SIMBA: SIngle-cell eMBedding Along with features

Huidong Chen[1, 2, 3], Jayoung Ryu[1, 2, 4], Michael E. Vinyard[1, 2, 3, 5], Adam Lerer[6]+, Luca Pinello[1, 2, 3]+

1. Molecular Pathology Unit, Center for Cancer Research, Massachusetts General Hospital, Charlestown, MA, USA
2. Department of Pathology, Harvard Medical School, Boston, MA, USA
3. Broad Institute of Harvard and MIT, Cambridge, MA, USA
4. Department of Biomedical Informatics, Harvard Medical School, Boston, MA, USA
5. Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA, USA
6. Facebook AI Research

+Corresponding author

Abstract

Recent advances in single-cell omics technologies enable the individual and joint profiling of cellular measurements including gene expression, epigenetic features, chromatin structure and DNA sequences. Currently, most single-cell analysis pipelines are cluster-centric, i.e., they first cluster cells into non-overlapping cellular states and then extract their defining genomic features. These approaches assume that discrete clusters correspond to biologically relevant subpopulations and do not explicitly model the interactions between different feature types. In addition, single-cell methods are generally designed for a particular task as distinct single-cell problems are formulated differently. To address these current shortcomings, we present *SIMBA*, a graph embedding method that jointly embeds single cells and their defining features, such as genes, chromatin accessible regions, and transcription factor binding sequences into a common latent space. By leveraging the co-embedding of cells and features, SIMBA allows for the study of cellular heterogeneity, clustering-free marker discovery, gene regulation inference, batch effect removal, and omics data integration. SIMBA has been extensively applied to scRNA-seq, scATAC-seq, and dual-omics data. We show that SIMBA provides a single framework that allows diverse single-cell analysis problems to be formulated in a unified way and thus simplifies the development of new analyses and integration of other single-cell modalities.

Introduction

Recent progress in single-cell molecular profiling technologies have dramatically advanced our ability to define cell types and states as well as discover key genes and regulatory regions in development and disease. We can now measure multiple modalities of individual cells. At the same time, the number of cells we are able to profile has increased rapidly. The emergence of single-cell multi-omics technologies allows for the measurements of multiple cellular layers, including genomics, epi-genomics, transcriptomics, and proteomics. Such assays have pioneered an avenue to better understand the interplay between these ‘omic’ layers and cell

states based on diverse genomic and molecular features including genes, regulatory elements, transcription factors. While single-cell multi-omic assays have quickly evolved towards the incorporation of more modalities with increasing resolution, harnessing their full potential poses significant computational challenges.

Many computational methods have been developed for single-cell single-modality analysis (e.g., scRNA-seq or scATAC-seq analysis) [1-4]. Common to these methods is a workflow with several standard steps, which typically include feature selection, dimension reduction, clustering, and differential feature detection. These "cluster-centric" analysis methods rely on accurately defined clustering solutions to discover meaningful and informative marker features. Unfortunately, clustering solutions may range widely within the space of the user-defined clustering resolution (number of clusters) and the chosen clustering algorithm. These parameters may markedly influence the resulting cluster assignment and clusters may not always correspond to the correct cell populations, thereby leading to inconsistent and potentially misleading biological annotations[5]. Although initial efforts have been made recently to develop clustering-free approaches to discover informative genes, they are specifically designed for extracting gene signatures [6, 7] or identifying perturbations between experimental conditions[8] from scRNA-seq data, and are therefore limited to single-modality and single-task analysis.

In addition to single-batch/modality analysis, approaches have also been proposed for multi-batch and cross-modality analysis, such as multimodal analysis (distinct cellular parameters are measured in the same cell)[9], batch correction (the same cellular parameter is measure in different batches) [10-12], and integration of multi-omics datasets (distinct cellular parameters are measured in different cells)[11, 12]. These approaches play a critical role in removing batch effects that confound true biological variation, improving the characterization of cell states by leveraging the unique strengths of each assay, and providing insights into the complex mechanisms of gene regulation. However, these tasks are formulated differently from those in single-batch/modality settings and thus require development of new dedicated analysis techniques. Also, while multiple types of cellular parameters (features) might be present, the relation between features cannot be exploited directly by most current methods. Furthermore, similar to single-batch/modality analysis methods, these methods identify marker features based on groups of cells obtained by clustering and therefore are limited to clustering solutions. Additionally, instead of directly identifying marker features in the integrated space, most batch correction/multi-omics integration methods need to first detect marker features in the uncorrected/unintegrated original space of each batch/modality independently, and then combine them, thus resulting in potentially inconsistent interpretations between batches/modalities.

To overcome the limitations in both single-batch/modality analysis and multi-batch/cross-modality analysis, we propose SIMBA (**SIngle-cell eMBedding Along with features**), a versatile single-cell embedding method that co-embeds cells and features into a shared latent space, in which the relation between cells and features or between two types of features (e.g., genes, peaks, or DNA sequences) can be assessed based on their proximity in the resulting embedding.

By formulating single-cell analyses as multi-entity graph embedding problems, SIMBA can be used to solve popular single-cell tasks in a unified framework that would otherwise require the development of distinct specialized approaches for each task. For each task, SIMBA constructs a graph, wherein differing entities (i.e., cells and features) are represented as nodes and relations between these entities are encoded as edges. Once the graph is constructed, SIMBA then applies a multi-entity graph embedding algorithm derived from social networking technologies as well as a Softmax-based transformation to embed the nodes/entities of the graph into a common low-dimensional space wherein cells and features can be analyzed based on their distance. We show that the SIMBA framework can perform many common and important single-cell analysis tasks, including: 1) dimensionality reduction techniques for studying cellular states; 2) clustering-free marker detection based on the similarity between single cells and features; 3) Single-cell multimodal analysis and the study of gene regulation; 4) batch correction and omics integration analysis as well as the simultaneous identification of marker features. SIMBA is adapted to these diverse analysis tasks by simply modifying how the input graph is constructed from the relevant single-cell data.

We extensively tested *SIMBA* in multiple scRNA-seq, scATAC-seq and dual-omics datasets covering popular single-cell tasks including scRNA-seq analysis, scATAC-seq analysis, multimodal analysis, batch correction, and multi-omics integration. We demonstrate that *SIMBA* jointly learns low-dimensional representations of both cells and features and thus enables the ability to simultaneously study cellular heterogeneity as well as proximity-based marker feature detection or gene regulation inference in a clustering-free way. We also demonstrate that SIMBA performs comparably to or better than current state-of-the-art methods specifically developed for each task.

We believe that SIMBA, as a broadly applicable approach for single cell omics study, will simplify the burden of developing methods for new single-cell tasks and measurements, while increasing interpretability of cellular mechanisms and functions.

Results

Overview of SIMBA

SIMBA is a single-cell embedding method with support for single- or multi- modality analyses that embeds cells and their associated genomic features into a shared latent space, generating interpretable and comparable embeddings of cells and features. It leverages recent graph embedding techniques that have been successful in modeling complex and hierarchical information present in natural languages, social networks, and other domains, as “knowledge graphs”. In our case, the graph encodes different components of cellular regulatory circuits and the relations between them.

SIMBA first encodes different types of entities such as cells, genes, open chromatin regions (peaks or bins), transcription factor (TF) motifs, and *k*-mers (short sequences of a specific length, *k*), into a single graph (**Fig. 1**, **Methods**) where each node represents an individual entity

and edges indicate relations between entities. For example, if a gene is expressed in a cell, an edge is created between the gene and cell. The weight of this edge is determined by the gene expression level. Similarly, an edge is added between a cell and a chromatin region if the region is open in this cell, or between a chromatin region and a TF motif if the TF motif is found in the region.

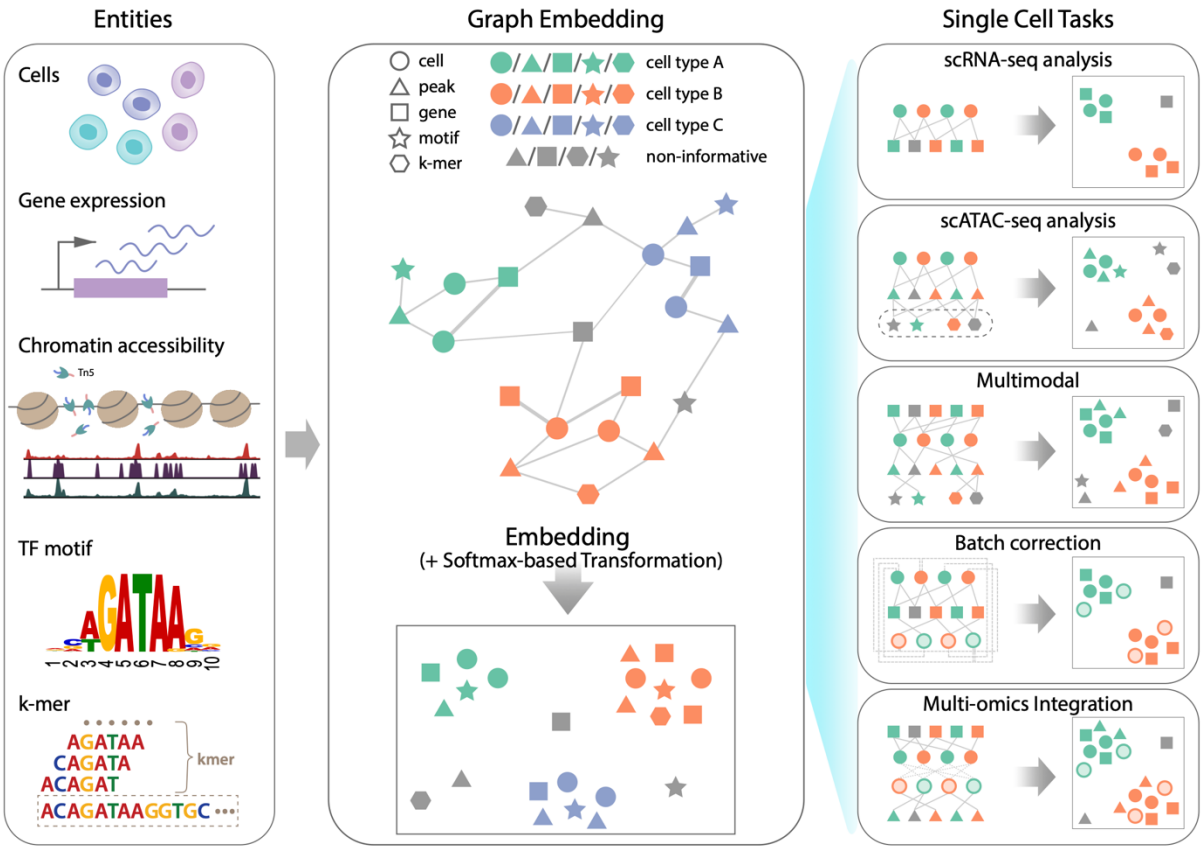

Figure1. SIMBA framework overview. SIMBA co-embeds cells and various features measured during single-cell experiments into a shared latent space to accomplish both common tasks involved in single-cell data analysis as well as tasks, which remain as open problems in single-cell genomics. **(Left)** Examples of possible biological entities may be encoded by SIMBA including cells, gene expression measurements, chromatin accessible regions, TF motifs, and k-mer sequences found in reads. **(Middle)** SIMBA embedding plot with multiple types of entities into a low-dimensional space. All entities represented as shapes (cell = circle, peak = triangle, gene = square, TF motif = star, k-mer = hexagon) are colored by relevant cell type (green, orange, and blue in this example). Non-informative features are colored dark grey. Within the graph, each entity is a node, and an edge indicates a relation between entities (e.g., a gene is expressed in a cell, a chromatin region is accessible in a cell, or a TF motif/k-mer is present within an open chromatin region, etc.). Once connected in a graph, these entities may be embedded into a shared low-dimensional space, with cell-type specific entities embedded in the same neighborhood and non-informative features embedded elsewhere. **(Right)** Common single-cell analysis tasks that may be accomplished using SIMBA.

Once the input graph is constructed, a low-dimensional representation of the graph nodes is then computed using an unsupervised graph embedding method (**Methods**). This graph embedding procedure leverages the PyTorch-BigGraph framework [13], which allows SIMBA to scale to millions of cells (**Methods**). The resulting joint embedding of cells and features not only reconstructs the heterogeneity of cells but also allows for the discovery of the defining features for each single cell without relying on a clustering solution, separating cell-type specific features from the non-informative features. In fact, the proximity between entity embeddings is informative of the potential importance of a feature to a cell and the interplay between features. When multiple types of features (e.g., transcriptomic and epigenetic features) are co-embedded, SIMBA provides an intuitive way to study gene regulation and the regulatory mechanisms underlying cellular differentiation and cell type formation.

Graph construction is inherently flexible, enabling SIMBA to be applied to a wide variety of single-cell tasks. In the following sections, we demonstrate the application of SIMBA to several popular single-cell tasks including scRNA-seq, scATAC-seq, multimodal analysis, batch correction and multi-omics integration (**Fig. 1**).

SIMBA enables simultaneous learning of cellular heterogeneity and individual-cell-level marker genes in scRNA-seq analysis

Single-cell RNA sequencing (scRNA-seq) is the most robust and widely used measurement to profile single cells. **Figure 2a** provides an illustrative overview of the SIMBA graph construction and the resulting low-dimensional embedding matrix of both cells and genes. Here we show how SIMBA enables simultaneous dimensionality reduction and clustering-free marker gene detection in scRNA-seq analysis. We applied SIMBA to a popular PBMCs dataset from 10x Genomics (**Supplementary Table 1**), which is used in the tutorials of both Scanpy[2]and Seurat[1] to illustrate their workflows. After the standard preprocessing steps including normalization and log-transformation, SIMBA discretizes the gene expression matrix into multiple gene expression levels (five levels, by default). The input graph is then constructed wherein two types of nodes – cells and genes are connected by edges that embody the relation between them and are weighted according to the corresponding multiple levels of gene expression. SIMBA then generates embeddings of these nodes through a graph embedding procedure (**Fig. 2a; Methods**).

We first visualized the SIMBA embeddings of cells using UMAP, which is a comparable output to other current single-cell analysis methods. Each of the previously assigned eight cell types, including B cells, megakaryocytes, CD14 monocytes, FCGR3A monocytes, dendritic cells, NK cells, CD4 T, and CD8 T cells, was clearly separated (**Fig. 2b**). We next applied UMAP to visualize the SIMBA embeddings of cells and genes together (**Fig. 2c**). The same set of marker genes used to annotate these cells from Scanpy[2] was highlighted on the UMAP plot. In addition, as a control, we also show the locations of two housekeeping genes *GAPDH* and *B2M*, which would not be expected to associate with any particular cell type. From the UMAP plot, we can see that SIMBA not only was able to embed major-cell-group specific genes to the correct locations

(e.g., *IL7R* was embedded into CD4T cells and *MS4A1* was embedded into B cells), but also was robust to rare-cell-group specific genes (e.g., *PPBP* was embedded into megakaryocytes). On the contrary, non-informative or non-cell-type specific genes such as the aforementioned housekeeping genes were embedded in the middle of all cell groups.

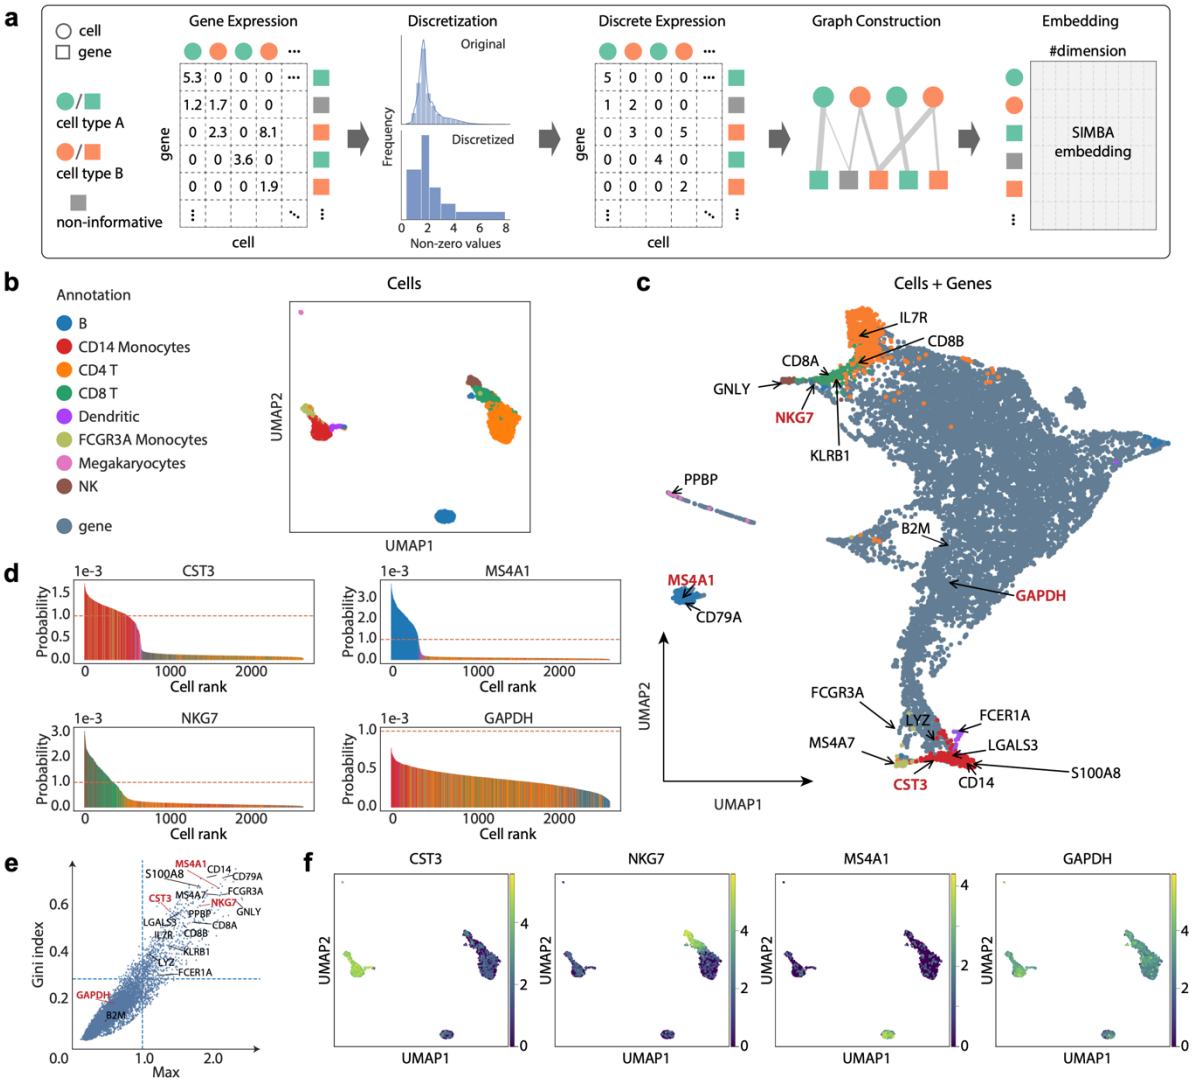

Figure 2. Single-cell RNA-seq analysis of the 10x PBMCs dataset using SIMBA. **(a)** SIMBA graph construction and embedding in scRNA-seq analysis. Biological entities including cells and genes are represented as shapes and colored by relevant cell types (green and orange). Non-informative genes are colored dark grey. Gene expression measurements for each cell are organized into a cell-by-gene matrix. These normalized non-negative observed values undergo discretization into five gene expression levels. Cells and genes are then assembled into a graph with nodes representing cells and genes, and edges between them representing different gene expression levels. This graph may then be embedded into a lower dimensional space resulting in a #entities x #dimension (by default, 50) SIMBA embedding matrix. **(b)** UMAP visualization of SIMBA embeddings of cells colored by cell type. **(c)** UMAP visualization of SIMBA embeddings of

cells and genes. Cells are colored according to cell type as defined in b. Genes are colored slate blue. Cell-type-specific marker genes and housekeeping genes recovered by Scanpy are indicated with text and arrows. Genes highlighted in red are shown in **d**, **e**, and **f**. **(d)** SIMBA barcode plots of genes *CST3*, *MS4A1*, *NKG7*, and *GAPDH*. The x-axis indicates the ordering of a cell as ranked by the probability for each cell to be associated with a given gene. The y-axis describes the probability. The sum of probability over all cells is equal to 1. Each cell is one bar and colored according to cell type as defined in b. **(e)** SIMBA ranking of genes based on the proposed metrics. All the genes are plotted according to the Gini index against max score. The same set of genes as in **c** are annotated. **(f)** UMAP visualization of SIMBA embeddings of cells colored by gene expression of (left to right): *CST3*, *NKG7*, *MS4A1*, and *GAPDH*.

In addition to visualizing all the entities at once using UMAP, SIMBA also provides a feature-specific 'barcode plot' to visualize the estimated probability of assigning a feature to a cell (**Fig. 2d, Methods**). The barcode plots in **Figure 2d** offer a rank-ordered probability of a given gene being associated with each cell (colored by cell type) where the total probability over all cells sums to one. An imbalance in probability indicates cell-type-specific association of a gene to a sub-population of cells, whereas a uniform probability distribution indicates a non-cell-type-specific gene. **Figure 2d** displays barcode plots for four genes, which are correspondingly highlighted in red, in **Figure 2c**. Three of these genes are commonly used marker genes for identifying subpopulations within PBMCs datasets (*CST3* for monocytes and dendritic cells, *MS4A1* for B cells, and *NGK7* for NK and CD8T cells). On the contrary, *GAPDH* is a housekeeping gene expressed in all cell types. For marker genes, we observed a clear excess in the probability of assigning each gene to their respective cell types. Conversely, for *GAPDH*, we observed a more uniform distribution with much lower probability of associating that gene with the top-ranked cells. Hence, SIMBA barcode plots serve as an informative way of visualizing gene expression patterns by showing the cell assignment probability distribution.

In addition, SIMBA also provides several quantitative metrics, including max value, Gini index, standard deviation, and entropy, to assess cell type specificity of various features (**Methods**). As an example, by inspecting the gene metric plot of max value (a higher value indicates higher cell-type specificity) vs Gini index (a higher value indicates higher cell-type specificity), we observed that the marker genes (e.g., *CST3*, *NKG7*, *MS4A1*) fall in the upper right corner, as opposed to housekeeping genes (e.g., *GAPDH*) in the lower left corner (**Fig. 2e**). Similar separation is observed with other metrics (**Supplementary Fig. 1b**). The cell type specificity of the selected marker genes was further confirmed by visualizing their expression pattern on UMAP plots (**Fig. 2f and Supplementary Fig. 1c**), accompanied by SIMBA barcode plots (**Supplementary Fig. 1d**).

To demonstrate that SIMBA provides a more robust means of detecting marker genes than the statistical-tests-based methods implemented by tools such as Scanpy and Seurat, we compared top 600 marker genes identified by SIMBA (based on max value and Gini index) with those identified by Scanpy (based on the Wilcoxon rank-sum test, one of the statistical tests recommended in Scanpy's tutorial) (**Supplementary Fig. 2a**). Upon comparison, we observed

that nearly half of the marker genes discovered by SIMBA overlap with the marker genes identified by Scanpy (**Supplementary Fig. 2a**). However, on inspection of the top non-overlapping marker genes, genes identified by SIMBA are found to be enriched only within certain groups of cells (**Supplementary Figs. 2b and 2c**) while genes identified by Scanpy but not by SIMBA include the housekeeping gene *B2M* and multiple ribosomal protein genes (e.g., *RPS3* and *RPS6*) that are expressed ubiquitously in all cell types (**Supplementary Figs. 2b and 2d**). In addition, we noticed that some of the identified genes by SIMBA can be potentially recovered using different statistical tests proposed in Scanpy. For example, IL7R was identified only by using the t-test and FCER1A was identified only by using the Wilcoxon rank-sum test as also noted in the Scanpy's tutorial. This example illustrates the limitations of a cluster-centric approach combined with multiple statistical testing methods. In contrast, SIMBA successfully identified both *IL7R* and *FCER1A* as informative genes with a single procedure and without clustering the cells (**Fig. 2e and Supplementary Fig. 1b**).

