs) is a class of deep-learning neural networks, which have gained popularity in image processing and bioinformatics (Angermueller et al. 2016b; Alipanahi et al. 2015; Angermueller et al. 2017). Our CNN architecture resembles the method described in (Singh et al. 2016) for predicting gene expression using bulk histone mark ChIP-seq data. However, instead of multiple histone marks, we used two classes of features, namely, CpG dinucleotide frequencies and CpG methylation levels for genomic bins surrounding the TSS (Fig. 1B, Methods). Elastic Net (EN) is a regularized regression method that combines the least absolute shrinkage and selection operator (LASSO) and ridge regression models. Random Forest (RF) consists of multiple decision trees, each of which is trained with random subsamplings from the training data and the result is obtained by combining the outputs of all decision trees in a democratized manner. As the baseline method, we computed the mean promoter de-methylation level (MPD) (ratio of unmethylated CpGs to all CpG calls in the promoter) as a predictor of gene activity level.

Supervised learning improves accuracy of gene activity prediction

We benchmarked the performance of MAPLE using four published single-cell multi-omics datasets generated using two different protocols (Supplemental Table S1). The datasets of Angermueller and colleagues (Angermueller et al. 2016a) and Hernando-Herraez and colleagues (Hernando-Herraez et al. 2019) were generated using scM&T-

seq (single-cell genome-wide Methylome and Transcriptome sequencing), while the datasets of Clark and colleagues (Clark et al. 2018) and Argelaguet and colleagues (Argelaguet et al. 2019) were generated using scNMT-seq (single-cell Nucleosome, Methylation and Transcription sequencing). For all datasets, the transcriptome and DNA methylome were jointly profiled for the same single cells. To evaluate the overall performance of the ensemble approach, we performed both internal and external cross validation, i.e. training a predictor with one dataset and predicting the gene activity levels using the remaining three datasets.

We evaluated the prediction accuracy using two metrics, Spearman's correlation and median squared error. We chose Spearman's correlation because it can capture both linear and nonlinear relationships in the data. We observed that the performance of the individual learners was sensitive to the selection of the training and test set data, especially for CNN and RF (Supplemental Fig. S3). These two methods performed well on the dataset of Angermueller et al. based on internal cross-validation, but poorly based on external cross-validation, indicating overfitting. On the other hand, their performance was low based on internal cross-validation using the dataset of Argelaguet et al., but higher based on external cross-validation using the same dataset as the training set. This is possibly due to the fact that the Argelaguet data set has the largest number of cells, which captured the variation in the data better and alleviated the overfitting to some extent. Hence, the performance of these two methods is dependent on the training set and not stable. Finally, the overall performance of EN was the lowest, possibly due to the relative simplicity of the model which is unable to capture the

complexity in the data. Therefore, to achieve a consistently improved performance independent of training and test data, we applied an ensemble approach by combining multiple predictors (Mendes-Moreira et al. 2012).

There are different ways for combining component predictors to build an ensemble predictor, such as unweighted average, weighted average and stacking (Methods). We evaluated all three approaches and found that they have similar performance across all combinations of training and test datasets (Supplemental Fig. S4). Therefore we chose the unweighted average approach due to its simplicity.

Ensemble based MAPLE clearly outperformed MPD in terms of correlation for the internal cross-validation tests (Supplemental Fig. S5). Regarding cross-dataset tests, across the twelve training-test dataset pairs, MAPLE achieved an average global Spearman's correlation of 0.62 across all genes and cells (Fig. 2A). In comparison, the baseline predictor using MPD as the feature gave an average global Spearman's correlation of 0.5, which means MAPLE had 24% average performance gain over MPD ($p=0.003$, one-sided t -test), reaching over 60% improvement for certain data sets (Supplemental Table S3). The Spearman's correlation across genes is also higher for MAPLE compared to the baseline predictor (Fig. 2B, Supplemental Fig. S6). We observed the same trend when using a median squared error as an alternative metric (Supplemental Fig. S7, S8). Taken together, these results demonstrate that a supervised predictor using both CpG frequency and methylation level can substantially enhance the accuracy of predicting gene expression activity at the single-cell level.

Next, we investigated the predictive power of individual features by computing a feature importance score using the random forest model. We found that feature importance score distributions are similar across all four datasets (Supplemental Fig. S9). For both DNA sequence based features (CpG ratio) and methylation rate features, the genomic bins that are closer to the TSS had higher importance. The importance score distribution of sequence based features was symmetric around TSS, with the importance score attenuating after \pm 2kb. In contrast, importance scores of the methylation rate features showed asymmetric distribution. Methylation rate of two bins immediately downstream of TSS had the highest predictive power, followed by the bin immediately preceding TSS. The importance score decreased after 1kb for the upstream region and 2kb for the downstream region.

Predicted gene activity improves identification of cell types

Since one of the most important utilities of single-cell data is to identify different cell types/states in a heterogeneous population, we next evaluated our trained predictor on a single-cell DNA methylation dataset containing 3,377 mouse neurons which were sequenced with the single-nucleus methylcytosine sequencing (snmC-seq) protocol (Luo et al. 2017). In this study, the authors identified different neuronal subtypes using gene body non-CpG (mCH) methylation, which is known to be inversely correlated with the expression level in adult neurons (Mo et al. 2015).

When we used the CpG (instead of non-CpG in order to be general) MPD values of all genes for dimensionality reduction, we observed no separation of cell types in the data and the MPD values of the marker genes for the neuronal subtypes were indistinguishable across the cell populations (Fig. 3A, B, Supplemental Fig. S10). In contrast, when we clustered the cells using gene activity levels of all genes predicted by MAPLE (trained on data from (Clark et al. 2018)), we identified two main clusters that were clearly separated (Fig. 3C). The larger cluster showed higher gene activity for excitatory neuron markers, such as *Tyro3*, *Slc17a7*, *Tbr1*, *Itpka* and the smaller cluster showed higher gene activity for inhibitory neuron markers, such as *Slc6a1*, *ErbB4* (Fig. 3D). The clustering result was consistent with that reported by the authors, with over 99% of the cells in the excitatory neuron cluster and 95% of the cells in the inhibitory cluster matching those reported by the original study, respectively. We observed a similar trend in a larger set of excitatory and inhibitory marker genes (Supplemental Fig. S11, S12, S13, S14).

To further evaluate the utility of predicted gene activity for cell type identification, we applied MAPLE to identify cell clusters in the embryoid body (EB) population in the multi-omics dataset by Clark et al. (Clark et al. 2018) using data from Angermueller et al. (Angermueller et al. 2016a) as the training set. In this study, the authors differentiated mouse embryonic stem cells to EB. Using expression data, they reported that the differentiated population clusters into two main clusters, one is labeled as pluripotent due to high expression of pluripotency genes and the other represents differentiated cells with the opposite pattern.

Using the MPD levels of promoters as the input for clustering did not recapitulate the heterogeneity observed using the transcriptome data. On the other hand, the same analysis using MAPLE predicted activities of all genes as the input resulted in two clusters that have similar sizes to the clusters defined using transcriptome data alone. Moreover, *Esrrb*, a pluripotency marker gene showed high levels of activity in one cluster whereas transcription factor *T*, a differentiation marker, displayed the opposite pattern (Fig. 4A, B), consistent with the pattern of heterogeneity based on the transcriptome data. Similar separation was evident for many additional marker genes for the two cell types (Supplemental Fig. S15, S16). To test the reliability of the separation of the cell types in the MAPLE input, we computed the dimensionality reduction with different parameter settings and we observed that the separation was evident across all parameter settings (Supplemental Fig. S17, S18).

In summary, these results demonstrate that use of a predictive modeling approach can significantly improve our ability to identify cell types and reveal heterogeneity in single-cell methylation data. Notably, even in this case of extremely small number of cells, using the meta-cell approach to predict gene activity did not lead to loss of resolution which could impede analysis of heterogeneity.

