Identification of Distinct Topological Structures From High-Dimensional Data

Bingxian Xu[1,2] and Rosemary Braun[1,2,3,4,5,6,*]

[1]Department of Molecular Biosciences, Northwestern University, Evanston, IL 60208, USA
[2]NSF-Simons National Institute for Theory and Mathematics in Biology, Chicago, IL 60611, USA
[3]Department of Engineering Sciences and Applied Mathematics, Northwestern University, Evanston, IL 60208, USA
[4]Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208, USA
[5]Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL 60208, USA
[6]Santa Fe Institute, Santa Fe, NM 87501, USA
[*]To whom correspondence should be addressed. Email: rbraun@northwestern.edu

Abstract

Single-cell RNA sequencing allows the direct measurement of the expression of tens of thousands of genes, providing an unprecedented view of the transcriptomic state of a cell. Within each cell, different biological processes such as differentiation or cell cycle take place simultaneously, each providing a different characterization of cell state. To identify gene sets that govern these processes for the purpose of disentangling convolved biological processes, we present "Identification of Distinct topological structures" (ID). ID works by constructing an alternative low-dimensional parametrization of the high-dimensional system, applying a finite perturbation to this alternative parametrization, and looking for genes that respond similarly. With this approach, we demonstrate that ID is capable of identifying structures within the data that will otherwise be missed. We further demonstrate the utility of ID in scRNA-seq datasets collected under various backgrounds, delineating cellular differentiation, characterizing cellular response to external perturbation, and dissecting the effect of genetic knock-outs.

1 Introduction

Life is composed of dynamic processes such as differentiation [1, 2, 3, 4], circadian rhythm [5, 6, 7, 8, 9], cell cycle [10, 11, 12, 13] and so on. Together, these processes define cellular state in a multi-faceted way. One may characterize the gene expression patterns associated with these processes by the topology of cells in the gene expression space. For example, the topology of the cell cycle should be a ring, and that of differentiation may look like a tree. Because of the large number of features assayed by state-of-the-art transcriptomic profiling technologies, a smaller set of feature genes is selected first on the basis of their signal-to-noise ratios. By defining cell-cell distances using feature genes, methods such as Uniform Manifold Approximation and Projection [14] (UMAP) are employed to map the data to a lower dimensional space that allows visualization of the topology/structure of the data [15, 16, 17, 18, 19]. However, different biological processes may give rise to different notions of "distance"; for instance, two cells of the same cell type (close on one set of axes) may not be in the same phase of the cell cycle (distant on another). This simple observation suggests that distances between cells defined by different gene sets need not agree with each other. When this happens, the distance matrix one obtains from a generic set of feature genes will not represent the true relationship between the cells. To overcome this, one must identify the genes that define these different biological processes and use them individually to unravel the topology of the data.

This problem is traditionally "solved" by some form of matrix factorization, where the original gene expression matrix is decomposed into a sum of simpler components. Constraints have been imposed on the factorized data to enhance interpretability. For example, by requiring the factorized matrix to be strictly positive, as is done in non-negative matrix factorization (NMF) [20, 21, 22, 23, 24], the gene expression of each cell can be viewed as the sum of the different gene programs it activates. However,

gene programs identified by NMF can still lack interpretability [25, 26, 27]. Additional constraints have been imposed by using prior knowledge of the system. For example, Spectra [27] uses user-defined gene sets to guide the factorization of the gene expression matrix. Similarly, CellUntangler [28] uses a variational auto-encoder to construct a latent space that best captures the topology of a user-defined gene set. In real-world settings, the best–known marker genes of a biological process may not have been detected in the data due to low read depth, which limits the applicability of the aforementioned methods. However, as these markers often drive the regulation of other genes, the data can retain signatures of the relevant topology even as the marker genes themselves remain undetected. Following this logic, scPrisma [29] exploits prior information about the topology of the data (e.g., the existence of periodic or linear signals) by conducting spectral template matching between the spectrum of the data and the theoretical covariate spectrum expected from a user-defined topology. Still, especially during the data exploration phase, neither the topology nor the relevant gene programs may known ahead of time. Because of this, an *unsupervised* algorithm is needed to identify gene sets that define the various aspects of the data. GeneTrajectory [30] is built for this exact purpose. Through the construction of a gene-gene distance matrix via conducting optimal transport on a cell-cell graph, GeneTrajectory constructs an embedding for genes that allows direct visualization of the topological structures present in the data and identifies genes responsible for each structure. However, the construction of a cell-cell graph implies that the data contain only a single underlying biological process. Imagine that a population of cells is both undergoing division (cell cycle) and differentiation, the cell-cell graphs defined by the cell cycle genes, the differentiation genes, and the union of all genes will be different. In this situation, it seems to be more reasonable to define the relationship between cells with multiple cell-cell graphs. For this reason, when GeneTrajectory is applied to real scRNA-seq data, it typically constructs a continuous landscape of genes, where the boundary between different sets is blurred by the use of a single cell-cell graph and potentially the noisiness of the data. Consequently, this prevents the correct classification of genes[30].