Lastly, we showed that SIMBA does not require variable gene selection, which is an essential step in standard scRNA-seq pipelines such as Seurat or Scanpy. SIMBA produces very similar embeddings for cells and genes with and without variable gene selection (**Supplementary Fig. 2e**), though we observed that variable gene section does improve efficiency of the training procedure.

SIMBA enables simultaneous characterization of cellular states and cis-regulatory elements by jointly modeling accessible sites and DNA sequences in scATAC-seq analysis

As one of the most popular single-cell epigenomic techniques, single-cell assay for transposase-accessible chromatin using sequencing (scATAC-seq) has been widely used to profile regions of open chromatin and identify functional *cis*-regulatory elements such as enhancers and active promoters. In scATAC-seq, cells are characterized by different types of features [14], such as regions of accessible chromatin ("peaks" or "bins") and *cis*-regulatory elements (DNA sequences) within these accessible regions including transcription factor (TF) motifs or *k*-mers. Unlike existing methods that use only peaks/bins or the sequence content found within them, in SIMBA both types of features can be encoded into a single graph. When only peaks/bins are used, SIMBA constructs a graph with nodes representing cells and chromatin regions (peaks or bins) and edges indicating the accessibility of the chromatin regions in cells (**Fig. 3a**). When the DNA sequences for chromatin regions are available, SIMBA can also encode DNA sequences including TF motifs and *k*-mers into the graph by adding edges between these entities as nodes and the existing chromatin region nodes. The edges in this case indicate the presence of TF motifs/*k*-mers within these chromatin-accessible regions. Through the embedding procedure, SIMBA generates embeddings of cells along with peaks and DNA sequences (**Methods**). Therefore, SIMBA enables dimensionality reduction by leveraging both chromatin accessible regions and cis-regulatory sequences. Simultaneously, it highlights the cell-type-specific open chromatin regions and regulatory DNA sequences in a clustering-free way.

To demonstrate the value of SIMBA embeddings for scATAC-seq analysis, we first applied SIMBA to a scATAC-seq data of 2,034 human hematopoietic cells with FACS-characterized cell

types[15](**Supplementary Table 1**). For the embeddings of cells alone, as shown using UMAP (**Fig. 3b**), SIMBA accurately embeds cells such that cells belonging to distinct cell types are visually separated. In addition to cells, SIMBA can also embed various types of features. The UMAP plot in **Figure 3c** highlights how distinct features from both positional (peaks) as well as sequence-content (TF motifs and k-mers) information are embedded together based on their biological relations. Notably, these highlighted features that are embedded within each cell type all have high cell-type specificity scores (shown in the upper right part of SIMBA metric plots in **Figure 3d**).

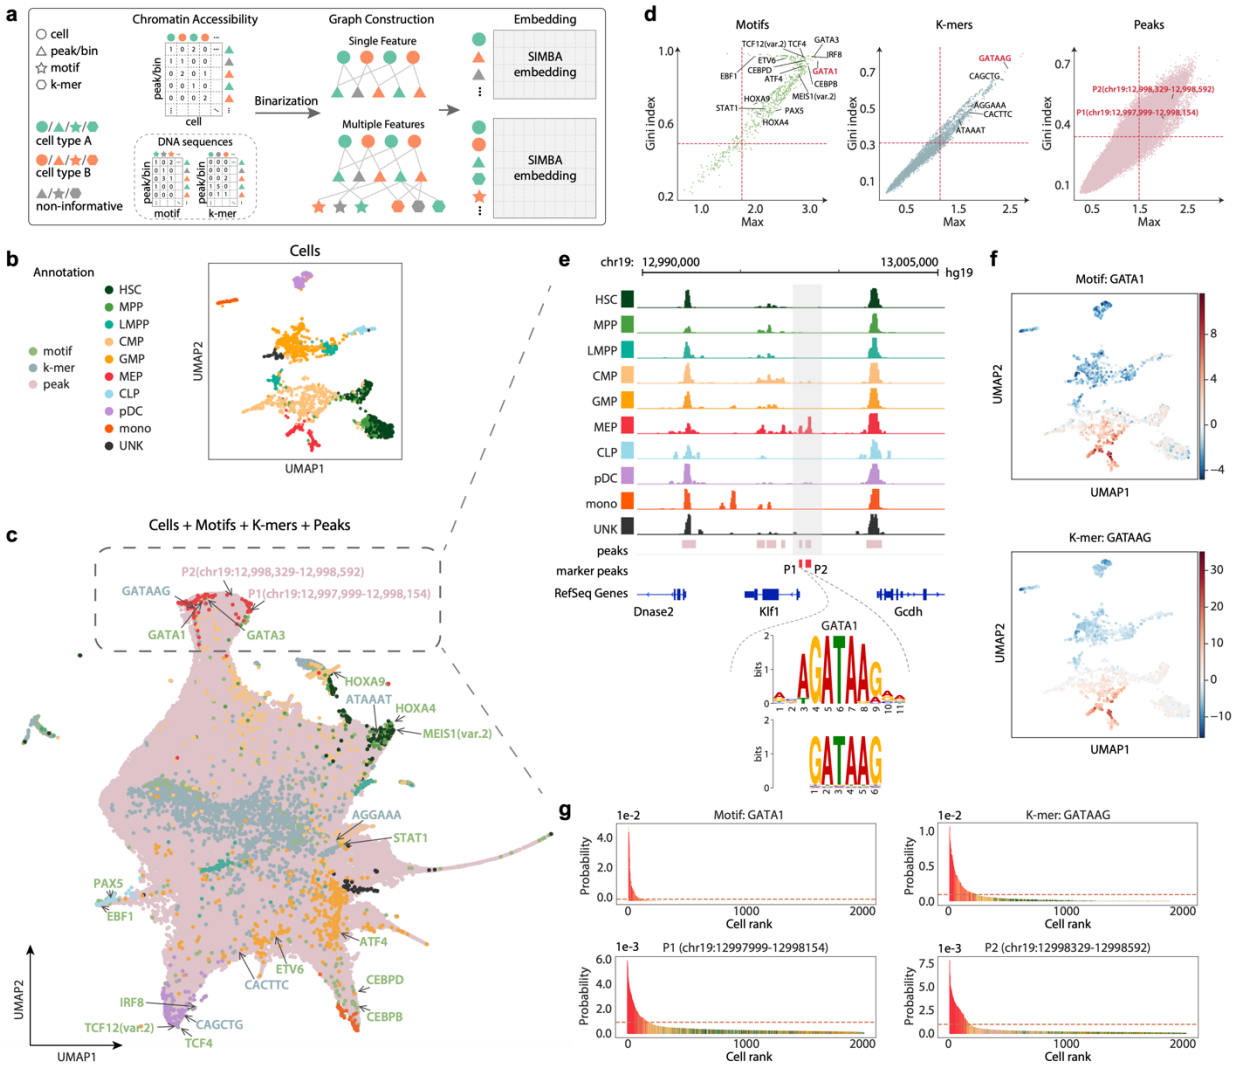

Figure 3. Single-cell ATAC-seq analysis of the human hematopoiesis dataset using SIMBA. **(a)** SIMBA graph construction and embedding in scATAC-seq analysis. Biological entities including cells, peaks/bins, TF motifs, k-mers are represented as shapes and colored by relevant cell types (green and orange). Non-informative features are colored dark grey. Cells and chromatin accessible features (peaks / bins) are organized into a cell x peaks / bins matrix. When sequence information (TF motif or *k*-mer sequence) within these regions is available, they can be organized

into two sub-matrices to associate a TF motif or k-mer sequence with each peak/bin. These constructed feature matrices are then binarized and assembled into a graph. When a single feature (chromatin accessibility) is used, the graph encodes cells and peaks/bins as nodes. When multiple features (both chromatin accessibility and DNA sequences) are used, this graph may then be extended with the addition of TF motifs and k-mer sequences as nodes connected. Finally, SIMBA embeddings of these entities are generated through a graph embedding procedure. **(b)** UMAP visualization of SIMBA embeddings of cells colored by cell type. **(c)** UMAP visualization of SIMBA embeddings of cells and features including TF motifs, k-mers, and peaks. Cells are colored by cell type while motifs, k-mers, and peaks are colored green, blue, and pink, respectively. Cell type specific features that are embedded near their corresponding cell types are indicated as the text labels (colored according to feature type) with arrows. **(d)** SIMBA metric plots of TF motifs, k-mers, and peaks. Cell-type specific features annotated in **(c)** are highlighted. **(e)** Genomic tracks of aligned scATAC-seq fragments, separated and colored by cell type. Two marker peaks P1 and P2 in red are shown beneath the alignment. Within the peak P1, k-mer GATAAG and its resembling GATA1 motif logo are highlighted. **(f)** UMAP visualization of SIMBA embeddings of cells colored by TF activity scores of the GATA1 motif and k-mer GATAAG enrichment. **(g)** SIMBA barcode plots of the GATA1 motif, k-mer GATAAG, and the two peaks P1 and P2. Cells are colored according to cell type labels described above. Dotted red line indicates the same cutoff used in all four plots.

This analysis led to several key findings in human hematopoietic differentiation.

First, SIMBA analysis identified key master regulators of hematopoiesis. As highlighted in **Fig. 3c**, we observed that motifs of previously reported TFs were embedded near their respective cell types in UMAP. For example, the GATA1 and GATA3 motifs are close to megakaryocyte-erythroid progenitor (MEP) cells[16]; the PAX5 and EBF1 motifs are close to common lymphoid progenitor (CLP) cells[17]; the CEBPB and CEBPD motifs are close to monocytes (mono)[18].

Second, SIMBA analysis identified an unbiased set of DNA sequences, i.e., *k*-mers, that represent important TF binding motifs involved in hematopoiesis, enabling *de novo* motif discovery. We observed that these *k*-mers were embedded near their resembling TF binding motifs and relevant cell subpopulations (**Fig. 3c** and **3e**, **Supplementary Fig. 3b**). For example, the DNA sequence CAGCTG is embedded in plasmacytoid dendritic cells (pDCs), and this sequence matches the TCF12 binding motif, which controls dendritic cell lineage specification. To further illustrate the interpretability of the SIMBA embeddings of TF motifs and k-mers, we calculated per-cell TF activity scores[19] (high-variance TF motifs/*k*-mers) and visualized them on SIMBA embeddings of cells . As shown in **Figure 3f**, the GATA1 TF motif and k-mer GATAAG that were both embedded in MEP cells by SIMBA, also showed high-level activity in MEP cells. The consistency between SIMBA embedding and TF activity was observed for most of other TF motifs and *k*-mers as well (**Supplementary Fig. 3a, 3b**).

Third, SIMBA analysis identified differential accessible regions that may mediate cell-type specific gene regulation. For example, the two peaks at chr19:12997999-12998154 (P1) and chr19:12998329-12998592 (P2) that were embedded within MEP cells were almost exclusively

observed in MEP cells on KLF1 genome track (**Fig. 3e**). Interestingly, P1, upstream of *KLF1*, contains the *k*-mer GATAAG that matches the GATA1 binding motif, while transcription factor GATA1 is known to regulate the gene *KLF1* and plays a pivotal role in erythroid cell and megakaryocyte development[20]. Therefore, by embedding these MEP-cell-related regulatory elements into the neighborhood of MEP cells, SIMBA demonstrates a new analytic way for studying cell-type-specific epigenetic landscape of the genome. To further validate the differential accessible regions identified by SIMBA, we randomly picked 100 peaks within each cell type in SIMBA co-embedding space. From the heatmap of chromatin accessibility, we clearly see that the peaks embedded in the neighborhood of each cell type by SIMBA shows strong cell-type specificity and this is robust to the number of cells within each cell type (**Supplementary Fig. 3c**).

Although SIMBA strongly differs from existing computational methods for scATAC-seq analysis in that SIMBA enables the co-embedding of cells and features, we still compared the resulting SIMBA embeddings of cells with these specialized methods by their ability to distinguish cell types. Our analyses show that SIMBA performs as well as or better than the methods evaluated, further demonstrating the wide utility of SIMBA (**Supplementary Figs. 4 and 5; Supplementary Note 1**).

SIMBA enables simultaneous learning of cellular heterogeneity and gene regulatory circuits from integrated analysis of single-cell multimodal data

scRNA-seq and scATAC-seq are two of the most widely adopted single-cell sequencing technologies, but they are limited to measuring only a single aspect of cell state at a time. To improve our ability to interrogate cellular states, several single-cell dual-omics technologies have been recently developed [21-24] to jointly profile transcriptome and chromatin accessibility within the same individual cells, therefore providing the potential to correlate gene expression with accessible regulatory elements and further delineate the yet elusive principles of gene regulation.

In this section, we demonstrate how SIMBA enables simultaneous learning of cellular heterogeneity and gene regulatory circuit from single-cell multiome data. We applied SIMBA to three recent single-cell dual-omics technologies: SHARE-seq[22], SNARE-seq[21], and a multiome PBMCs dataset from 10x Genomics (**Supplementary Table 1**).

Figure 4a illustrates the procedure of graph construction and generation of the final SIMBA embedding matrix. Briefly, for scRNA-seq, the gene expression matrix is discretized to generate different levels of gene expression. For scATAC-seq, both the chromatin accessibility matrix and motif/*k*-mer match matrix are binarized. In this graph, there are five entity (node) types, including cells, genes, peaks, motifs, and *k*-mers. For scRNA-seq, an edge indicates whether a gene is expressed in a cell and its weight indicates the gene expression level (five levels, by default). For scATAC-seq, an edge indicates whether a peak is present in a cell or if a TF motif/*k*-mer is present within a peak.

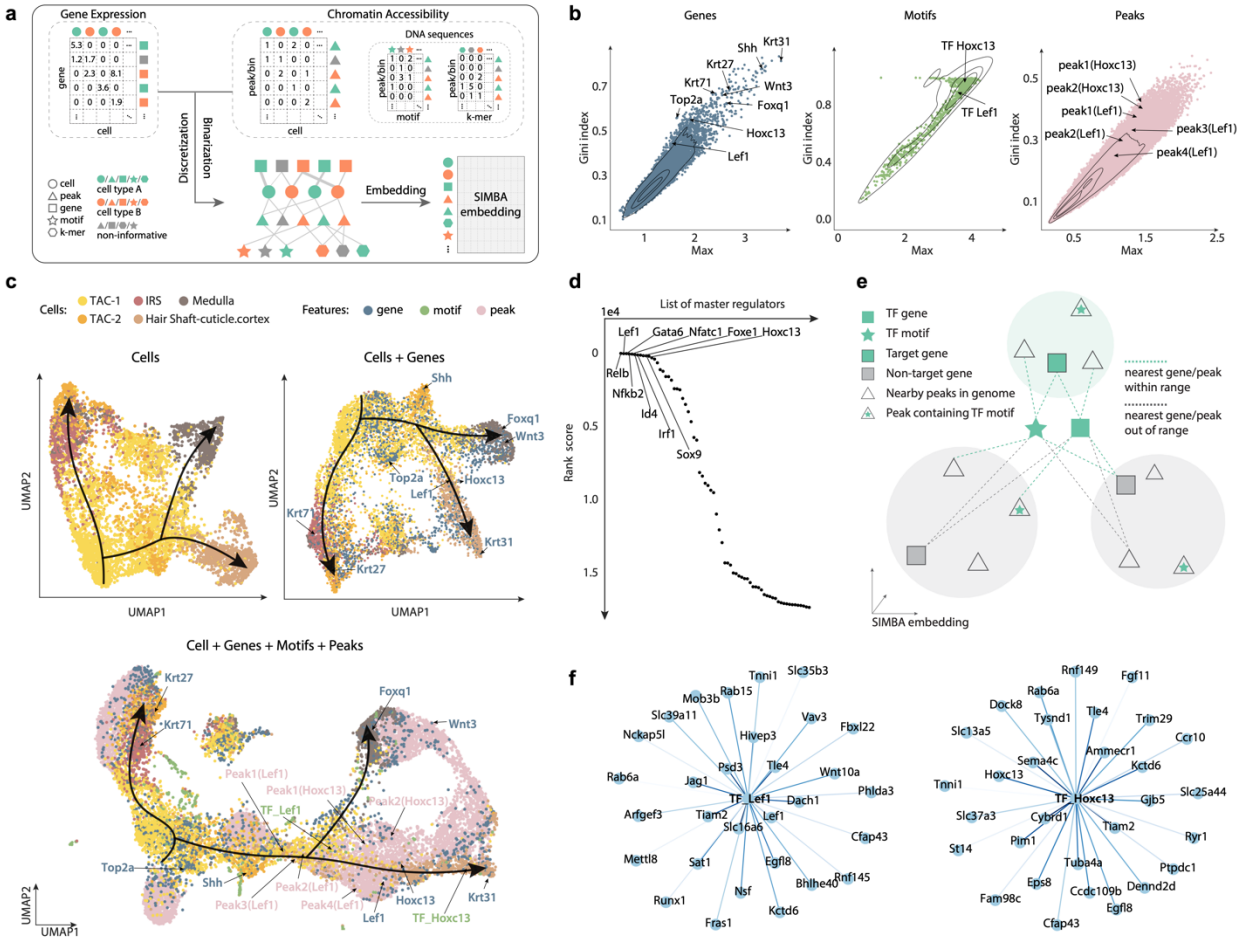

Figure 4. Multimodal analysis of the SHARE-seq hair follicle dataset using SIMBA. **(a)** SIMBA graph construction and embedding in multimodal analysis. Overview of SIMBA's approach to multimodal (scRNA-seq + scATAC-seq) data analysis. **(b)** SIMBA metric plots of genes, TF motifs, and peaks. All these features are plotted according to the Gini index against max score. Cell-type specific genes, TF motifs, and peaks are highlighted. **(c)** UMAP visualization of SIMBA embeddings of cells (Top-left), cells and genes (Top-right), and cells along with genes, TF motifs, and peaks (Bottom). **(d)** Ranked scatter plot of candidate master regulators as identified by SIMBA. **(e)** Schematic description of SIMBA's strategy for identifying target genes given a master regulator. **(f)** Top 30 target genes of transcription factors Lef1 and Hoxc13 as inferred by SIMBA.

To demonstrate the usefulness and versatility of the SIMBA embeddings, we analyzed the cell populations undergoing hair follicle differentiation from mouse skin profiled with SHARE-seq.

First, we used SIMBA to assess the cell-type specificity of different types of features, including genes, TF motifs, and peaks (**Fig. 4b, Methods**). As shown in **Figure 4b**, genes (e.g., *Lef1* and *Hoxc13*), which are associated with hair follicles each have relatively higher max values and Gini index scores. Similarly, TF motifs and peaks proximal to the genomic loci of these genes also

410 score in the upper right quadrant of the metric plots. SIMBA's cell-type specificity metrics
411 successfully revealed the key genes and regulatory factors important to the hair follicle
412 differentiation process.

413 Next, we visualized and interrogated the SIMBA embeddings of 1) cells; 2) cells and genes; and
414 3) cells together with genes, TF motifs, and peaks. **Figure 4c** shows the UMAP visualization of
415 the SIMBA embeddings of cells and informative features selected based on SIMBA metric plots.
416 The UMAP visualization of SIMBA embeddings of cells and the full set of features was also
417 performed (**Supplementary Fig. 6a**). However, we observed that with the SIMBA embeddings of
418 cells and the informative features selected by SIMBA, UMAP plots show more visible cell
419 subgroups without obstruction by entities that may dominate the space given their larger
420 number (e.g., peaks). SIMBA embeddings of cells were able to reveal the three fate decisions
421 from transit-amplifying cells (TACs), including inner root sheath (IRS), medulla, and
422 cuticle/cortex. SIMBA embeddings also uncovered important genes and regulatory factors
423 along the hair follicle differentiation trajectories. For example, the marker genes *Krt71*, *Krt31*,
424 and *Foxq1* were embedded into their corresponding cell types: IRS, cuticle/cortex, and medulla,
425 respectively. The Lef1 motif was embedded into the beginning of medulla and cuticle/cortex
426 lineages while the Hoxc13 motif was embedded into the late stage of cuticle/cortex
427 differentiation. Peaks near the *Lef1* and *Hoxc13* loci were also embedded into the nearby
428 regions of these genes and motifs, as expected. To show the robustness of SIMBA embedding,
429 we also performed the single-modality analyses within the SHARE-seq dataset, separating the
430 scRNA-seq and scATAC-seq components. With the consistent embedding results of cells and
431 features as in multimodal analysis, we further demonstrated that SIMBA embedding procedure
432 is robust to the type and number of features encoded in the input graph (**Supplementary Fig.
433 6b**,**6c**). Each marker gene was further validated using the UMAP plots with cells colored by gene
434 expression as well as using the SIMBA barcode plots. The two aforementioned TF motifs and
435 their respective peak sets were also validated visually using SIMBA barcode plots. As expected,
436 in SIMBA barcode plots these marker features clearly showed an imbalanced distribution with
437 much higher probabilities in the correct cell types (**Supplementary Fig. 7a-d**).

438 Further, we demonstrated that the SIMBA co-embedding space of cells and features provides
439 the potential to identify master regulators of differentiation and infer their target regulatory
440 genes. To define a master regulator *a priori*, we postulate that both its TF motif and TF gene
441 should be cell-type specific and given that active gene regulation involves the expression of TF
442 and accessibility of its binding sites, TF motif and TF gene should be embedded closely in the
443 shared latent space. Extending this logic to identify putative master regulators, we assessed the
444 cell-type-specificity of TF motifs and genes based on SIMBA metrics and ranked all potential
445 master regulators based on the distance between the TF motif and the respective TF gene in
446 the shared SIMBA embedding space (**Methods**). SIMBA successfully identified previously
447 described master regulators such as Lef1, Gata6, Nfatc1, and Hoxc13. as the top master
448 regulators related to lineage commitment in mouse skin (**Fig. 4d, Supplementary Table 2**). To
449 infer the target genes of a given master regulator, we postulate that in the shared SIMBA
450 embedding space, 1) the target gene is close to both the TF motif and the TF gene; 2) the
451 accessible regions (peaks) near the target gene loci must be close to both the TF motif and the

452 target TF gene. Resting on these assumptions of *cis*-regulatory dynamics, the inference of target
453 genes was performed by calculating the distance between target gene candidates and the
454 respective TF motif and gene. In addition, nearby peaks around the target gene's locus and the
455 presence of TF motif in these nearby peaks are also considered (**Fig. 4e, Methods**). The top 30
456 target genes of TF Lef1 and TF Hoxc13 inferred by SIMBA are shown respectively (**Fig. 4f,**
457 **Supplementary Fig. 7e**). The full list of ranked target genes is provided in **Supplementary Table**
458 **3**. Notably, our approach recovered targets genes that were also reported in the original
459 study[22]. For example, genes *Lef1*, *Jag1*, *Hoxc13*, *Gtf2ird1* are regulated by the TF Lef1, while
460 genes *Cybrd1*, *Hoxc13*, *St14* are regulated by the TF Hoxc13.