Predicted gene activity enhances integration with scRNA-seq data

We reasoned that accurate prediction of gene activity level can significantly enhance the integration of single-cell methylome data with transcriptome data. To test this

hypothesis, we took advantage of published true multi-omics datasets in which DNA methylome and transcriptome were measured for the same single cells. Therefore, we know the ground truth about the matching of the two data types for a given cell. We first used the dataset by Clark et al. that is composed of differentiated embryoid body cells and undifferentiated embryonic stem cells (ESCs) (Clark et al. 2018). Using predicted gene activity levels by either MAPLE or the MPD method, we integrated the methylome and transcriptome data using Seurat (Stuart et al. 2019). We then clustered the cells using the co-embedded transcriptome and methylome data produced by Seurat. We computed the fraction of cells in the methylome data that were assigned to the correct cluster, using a k-nearest neighbor (kNN) classifier and cells from single cell RNA data. In contrast to MPD, clusters computed based on MAPLE-predicted gene activity showed higher homogeneity (i.e. greater fraction of the cells in each cluster belongs to the same type, Fig. 4C-H, Supplemental Fig. S19), suggesting improved accuracy in data integration.

We further evaluated the performance of data integration using a larger dataset on primary tissue. The dataset includes 850 cells from four different embryonic time points (Argelaguet et al. 2019) during mouse gastrulation. Each cell was sequenced with multiple modalities, including chromatin accessibility, DNA methylation and RNA expression. We integrated the single-cell methylome data with the expression data, as described above and computed the fraction of cells in the methylome data that were assigned to the correct cluster. Using gene activity predicted by MPD, neither linear (PCA) nor non-linear (UMAP) dimensionality reduction methods on the integrated data

resulted in accurate matching of cells from the same developmental stage (Fig. 5A,B, Supplemental Fig. S20). In comparison, using MAPLE data as the input for integration, the resulting integrated data showed much higher fractions of matched cells based on the two data types for all four developmental stages (Fig. 5C,D, Supplemental Fig. S20, Supplemental Fig. S21). Although part of the populations from E6.5 and E7.5 has some overlap in the integrated data based on MAPLE input, this outcome is due to the biological nature of the data, as it is observed in all data modalities including gene expression (Argelaguet et al. 2019).

Discussion

While recent advances in single-cell DNA methylome sequencing technologies improve our ability to study epigenetic heterogeneity, data analysis poses unique challenges, especially the issue of connecting the methylome and transcriptome data. Here, we addressed this challenge by developing a supervised learning approach to inferring gene activity based on DNA methylation data. The inferred gene activity score acts as an intermediate input for integrating the two types of data. A similar approach is commonly used to integrate single-cell ATAC-seq data with transcriptome data where summed chromatin accessibility signal in the promoter and gene body is used as a proxy to gene activity and intermediate input for data integration. However, simply using summed DNA methylation signal in the promoter region does not work as we demonstrated using the mean promoter demethylation method, due to the more complex relationship between DNA methylation and gene expression. Instead, we showed that predictive modeling better captures the relationship between DNA

methylation and gene expression. As a result, the predicted gene activity score helps improve cell type identification in the methylome data as well as integration with transcriptome data.

For constructing meta-cells, we used the PCA that is computed with methylation levels of the variably methylated promoters (VMPs) as the input. It might be argued that the presence of CpG-poor promoters can affect quality of constructed meta-cells. We evaluated this issue by computing the overlap between the VMPs and CpG-poor promoters (10% of all promoters having the lowest number of CpG sites). We found that a very small fraction of VMPs are CpG-poor (0.6%) and their removal has little effect (7.5%) on the adjacency matrix (Supplemental Table S4), demonstrating reliability of meta-cells.

Despite the fact that using information from neighboring cells may lead to minor loss of resolution, the reliability gained by using meta-cell for predicting gene activity outweighs its cost. The ability of MAPLE to identify subpopulations with a very small number of EB cells (Clark et al. 2018) shows that the use of meta-cells does not significantly impede the ability of our method for detecting heterogeneity.

Although MAPLE consistently outperforms MPD for all datasets we studied, the performance gain varied across the training and test data sets used. It has been reported that expression of genes having CpG island promoters have low sensitivity to DNA methylation change (Antequera 2003; Fan and Zhang 2009). We therefore

investigated whether different gene sets contributed to the performance differences by MAPLE by categorizing genes into four groups: genes having promoters overlapping with CpG islands, CpG shores, CpG shelves and open sea (all remaining genes). We found the performance of MAPLE is highest using genes in the open sea and lowest using genes with CpG islands (Supplemental Fig. S22). This result suggests that the expression of genes located in CpG islands, shores and shelves has low sensitivity to changes in methylation signal. In contrast, genes in the open sea have a more dynamic expression pattern in response to methylation change.

Our analysis revealed some interesting biological insights. First, for many genes, promoter methylation and gene expression are positively correlated, which is inconsistent with the observation based on bulk data. This result suggests a more nuanced relationship between gene expression and promoter methylation that is dependent upon cell type or state and not previously captured by bulk data. Second, we found that the importance of CpG methylation in predicting gene expression is not symmetric around TSS. Overall, the downstream region has higher predictive power than the upstream region. Furthermore, CpG methylation of the 500bp region immediately downstream of the TSS has the greatest predictive power.

In our framework, we used promoter methylation since it can be directly associated with genes. It has been reported that gene-distal transcriptional enhancers also have unique DNA methylation signature (Lister et al. 2009; Stadler et al. 2011; Hon et al. 2013) . We anticipate that inclusion of enhancer DNA methylation signals can further improve the

accuracy of the predictive model. To this end, a major challenge is associating distal enhancers with their target genes. This question has been addressed extensively for bulk data using both experimental and computational approaches (Jung et al. 2019; He et al. 2014; Javierre et al. 2016; Cao et al. 2017). However, due to unique challenges in single-cell data, additional research is warranted to address this important question. Finally, we used an ensemble learning approach in our method. Although we chose three representative predictors as components of the ensemble predictor, benefits of adding more predictors can be explored in future work.

Methods

Processing of public datasets

Datasets used in this study are listed in Supplemental Table S1. Briefly, we used four single-cell multi-omics datasets generated by two different experimental protocols, scM&T-seq (Angermueller et al. 2016a) and scNMT-seq (Clark et al. 2018). In addition, we used the snmC-seq (Luo et al. 2017) dataset as the methylation-only data. We used Bismark methylation call files (COV files) when available from the authors, or converted the methylation calls into Bismark COV format (Krueger and Andrews 2011). For each dataset, we only used cells for which both scRNA-seq and sc-Methylome data were available.

For the single-cell RNA-seq data, we normalized the number of reads by the total number of reads per cell and obtained CPM (counts per million). CPM values were log

transformed and were further normalized by the maximum log expression value in the dataset in order to fit all datasets into the same range and make training and test sets compatible.

Computing meta-cells

In order to alleviate the problem of data sparsity, we combined DNA methylation data from neighboring cells into meta-cells as follows. We first counted the CpG methylation calls in the +/- 5kbp region flanking the transcription start site (TSS) of the genes and computed the methylation level as the ratio of methylated CpGs to all CpG calls in those regions. For regions with no cytosine calls for a particular cell, we used the mean methylation level across all promoters of that cell. Next, we performed principal component analysis (PCA) on this methylation level matrix using the top 5,000 promoters with highest variance of methylation across the cells. We then used the top $d=10$ (d is an adjustable parameter in our method) principal components as the feature space to compute the Euclidean distance between each cell pair, since the total variance explained after the 10th principle component was minimal (Supplemental Fig. S23). Based on the distance, we defined the local neighborhood of each cell as the k nearest cells. Note that each meta-cell corresponds to an actual cell in the dataset, i.e. the number of meta-cells is the same as the number of single cells in the data. Since we observed that there is minimal improvement in terms of non-empty bins for k larger than 20 (Supplemental Fig. S24) for all datasets in this study, we used $k=20$.

We evaluated the effect of CpG-poor promoters for meta-cell as follows. First, we determined the 20 nearest neighbors of each cell using the top 10 dimensions of PCA and built an adjacency matrix. Next, we repeated this procedure when CpG-poor promoters are removed from the input gene set when computing the PCA. Then we determined the number of changes in the two matrices with and without the CpG-poor promoters. Finally, we calculated the ratio of the number of changed edges to the total number of edges in the original adjacency matrix (with CpG-poor promoters) as the percentage of adjacency difference.

Calculation of DNA methylation rate

Transcription start sites (TSS) were defined based on the GENCODE annotation (release vm23) for mouse genome (release GrCM38). Promoter was defined as the region spanning 5kbp up and downstream of the TSS.

To compute CpG methylation rate, we first divided each promoter into 20 bins of 500bp length. For each meta-cell, the numbers of methylated and unmethylated CpG sites for each bin were counted. The methylation rate of each bin was calculated by dividing the total number of methylated cytosines to all cytosines in that bin, considering all the cytosine calls for all cells in the corresponding meta-cell.

Selection of the bin size may have an effect on the performance. Very large bin sizes can lead to loss of resolution whereas too small bin sizes will result in many bins with few CpG calls, leading to inaccurate estimates of methylation rates. To evaluate this

issue, we performed external cross validation using different bin sizes. We found that there is a minor difference in performance using different bin sizes (Supplemental Fig. S25).

Feature Set

Two classes of features were computed, gene-dependent and cell-independent feature and cell- and gene-dependent feature. For each gene, we extracted the sequence information for the +/- 5kbp region around the TSS. For each 500bp bin, we computed the CpG frequency, considering both strands, which resulted in 20 gene-dependent, cell-independent features. Cell- and gene-dependent features are CpG methylation rates of the 20 promoter bins of a gene in a given cell. They were calculated by using all methylation calls for all cells in the corresponding meta-cell.

Training and testing the predictor

We randomly subsampled 100,000 cell-gene pairs from each training dataset in order to speed up the training process. Three supervised learners, convolutional neural network (CNN), elastic net (EN), and random forest (RF) were trained with each of the four multi-omics datasets. We evaluated the performance with different parameter settings for individual learners using 5-fold internal cross validation (Supplemental Fig. S26, S27, S28). To avoid bias due to different experimental protocols and/or datasets, we trained and tested the predictors using external cross validation, i.e. training and testing using different datasets.

We trained the CNN models with 50 filters using the ReLU activation function for hidden layers and linear activation for the output layer. Kernel size was set to 5 and max pooling was set to size of 4. Mean Squared Error (MSE) was used as the loss function and patience for early stopping was set to 10 epochs. Models were regularized by setting the dropout rate to 0.2 to avoid overfitting (Srivastava et al. 2014). We used the R keras package for the implementation of the model (Falbel et al. 2019).

We trained the EN predictor (Friedman et al. 2010) by setting the value of *alpha* to 0.5.

The value of the λ was determined using 10-fold internal cross validation. We used the R glmnet package to train the predictor with cv.glmnet function (Friedman et al. 2009).

Random Forest (RF) predictors were constructed with 500 trees, where each tree was grown with random subsampling of training data using 80% of the training samples. We used the R randomForest package (Liaw and Wiener 2018) to build the RF models.

There are different rules for building an ensemble from the individual predictors. The unweighted average approach employed in this study is the mean of the outputs of the individual predictors and does not depend on any prior assumptions on the underlying models. As an alternative, weighted average prioritizes some of the predictors over others based on some prior information. We first calculated the median correlations between predicted gene activity and observed gene expression from the cross validation results and used them as weights for combining the outputs of the predictors. In other words, the predictor that was associated with higher correlation in the training

data received higher weight. As another alternative, we used median accuracy (1 - error) as the weight for each predictor.

A more complicated combination rule is the stacked learning (stacking) approach, in which a second layer predictor is trained on the outputs of the first layer predictors (Wolpert 1992; Breiman 1996) using training data. We tested stacking approach with three different second level predictors (belonging to the same family of learners used for the first layer): elastic net regression, artificial neural network and random forest. This approach gave mixed results on the different datasets, offsetting the advantage of stability of ensemble approach. As a result, we decided to employ the unweighted averaging approach, which is the simplest ensemble model and does not rely on any prior information. We compared the correlations between MAPLE and MPD using *t*-test, and we replicated the correlation coefficients achieved using MPD, so that there are the same number (12) of values for comparison of two matrices.

Feature Importance

We computed feature importance scores with the random forest model using the “importance” function in R random forest package (Liaw and Wiener 2018). We used the MSE difference as the importance measure. It was calculated by training a model using all features, then iteratively permuting the values for one feature at a time and computing the difference between overall MSE and MSE obtained with the permuted feature.

Integration of methylome and transcriptome data using predicted gene activity

We integrated single-cell DNA methylation data with single-cell RNA sequencing data using Seurat version 3 (Stuart et al. 2019). Only cells that have both expression and methylation data were used for the integration. After normalization of both data types, the top 3000 integration features (genes) were selected using the “SelectIntegrationFeatures” function, followed by integration anchors (cells) selected by using the “FindIntegrationAnchors” function, the selected features. The k.filter (number of neighbors) was set to 100 and normalization was set to SCT (Hafemeister and Satija 2019). The two data types were integrated by using those anchors and the “IntegrateData” function. Finally, we run PCA,UMAP and clustering on the integrated (co-embedded) data including cells from both data types.

Software availability

MAPLE software is implemented in R (R Core Team 2019) and is freely available under the MIT license. Source code is available as Supplemental Code and at GitHub (<https://github.com/tanlabcode/MAPLE.1.0>).

Acknowledgements

We thank the Research Information Services at the Children’s Hospital of Philadelphia for providing computing support. This work was supported by National Institutes of Health of United States of America grants CA226187, CA233285, and HL156090 (to K.T.), a grant from the Leona M. and Harry B Helmsley Charitable Trust (2008-04062 to K.T.) and a grant from the Alex’s Lemonade Stand Foundation (to K.T.)

Disclosure Declaration

The authors have no competing interests.

Figure Legend

Figure 1. Schematic overview of the computational framework. **A)** Generation of meta-cells from single-cell DNA methylation data. The gene-by-cell DNA methylation matrix is used for principle component analysis (PCA) to reduce the dimensionality. Each point in the PCA plot is a single cell. A meta-cell is the set of k nearest cells to an individual cell in the PCA space. **B)** Prediction of gene activity level by combining DNA methylation and sequence information in promoter. TSS, Transcription Start Site. ML, Methylation Level. Each horizontal box represents a genomic bin. Each circle represents a CpG site and filled circle represents methylated cytosine and empty circle represents unmethylated cytosine in a meta-cell. CpG ratio is defined as the percentage of CpG dinucleotides in a genomic bin. CpG methylation level is defined as the ratio of the number of methylated CpG calls to all CpG calls. **C)** Integration of single-cell methylation and single-cell RNA-seq data based on predicted gene activity scores using single-cell DNA methylation data.

Figure 2. Prediction accuracy of gene expression using DNA methylation data and ensemble learning. **A)** Heatmap showing global Spearman's correlation coefficients between observed gene expression and predicted gene activity for all genes across all cells in a dataset. Rows represent training datasets and columns represent test datasets. Row and column Roman numbers correspond to the datasets shown. **B)** Distribution of Spearman's correlation coefficients across genes using the dataset of Clark et al. as the training set. Each data point represents one cell. MAPLE (EL) correlations are significantly higher than those of MPD ($p < 10^{-16}$, one-sided t -test for all three comparisons).

Figure 3. Predicted gene activity using methylome data improves cell subtype identification for neurons. **A)** UMAP of clustering result generated using mean promoter de-methylation (MPD) as the input. **B)** Violin plot of MPD values for marker genes for excitatory (*Tyro3*, *Slc17a7*, *Tbr1*, *Itpka*) and inhibitory (*Slc6a1*, *ErbB4*) neurons. **C)** Same as panel A) but using predicted gene activity (PGA) as the input. **D)** Same as panel B) but using predicted gene activity as the input.