461 In addition to SHARE-seq, we also applied SIMBA to another two dual-omics datasets, the
462 mouse cerebral cortex dataset profiled by SNARE-seq[21] (**Supplementary Fig. 8**) and the
463 multiome PBMCs dataset from 10x Genomics (**Supplementary Fig. 9**). By validating the
464 embeddings of cells and features with given cell type labels (**Supplementary Fig. 8a and Fig.9a**),
465 marker genes from the original study (**Supplementary Fig. 8a,b,d and Fig.9a,b,d**), and
466 differentially accessible chromatin regions (**Supplementary Fig. 8c and Fig.9c**), we further
467 demonstrate the suitability of SIMBA for multimodal analysis.

468

469 **SIMBA enables simultaneous batch correction and clustering-free marker gene detection**

470

471 Efforts to collect data from single cells has grown to the level of consortia that span multiple
472 institutions with the hopes of finely mapping and characterizing specific tissues. This has
473 brought with it an increased demand for analysis methods that are capable of negating
474 technical covariates inherent to multi-batch data collection, including experimental replicate
475 identity, sample preparation, and sequencing platform. Batch correction that removes the
476 effects of technical covariation while preserving true biological signals is required prior to
477 downstream analysis [25, 26]. Unlike existing methods that require clustering and the merging of
478 marker genes detected in the original gene expression space of each batch, SIMBA enables
479 simultaneous batch effect removal and cell-type-specific marker gene detection in the same
480 integrated space without clustering.

481

482 We first demonstrate that SIMBA readily corrects batch effects and produces joint embeddings
483 of cells and genes across multiple datasets with different sequencing platforms and cell type
484 compositions. Thus, while previous methods primarily rely on specialized tools for batch
485 correction before downstream clustering and marker detection, SIMBA works as a stand-alone
486 package obviating the need for prior input data correction when applied to multi-batch scRNA-
487 seq dataset. In SIMBA, batch correction is accomplished by encoding multiple scRNA-seq
488 datasets into a single graph (**Fig. 5a**). Cells in different batches are linked to genes as in the
489 previously described scRNA-seq graph construction. Here, the gene nodes are shared between
490 the cell nodes of different batches. In addition to the experimentally measured edges, batch
491 correction is further enhanced through computationally inferred edges drawn between similar
492 cell nodes across datasets using a truncated randomized singular value decomposition (SVD)-
493 based procedure. SIMBA then generates the embeddings of all nodes including cells of each

batch and genes from the resulting graph, which may then be visualized using tools such as UMAP as in the analyses presented in the sections above (**Methods**).

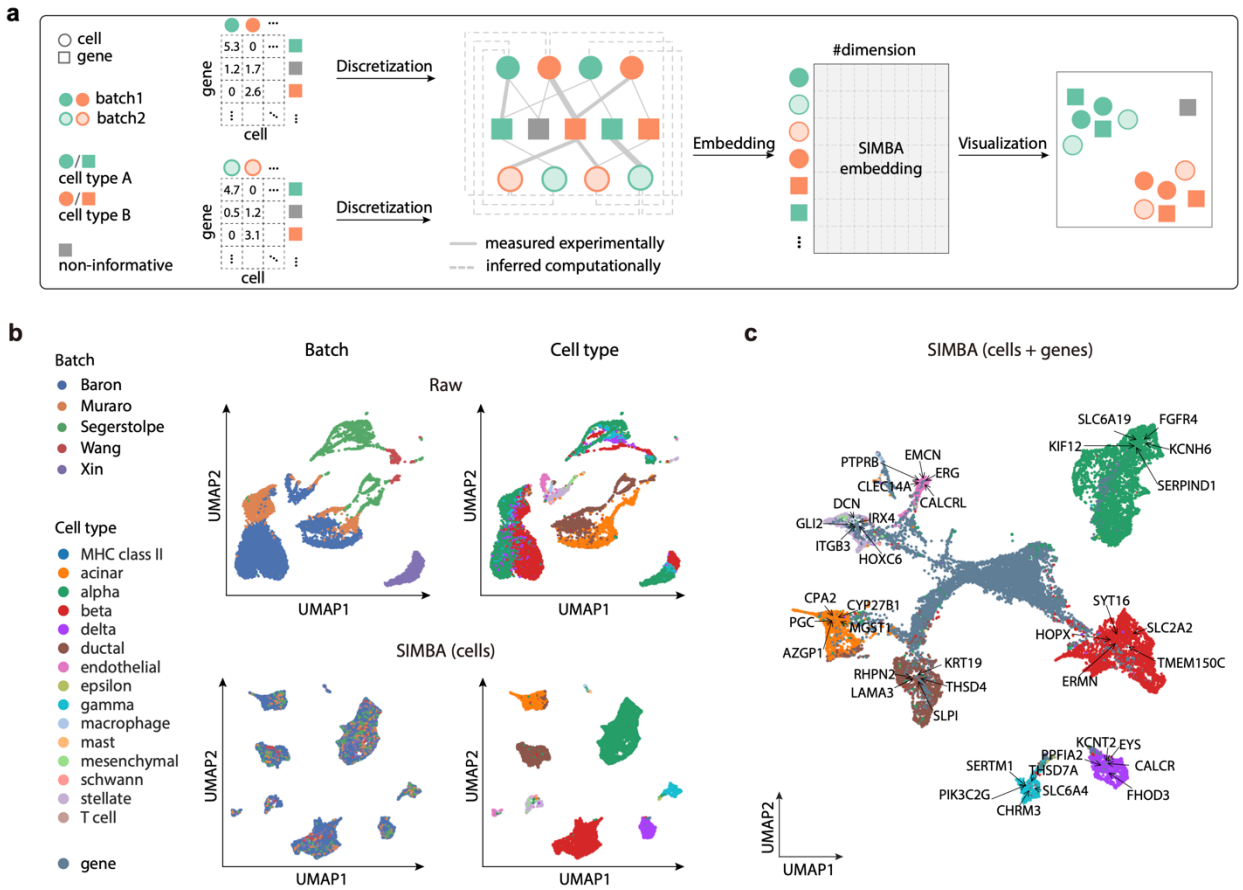

Figure 5. Batch correction analysis of scRNA-seq data using SIMBA. **(a)** SIMBA graph construction and embedding in batch correction analysis. Overview of SIMBA's approach to batch correction across scRNA-seq datasets. Distinct shapes indicate the type of entity (cell or gene). Colors distinguish batches or cell types. **(b)** UMAP visualization of the scRNA-seq human pancreas dataset with five batches of different studies before and after batch correction. Cells are colored by scRNA-seq data source and cell type respectively. Top: UMAP visualization before batch correction; Bottom: UMAP visualization after batch correction with SIMBA; **(c)** UMAP visualization of SIMBA embeddings of cells and genes, with batch effect removed and known marker genes highlighted.

SIMBA was applied to a mouse atlas dataset composed of two batches and a human pancreas dataset that spans several batches used in a recent benchmark study[25] (**Supplementary Table 1**). The mouse atlas dataset contains two scRNA-seq datasets with shared cell types from different sequencing platform. The human pancreas dataset contains five samples pooled from five sources using four different sequencing techniques, in which not all cell types are shared across each sample. For both datasets, SIMBA successfully corrected the batch effects where

515 the resulting embeddings of cells are clustered by the cell types while each batch is mixed
516 evenly, indicating the preservation of biological signal and the elimination of confounding
517 technical covariates (**Fig. 5b, Supplementary Fig. 12b**). It is important to note that the mouse
518 atlas dataset was collected from nine different organ systems, so there exists some expected
519 heterogeneity within the same cell type labels. Conversely, the human pancreas datasets are
520 curated from a single organ and SIMBA sufficiently separated cell types into transcriptionally-
521 distinct homogeneous cell clusters (**Fig. 5b**).
522
523 In addition to batch effect removal, SIMBA also simultaneously identifies cell-type-specific
524 marker genes (**Fig. 5c**). Having effectively eliminated differences between cells due to technical
525 covariates, marker genes are discoverable across multiple samples by querying according to
526 cells of each cell type within the batch-corrected SIMBA embedding. As shown in the SIMBA co-
527 embedding of cells and variable genes, known marker genes were correctly placed within each
528 cell type while non-marker genes were embedded away from any cell type (**Supplementary Fig.**
529 **10, 11**). The resulting marker genes recapitulated the clustering-based differential expression
530 (DE) analysis results for each dataset[27-32] (e.g. *Cdh5*, *Tie1*, *Myct1* for endothelial cell and *C1qc*,
531 Fcgr1 for macrophage, *S100a8*, *Trem3* for Neutrophil in the mouse atlas dataset and *KIF12* for
532 alpha cell and *KRT19* for ductal cell in the human pancreas dataset) and are shown to be
533 expressed specifically in the queried cell types (**Supplementary Fig. 10, 11**). This distinguishes
534 SIMBA from other batch correction methods, in which clustering is performed first in the batch-
535 corrected space and then marker genes of each batch are identified through DE analysis in the
536 original, uncorrected space of the corresponding batch.
537
538 While SIMBA is a generalizable graph embedding method capable of performing multiple tasks
539 outside of batch correction, we still evaluated its performance for this task with methods that
540 were specifically designed for batch correction. We considered three widely adopted batch
541 correction methods that demonstrated top-tier performance based on a recent benchmark
542 study[25]: Seurat3, LIGER and Harmony. Our results show that SIMBA achieved comparable batch
543 correction performance both qualitatively and quantitatively (**Supplementary Note 2,**
544 **Supplementary Figure 12**).
545
546 **SIMBA enables simultaneous multi-omics integration and clustering-free multi-type marker**
547 **feature detection**
548
549 Single-cell assays are now capable of measuring a broad range of cellular modalities including
550 mRNA, chromatin accessibility, DNA methylation and cell-surface proteins. Thus, data is being
551 generated that describes cells by varying features sets, which has motivated the need for
552 methods that leverage these features to perform multi-omics integration such that a more
553 comprehensive description of cell state may be learned. This is different from multi-modal
554 analysis because the correspondence between individual cells is unknown. We demonstrate
555 that SIMBA enables simultaneous multi-omics integration and clustering-free detection of
556 distinct marker features, specifically as it is applied to datasets comprised of scRNA-seq and
557 scATAC-seq.
558

SIMBA accomplishes this integration by first building one graph for scRNA-seq data and another graph for scATAC-seq data independently as described in previous sections (**Fig. 6a**). To connect these two graphs, SIMBA then calculates gene activity scores by summarizing accessible regions from scATAC-seq data and then infers edges between cells of different assays based on their shared gene expression modules as previously described in the batch correction section. Finally, SIMBA embeds the graph of cells, genes, and peaks into a common, low-dimensional space. The SIMBA embeddings of these multi-omics entities can be visualized using UMAP or similar visualization tools (**Methods**).

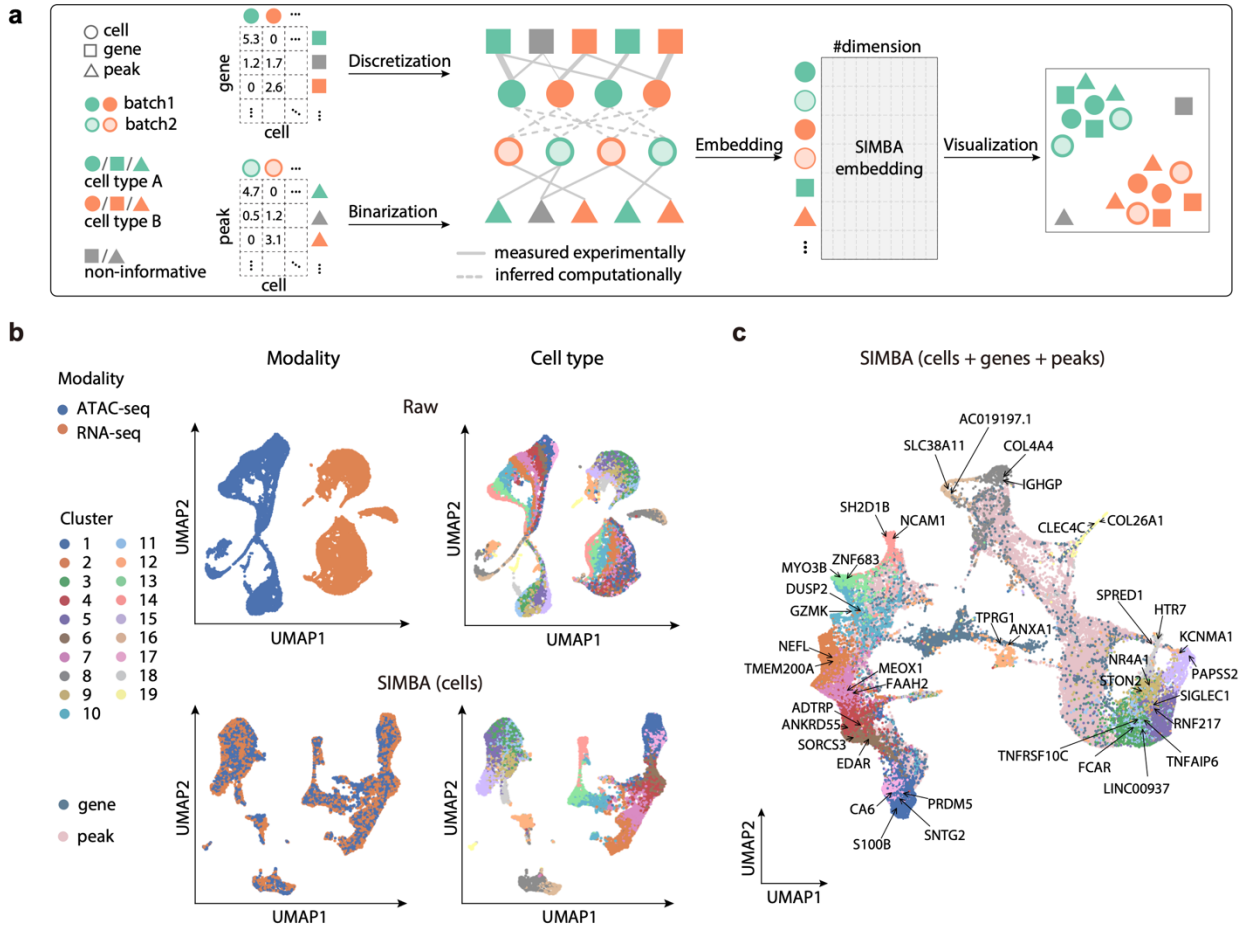

Figure 6. Multi-omics integration of scRNA-seq + scATAC-seq data using SIMBA. **(a)** SIMBA graph construction and embedding in multi-omics integration. Overview of SIMBA's approach to data integration across scRNA-seq and scATAC-seq. Distinct shapes indicate the type of entity (cell, gene, or peak). Colors distinguish batches or cell types. **(b)** UMAP visualization of the integrated scRNA-seq and scATAC-seq data manually created from the 10x human PBMCs dataset before and after data integration. Cells are colored by single-cell modality and cell type respectively. Top: UMAP visualization before integration; Bottom: UMAP visualization after integration with SIMBA. **(c)** UMAP visualization of SIMBA embeddings of cells, genes, and peaks with two cell modalities integrated and known marker genes highlighted.

To facilitate the evaluation of data integration performance, we created datasets with ground-truth labels by manually splitting the dual-omics datasets into two single-modality datasets (i.e., scRNA-seq and scATAC-seq), in which we know the true matching between cells across the two modalities. We then applied SIMBA to the integration analysis of two case studies where scRNA-seq and scATAC-seq datasets are generated from the SHARE-seq mouse skin dataset and the 10x Genomics multiome human PBMCs dataset, respectively (**Supplementary Table 1**). We observed that SIMBA was able to preserve cellular heterogeneity while evenly mixing the two modalities (**Fig. 6b, Supplementary Fig. 15b**).

Importantly, SIMBA simultaneously identified marker genes and peaks at single-cell resolution. In the co-embedding space, we observed that the neighbor genes of cells (highlighted in UMAP plots), are each exclusively expressed in their corresponding cell types (**Supplementary Figs. 13a-e, 14a-c,e**). For example, in the SHARE-seq mouse skin dataset, *Foxq1* and *Shh* are located within medulla and TAC-2, respectively; in the 10x PBMCs dataset, *PAPSS2* and *KCNMA1*, which are the marker genes of blood monocytes, are embedded close to each other. Similarly, we observed that the neighbor peaks of cells show a clear cell-type-specific accessibility pattern that is robust to the cluster size of a given cell type (**Supplementary Figs. 13f and 14d).**

While SIMBA distinguishes itself from other multi-omics integration methods by generating joint embeddings of cells and features and thus enabling simultaneous multi-omics integration and marker feature detection, we still compared SIMBA with two leading and widely-adopted single-cell multi-omics integration methods, Seurat3 and LIGER, based on their ability to integrate single-cell modalities while persevering cellular heterogeneity (**Supplementary Note 3**). We observed that SIMBA achieved overall best performance on the mouse skin SHARE-seq dataset and 10x PBMCs multiome dataset.

Discussion

Multimodal measurements of individual cells offer new and unexplored opportunities to investigate complex interactions between components involved in establishing transcriptional programs. Despite the exciting potential for discovery associated with these datasets, methods to delineate cell states and cellular processes using multiple genomic features across modalities remains under development.

As presented in this manuscript, SIMBA models cells and measured features as nodes encoded in a graph and employs a scalable and efficient graph embedding procedure to embed nodes of cells and features into a shared latent space. We demonstrate that direct graph representations of single-cell data capture not only the relations between cells and the quantified features of the experiment (e.g., gene expression or chromatin accessibility) but also the hierarchical relations between features. An example of such a hierarchical relation is the coordinate-level description of an ATAC-seq peak and the corresponding TF motifs and/or *k*-mer sequences contained within that region. In the resulting joint embedding, proximity-based queries can be performed to discover cell-type-specific regulatory mechanisms and the respective features integral to such mechanisms. Therefore, SIMBA enables simultaneous learning of cellular

heterogeneity and cell-type-specific multimodal features, and complements the current gene regulatory network analyses. SIMBA also circumvents the typical workflow led by cell clustering and followed by differential feature detection, and thus avoids relying on user-defined clustering resolution that may lead to artifactual discovery or false negative results.

SIMBA has been extensively benchmarked across single-cell modalities and tasks and obtained comparable or better performance compared to current state-of-the-art methods developed for the respective task. These results suggest a wide applicability of SIMBA's graph-based framework and therefore obviates the need to optimally combine multiple analysis tools. This contrasts with task-specific methods, which have been specially developed for a given data type and modality.

A foreseeable extension of SIMBA would include encoding increasingly complex measurements of single-cells such as spatial transcriptomics wherein transcriptomic and (real) proximity data should be considered [33]. We also envision extending this framework to chromatin looping data wherein the interaction between DNA segments can be encoded to represent how regulatory regions are linked to genes. Another potential extension of SIMBA could consider single-cell lineage-tracing datasets[34] wherein both cellular lineage information and gene expression measurements are captured and can be potentially encoded into a SIMBA graph to represent their longitudinal relations. In general, we are interested in the further incorporation of external information and hierarchical relationships between features in the graph. While extensions to different experimental designs is relatively straightforward, we note that the interpretation of the output embedding in each case may vary as a function of the input graph construction and training process thus requiring some level of domain-specific expertise.

Graph embedding methods hold significant promise for the analysis of biological data. Previous applications of graph embedding include functional annotation of genes [35], transcription factor binding to DNA motifs [36] and more recent single-cell RNA-seq analyses [37, 38]. The graph encoding and embedding procedures we have outlined may be potentially improved and extended to better represent biological entities and capture their respective relations. We envision employing continuous edge weights as well as the consideration of both node attributes and edge information to determine the proximity of node pairs [39]. Considering embedding distances as a function of the node attribute may better equip our approach to extend to time series data as well as to better align cells stratified by positions on a developmental axis. In addition, to improve interpretability it may be possible to provide a metric of uncertainty alongside the resulting embedding and generalization to the unseen nodes[40]. This could enable the mapping of new cells in the same manifold and a straightforward and automatic cell type annotation.

The ability to generate high-quality, complex datasets that capture and quantify multiple classes of molecules (omics) will undoubtedly continue to increase. Technological advances have already outpaced our ability to gain integrative insights from such data, highlighting a need for methods that break through previous limitations as well as extend easily to new tasks. We believe SIMBA satisfies these conditions for such a comprehensive yet accessible method for exploring cellular heterogeneity and investigating the regulatory mechanisms that drive

667 cellular diversity while laying a groundwork for the development of new non-cluster-centric
668 single cell omics computational methods.
669
670
671 ## Methods
672
673 **Single-cell data preprocessing**
674
675 a. Single-cell RNA-seq
676 Genes expressed in fewer than three cells were filtered. Raw counts were library size-
677 normalized and subsequently log-transformed. Optionally, variable gene selection [12] (a
678 python version is implemented in SIMBA that is inspired by Scanpy[2]) may be performed
679 to remove non-informative genes and accelerate the training procedure. Notable
680 differences in the resulting cell embeddings were not observed upon limiting feature
681 input to those identified by variable gene selection but SIMBA embeddings of non-
682 variable genes will not be generated as they are not encoded in the graph.
683
684 b. Single-cell ATAC-seq
685 Peaks present in fewer than three cells were filtered. Optionally, we implemented a
686 scalable truncated-SVD-based procedure to select variable peaks as a preliminary step
687 to additionally filter non-informative peaks and accelerate the training procedure. First
688 the top k principal components (PCs) were selected, with k chosen based on the elbow
689 plot of variance ratio. Then for each of the top *k* PCs, peaks were automatically selected
690 based on the loadings using a knee point detection algorithm implemented by 'kneed'[41].
691 Finally, peaks selected for each PC were combined and denoted as "variable peaks".
692 Similar to the observation made with scRNA-seq data, the optional step of variable peak
693 selection has a negligible effect on the resulting cell embedding. Despite this minimal
694 impact on the resulting embedding, this feature selection step imparts a significant
695 practical advantage in reducing training procedure time.
696
697 *k*-mer and motif scanning was performed using packages 'Biostrings' and 'motifmatchr'
698 with JASPAR2020[42]. Included in the implementation of SIMBA is a convenient R
699 command line script "scan_for_kmers_motifs.R" , which will convert a list of peaks
700 (formatted in a bed file) to a sparse peaks-by-*k*-mers/motifs matrix, which is stored as
701 an hdf5-formated file.
702
703 **Graph construction (five scenarios)**
704
705 i. Single-cell RNA-seq analysis
706 The distribution of non-zero values in the normalized gene expression matrix was first
707 approximated using a *k*-means clustering-based procedure. First, the continuous non-
708 zero values were binned into n intervals (by default n=5). Bin widths were defined using
709 1-dimensional *k*-means clustering wherein the values in each bin are assigned to the
710 same cluster center. The continuous matrix is then converted into a discrete matrix

wherein$1, \ldots, n$ are used to denote n levels of gene expression. Zero values are retained in this matrix. Then the graph was constructed by encoding two types of entities, cells and genes, as nodes and relations with n different weights between them, i.e., n levels of gene expression, as edges. These n relation weights range from 1.0 to 5.0 with a step size of $5/n$ denoting gene expression levels (lowest: 1.0, highest: 5.0), such that edges corresponding to high expression levels affect embeddings more strongly than those with intermediate or low expression levels. This discretization is implemented in the SIMBA package using the function, "si.tl.discretize()".

ii. Single-cell ATAC-seq analysis

Peak-by-cell matrices were binarized, with "1" indicating at least one read within a peak and "0" otherwise. The graph was constructed by encoding two types of entities, cells and peaks, as nodes and the relation between them, denoting the presence of a given peak in a cell, as edges. The single relation type was assigned with a weight of 1.0. When the DNA sequence features were available, they were encoded into the graph using *k*-mer and motif sequence entities as nodes. This was performed by first binarizing the peak-by-*k*-mer/motif matrix then constructing an extension to the original peak/cell graph using the peaks, *k*-mers, and motifs as nodes and the presence of these entities within peaks as edges between these additional nodes and the peak nodes. The relation between k-mers and peaks was assigned a weight of 0.02 while the relation between TF motifs was assigned a weight of 0.2. Of note, *k*-mers and motifs may be used independently of each other as node inputs to the graph, depending on the specific analysis task.

iii. Multimodal analysis

Combination of the above outlined strategies for graph construction of scRNA-seq and scATAC-seq data was used to construct a multi-omics graph.

iv. Batch correction

A graph for each batch was constructed as described in i). Edges between cells of different batches were inferred through a procedure based on truncated randomized singular value decomposition (SVD) to link disjoint graphs of different batches. More specifically, in the case of scRNA-seq data, consider two gene expression matrices $X1_{n_1 \times m}$ and $X2_{n_2 \times m}$, where n_1, n_2 denotes the number of cells and m denotes the number of the shared features, i.e., variable genes, between datasets. The matrix $X_{n_1 \times n_2}$ was then computed by multiplying $X1$ and $X2$:

$$X = X1 \times X2^T$$

Truncated randomized SVD was subsequently performed on X:

$$X \approx U \times \Sigma \times V^T$$

where U is an $n_1 \times d$ matrix, Σ is an $d \times d$ matrix, and V is an $n_2 \times d$ matrix (by default $d = 20$).