Figure 4. Predictive modeling improves integration with transcriptome data of cell lines. **A)** Cell heterogeneity based on transcriptome and DNA methylome data. Left

panel, UMAP using RNA-seq data as the input. Color scale represents the log normalized (using Seurat) expression level (read counts) of *Esrrb* for EBs. Middle panel, UMAP using mean promoter demethylation as the input. Color scale represents the MPD (1 - mean methylation level) level of the *Esrrb* gene. Right panel, UMAP using MAPLE predicted gene activity based on DNA methylation data as the input. Color scale represents the MAPLE predicted gene activity levels of *Esrrb*. **B)** Same as panel A) but for the *T* gene. **C)** UMAP based on integrated RNA-seq and DNA methylation data. Mean promoter de-methylation (MPD) was used as the input for data integration using Seurat. EB, embryoid body. ESC, embryonic stem cell. **D)** Density clustering of the data shown in the UMAP in panel C). **E)** Confusion matrix plot based on the clustering result shown in panel D), illustrating the agreement between cell type assignment based on clustering and true cell type. Size of each quadrant is proportional to the number of cells classified. **F)** Same as panel C), but using predicted gene activity as the input. **G)** Same as panel D), but using predicted gene activity as the input. **H)** Same as panel E), but using predicted gene activity as the input. Chi-square test p-value for the confusion matrices in g and h is 0.002.

Figure 5. Predictive modeling improves integration with transcriptome data of primary tissues. **A)** UMAP plots for integrated expression and DNA methylation data. Mean promoter de-methylation (MPD) was used as the input for data integration using Seurat. **B)** Same as panel A), but using MAPLE predicted gene activity as the input. **C)** Pie chart showing the percentage of correctly and mis-classified cells using scDNA-methylation data, based on k-nearest neighbor (kNN) classification on the scRNA-seq cells for the MPD based UMAP in A) . **D)** Same as panel C), but for the MAPLE based UMAP in B). Chi-square test p-value for the comparison between correct and mis-classifications in panels C and D is 3.6×10^{-10} .

References

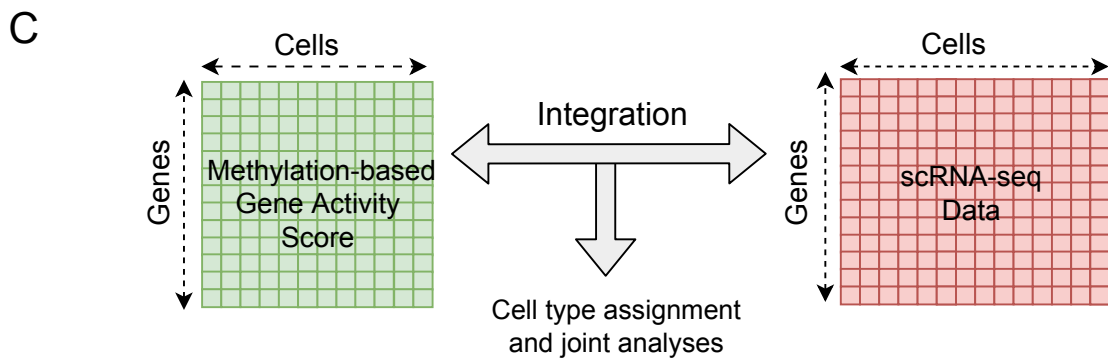
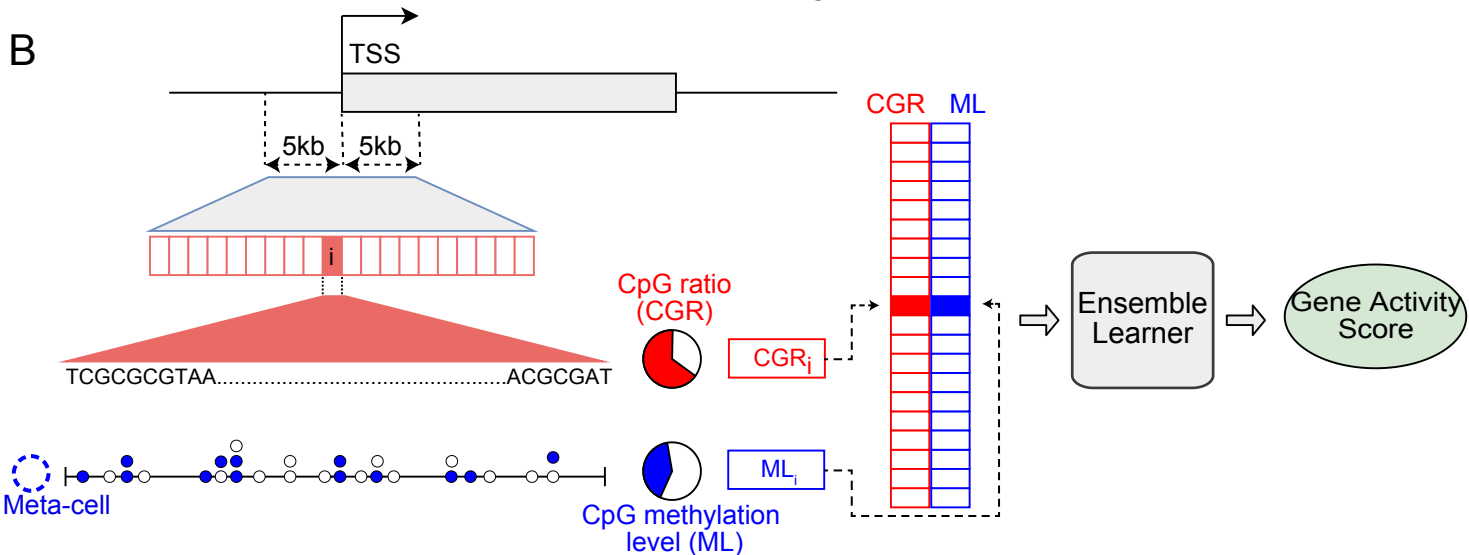
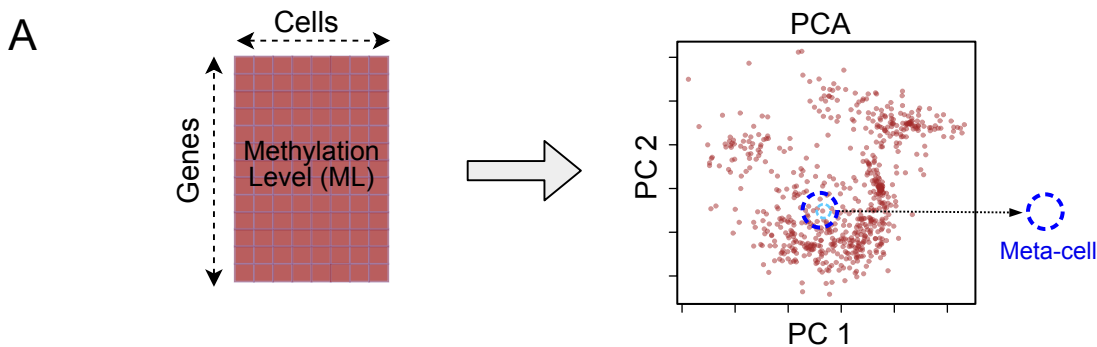
- Alipanahi B, Delong A, Weirauch MT, Frey BJ. 2015. Predicting the sequence specificities of DNA- and RNA-binding proteins by deep learning. *Nature Biotechnology* **33**: 831–838. <http://dx.doi.org/10.1038/nbt.3300>.
- Angermueller C, Clark SJ, Lee HJ, Macaulay IC, Teng MJ, Hu TX, Krueger F, Smallwood S, Ponting CP, Voet T, et al. 2016a. Parallel single-cell sequencing links transcriptional and epigenetic heterogeneity. *Nat Methods* **13**: 229–232.
- Angermueller C, Lee HJ, Reik W, Stegle O. 2017. DeepCpG: accurate prediction of single-cell DNA methylation states using deep learning. *Genome Biol* **18**: 67.
- Angermueller C, Pärnamaa T, Parts L, Stegle O. 2016b. Deep learning for computational biology. *Mol Syst Biol* **12**: 878.

- Antequera F. 2003. Structure, function and evolution of CpG island promoters. *Cell Mol Life Sci* **60**: 1647–1658.
- Argelaguet R, Clark SJ, Mohammed H, Carine Stapel L, Krueger C, Kapourani C-A, Imaz-Rosshandler I, Lohoff T, Xiang Y, Hanna CW, et al. 2019. Multi-omics profiling of mouse gastrulation at single-cell resolution. *Nature* **576**: 487–491. <http://dx.doi.org/10.1038/s41586-019-1825-8>.
- Breiman L. 1996. Stacked regressions. *Mach Learn* **24**: 49–64.
- Buenrostro JD, Wu B, Litzenburger UM, Ruff D, Gonzales ML, Snyder MP, Chang HY, Greenleaf WJ. 2015. Single-cell chromatin accessibility reveals principles of regulatory variation. *Nature* **523**: 486–490.
- Cao Q, Anyansi C, Hu X, Xu L, Xiong L, Tang W, Mok MTS, Cheng C, Fan X, Gerstein M, et al. 2017. Reconstruction of enhancer–target networks in 935 samples of human primary cells, tissues and cell lines. *Nat Genet* **49**: 1428–1436.
- Chen X, Miragaia RJ, Natarajan KN, Teichmann SA. 2018. A rapid and robust method for single cell chromatin accessibility profiling. *Nat Commun* **9**: 5345.
- Clark SJ, Argelaguet R, Kapourani C-A, Stubbs TM, Lee HJ, Alda-Catalinas C, Krueger F, Sanguinetti G, Kelsey G, Marioni JC, et al. 2018. scNMT-seq enables joint profiling of chromatin accessibility DNA methylation and transcription in single cells. *Nat Commun* **9**: 781.
- Deaton AM, Bird A. 2011. CpG islands and the regulation of transcription. *Genes Dev* **25**: 1010–1022.
- Falbel D, Allaire JJ, Chollet F, RStudio, Google, Tang Y, Van Der Bijl W, Studer M, Keydana S. 2019. *R interface to keras*. CRAN <https://cran.r-project.org/web/packages/keras/>.
- Fan S, Zhang X. 2009. CpG island methylation pattern in different human tissues and its correlation with gene expression. *Biochem Biophys Res Commun* **383**: 421–425.
- Friedman J, Hastie T, Tibshirani R. 2009. glmnet: Lasso and elastic-net regularized generalized linear models. *R package version 1*.
- Friedman J, Hastie T, Tibshirani R. 2010. Regularization paths for generalized linear models via coordinate descent. *J Stat Softw* **33**: 1.
- Gong W, Kwak I-Y, Pota P, Koyano-Nakagawa N, Garry DJ. 2018. Drlmpute: imputing dropout events in single cell RNA sequencing data. *BMC Bioinformatics* **19**: 220.
- Greenberg MVC, Bourc'his D. 2019. The diverse roles of DNA methylation in mammalian development and disease. *Nat Rev Mol Cell Biol* **20**: 590–607.

- Grosselin K, Durand A, Marsolier J, Poitou A, Marangoni E, Nemati F, Dahmani A, Lameiras S, Reyat F, Frenoy O, et al. 2019. High-throughput single-cell ChIP-seq identifies heterogeneity of chromatin states in breast cancer. *Nat Genet* **51**: 1060–1066.
- Hafemeister C, Satija R. 2019. Normalization and variance stabilization of single-cell RNA-seq data using regularized negative binomial regression. *Genome Biology* **20**: 296.
- He B, Chen C, Teng L, Tan K. 2014. Global view of enhancer-promoter interactome in human cells. *Proc Natl Acad Sci U S A* **111**: E2191–9.
- Hernando-Herraez I, Evano B, Stubbs T, Commere P-H, Jan Bonder M, Clark S, Andrews S, Tajbakhsh S, Reik W. 2019. Ageing affects DNA methylation drift and transcriptional cell-to-cell variability in mouse muscle stem cells. *Nat Commun* **10**: 4361.
- Hon GC, Rajagopal N, Shen Y, McCleary DF, Yue F, Dang MD, Ren B. 2013. Epigenetic memory at embryonic enhancers identified in DNA methylation maps from adult mouse tissues. *Nat Genet* **45**: 1198–1206.
- Hu Y, Huang K, An Q, Du G, Hu G, Xue J, Zhu X, Wang C-Y, Xue Z, Fan G. 2016. Simultaneous profiling of transcriptome and DNA methylome from a single cell. *Genome Biol* **17**: 88.
- Javierre BM, Burren OS, Wilder SP, Kreuzhuber R, Hill SM, Sewitz S, Cairns J, Wingett SW, Varnai C, Thiecke MJ, et al. 2016. Lineage-Specific Genome Architecture Links Enhancers and Non-coding Disease Variants to Target Gene Promoters. *Cell* **167**: 1369–1384 e19.
- Jones PA. 2012. Functions of DNA methylation: islands, start sites, gene bodies and beyond. *Nat Rev Genet* **13**: 484–492.
- Jung I, Schmitt A, Diao Y, Lee AJ, Liu T, Yang D, Tan C, Eom J, Chan M, Chee S, et al. 2019. A compendium of promoter-centered long-range chromatin interactions in the human genome. *Nat Genet* **51**: 1442–1449.
- Kapourani C-A, Sanguinetti G. 2016. Higher order methylation features for clustering and prediction in epigenomic studies. *Bioinformatics* **32**: i405–i412.
- Korsunsky I, Millard N, Fan J, Slowikowski K, Zhang F, Wei K, Baglaenko Y, Brenner M, Loh P-R, Raychaudhuri S. 2019. Fast, sensitive and accurate integration of single-cell data with Harmony. *Nat Methods* **16**: 1289–1296.
- Krueger F, Andrews SR. 2011. Bismark: a flexible aligner and methylation caller for Bisulfite-Seq applications. *Bioinformatics* **27**: 1571–1572.
- Lee D-S, Luo C, Zhou J, Chandran S, Rivkin A, Bartlett A, Nery JR, Fitzpatrick C,