Both U and V were further $L2$ normalized. For each cell in U, we searched for k nearest neighbors in V and vice versa (by default, $k = 20$). Eventually, only the mutual nearest neighbors between U and V were retained as inferred edges between cells (represented as dashed lines in **Fig. 5a**). The procedure of inferring edges between cells of different batches is implemented in the function "si.tl.infer_edges()" in the SIMBA package.

For multiple batches, SIMBA can flexibly infer edges between any pair of datasets. In practice, however edges are inferred between the largest dataset(s) or the dataset(s) containing the most complete set of expected cell types and other datasets.

v. Multi-omics integration

scRNA-seq and scATAC-seq graphs were constructed following steps i) and ii), respectively. To infer the edges between cells of scRNA-seq and scATAC-seq, gene activity scores were first calculated for scATAC-seq data[3]. More specifically, for each gene, peaks within 100kb upstream and downstream of the TSS were considered. Peaks overlapping gene body region or within 5kb upstream of gene bodies were given the weight of 1.0. Otherwise, peaks were weighted based on their distances to TSS using the exponential decay function: $e^{\frac{-distance}{5000}}$. Subsequently, the gene score of each gene was computed as a weighted sum of the considered peaks. These gene scores were then scaled to respective gene size. These steps are implemented by the function "si.tl.gene_scores()" in SIMBA. For user convenience, the SIMBA package curates the gene annotations of several commonly used reference genomes, including hg19, hg38, mm9, and mm10. Once gene scores were obtained, the same procedure described in iv) was performed to infer edges between cells profiled by scRNA-seq and scATAC-seq using the function, "si.tl.infer_edges()" in SIMBA.

The procedure of generating constructed graphs is implemented in the function, "si.tl.gen_graph()" in the SIMBA package.

Graph Embeddings with Type Constraints

Following the construction of a multi-relational graph between biological entities, we adapted graph embedding techniques from the knowledge graph and recommendation systems literature to construct unsupervised representations for these entities.

We provide as input a directed graph $G = (V, E)$, where V is a set of entities (vertices) and E is a set of edges, with a generic edge $e = (u, v)$ between a source entity u and destination entity v. We further assume that each entity has a distinct known type (e.g., cell, peak, etc.).

Graph embedding methods learn a D-dimensional embedding vector for each $v \in V$ by optimizing a link prediction objective via stochastic gradient descent, with D=50 used for our experiments. We will denote the full embedding matrix as $\theta \in R^{|V| \times D}$ and the embedding for an entity v as θ_v.

For an edge $e = (u, v)$, we denote $s_e = \theta_u * \theta_v$ as the score for e, and optimize a multi-class log loss

$$\mathcal{L} = -log \frac{\exp(s_e)}{\sum_{e' \in \mathcal{N}} \exp(s'_e)}$$

Where $\mathcal{N}$ is a set of "negative sampled" candidate edges generated by corrupting e [43]. This log loss objective attempts to maximize the score for all $(u, v) \in E$ and minimize it for $(u, v) \notin E$.

Negative samples are constructed by replacing either the source or target entity in the target edge $e = (u, v)$ with a randomly sampled entity. However, in graphs like ours where only edges between certain entity types are possible, previous work has shown that it is beneficial to optimize the loss only over candidate edges that satisfy the type constraints[44]. Thus, for e.g., a cell-peak edge we only sample negative candidates between cell and peak entities. This modification is crucial in our setting since most randomly selected edges will be of invalid type (e.g., peak-peak), forcing the embeddings to primarily be optimized for irrelevant tasks (e.g., having low dot product between every pair of peaks).

Furthermore, it has been frequently observed that in graphs with wide distribution of node degrees, it is advantageous to sample negatives proportional to some function of the node degree to produce more informative embeddings that don't merely capture the degree distribution [13, 45]. For each graph edge in the dataset encountered in a training batch, we produce 100 negatives by corrupting the edge with a source or destination sampled uniformly from the nodes with the correct types for this relation and 100 by corrupting the edge with a source or destination node sampled with probability proportional to its degree[13].

As with many ML methods, graph embeddings are prone to overfitting in a low-data regime (i.e., low ratio of edges to parameters). We observed overfitting measurable as a gap between training and validation loss on the link prediction task, which we addressed with $L2$ regularization on the embeddings θ,

$$\mathcal{L}_{reg} = \mathcal{L} + \lambda \sum_{u \in N} \sum_{d=1}^{D} \theta_{ud}^2.$$

with λ =$wd * wd_interval$. For weight decay parameter (wd), by default it is calculated automatically as $\frac{C}{N_e}$, where N_e is the training sample size (i.e., the total number of edges) and C is a constant. For weight decay interval ($wd_interval$), we set it to 50 for all experiments.

We use the PyTorch-BigGraph framework, which provides efficient computation of multi-relation graph embeddings over multiple entity types and can scale to graphs with millions or billions of entities[13].

The resulting graph embeddings have two desirable properties that we will take advantage of:

1. First-order similarity: for two entity types T_1, T_2 with a relation between them, edges with high likelihood should have higher dot product; specifically, for any $u \in T_1$, the predicted probability distribution over edges to T_2 originating from u is approximated as $\frac{e^{x_u * x_v}}{\sum_{v' \in T_2} e^{x_u * x_{v'}}}$.
2. Second-order similarity: within a single entity type, entities that have 'similar contexts', i.e., a similar distribution of edge probabilities, should have similar embeddings. Thus, the embeddings of each entity type provide a low-rank latent space that encodes the similarity of those entities' edge distributions.

Evaluation of the model during training

During the PyTorch-BigGraph training procedure, a small percent of edges is held out (by default, the evaluation fraction is set to 5%) to monitor overfitting and evaluate the final model. Five metrics are computed on the reserved set of edges, including mean reciprocal rank (MRR, the average of the reciprocal of the ranks of all positives), R1 (the fraction of positives that rank better than all their negatives, i.e., have a rank of 1), R10 (the fraction of positives that rank in the top 10 among their negatives), R50 (the fraction of positives that rank in the top 50 among their negatives), and AUC (Area Under the Curve). By default, we show MRR along with training loss and validation loss while other metrics are also available in SIMBA package (**Supplementary Fig. 1a**). The learning curves for validation loss and these metrics can be used to determine when training has completed. The relative values of training and validation loss along with these evaluation metrics can be used to identify issues with training (underfitting vs overfitting) and tune the hyperparameters weight decay, embedding dimension, and number of training epochs appropriately. For example, in **Supplementary Figure 1a** training can be stopped once the validation loss plateaus. However, for most datasets we find that the default parameters do not need tuning.

Softmax transformation

PyTorch-BigGraph training provides initial embeddings of all entities (nodes). However, entities of different types (e.g., cells vs peaks, cells of different batches or modalities)

have different edge distributions and thus may lie on different manifolds of the latent space. To make the embeddings of entities of different types comparable, we transform the embeddings of features with the Softmax function by utilizing the first-order similarity between cells (reference) and features (query). In the case of batch correction or multi-omics integration, the Softmax transformation is also performed based on the first-order similarity between cells of different batches or modalities.

Given the initial embeddings of cells (reference) $(v_{c_1}, \dots, v_{c_n})$ and features $(v_{f_1}, \dots, v_{d_m})$, the model-estimated probability of an edge (c_i, f_j) obeys

$$P\left(v_{c_i,f_j}\right) \propto \exp\left(v_{c_i} \cdot v_{f_j}\right)$$

Therefore, if a random edge was sampled from feature f_j to a cell, the model would estimate the distribution over such edges as

$$p_{c_i,f_j} = \frac{\exp\left(v_{c_i} \cdot v_{f_j}\right)}{\sum_{k=1}^{n} \exp\left(v_{c_k} \cdot v_{f_j}\right)}$$

i.e., the Softmax weights between all cells $\{c_i\}$ and the feature f_j. We can then compute new embeddings for features as a linear combination of the cell embeddings weighted by the edge probabilities raised to some power.

$$\hat{v}_{f_j} = \frac{\sum_{i=1}^{n} p_{c_i,f_j}^{T^{-1}} v_{c_i}}{\sum_{i=1}^{n} p_{c_i,f_j}^{T^{-1}}}$$

T is a temperature hyperparameter that controls the sharpness of the weighting over cells. At $T = 1$, the cell embeddings are weighted by their estimated edge probabilities; at $T \to 0$, each feature embedding is assigned the cell embedding of its nearest neighbor; at $T \to \infty$, it becomes a discrete uniform distribution, and each query becomes the average of reference embeddings. We set $T = 0.5$ for all the analyses.

These steps are implemented in the function “si.tl.embed()” in the SIMBA package.

Metrics to assess cell-type specificity

Four metrics are proposed to assess the cell type specificity of each feature from different aspects, including max value (a higher value indicates higher cell-type specificity), Gini index (a higher value indicates higher cell-type specificity), standard deviation (a higher value indicates higher cell-type specificity), and entropy (a lower value indicates higher cell-type specificity). We observe these four metrics generally give consistent results. For SIMBA metric plot, by default, Gini index is plotted against max value. For feature f_j :

The max value is defined as the average normalized similarity of top k cells (by default, k=50). The similarity normalization function is defined as:

$$norm(x_i) = x_i - log\frac{\sum_{j=1}^{n}\exp(x_j)}{n}$$

Where $i = 1, \dots, n$. n is the number of cells and x_i represents the dot product of $\hat{v}_{f_j}$ and the embedding of cell i.

The max value is computed as:

$$\max(f_j) = \frac{\sum_{i=1}^{k} norm(x_i)}{k}$$

The Gini index is computed as:

$$\mathrm{gini}(\mathrm{f_j}) = \frac{\sum_{i=1}^{n}(2i - n - 1) * p_{c_i,f_j}}{n\sum_{i=1}^{n} p_{c_i,f_j}}$$

The standard deviation is computed as:

$$\mathrm{std}(\mathrm{f_j}) = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(p_{c_i,f_j} - \mu)^2}$$

Where $\mu = \frac{1}{n}\sum_{i=1}^{n} p_{c_i,f_j}$.

Entropy is computed as:

$$\mathrm{entropy}(\mathrm{f_j}) = -\sum_{i=1}^{n} p_{c_i,f_j}\log(p_{c_i,f_j})$$

Identification of master regulators

To identify master regulators, we take into consideration both the cell type specificity of each pair of TF motif and TF gene and the distance between them. More specifically, for each TF motif, first its distances (Euclidean distance by default) to all the genes are calculated in the SIMBA embedding space. Then the rank of this TF gene among all these genes is computed. In addition, we also assess the cell type specificity of this pair of TF motif and TF gene based on SIMBA metrics (by default, max value and Gini index are used). The same procedure is performed for all TFs. Finally, we identify master regulators by filtering out TFs with low cell-type specificity and scoring them based on

TF gene rank. This procedure is implemented in the function "st.tl.find_master_regulators()" in SIMBA package.

Identification of TF target genes

Given a master regulator, its target genes are identified by comparing the locations of the TF gene, TF motif, and the peaks near the genomic loci of candidate target genes in the SIMBA co-embedding space (**Fig. 4e**). More specifically we first search for k nearest neighbor genes around the motif (TF motif) and the gene (TF gene) of this master regulator, respectively (k = 200 by default). The union of these neighbor genes is the initial set of candidate target genes. These genes are then filtered based on the criterion that open regions (peaks) within 100kb upstream and downstream of the TSS of a putative target gene must contain the TF motif.

Next, for each candidate target gene, we compute four types of distances in SIMBA embedding space: distances between the embeddings of 1) the candidate target gene and TF gene; 2) the candidate target gene and TF motif; 3) peaks near the genomic locus of the candidate target gene and TF motif; 4) peaks near the genomic locus of the candidate target gene and the candidate gene. All the distances (Euclidean distances by default) are converted to ranks out of all genes or all peaks to make the distances comparable across different master regulators.

The final list of target genes is decided using the calculated ranks based on two criteria: 1) at least one of the nearest peaks to TF gene or TF motif is within a predetermined range (top 1,000 by default); 2) the average rank of the candidate target gene is within a predetermined range (top 5,000 by default). This procedure is implemented in the function "st.tl. find_target_genes ()" in SIMBA.

Benchmarking scATAC-seq computational methods

To compare SIMBA to other scATAC-seq computational methods including SnapATAC [4], Cusanovich2018[46], and cisTopic[47], we employed the previously developed benchmarking framework from Chen et al[14](**Supplementary Table 1**). This framework evaluates different methods based on their ability to distinguish cell types. We applied three clustering algorithms: k-means clustering, hierarchical clustering, and Louvain on the feature matrix derived from each method.

For datasets with ground-truth (FACS-sorted labels or known tissue labels), including simulated bone marrow data, Buenrostro 2018, and sci-ATAC-seq subset, three metrics including Adjusted Rand Index (ARI), Adjusted Mutual Information (AMI), and Homogeneity are applied to evaluate the performance. ARI measures the similarity between two clusters, comparing all pairs of samples assigned to matching or different clusters in the predicted clustering solution vs the true cluster/cell type label. AMI

describes an observed frequency of co-occurrence compared to an expected frequency of co-occurrence between two variables, informing the mutual dependence or strength of association of these two variables. Homogeneity measures whether a clustering algorithm preserves cluster assignments towards samples that belong to a single class. A higher metric value indicates a better clustering solution.

For 10x PBMCs dataset with no ground truth, the Residual Average Gini Index (RAGI) proposed in the benchmarking study[14] is used as the clustering evaluation metric. RAGI measures the relative exclusivity of marker genes to their corresponding clusters in comparison to housekeeping genes, which should demonstrate low specificity to any given cluster. In brief, the mean Gini Index is computed for both marker genes and housekeeping genes. The difference between the means is computed to obtain the average residual specificity (i.e., RAGI) of a clustering solution with respect to marker genes. A higher RAGI indicates a better separation of biologically distinct clusters.

Benchmarking single-cell batch correction methods

The batch correction performance of SIMBA was compared to Seurat3[12], LIGER[11] and Harmony[10] in two benchmark datasets: the mouse atlas dataset and the human pancreas dataset (**Supplementary Table 1**). For Seurat3, LIGER and Harmony, the batch correction was done with the same parameters used in a previous benchmark study[25].

To evaluate the batch integration performance, average Silhouette width (ASW), adjusted Rand index (ARI), and local inverse Simpson's index (LISI)[10] were calculated for the batches and cell types using the Euclidean distance as described in a previous benchmark[25]. To make a fair evaluation, only the cell types that are present in all batches were considered. We used the same number of dimensions (50) for these methods and all other parameters were set as in the benchmark.

Average Silhouette width (ASW)

Average Silhouette width is the mean value of Silhouette scores calculated from each cell. Silhouette width measures the relative closeness of cells with the same label compared to the cells with the different label and ranges from -1 to +1. Silhouette score for a data point with a label is calculated as

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$$

where $a(i)$ is the distance to the closest point with the same label, and $b(i)$ is the distance to the closest point with different labels. A high Silhouette score means the point is located more closely with the same label, where a low Silhouette score closer to -1 means the point is located closer with different labels than that of itself. The ideal batch correction result will give a low ASW score for batch labels as the point is well

mixed with other batches and a high ASW score for the cell type labels as the cells of the same cell type should cluster together after the batch correction. The final score is calculated as the median ASW scores from 20 subsets of randomly sampled 80% cells.

Average Rand Index (ARI)

To evaluate the cell type purity, the true cell type labels and the k-means clustering solution were used to calculate the cell type ARI. To evaluate the batch correction performance, the true batch labels and the k-means clustering solution were used to calculate the batch ARI. The final ARI was calculated as the median ARI scores of 20 subsets comprised of randomly sampled 80% cells for batches and cell types, respectively. A superior batch correction will have a high cell type ARI (high agreement between the clustering solution and the true cell type labels), and a low batch ARI (the clustering solution is not mainly driven by batches and clusters contain cells with well-mixed batch labels).

Local Inverse Simpson's Index (LISI)

Local Inverse Simpson's Index (LISI) [10] measures the local batch and cell type mixing. For each data point, it considers the Gaussian kernel weighted distribution of labels in its neighborhood with a perplexity argument. We set perplexity to 50 40 as in the previous benchmark study. Using the weighted neighborhood label distribution, the inverse Simpson's index is calculated as $\frac{1}{\sum_l p(l)}$ where l is the batch or cell type labels and $p(l)$ is the probability of each label in the local neighborhood obtained with the kernel. For each cell, the LISI is the expected number of cells to be sampled locally before a cell of the same label is sampled. A perfect batch correction will have a cell type LISI (cLISI) of 1 and a batch LISI (integration LISI, iLISI) close to the number of batches. The final LISI score was calculated as the average LISI scores of all cells.

Further details are described in **Supplementary Note 2**.

Benchmarking single cell multi-omics integration methods

Two pairs of scRNA-seq and scATAC-seq datasets manually split from the dual-omics SHARE-seq mouse skin dataset and 10X PBMCs dataset respectively were used for the modality integration task. For Seurat3 and LIGER, the parameters and preprocessing were done as described in their documentations. However for the LIGER analysis of the SHARE-seq mouse skin dataset the parameter 'lambda' was set to 30 and the 'ref_dataset' was set to scATAC-seq to get a better alignment. For the Raw results, the activity matrix of scATAC-seq was constructed using Seurat3 and the first 20 PCs of the scRNA-seq count matrix and the activity matrix were used for the comparison. The integration results generated by each method were evaluated with four metrics—

Anchoring distance, anchoring distance rank, Silhouette index, and cluster agreement—as described below.

Anchoring distance

The Anchoring distance was proposed in Dou et al., 2020[48] and is the normalized distance between the matched cells of two modalities (e.g. RNA and ATAC). Here we considered the Euclidean distance and normalized the distance by the mean of the distances calculated between random pairs of cells. The number of pairs randomly sampled was set to 10% of the total number of cells.

Anchoring distance rank

Given that the anchoring distance does not account for the local density of cells, we propose a new metric entitled *anchoring distance rank* (ADR). The ADR is based on the normalized rank of the distance between the matched cells of two modalities. For each cell x_{ij} with cell identity i and modality j, the distance between the cell and all the other cells of the other modality j', $d(x_{ij}, x_{kj'}), k = 1, \dots, N$ is calculated, where N is the total number of cells. Then the rank of $r_i = d(x_{ij}, x_{ij'})$ within the calculated distances is normalized by the number of pairs $N - 1$ to obtain the final anchoring rank $m_i = \frac{r_i - 1}{N - 1}$. For each cell, an anchoring rank of 0 indicates an ideal modality integration performance as the matched cells are closest to each other in the embedding.

Silhouette index

The silhouette index was calculated as described in 10) based on the cluster assignment wherein each cluster consists of two cells, one cell from a scRNA-seq dataset and one cell from a scATAC-seq dataset.

Fraction in the same cluster

Fraction in the same cluster was calculated as the fraction of the matched cells from two modalities in the same cluster. The clusters of cells were generated using Louvain algorithm and the number of clusters is equal to the number of cell types in the dataset.

Further details are described in **Supplementary Note 3**.

Data availability:

All the datasets used in this study (eight scRNA-seq datasets, four scATAC-seq datasets, and three dual-omics datasets) are summarized in **Supplementary Table 1**. All these datasets are curated in the SIMBA package, and they can be easily downloaded and imported directly to reproduce the analyses presented in this manuscript.

Code availability:

We provide a comprehensive Python package 'simba' available at https://anaconda.org/bioconda/simba and https://github.com/pinellolab/simba. All the proposed procedures are implemented in the "simba" package. 'simba' can be easily installed with conda "*conda install simba*". We also built a website (https://simba-bio.readthedocs.io), providing a detailed introduction of the 'simba' software and several SIMBA tutorials for different types of single-cell analyses presented in this manuscript.

Acknowledgements

The authors thank Ledell (Yu) Wu, Facebook, for the helpful discussions about Starspace; Dr.Sai Ma and Dr. Jason Buenrostro, Broad Institute of MIT and Harvard, for sharing the SHARE-seq data and metadata. The authors also acknowledge members of the Pinello lab, Mass General Hospital/Harvard Medical School, for helpful comments and feedback. This project has been made possible in part by grant number 2019-202669 from the Chan Zuckerberg Foundation to LP. LP is also partially supported by the National Human Genome Research Institute (NHGRI) Genomic Innovator Award (R35HG010717). M.E.V. is supported by the National Institutes of Health (NIH) under the Ruth L. Kirschstein National Research Service Award (1F31CA257625-01) from the National Cancer Institute (NCI).

Author contributions

H.C. and L.P. conceived this project. H.C. designed SIMBA, wrote the SIMBA package, and performed SIMBA analysis on all datasets. A.L. contributed to the adaption of PyTorch-BigGraph to single cell analysis. J.R. and H.C. performed the comparison analysis on batch correction and data integration. M.E.V. and H.C. performed the comparison analysis on scATAC-seq data. L.P. and A.L. provided guidance and supervised this project. All the authors wrote and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

References:

1. Satija, R., Farrell, J.A., Gennert, D., Schier, A.F. & Regev, A. Spatial reconstruction of single-cell gene expression data. *Nat Biotechnol* **33**, 495-502 (2015).
2. Wolf, F.A., Angerer, P. & Theis, F.J. SCANPY: large-scale single-cell gene expression data analysis. *Genome Biol* **19**, 15 (2018).
3. Granja, J.M. et al. ArchR is a scalable software package for integrative single-cell chromatin accessibility analysis. *Nat Genet* **53**, 403-411 (2021).
4. Fang, R. et al. Comprehensive analysis of single cell ATAC-seq data with SnapATAC. *Nat Commun* **12**, 1337 (2021).
5. Kiselev, V.Y., Andrews, T.S. & Hemberg, M. Challenges in unsupervised clustering of single-cell RNA-seq data. *Nat Rev Genet* **20**, 273-282 (2019).

1167 6. Cortal, A., Martignetti, L., Six, E. & Rausell, A. Gene signature extraction and cell identity
1168 recognition at the single-cell level with Cell-ID. *Nat Biotechnol* (2021).
1169 7. Vandenbon, A. & Diez, D. A clustering-independent method for finding differentially
1170 expressed genes in single-cell transcriptome data. *Nat Commun* **11**, 4318 (2020).
1171 8. Dann, E., Henderson, N.C., Teichmann, S.A., Morgan, M.D. & Marioni, J.C. Differential
1172 abundance testing on single-cell data using k-nearest neighbor graphs. *Nat Biotechnol*
1173 (2021).
1174 9. Hao, Y. et al. Integrated analysis of multimodal single-cell data. *Cell* (2021).
1175 10. Korsunsky, I. et al. Fast, sensitive and accurate integration of single-cell data with
1176 Harmony. *Nat Methods* **16**, 1289-1296 (2019).
1177 11. Welch, J.D. et al. Single-Cell Multi-omic Integration Compares and Contrasts Features of
1178 Brain Cell Identity. *Cell* **177**, 1873-1887 e1817 (2019).
1179 12. Stuart, T. et al. Comprehensive Integration of Single-Cell Data. *Cell* **177**, 1888-1902
1180 e1821 (2019).
1181 13. Lerer, A. et al. Pytorch-biggraph: A large-scale graph embedding system. *arXiv preprint*
1182 *arXiv:1903.12287* (2019).
1183 14. Chen, H. et al. Assessment of computational methods for the analysis of single-cell
1184 ATAC-seq data. *Genome Biology* **20**, 241 (2019).
1185 15. Buenrostro, J.D. et al. Integrated Single-Cell Analysis Maps the Continuous Regulatory
1186 Landscape of Human Hematopoietic Differentiation. *Cell* **173**, 1535-1548 e1516 (2018).
1187 16. Ferreira, R., Ohneda, K., Yamamoto, M. & Philipsen, S. GATA1 function, a paradigm for
1188 transcription factors in hematopoiesis. *Molecular and cellular biology* **25**, 1215-1227
1189 (2005).
1190 17. Tijchon, E., Havinga, J., Van Leeuwen, F. & Scheijen, B. B-lineage transcription factors
1191 and cooperating gene lesions required for leukemia development. *Leukemia* **27**, 541-
1192 552 (2013).
1193 18. Friedman, A. Transcriptional control of granulocyte and monocyte development.
1194 *Oncogene* **26**, 6816-6828 (2007).
1195 19. Schep, A.N., Wu, B., Buenrostro, J.D. & Greenleaf, W.J. chromVAR: inferring
1196 transcription-factor-associated accessibility from single-cell epigenomic data. *Nat*
1197 *Methods* **14**, 975-978 (2017).
1198 20. Moriguchi, T. & Yamamoto, M. A regulatory network governing Gata1 and Gata2 gene
1199 transcription orchestrates erythroid lineage differentiation. *International journal of*
1200 *hematology* **100**, 417-424 (2014).
1201 21. Chen, S., Lake, B.B. & Zhang, K. High-throughput sequencing of the transcriptome and
1202 chromatin accessibility in the same cell. *Nat Biotechnol* (2019).
1203 22. Ma, S. et al. Chromatin Potential Identified by Shared Single-Cell Profiling of RNA and
1204 Chromatin. *Cell* (2020).
1205 23. Cao, J. et al. Joint profiling of chromatin accessibility and gene expression in thousands
1206 of single cells. *Science* **361**, 1380-1385 (2018).
1207 24. Zhu, C. et al. An ultra high-throughput method for single-cell joint analysis of open
1208 chromatin and transcriptome. *Nat Struct Mol Biol* (2019).
1209 25. Tran, H.T.N. et al. A benchmark of batch-effect correction methods for single-cell RNA
1210 sequencing data. *Genome Biol* **21**, 12 (2020).

26. Leek, J.T. et al. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nature Reviews Genetics* **11**, 733-739 (2010).
27. Han, X. et al. Mapping the mouse cell atlas by microwell-seq. *Cell* **172**, 1091-1107. e1017 (2018).
28. Tabula Muris, C. et al. Single-cell transcriptomics of 20 mouse organs creates a Tabula Muris. *Nature* **562**, 367-372 (2018).
29. Baron, M. et al. A single-cell transcriptomic map of the human and mouse pancreas reveals inter-and intra-cell population structure. *Cell systems* **3**, 346-360. e344 (2016).
30. Muraro, M.J. et al. A single-cell transcriptome atlas of the human pancreas. *Cell systems* **3**, 385-394. e383 (2016).
31. Wang, Y.J. et al. Single-cell transcriptomics of the human endocrine pancreas. *Diabetes* **65**, 3028-3038 (2016).
32. Segerstolpe, Å. et al. Single-cell transcriptome profiling of human pancreatic islets in health and type 2 diabetes. *Cell metabolism* **24**, 593-607 (2016).
33. Palla, G. et al. Squidpy: a scalable framework for spatial single cell analysis. *bioRxiv 2021.02.19.431994v2* (2021).
34. VanHorn, S. & Morris, S.A. Next-Generation Lineage Tracing and Fate Mapping to Interrogate Development. *Dev Cell* **56**, 7-21 (2021).
35. Ietswaart, R., Gyori, B.M., Bachman, J.A., Sorger, P.K. & Churchman, L.S. GeneWalk identifies relevant gene functions for a biological context using network representation learning. *Genome Biol* **22**, 55 (2021).
36. Yuan, H., Kshirsagar, M., Zamparo, L., Lu, Y. & Leslie, C.S. BindSpace decodes transcription factor binding signals by large-scale sequence embedding. *Nat Methods* (2019).
37. Li, H., Xiao, X., Wu, X., Ye, L. & Ji, G. scLINE: A multi-network integration framework based on network embedding for representation of single-cell RNA-seq data. *J Biomed Inform* **122**, 103899 (2021).
38. Buterez, D., Bica, I., Tariq, I., Andrés-Terré, H. & Liò, P. CELLVGAE: AN UNSUPERVISED SCRNA-SEQ ANALYSIS WORKFLOW WITH GRAPH ATTENTION NETWORKS. *bioRxiv 2020.12.20.423645v1* (2020).
39. Sheikh, N., Kefato, Z. & Montresor, A. gat2vec: representation learning for attributed graphs. *Computing* **101**, 187-209 (2018).
40. Bojchevski, A. & Günnemann, S. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. *arXiv preprint arXiv:1707.03815* (2017).
41. Satopaa, V., Albrecht, J., Irwin, D. & Raghavan, B. in 2011 31st international conference on distributed computing systems workshops 166-171 (IEEE, 2011).
42. Fornes, O. et al. JASPAR 2020: update of the open-access database of transcription factor binding profiles. *Nucleic acids research* **48**, D87-D92 (2020).
43. Kadlec, R., Bajgar, O. & Kleindienst, J. Knowledge base completion: Baselines strike back. *arXiv preprint arXiv:1705.10744* (2017).
44. Krompaß, D., Baier, S. & Tresp, V. in International semantic web conference 640-655 (Springer, 2015).
45. Mikolov, T., Chen, K., Corrado, G. & Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).

1255 46. Cusanovich, D.A. et al. The cis-regulatory dynamics of embryonic development at single-
1256 cell resolution. *Nature* **555**, 538-542 (2018).
1257 47. Bravo Gonzalez-Blas, C. et al. cisTopic: cis-regulatory topic modeling on single-cell ATAC-
1258 seq data. *Nat Methods* **16**, 397-400 (2019).
1259 48. Dou, J. et al. Unbiased integration of single cell multi-omics data. *bioRxiv*,
1260 2020.2012.2011.422014 (2020).
1261
1262

Supplementary Figures

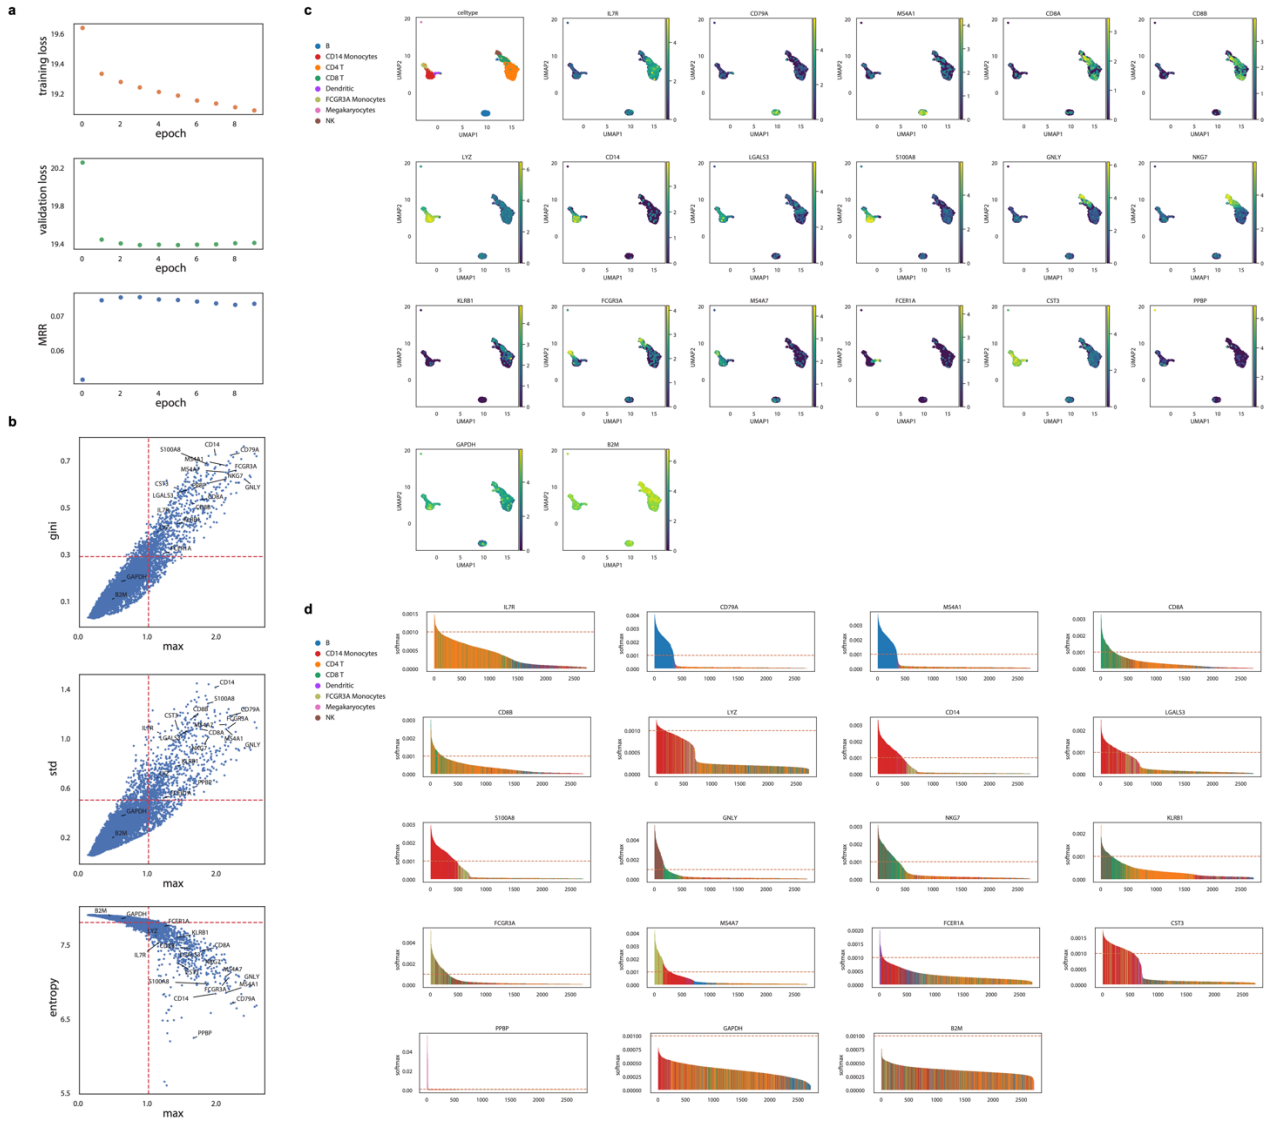

Supplementary Figure 1. SIMBA analysis of the scRNA-seq 10x PBMCs dataset.

- **a.** Three default metrics used to evaluate SIMBA training procedure, including training loss (top), validation loss (middle), mean reciprocal rank (MRR)
- **b.** SIMBA metric plots of genes. All the genes are plotted according to the Gini index against max score, standard deviation (std) against max score, and entropy against max score, respectively. The same set of genes as in Figure 2c are highlighted.
- **c.** UMAP visualization of SIMBA embeddings of cells colored by cell type or gene expression of those genes highlighted in (b).
- **d.** SIMBA barcode plots of the genes highlighted in (b).

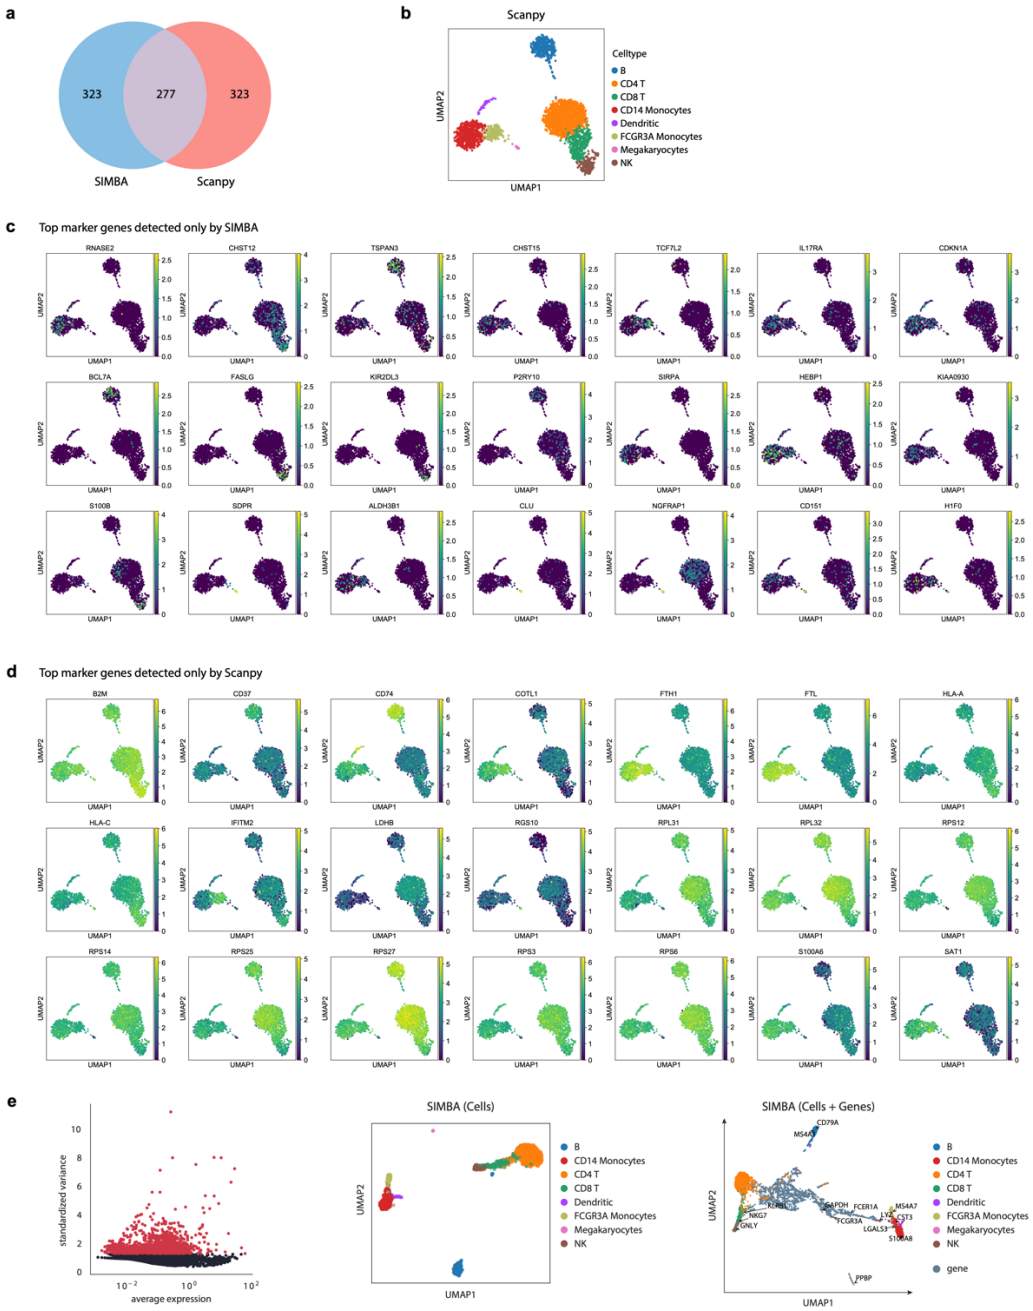

Supplementary Figure 2. Comparison of SIMBA with Scanpy on the scRNA-seq 10x PBMCs dataset.

- **a.** Venn diagram of top marker genes identified by SIMBA and Scanpy
- **b.** Scanpy-derived UMAP visualization of cells colored by cell type
- **c.** Top marker genes detected only by SIMBA. Colored by intensity of gene expression.
- **d.** Top marker genes detected only by Scanpy. Colored by intensity of gene expression.
- **e.** SIMBA embedding result after implementing variable gene selection. Left: variable gene selection step implemented in SIMBA. Middle: UMAP visualization of SIMBA embeddings of cells. Right: UMAP visualization of SIMBA embeddings of cells and variable genes.

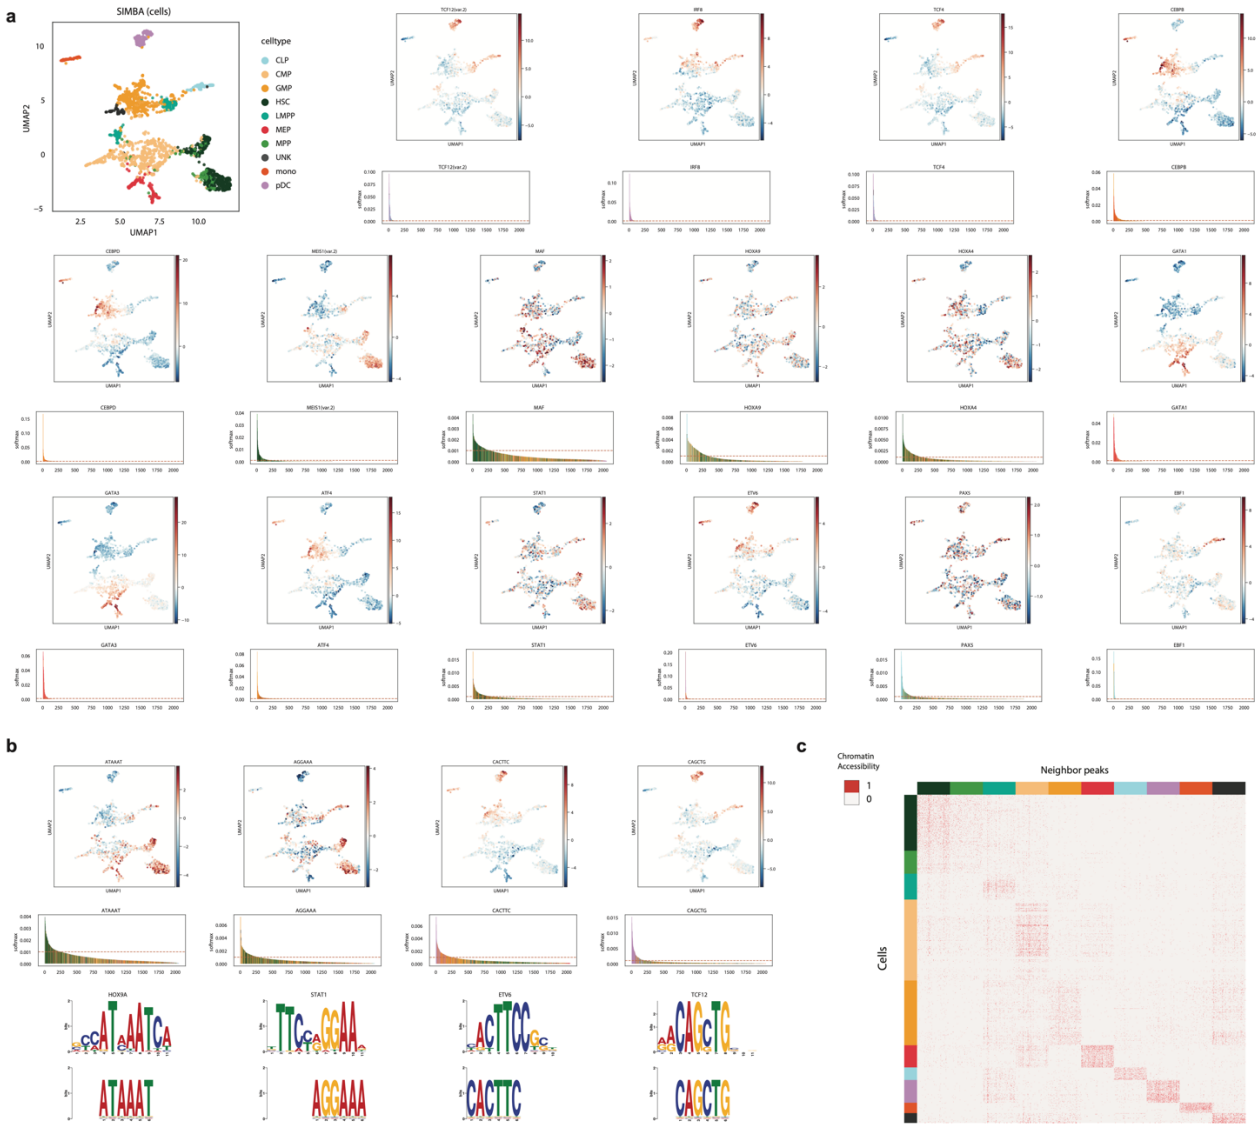

Supplementary Figure 3. SIMBA analysis of the *Buenrostro2018* dataset

- **a.** UMAP visualization of SIMBA embeddings of cells colored by cell type (top-left), and TF activity scores of TF motifs calculated with chromVAR, respectively. The SIMBA barcode plot of each TF motif is shown below the UMAP plot.
- **b.** Top: UMAP visualization of SIMBA embeddings of cells colored by TF activity scores of k-mers calculated with chromVAR. Middle: SIMBA barcode plots of the corresponding k-mers. Bottom: the matching known motif against the enriched k-mer sequences.
- **c.** Heatmap of cells against neighboring peaks of each cell type that are selected in the SIMBA co-embedding space. Chromatin accessibility is binary and colored accordingly.

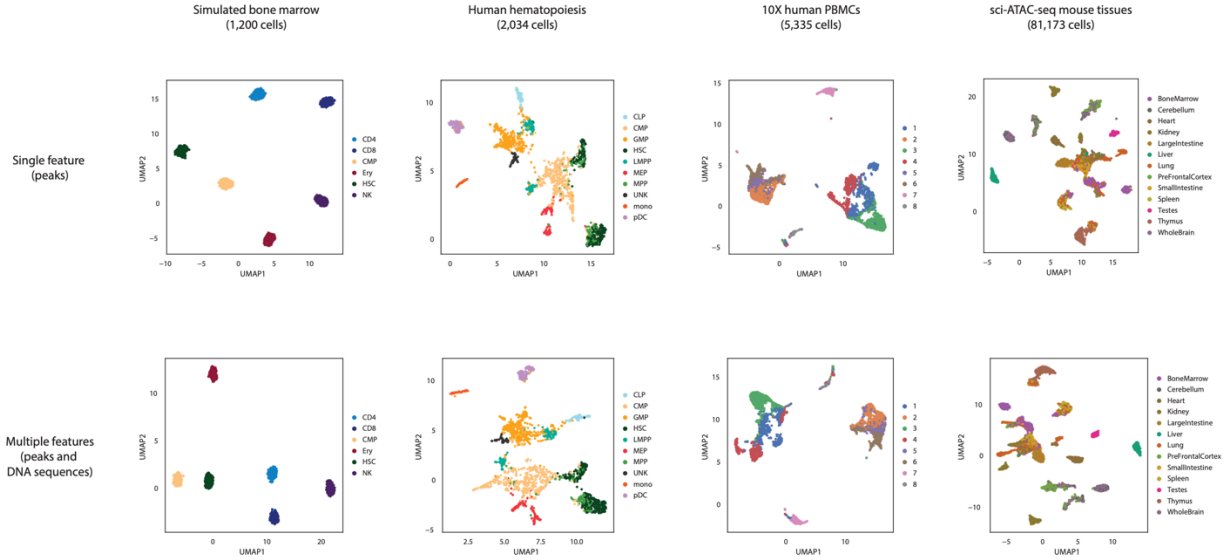

Supplementary Figure 4. Comparison of SIMBA performance using scATAC-seq peaks and DNA sequence content vs only scATAC-seq peaks. **Top**: UMAP visualization of SIMBA embeddings of cells for each indicated dataset generated from only scATAC-seq peak information. **Bottom**: UMAP visualization of SIMBA embeddings of cells for each indicated dataset generated from scATAC-seq peak information and DNA sequence content information.