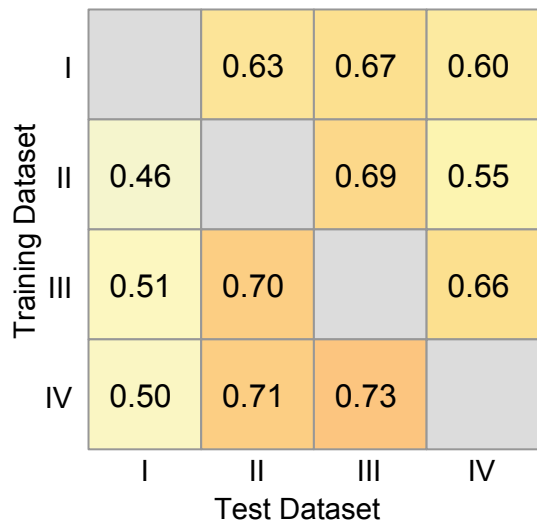
- O'Connor C, Dixon JR, et al. 2019. Simultaneous profiling of 3D genome structure and DNA methylation in single human cells. *Nat Methods* **16**: 999–1006.
- Liaw A, Wiener M. 2018. *Classification and regression based on a forest of trees using random inputs*. <https://cran.r-project.org/web/packages/randomForest/>.
- Lister R, Mukamel EA, Nery JR, Urich M, Puddifoot CA, Johnson ND, Lucero J, Huang Y, Dwork AJ, Schultz MD, et al. 2013. Global epigenomic reconfiguration during mammalian brain development. *Science* **341**: 1237905.
- Lister R, Pelizzola M, Downen RH, Hawkins RD, Hon G, Tonti-Filippini J, Nery JR, Lee L, Ye Z, Ngo Q-M, et al. 2009. Human DNA methylomes at base resolution show widespread epigenomic differences. *Nature* **462**: 315–322.
- Luo C, Keown CL, Kurihara L, Zhou J, He Y, Li J, Castanon R, Lucero J, Nery JR, Sandoval JP, et al. 2017. Single-cell methylomes identify neuronal subtypes and regulatory elements in mammalian cortex. *Science* **357**: 600–604.
- Luo C, Rivkin A, Zhou J, Sandoval JP, Kurihara L, Lucero J, Castanon R, Nery JR, Pinto-Duarte A, Bui B, et al. 2018. Robust single-cell DNA methylome profiling with snmC-seq2. *Nat Commun* **9**: 3824.
- Mendes-Moreira J, Soares C, Jorge AM. 2012. Ensemble approaches for regression: A survey. *Acm computing surveys*.
<https://dl.acm.org/doi/abs/10.1145/2379776.2379786>.
- Mo A, Mukamel EA, Davis FP, Luo C, Henry GL, Picard S, Urich MA, Nery JR, Sejnowski TJ, Lister R, et al. 2015. Epigenomic Signatures of Neuronal Diversity in the Mammalian Brain. *Neuron* **86**: 1369–1384.
- Mulqueen RM, Pokholok D, Norberg SJ, Torkenczy KA, Fields AJ, Sun D, Sinnamon JR, Shendure J, Trapnell C, O'Roak BJ, et al. 2018. Highly scalable generation of DNA methylation profiles in single cells. *Nat Biotechnol* **36**: 428–431.
- Nagano T, Lubling Y, Stevens TJ, Schoenfelder S, Yaffe E, Dean W, Laue ED, Tanay A, Fraser P. 2013. Single-cell Hi-C reveals cell-to-cell variability in chromosome structure. *Nature* **502**: 59–64.
- R Core Team. 2019. *R: A language and environment for statistical computing*.
<https://cran.r-project.org/>.
- Singh R, Lanchantin J, Robins G, Qi Y. 2016. DeepChrome: deep-learning for predicting gene expression from histone modifications. *Bioinformatics* **32**: i639–i648.
- Smallwood SA, Lee HJ, Angermueller C, Krueger F, Saadeh H, Peat J, Andrews SR, Stegle O, Reik W, Kelsey G. 2014. Single-cell genome-wide bisulfite sequencing for assessing epigenetic heterogeneity. *Nat Methods* **11**: 817–820.

- Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res* **15**: 1929–1958.
- Stadler MB, Murr R, Burger L, Ivanek R, Lienert F, Schöler A, van Nimwegen E, Wirbelauer C, Oakeley EJ, Gaidatzis D, et al. 2011. DNA-binding factors shape the mouse methylome at distal regulatory regions. *Nature* **480**: 490–495.
- Stevens TJ, Lando D, Basu S, Atkinson LP, Cao Y, Lee SF, Leeb M, Wohlfahrt KJ, Boucher W, O’Shaughnessy-Kirwan A, et al. 2017. 3D structures of individual mammalian genomes studied by single-cell Hi-C. *Nature* **544**: 59–64.
- Stuart T, Butler A, Hoffman P, Hafemeister C, Papalexi E, Mauck WM 3rd, Hao Y, Stoeckius M, Smibert P, Satija R. 2019. Comprehensive Integration of Single-Cell Data. *Cell* **177**: 1888–1902.e21.
- Tran HTN, Ang KS, Chevrier M, Zhang X, Lee NYS, Goh M, Chen J. 2020. A benchmark of batch-effect correction methods for single-cell RNA sequencing data. *Genome Biol* **21**: 12.
- Urich MA, Nery JR, Lister R, Schmitz RJ, Ecker JR. 2015. MethylC-seq library preparation for base-resolution whole-genome bisulfite sequencing. *Nat Protoc* **10**: 475–483.
- van Dijk D, Sharma R, Nainys J, Yim K, Kathail P, Carr AJ, Burdziak C, Moon KR, Chaffer CL, Pattabiraman D, et al. 2018. Recovering Gene Interactions from Single-Cell Data Using Data Diffusion. *Cell* **174**: 716–729.e27.
- Wagner F, Yan Y, Yanai I. 2017. K-nearest neighbor smoothing for high-throughput single-cell RNA-Seq data. *BioRxiv*.
<https://www.biorxiv.org/content/10.1101/217737v1.abstract>.
- Weber M, Hellmann I, Stadler MB, Ramos L, Pääbo S, Rebhan M, Schübeler D. 2007. Distribution, silencing potential and evolutionary impact of promoter DNA methylation in the human genome. *Nat Genet* **39**: 457–466.
- Welch JD, Kozareva V, Ferreira A, Vanderburg C, Martin C, Macosko EZ. 2019. Single-Cell Multi-omic Integration Compares and Contrasts Features of Brain Cell Identity. *Cell* **177**: 1873–1887.e17.
- Wolpert DH. 1992. Stacked generalization. *Neural Netw* **5**: 241–259.
- Zhu Q, Gao P, Tober J, Bennett L, Chen C, Uzun Y, Li Y, Howell ED, Mumau M, Yu W, et al. 2020. Developmental trajectory of prehematopoietic stem cell formation from endothelium. *Blood* **136**: 845–856.

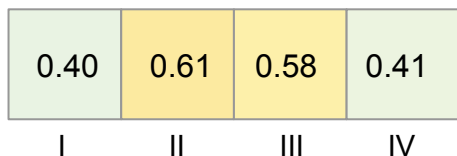


A

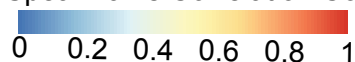
Ensemble Learner (EL)



Mean Promoter Demethylation (MPD)



Spearman's Correlation Coefficient



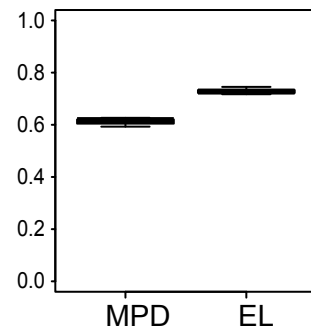
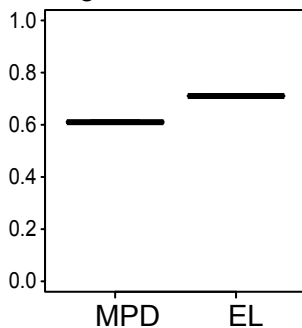
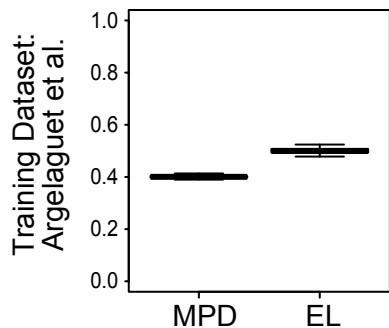
Datasets

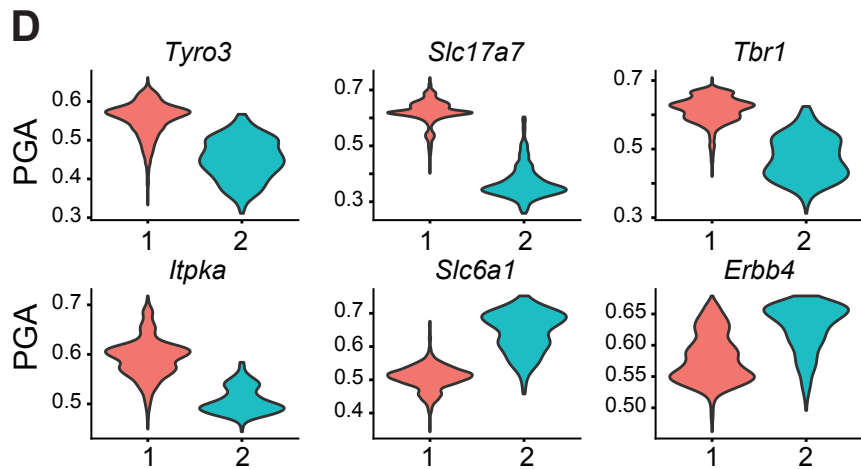
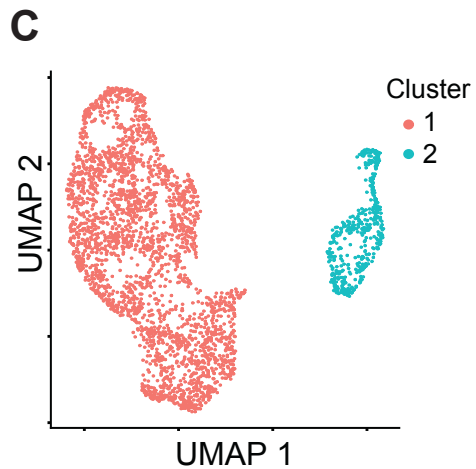
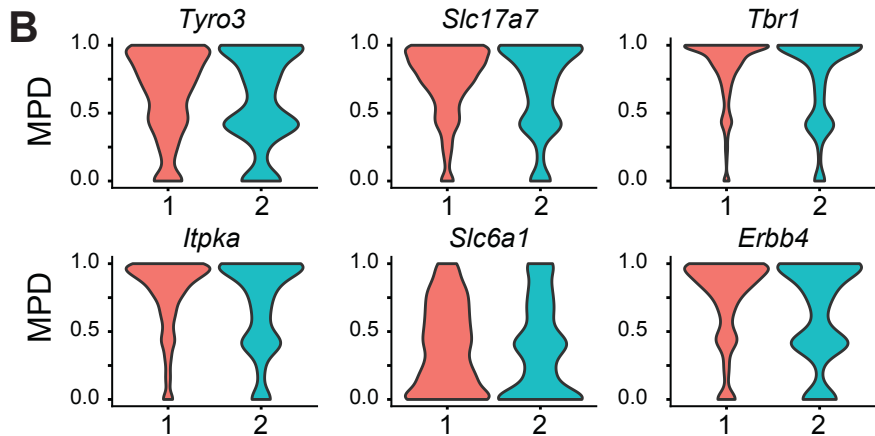
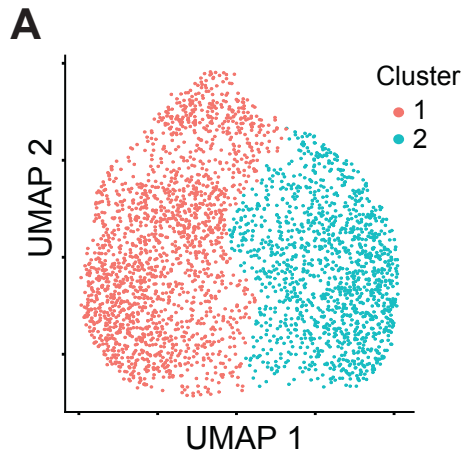
I Hernado-Herraez et al.