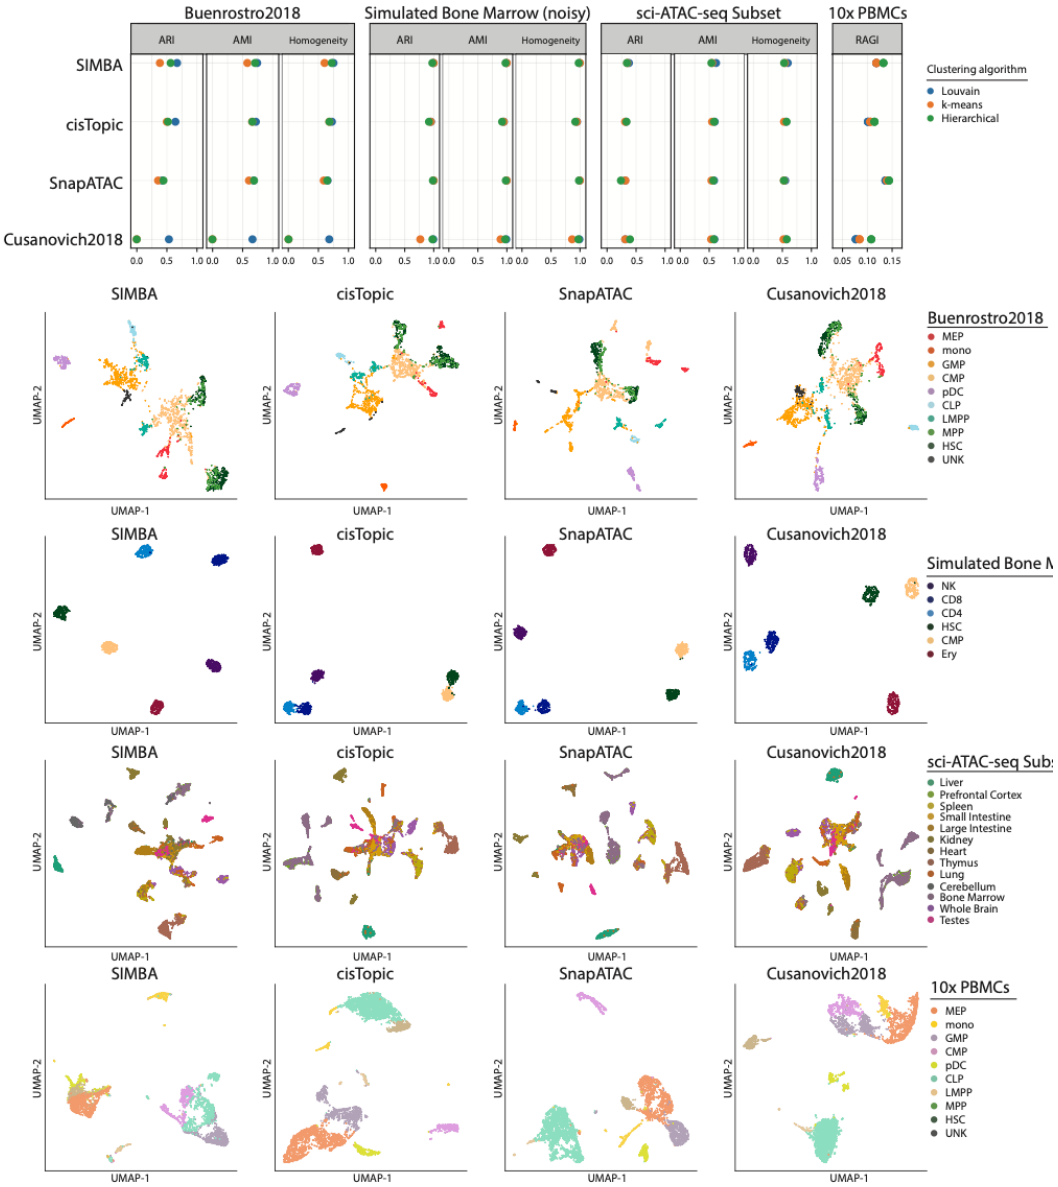

Supplementary Figure 5. Benchmark of SIMBA against top-performing scATAC-seq analysis methods.

Top: Evaluation of SIMBA and other methods including cisTopic, SnapATAC, *Cusanovich2018* for scATAC-seq analysis using metrics 1) ARI, AMI, and Homogeneity for datasets with ground truth cell type labels and 2) Residual Average Gini Index (RAGI) for the 10x PBMCs dataset without ground truth labels.

Bottom: UMAP visualization of feature matrices produced by each method on each dataset colored by cell type annotation or cluster label.

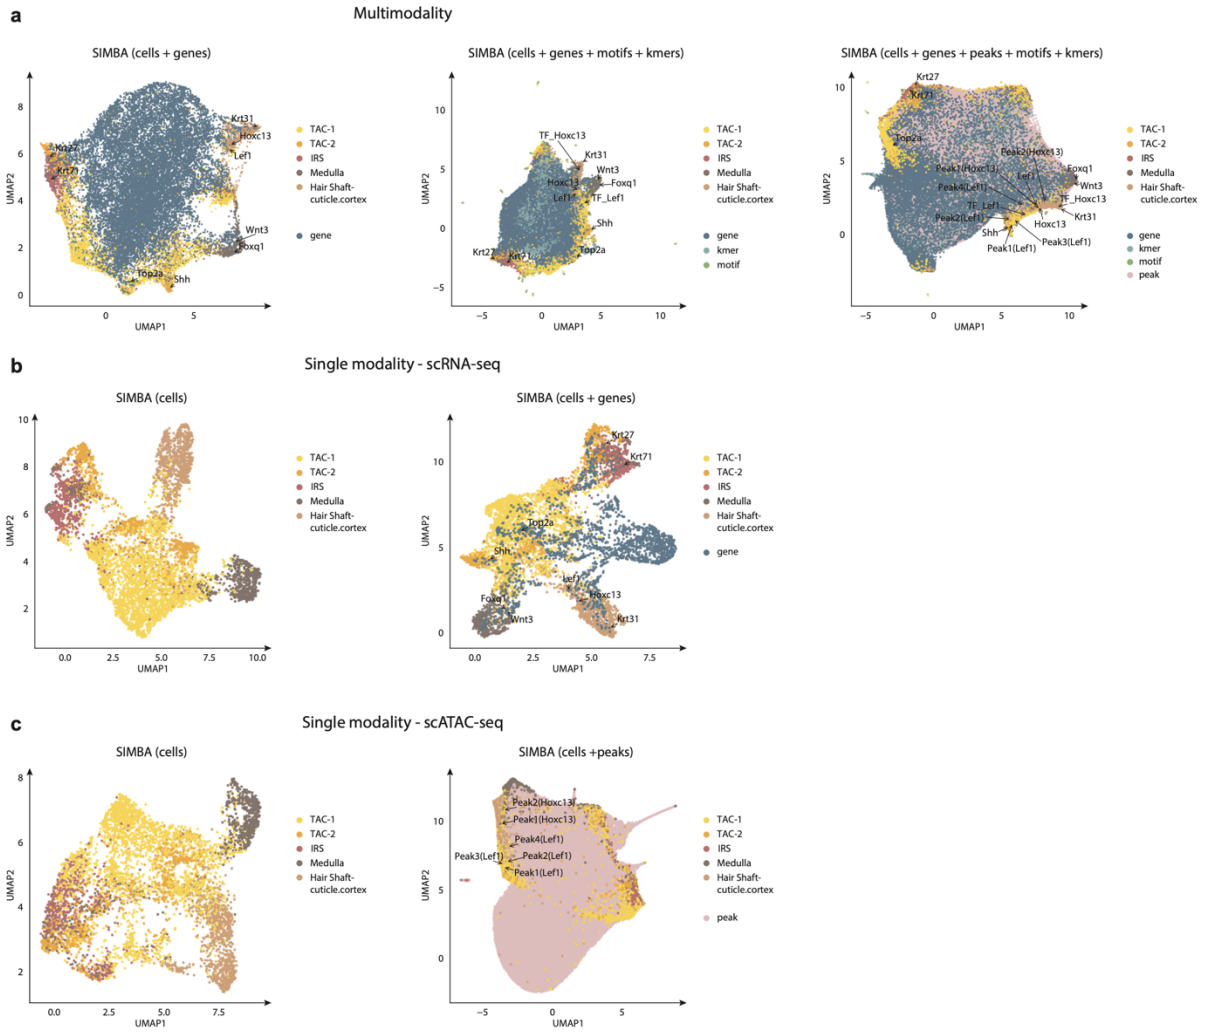

Supplementary Figure 6. SIMBA multimodal analysis of the SHARE-seq hair follicle dataset.

a. SIMBA embedding results when both gene expression and chromatin accessibility are encoded in the graph. Left: UMAP visualization of SIMBA embeddings of cells and genes. Middle: UMAP visualization of SIMBA embeddings of cells along with genes, TF motifs, and k-mers. Right: UMAP visualization of SIMBA embeddings of cells along with genes, peaks, TF motifs, and k-mers.

b. SIMBA embedding results when only gene expression is encoded in the graph. Left: UMAP visualization of SIMBA embeddings of cells. Right: UMAP visualization of SIMBA embeddings of cells and variable genes.

c. SIMBA embedding results when only chromatin accessibility is encoded in the graph. Left: UMAP visualization of SIMBA embeddings of cells. Right: UMAP visualization of SIMBA embeddings of cells and peaks.

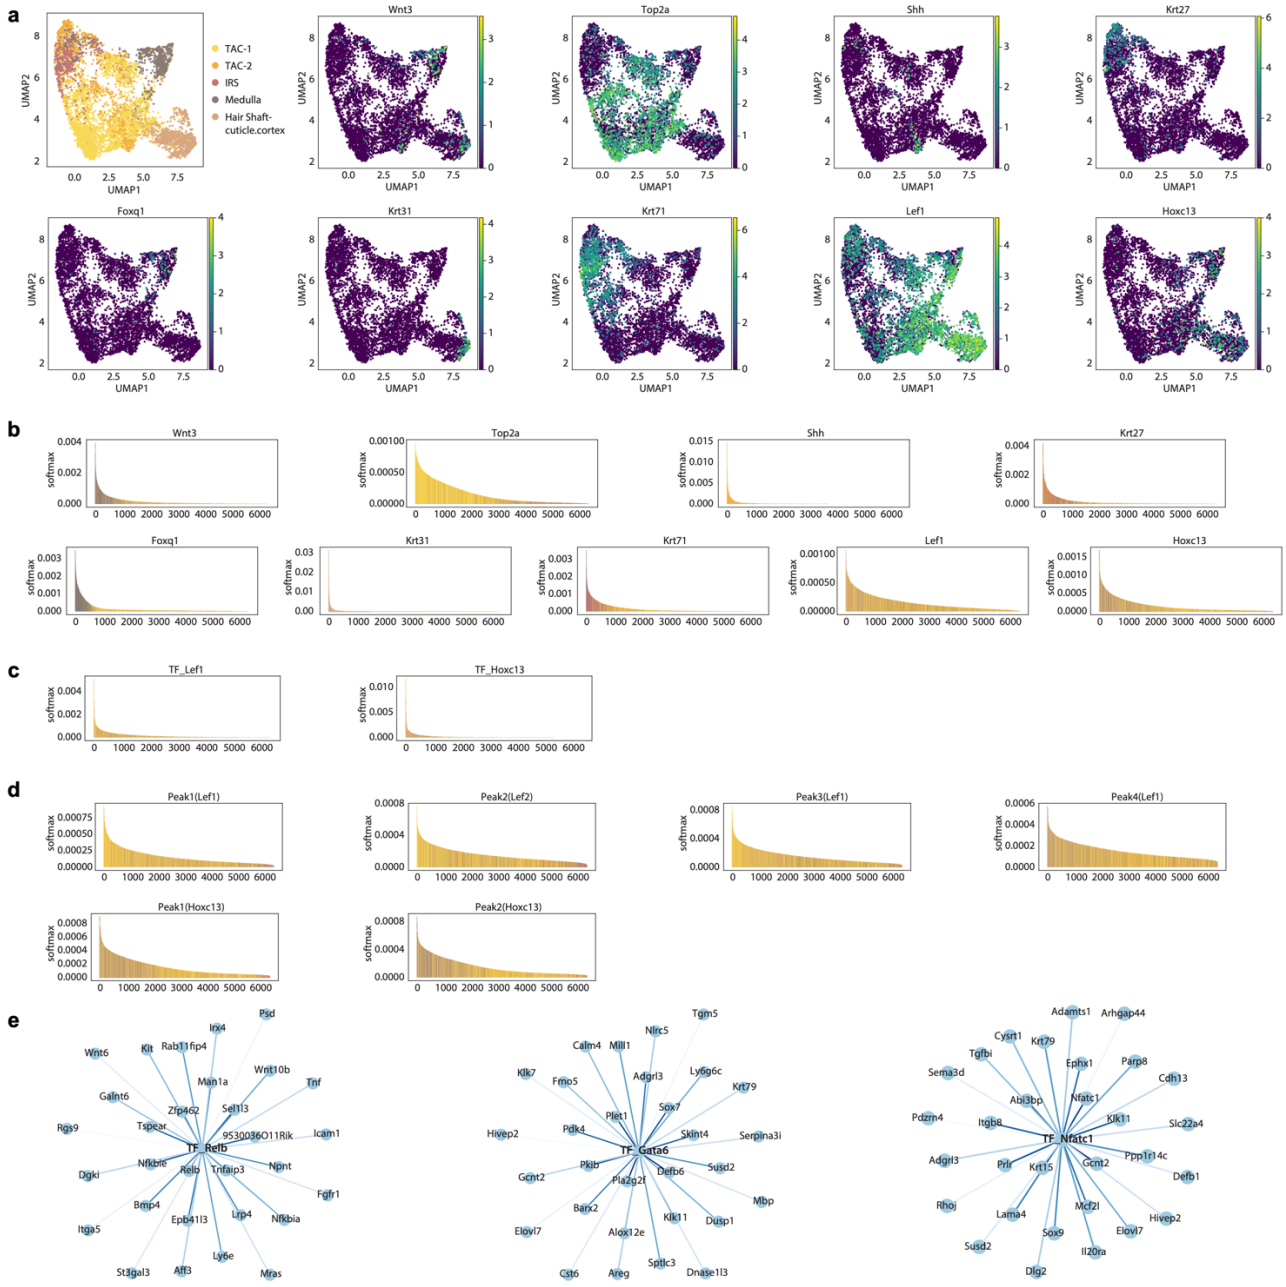

Supplementary Figure 7. Cell type specific marker genes and the target genes of master regulators identified by SIMBA in the SHARE-seq hair follicle subset dataset.

- **a.** UMAP visualization of SIMBA embeddings of cells colored by cell type and gene expression intensity.
- **b.** SIMBA barcode plots of each gene plotted above.
- **c.** SIMBA barcode plots of TF motifs *Lef1* and *Hoxc13.*
- **d.** SIMBA barcode plots of peaks near the loci of *Lef1* and *Hoxc13.*
- **e.** Top 30 target genes of the master regulators *Relb, Gata6,* and *Nfatc1* as inferred by SIMBA.

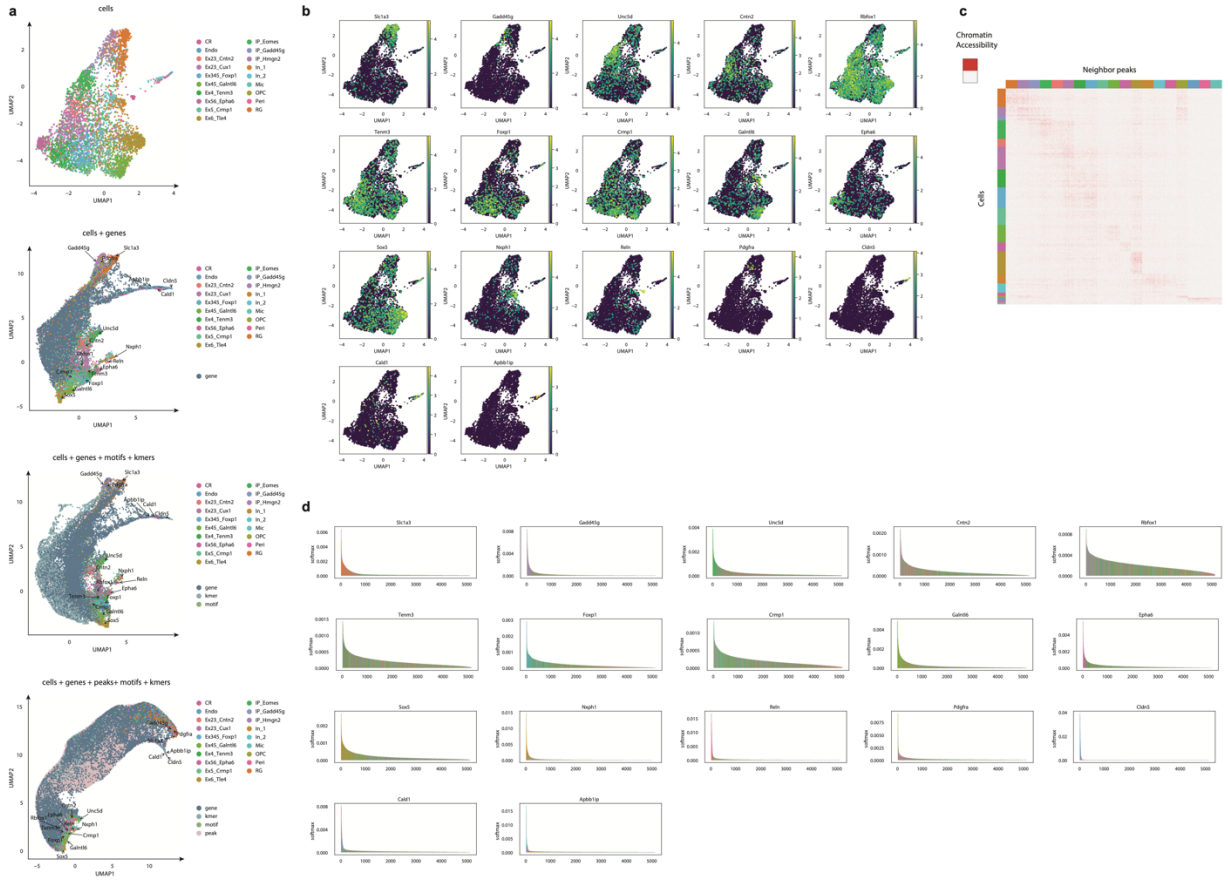

Supplementary Figure 8. SIMBA multimodal analysis of the SNARE-seq mouse cerebral cortex dataset.

- **a.** From top to bottom: UMAP visualization of SIMBA embeddings of (i) cells (ii) genes alongside cells (iii) genes, motifs, and k-mers alongside cells (iv) genes, peaks, motifs, and k-mers alongside cells.
- **b.** UMAP visualization of SIMBA embeddings of cells colored by indicated gene expression intensity.
- **c.** Heatmap of cells against neighboring peaks of each cell type that are selected in the SIMBA co-embedding space. Chromatin accessibility is binary and colored accordingly.
- **d.** SIMBA barcode plots of the genes highlighted in (a).

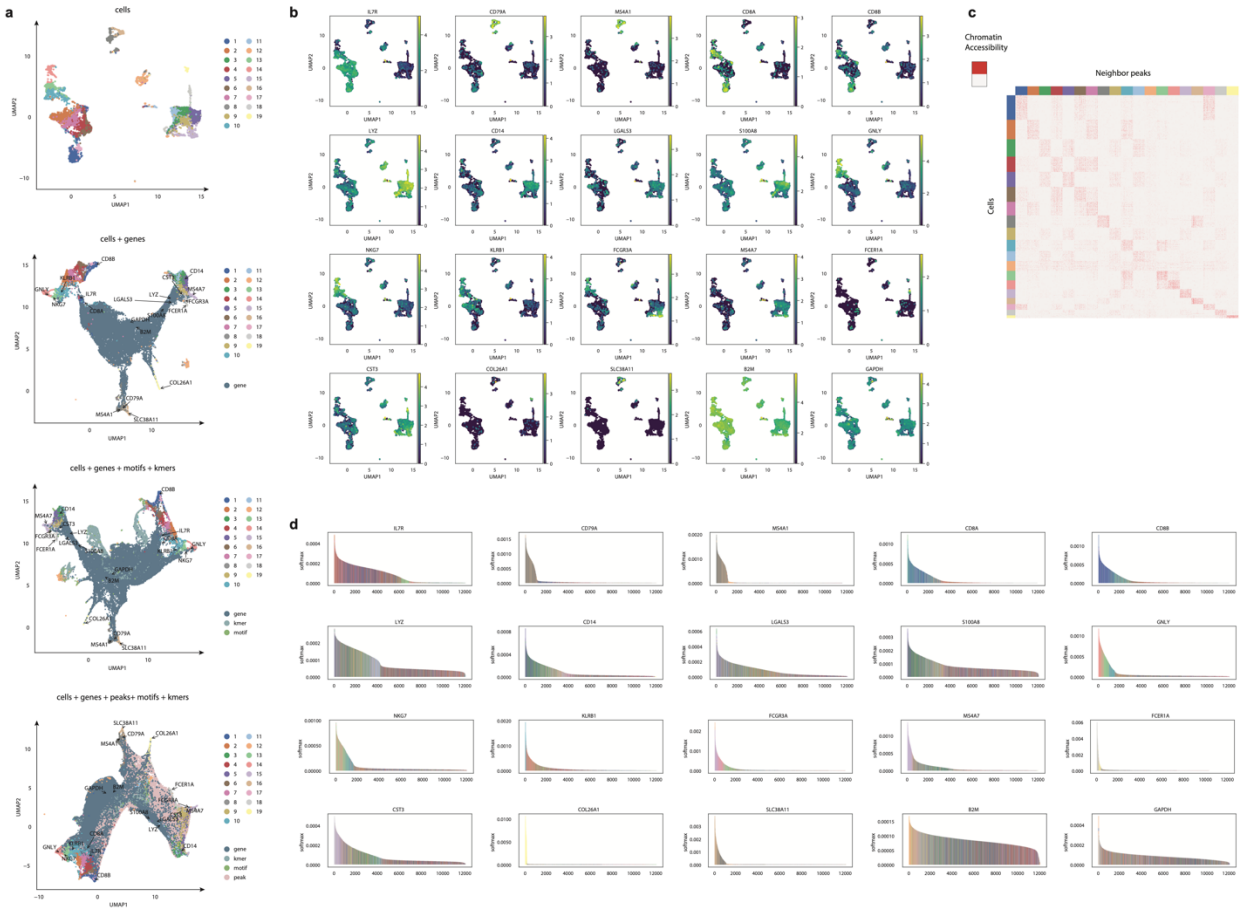

Supplementary Figure 9. SIMBA multimodal analysis of the 10x multiome PBMCs dataset.