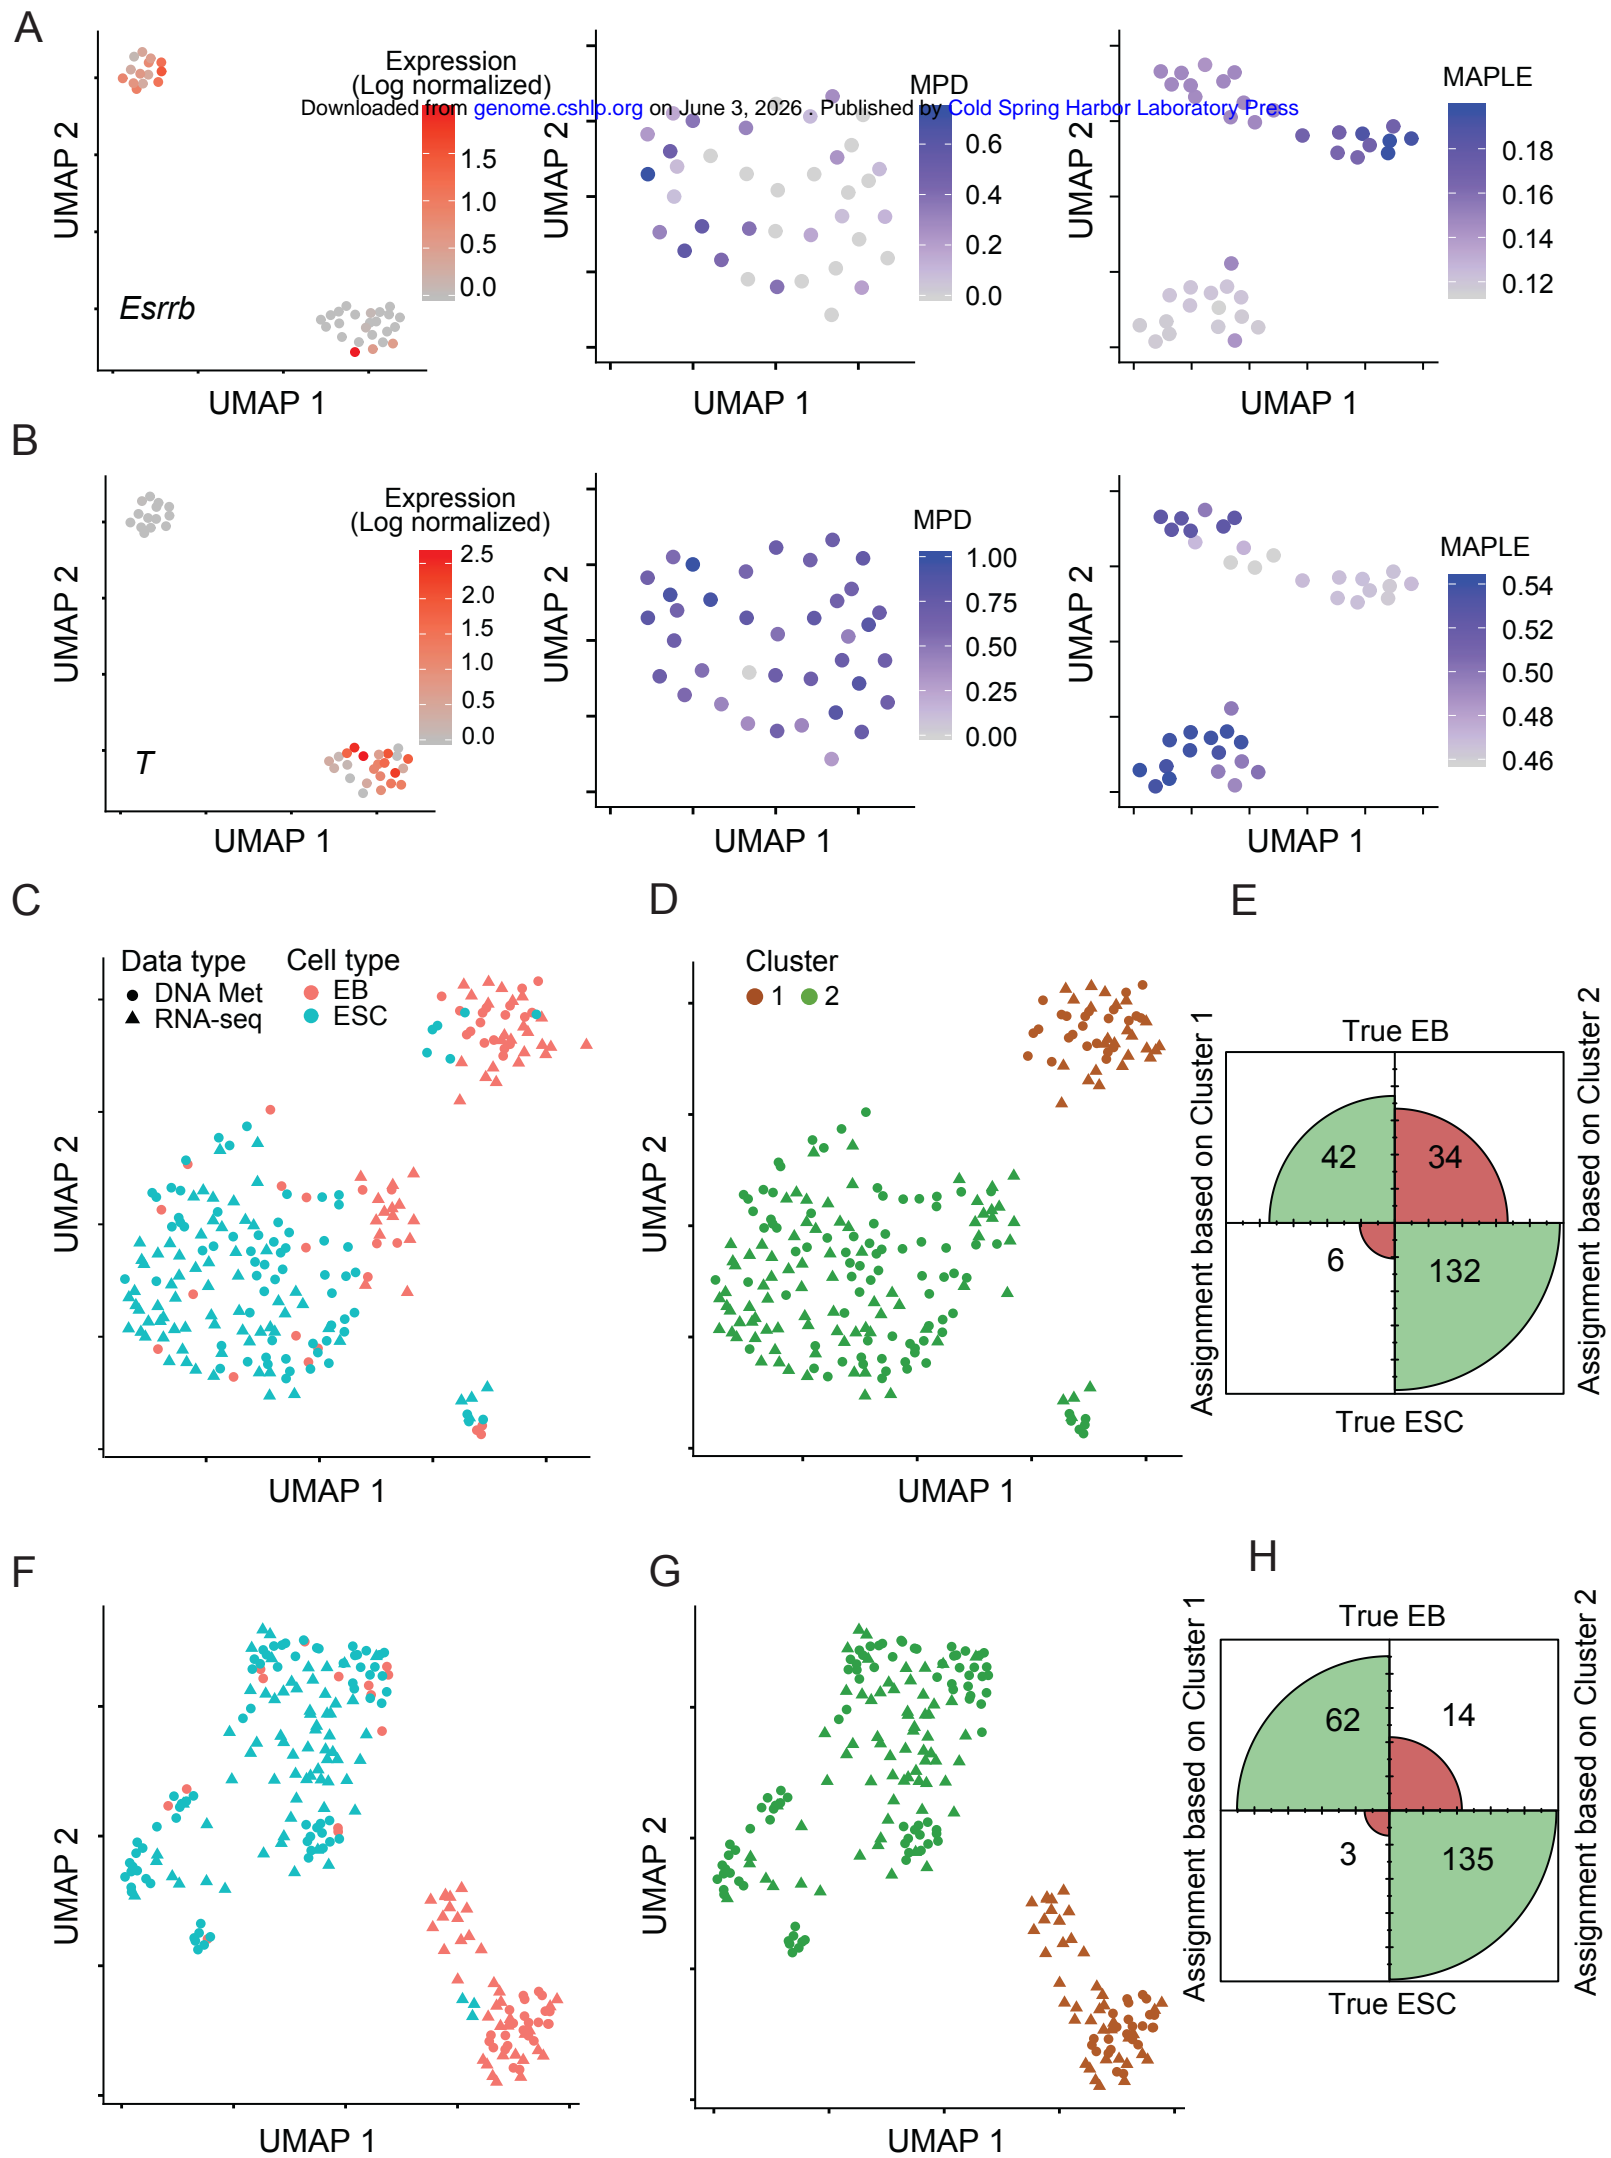
II Angermueller et al.

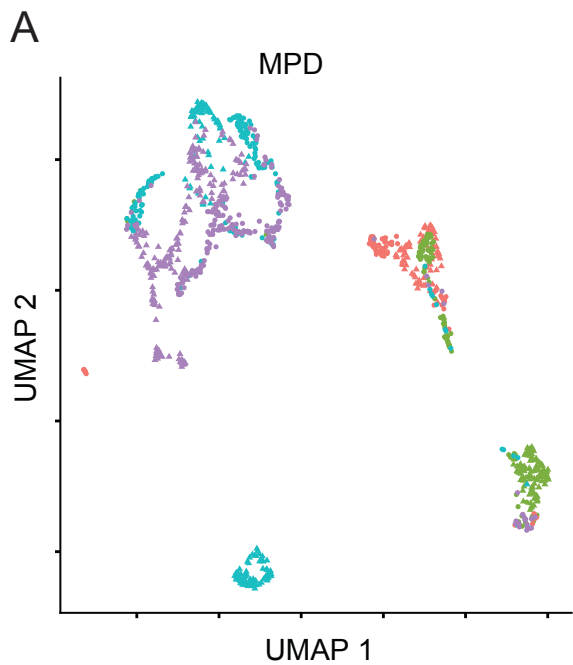
III Clark et al.

IV Argelaguet et al.

BTest Dataset:
Hernado-Herraez et al.Test Dataset:
Angermueller et al.Test Dataset:
Clark et al.







Embryonic Day

- E4.5
- E6.5
- E5.5
- E7.5

Data Type

- scDNA Meth
- ▲ scRNA