- **a.** From top to bottom: UMAP visualization of SIMBA embeddings of (i) cells (ii) genes alongside cells (iii) genes, motifs, and k-mers alongside cells (iv) genes, peaks, motifs, and k-mers alongside cells.
- **b.** UMAP visualization of SIMBA embeddings of cells colored by indicated gene expression intensity.
- **c.** Heatmap of cells against neighboring peaks of each cluster that are selected in the SIMBA co-embedding space. Chromatin accessibility is binary and colored accordingly.
- **d.** SIMBA barcode plots of the genes highlighted in (a).

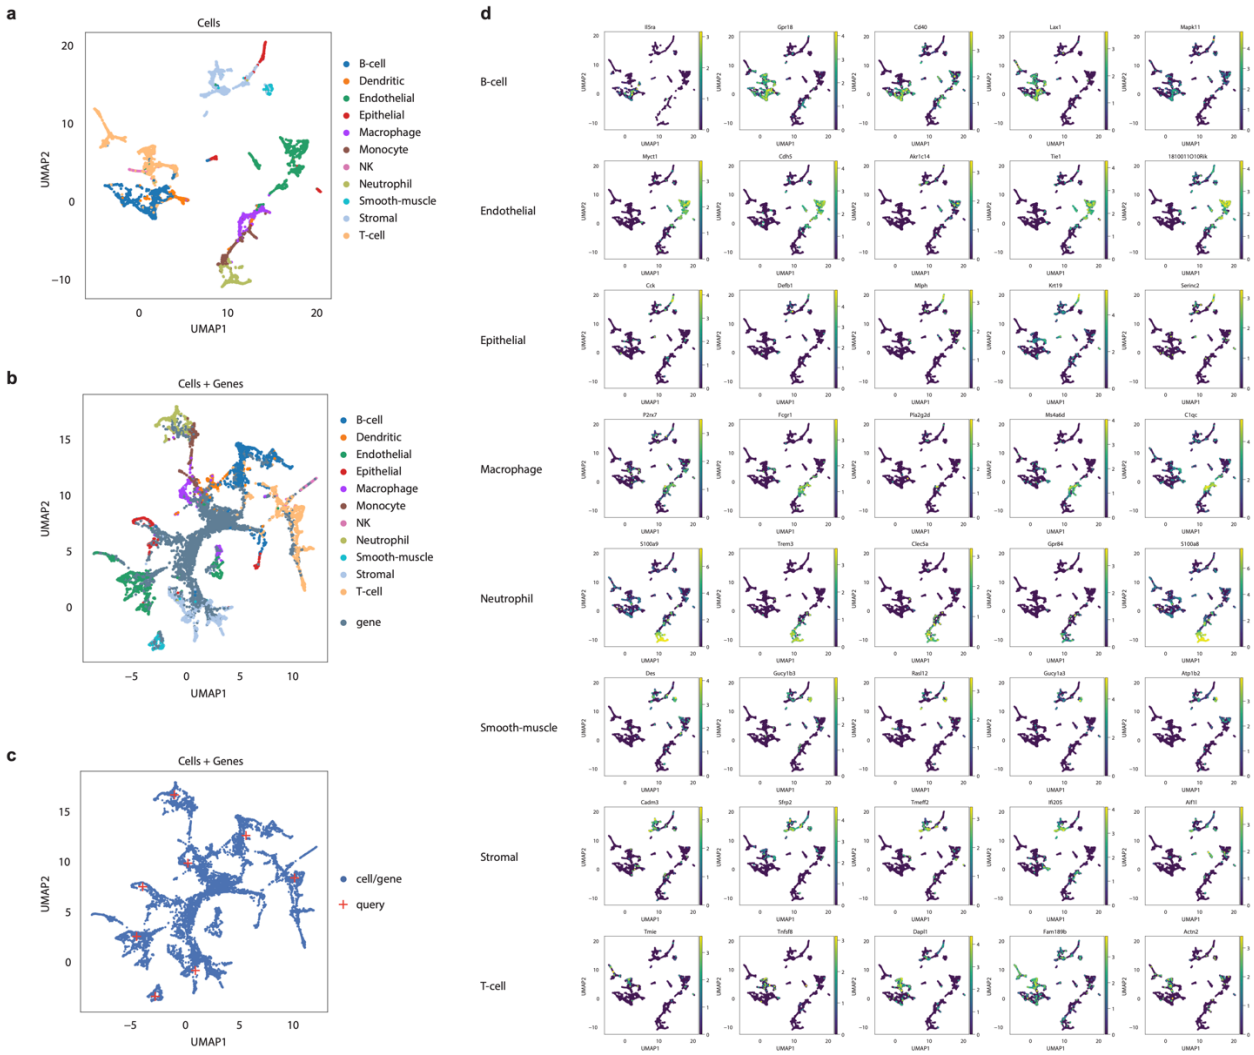

Supplementary Figure 10. SIMBA-inferred marker genes for the scRNA-seq mouse atlas dataset in batch correction analysis.

- **a.** UMAP visualization of SIMBA embeddings of cells colored by cell type.
- **b.** UMAP visualization of SIMBA embeddings of cells and genes.
- **c.** UMAP visualization of SIMBA embeddings of cells and genes. Biological "query" points are highlighted with a red "+". Nearby informative genes are colored accordingly.
- **d.** UMAP visualization of SIMBA embeddings of cells colored by indicated gene expression intensity, separated by cell type.

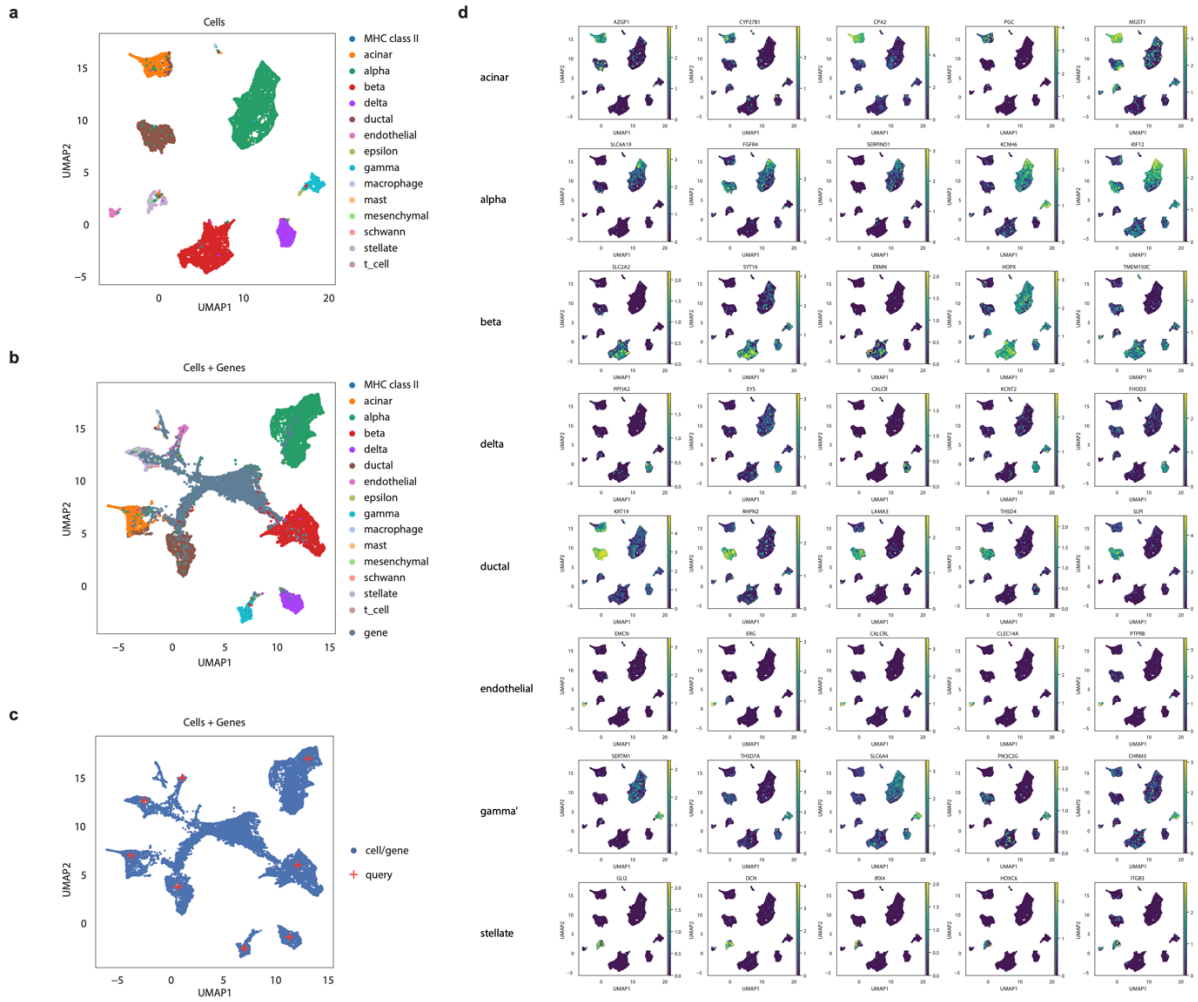

Supplementary Figure 11. SIMBA-inferred marker genes for the scRNA-seq human pancreas dataset in batch correction analysis.

- **a.** UMAP visualization of SIMBA embeddings of cells colored by cell type.
- **b.** UMAP visualization of SIMBA embeddings of cells and genes.
- **c.** UMAP visualization of SIMBA embeddings of cells and genes. Biological "query" points are highlighted with a red "+". Nearby informative genes are colored accordingly.
- **d.** UMAP visualization of SIMBA embeddings of cells colored by indicated gene expression intensity, separated by cell type.

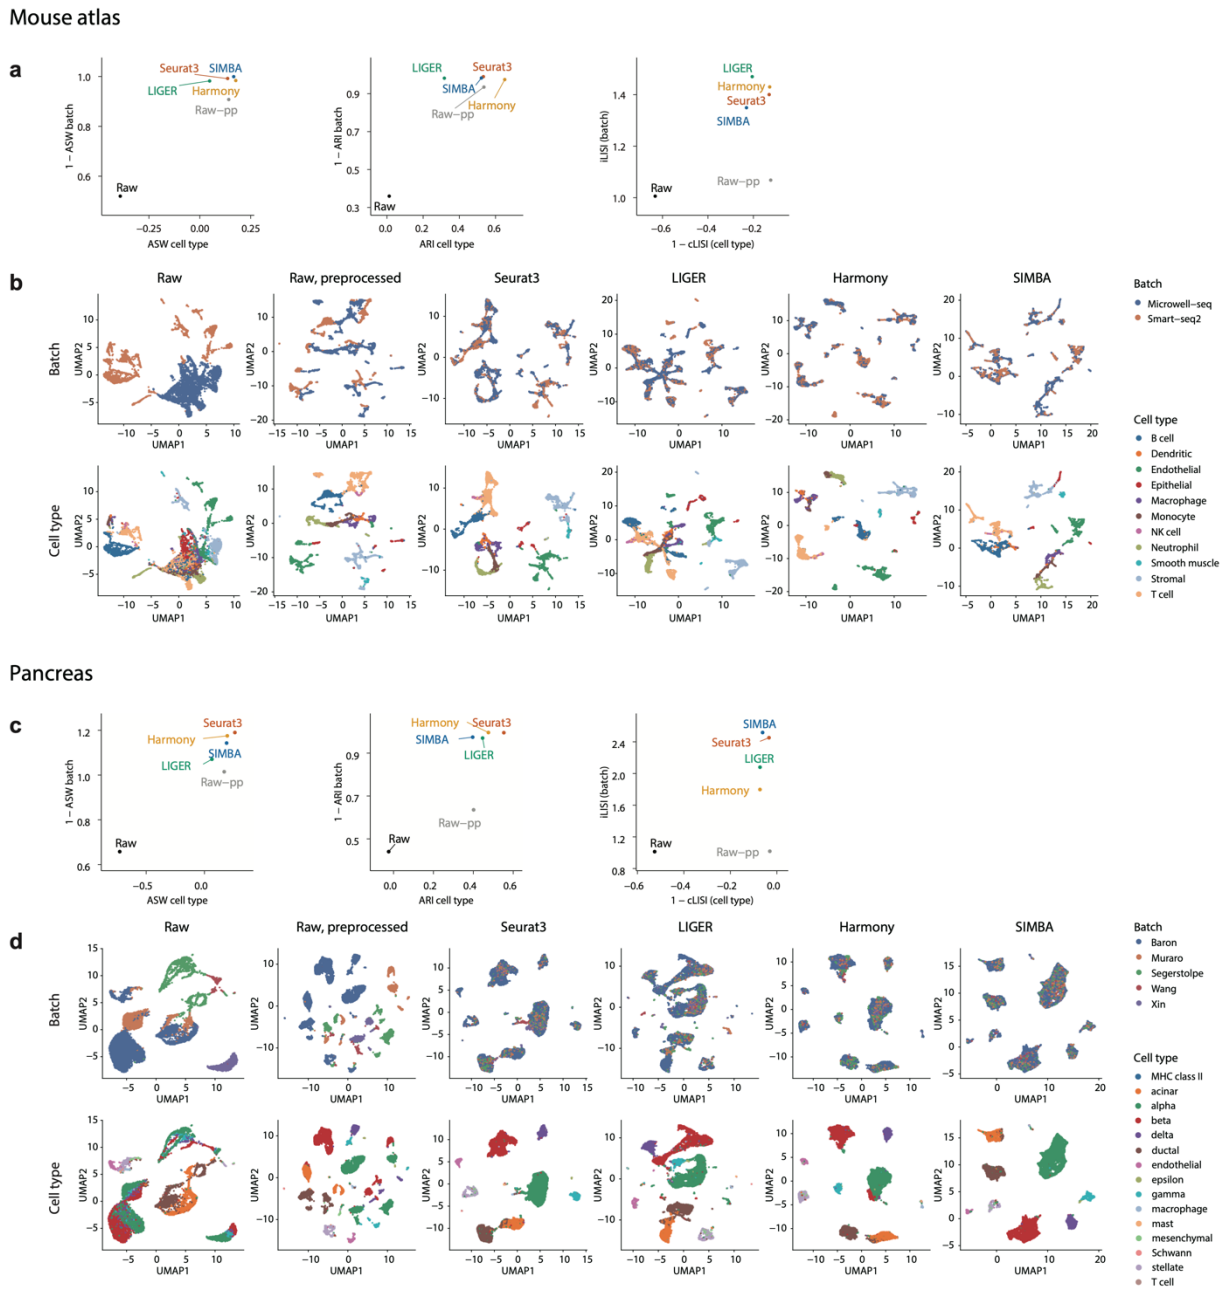

Supplementary Figure 12. Comparison of SIMBA to other methods for batch correction of the mouse atlas (a-b) and human pancreas scRNA-seq datasets (c-d).

a, c. Quantitative comparison of SIMBA with three other batch correction methods including Seurat3, LIGER and Harmony, using, left-to-right: average silhouette width (ASW), adjusted Rand index (ARI), and local inverse Simpson's index (LISI)

b, d. UMAP visualization of raw and preprocessed data alongside the batch corrected results produced by Seurat3, LIGER, Harmony, and SIMBA. Colored by technology (top) and cell type (bottom).

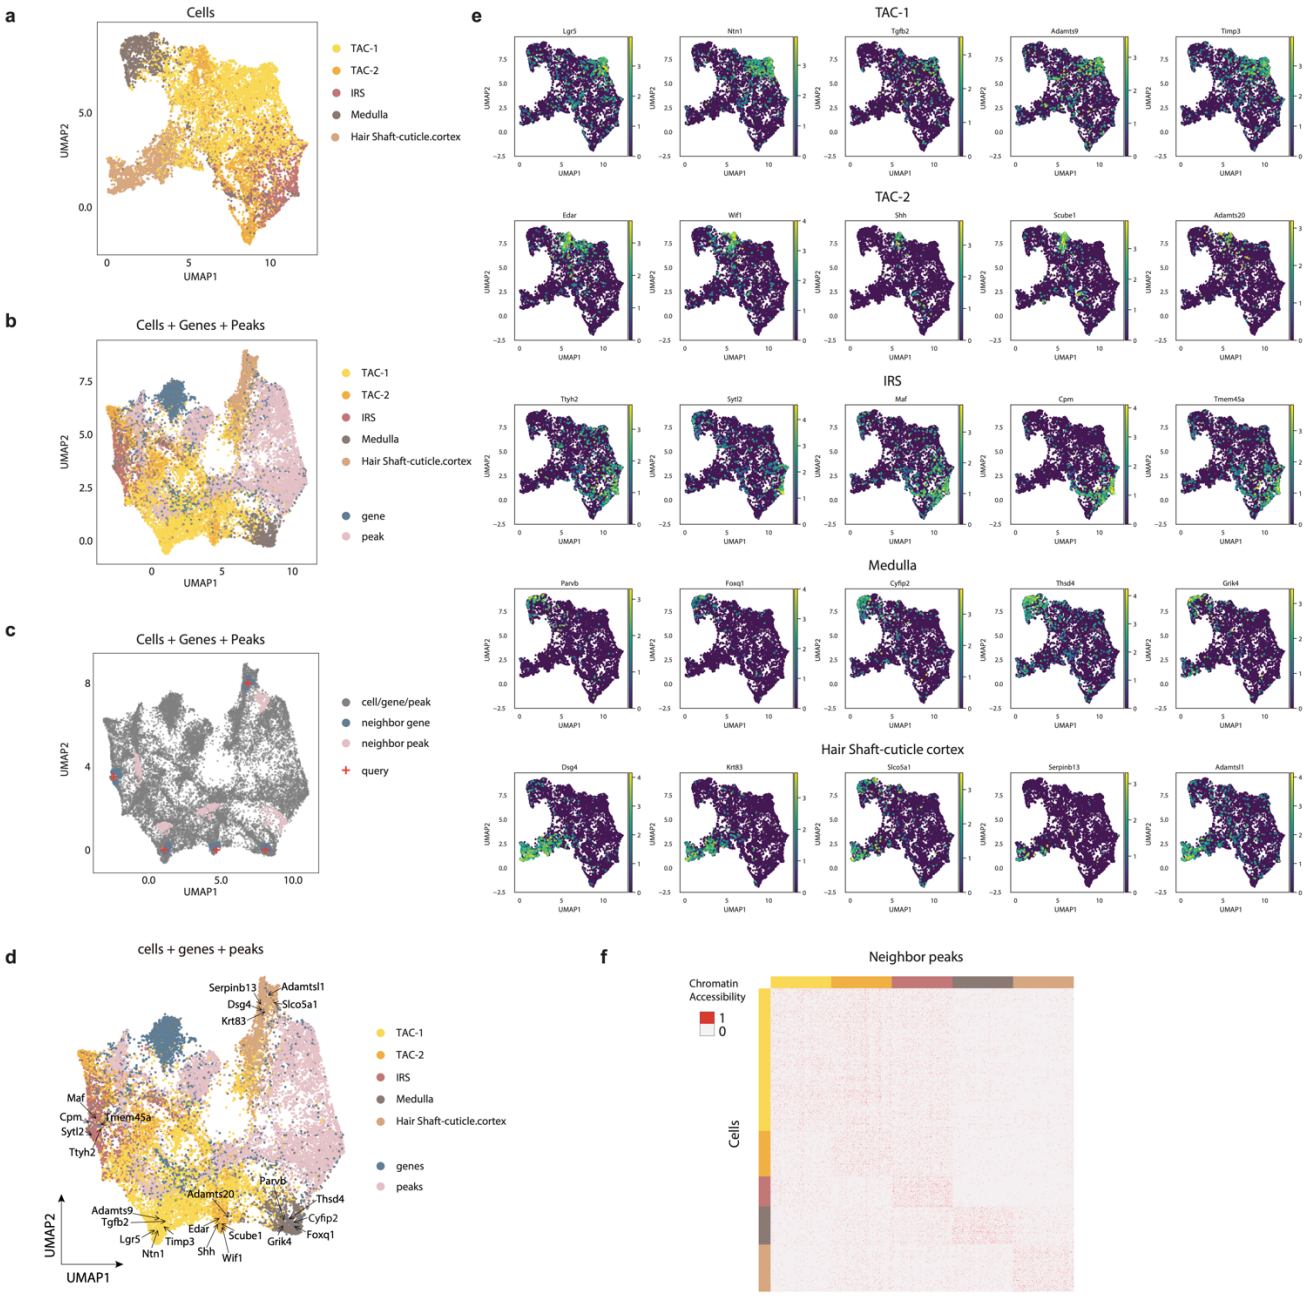

Supplementary Figure 13. SIMBA-inferred marker features for the SHARE-seq mouse skin dataset in multi-omics integration analysis.

- **a.** UMAP visualization of SIMBA embeddings of cells with two cellular modalities integrated.
- **b.** UMAP visualization of SIMBA embeddings of cells, genes, and peaks with two cellular modalities integrated.

c. UMAP visualization of SIMBA embeddings of cells, genes, and peaks with two cellular modalities integrated. Biological “query” points are highlighted with a red “+”. Nearby informative genes and peaks are colored accordingly.

d. UMAP visualization of SIMBA embeddings of cells, genes, and peaks with two cell modalities integrated and known marker genes highlighted.

e. UMAP visualization of SIMBA embeddings of cells colored by indicated gene expression intensity, separated by cell type.

f. Heatmap of cells against neighboring peaks of each cell type that are selected in the SIMBA co-embedding space. Chromatin accessibility is binary and colored accordingly.

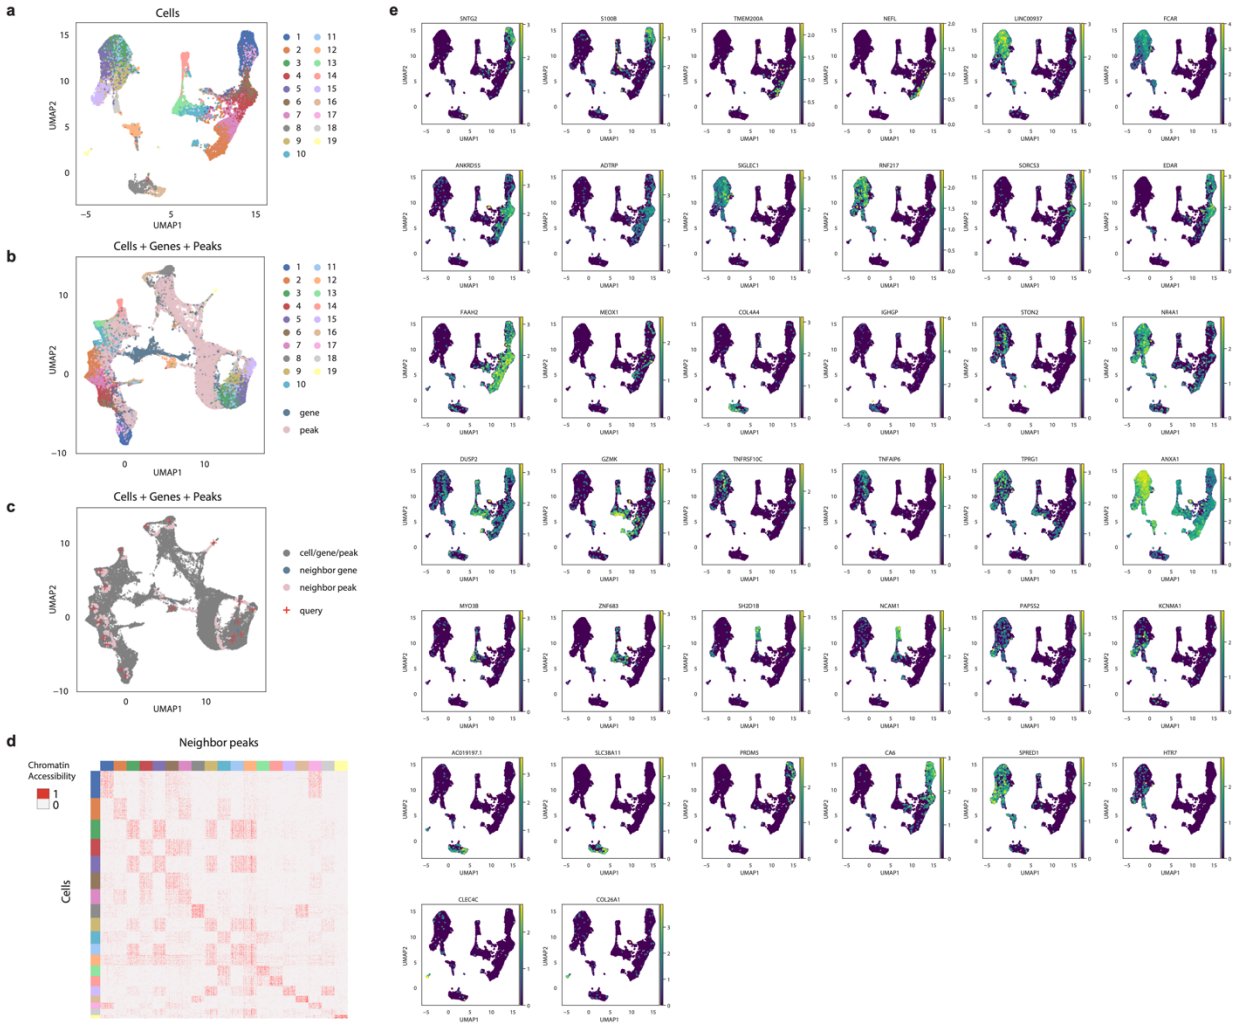

Supplementary Figure 14. SIMBA-inferred marker features for the 10x human PBMCs dataset in multi-omics integration analysis.

- **a.** UMAP visualization of SIMBA embeddings of cells with two cellular modalities integrated.
- **b.** UMAP visualization of SIMBA embeddings of cells, genes, and peaks with two cellular modalities integrated.
- **c.** UMAP visualization of SIMBA embeddings of cells, genes, and peaks with two cellular modalities integrated. Biological "query" points are highlighted with a red "+". Nearby informative genes and peaks are colored accordingly.
- **d.** Heatmap of cells against neighboring peaks of each cluster that are selected in the SIMBA co-embedding space. Chromatin accessibility is binary and colored accordingly.
- **e.** UMAP visualization of SIMBA embeddings of cells colored by indicated gene expression intensity, separated by cell type.

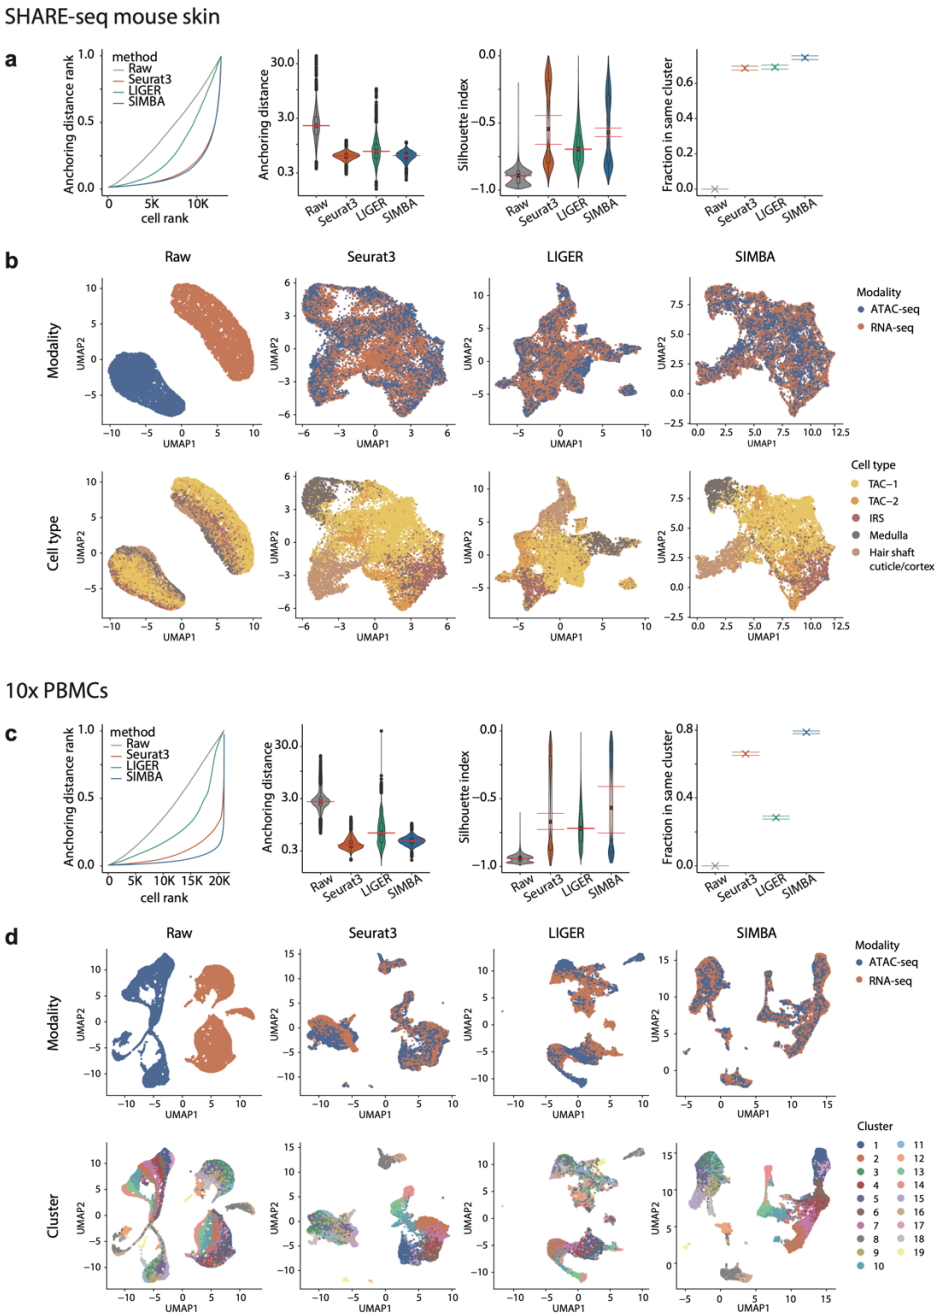

Supplementary Figure 15. Comparison of SIMBA to other methods for multi-omics integration of the SHARE-seq mouse skin (a-b) and 10x multiome human PBMCs (c-d) datasets.

a, c. Quantitative comparison of SIMBA with two other methods including Seurat3, LIGER for multi-omics integration, using, left-to-right: anchoring distance rank, anchoring distance, silhouette index, and Fraction in the same cluster.

b, d. UMAP visualization of the raw scRNA-seq and scATAC-seq data from the 10x multiome human PBMCs dataset alongside the integrated results produced by Seurat3, LIGER, and SIMBA. Colored by data modality (top) and cluster assignment (bottom). The red intervals of violin plot of Anchoring distance and Silhouette index shows the 95% of the mean.

Supplementary Notes

Supplementary Note 1: Comparison with scATAC-seq methods

To assess SIMBA's ability to cluster cell types based on scATAC-seq profiles, we compared SIMBA with specialized methods specifically designed for this task. We observed that SIMBA yields consistent embeddings of cells when using either a single feature (peaks) or multiple features (peaks and DNA sequences from within those peaks) as input to the graph. This comparison was performed across four scATAC-seq datasets of varying profiling technologies and organisms (**Supplementary Fig. 4**). Given these differences, to create a fair comparison we used the same set of features (i.e., peaks) for SIMBA as other methods. SIMBA's performance was compared against three of the top methods, including SnapATAC[4], *Cusanovich*2018[46], and cisTopic[47] recommended by our recent benchmark study[14]. This comparison was first made qualitatively based on UMAP visualization and then quantitatively based on clustering performance. SIMBA performed as well as or better than each of the methods evaluated. These results comparing SIMBA to scATAC-seq-specialized methods highlight SIMBA's wide utility for single-cell analyses (**Supplementary Fig. 5**).

Supplementary Note 2: Comparison with batch correction methods

Multiple methods have now been developed to correct for the technical effects of sample preparation and data collection in single cells. To assess SIMBA's performance in removing batch effects, we compared it to Seurat3[12], LIGER[11] and Harmony[10], three top-performing batch correction methods recommended in a recent benchmark study[25].

Two datasets, including a mouse atlas dataset and a human pancreas dataset (see **Supplementary Table 1**), were used for the evaluation. The mouse atlas dataset is composed of two scRNA-seq subsets with shared cell types from different sequencing platforms. The human pancreas dataset is composed of five samples pooled from five distinct sources using four different sequencing techniques wherein not all cell types are shared across each sample.

To qualitatively compare these methods, we visualized cells of each dataset before and after batch-correction in UMAP plots (**Supplementary Fig. 12b,d**). To quantitatively evaluate the performance of each method, using the benchmarking pipeline laid out in Tran *et al*[25], we measured the conservation of biological information and batch effect removal based on three different metrics: average silhouette width (ASW), adjusted Rand index (ARI), and local inverse Simpson's index (LSI)[10] as in the previously-mentioned benchmark study[25] (**Supplementary Fig. 12a,c**; **Methods**). Each metric measures the relative mixing of class labels, where optimal performance is associated with maximal mixing in the batch labels and minimal mixing in the cell type labels.

The "Raw" batch correction results are the first 50 principal components of the horizontally concatenated gene-by-cell expression count matrix using *stats::prcomp* in R package with

centering and scaling. The "Raw, preprocessed" batch correction used the preprocessed data with log normalization with scaling factor 10^4 and selection of 3000 highly variable genes with Seurat v3 with no restriction on the minimum number of cells and genes.

For batch correction using Seurat v3, default options are used for pancreas dataset whereas for mouse atlas dataset no cutoff was used for the minimum number of cells and genes as in Tran *et al*.[25]. The dimension of the batch corrected embedding is set as 50 dimensions following the default option for *Seurat::RunPCA* and for the consistency with SIMBA.

For batch correction using LIGER, the same arguments are used (lambda = 5, nrep = 3) are used for *liger::optimizeALS* in Tran *et al*. other than the number of factors k was set as 50 for consistency with other methods for both datasets.

For batch correction using Harmony, the same arguments are used as in Tran *et al*.[25] other than the number of dimensions of the output embedding was set to 50 instead of 20. We note that the output embedding of 20 dimensions would result in the similar result as when used 50 dimensions in these methods.

Supplementary Note 3: **Comparison with multi-omics integration methods**

Seurat3 and LIGER are two of the most widely-adopted methods for single-cell data integration. Here, we demonstrate that SIMBA outperforms these methods on two separate datasets, the recently published SHARE-seq mouse skin dataset and the similarly recent 10x PBMCs multiome dataset (**Supplementary Table 1**). We focus on Seurat3 and LIGER as they have explicit documentation for the task of integrating scRNA-seq and scATAC-seq data.

We first qualitatively evaluated these methods by inspecting UMAP visualization plots. For the SHARE-seq dataset, we observed that all three methods perform comparably well in mixing cells of two modalities though LIGER generated particularly small and noisy clusters (**Supplementary Fig. 15b)**. For the 10X PBMCs dataset, SIMBA resulted in the best mixing of cells belonging to each modality whereas other methods clustered cells separately within the originating modalities (**Supplementary Fig. 15d**). We next quantitatively assessed the integration performance of each method using four metrics that measure the distances between matched cells in the integrated space (**Methods**). In addition to the commonly-used metrics including anchoring distance, Silhouette index, and Fraction in the same cluster, we developed an additional metric, *anchoring distance rank* (ADR), which represents the normalized rank of the distance between matching cells. If two matching cells from scRNA-seq and scATAC-seq are mutually closest to one another, their ADR will be close to 0 (**Methods**) and thus a minimized ADR is ideal. Overall SIMBA showed the best performance according to ADR as well as cluster agreement while showing comparable or better performance according to the remaining metrics for both datasets (**Supplementary Fig. 15a,c**).

The modality integration procedure for Seurat v3 and LIGER follows the tutorial provided by the authors (Seurat v3: https://satijalab.org/seurat/archive/v3.1/atacseq_integration_vignette.html; LIGER: http://htmlpreview.github.io/?https://github.com/welch-lab/liger/blob/master/vignettes/Integrating_scRNA_and_scATAC_data.html).

Both Seurat v3 and LIGER formulate the modality integration task between scRNA-seq and scATAC-seq data as a batch correction task between scRNA-seq and gene activity matrix constructed from scATAC-seq. In Seurat v3, the gene activity score of a gene is calculated as the sum of the read counts in the peaks that falls within from 2kb upstream of the TSS to the end of the gene body. In LIGER, this score is calculated as the sum of all read counts that falls within 3kb upstream of the TSS to the end of the gene body.

The "Raw" results start from a scRNA-seq count matrix and a gene activity matrix calculated by Seurat v3. Filtering for the shared genes in both modalities resulted in 16738 genes for the SHARE-seq mouse skin dataset and 11045 genes for the 10X PBMCs dataset. Gene-by-cell gene expression matrix and gene activity matrix were horizontally concatenated along matching rows (genes). The output embedding is the first 20 principal components calculated by the R function *stats::prcomp* with centering and scaling.

For the modality integration using Seurat v3, the gene expression count was filtered using the default parameters *min.cells = 3* and *min.features = 200*. The co-embedding was created as described in the tutorial of the package using the scRNA-seq. The output embedding consists of the first 50 principal components, which is the default option of *Seurat::RunPCA*.

For the modality integration using LIGER, the gene expression count and gene activity matrices were normalized and filtered for the genes that are shared between both matrices. The values were then scaled according to the tutorial. In applying LIGER to the SHARE-seq mouse skin dataset, the function, *liger::optimizeALS* was used with the default parameters, k = 20 and lambda = 5. The scRNA-seq dataset was indicated as the reference in the function, *liger::quantile_norm* as described in the documentation. The scRNA-seq and scATAC-seq modalities of the 10X PBMC multiome dataset were unable to be aligned using the default parameters. Thus *lambda = 30* and *max.iters = 100* were used for the *liger::optimizeALS* function and the scATAC-seq dataset was indicated as the reference using the *liger::quantile_norm* function to ensure a better alignment.

References

1. Satija, R., Farrell, J.A., Gennert, D., Schier, A.F. & Regev, A. Spatial reconstruction of single-cell gene expression data. *Nat Biotechnol* **33**, 495-502 (2015).
2. Wolf, F.A., Angerer, P. & Theis, F.J. SCANPY: large-scale single-cell gene expression data analysis. *Genome Biol* **19**, 15 (2018).
3. Granja, J.M. et al. ArchR is a scalable software package for integrative single-cell chromatin accessibility analysis. *Nat Genet* **53**, 403-411 (2021).

1596 4. Fang, R. et al. Comprehensive analysis of single cell ATAC-seq data with SnapATAC. *Nat*
1597 *Commun* **12**, 1337 (2021).
1598 5. Kiselev, V.Y., Andrews, T.S. & Hemberg, M. Challenges in unsupervised clustering of
1599 single-cell RNA-seq data. *Nat Rev Genet* **20**, 273-282 (2019).
1600 6. Cortal, A., Martignetti, L., Six, E. & Rausell, A. Gene signature extraction and cell identity
1601 recognition at the single-cell level with Cell-ID. *Nat Biotechnol* (2021).
1602 7. Vandenbon, A. & Diez, D. A clustering-independent method for finding differentially
1603 expressed genes in single-cell transcriptome data. *Nat Commun* **11**, 4318 (2020).
1604 8. Dann, E., Henderson, N.C., Teichmann, S.A., Morgan, M.D. & Marioni, J.C. Differential
1605 abundance testing on single-cell data using k-nearest neighbor graphs. *Nat Biotechnol*
1606 (2021).
1607 9. Hao, Y. et al. Integrated analysis of multimodal single-cell data. *Cell* (2021).
1608 10. Korsunsky, I. et al. Fast, sensitive and accurate integration of single-cell data with
1609 Harmony. *Nat Methods* **16**, 1289-1296 (2019).
1610 11. Welch, J.D. et al. Single-Cell Multi-omic Integration Compares and Contrasts Features of
1611 Brain Cell Identity. *Cell* **177**, 1873-1887 e1817 (2019).
1612 12. Stuart, T. et al. Comprehensive Integration of Single-Cell Data. *Cell* **177**, 1888-1902
1613 e1821 (2019).
1614 13. Lerer, A. et al. Pytorch-biggraph: A large-scale graph embedding system. *arXiv preprint*
1615 *arXiv:1903.12287* (2019).
1616 14. Chen, H. et al. Assessment of computational methods for the analysis of single-cell
1617 ATAC-seq data. *Genome Biology* **20**, 241 (2019).
1618 15. Buenrostro, J.D. et al. Integrated Single-Cell Analysis Maps the Continuous Regulatory
1619 Landscape of Human Hematopoietic Differentiation. *Cell* **173**, 1535-1548 e1516 (2018).
1620 16. Ferreira, R., Ohneda, K., Yamamoto, M. & Philipsen, S. GATA1 function, a paradigm for
1621 transcription factors in hematopoiesis. *Molecular and cellular biology* **25**, 1215-1227
1622 (2005).
1623 17. Tijchon, E., Havinga, J., Van Leeuwen, F. & Scheijen, B. B-lineage transcription factors
1624 and cooperating gene lesions required for leukemia development. *Leukemia* **27**, 541-
1625 552 (2013).
1626 18. Friedman, A. Transcriptional control of granulocyte and monocyte development.
1627 *Oncogene* **26**, 6816-6828 (2007).
1628 19. Schep, A.N., Wu, B., Buenrostro, J.D. & Greenleaf, W.J. chromVAR: inferring
1629 transcription-factor-associated accessibility from single-cell epigenomic data. *Nat*
1630 *Methods* **14**, 975-978 (2017).
1631 20. Moriguchi, T. & Yamamoto, M. A regulatory network governing Gata1 and Gata2 gene
1632 transcription orchestrates erythroid lineage differentiation. *International journal of*
1633 *hematology* **100**, 417-424 (2014).
1634 21. Chen, S., Lake, B.B. & Zhang, K. High-throughput sequencing of the transcriptome and
1635 chromatin accessibility in the same cell. *Nat Biotechnol* (2019).
1636 22. Ma, S. et al. Chromatin Potential Identified by Shared Single-Cell Profiling of RNA and
1637 Chromatin. *Cell* (2020).
1638 23. Cao, J. et al. Joint profiling of chromatin accessibility and gene expression in thousands
1639 of single cells. *Science* **361**, 1380-1385 (2018).

1640 24. Zhu, C. et al. An ultra high-throughput method for single-cell joint analysis of open
1641 chromatin and transcriptome. *Nat Struct Mol Biol* (2019).
1642 25. Tran, H.T.N. et al. A benchmark of batch-effect correction methods for single-cell RNA
1643 sequencing data. *Genome Biol* **21**, 12 (2020).
1644 26. Leek, J.T. et al. Tackling the widespread and critical impact of batch effects in high-
1645 throughput data. *Nature Reviews Genetics* **11**, 733-739 (2010).
1646 27. Han, X. et al. Mapping the mouse cell atlas by microwell-seq. *Cell* **172**, 1091-1107. e1017
1647 (2018).
1648 28. Tabula Muris, C. et al. Single-cell transcriptomics of 20 mouse organs creates a Tabula
1649 Muris. *Nature* **562**, 367-372 (2018).
1650 29. Baron, M. et al. A single-cell transcriptomic map of the human and mouse pancreas
1651 reveals inter-and intra-cell population structure. *Cell systems* **3**, 346-360. e344 (2016).
1652 30. Muraro, M.J. et al. A single-cell transcriptome atlas of the human pancreas. *Cell systems*
1653 **3**, 385-394. e383 (2016).
1654 31. Wang, Y.J. et al. Single-cell transcriptomics of the human endocrine pancreas. *Diabetes*
1655 **65**, 3028-3038 (2016).
1656 32. Segerstolpe, Å. et al. Single-cell transcriptome profiling of human pancreatic islets in
1657 health and type 2 diabetes. *Cell metabolism* **24**, 593-607 (2016).
1658 33. Palla, G. et al. Squidpy: a scalable framework for spatial single cell analysis. *bioRxiv*
1659 *2021.02.19.431994v2* (2021).
1660 34. VanHorn, S. & Morris, S.A. Next-Generation Lineage Tracing and Fate Mapping to
1661 Interrogate Development. *Dev Cell* **56**, 7-21 (2021).
1662 35. Ietswaart, R., Gyori, B.M., Bachman, J.A., Sorger, P.K. & Churchman, L.S. GeneWalk
1663 identifies relevant gene functions for a biological context using network representation
1664 learning. *Genome Biol* **22**, 55 (2021).
1665 36. Yuan, H., Kshirsagar, M., Zamparo, L., Lu, Y. & Leslie, C.S. BindSpace decodes
1666 transcription factor binding signals by large-scale sequence embedding. *Nat Methods*
1667 (2019).
1668 37. Li, H., Xiao, X., Wu, X., Ye, L. & Ji, G. scLINE: A multi-network integration framework
1669 based on network embedding for representation of single-cell RNA-seq data. *J Biomed*
1670 *Inform* **122**, 103899 (2021).
1671 38. Buterez, D., Bica, I., Tariq, I., Andrés-Terré, H. & Liò, P. CELLVGAE: AN UNSUPERVISED
1672 SCRNA-SEQ ANALYSIS WORKFLOW WITH GRAPH ATTENTION NETWORKS. *bioRxiv*
1673 *2020.12.20.423645v1* (2020).
1674 39. Sheikh, N., Kefato, Z. & Montresor, A. gat2vec: representation learning for attributed
1675 graphs. *Computing* **101**, 187-209 (2018).
1676 40. Bojchevski, A. & Günnemann, S. Deep gaussian embedding of graphs: Unsupervised
1677 inductive learning via ranking. *arXiv preprint arXiv:1707.03815* (2017).
1678 41. Satopaa, V., Albrecht, J., Irwin, D. & Raghavan, B. in 2011 31st international conference
1679 on distributed computing systems workshops 166-171 (IEEE, 2011).
1680 42. Fornes, O. et al. JASPAR 2020: update of the open-access database of transcription
1681 factor binding profiles. *Nucleic acids research* **48**, D87-D92 (2020).
1682 43. Kadlec, R., Bajgar, O. & Kleindienst, J. Knowledge base completion: Baselines strike back.
1683 *arXiv preprint arXiv:1705.10744* (2017).

44. Krompaß, D., Baier, S. & Tresp, V. in International semantic web conference 640-655 (Springer, 2015).
45. Mikolov, T., Chen, K., Corrado, G. & Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
46. Cusanovich, D.A. et al. The cis-regulatory dynamics of embryonic development at single-cell resolution. *Nature* **555**, 538-542 (2018).
47. Bravo Gonzalez-Blas, C. et al. cisTopic: cis-regulatory topic modeling on single-cell ATAC-seq data. *Nat Methods* **16**, 397-400 (2019).
48. Dou, J. et al. Unbiased integration of single cell multi-omics data. *bioRxiv*, 2020.2012.2011.422014 (2020).