To address the difficulty of conducting unsupervised identification of gene sets that define different topological structures, we developed *Identification of Distinct topological structures (ID)*, which takes only the count matrix as input to generate an embedding for genes with which they can be clustered. ID achieves state-of-the-art performance by first constructing an alternative low–dimensional representation of the system and then identifying genes that respond similarly when the alternative parametrization is perturbed. The intuition behind ID is that genes that are related to the same biological process would respond similarly to the same perturbation. We tested the performance of ID in multiple toy datasets and evaluated how it compared to that of existing unsupervised approaches. We also applied ID on multiple scRNA-seq datasets to unravel structures hidden within the data, and uncover how cells respond to both environmental and genetic perturbation. Lastly, using two independent human lung datasets, we "IDed" a conserved gene set that robustly defines lung epithelial differentiation.

2 Methods

The state of a cell can be defined by the genes it expresses, and the distance between two cells is conveniently defined by the difference between their states. These distances can be used to visualize the cell state manifold and deduce cellular trajectories [31, 17, 32]. However, if the distances defined by different subsets of genes disagree with one another, the cell state manifold they construct would also differ; what, then, is the "correct" cell state manifold?

To reconcile the disagreement between cell–cell distances defined by different gene sets, we developed *Identification of Distinct topological structures (ID)* to identify genes that define cell–cell distances in similar ways. The key output of ID is an embedding of genes, where each cluster of genes defines similar cell–cell distances. In the following sections, we describe the ID algorithm in detail.

2.1 The ID algorithm

Given a raw data matrix, we first standardize the mean and variance of each feature by Z-scoring across samples to obtain the final data matrix we use as input, $\mathbf{X} \in \mathbb{R}^{s \times n}$, where s and n are the number of samples and features respectively. For scRNA-seq count matrices, we would first normalize by the library size of the cell before Z-scoring by each feature. Then, we train a variational autoencoder [33]

to map x_s (the data vector for sample s) to a latent space z and then back to x_s. This VAE is the core of ID and we describe its optimization in the next section. Next, we construct a response matrix, $\mathbf{M} \in \mathbb{R}^{n \times t}$, where t is the number of computational perturbations. Each entry of $\mathbf{M}$, $\mathbf{M}_{i,j}$, denotes how feature i responded to the jth perturbation and is the key to our identification of related features. The ith column of $\mathbf{M}$, denoted by $\mathbf{M}_{(i)}$ is computed by:

$$\mathbf{M}_{(i)} = |f(z) - f(z + \delta z)| \tag{1}$$

where $f(\cdot)$ is the decoder network, z is a point in the latent space and δz a random direction with a fixed magnitude, which we have set to $|\delta z| = 0.1$ by default. It is easy to see that $\mathbf{M}_{(i)}$ represents how the feature space would respond to minor perturbation to its latent representation, as illustrated in Figure 1A and B. By sampling multiple points in the latent space as well as random directions, and plugging them into Eq 1, we assemble the response matrix, $\mathbf{M}$. Finally, we scale each column of $\mathbf{M}$ via Z-scoring to obtain the scaled response matrix, $\hat{\mathbf{M}}$.

By default, we set the number of trials, t, to $50,000$. To obtain the final cluster assignment for each gene, we first use principal component analysis to extract the top 20 PCs of $\hat{\mathbf{M}}$ and subsequently map it to a 2D plane with UMAP (**n_neighbour** $= 5$) for clustering.

2.2 VAE optimization

By default, ID uses a 2–layer encoder neural network with 300 fully connected neurons per layer and ReLU activation to project count data onto a 50–dimensional latent space (z). Subsequently, z_s are used as input to a decoder network to reconstruct the real data. The loss function was simply the mean squared error between the input and the reconstructed data. During each training loop, we focus only on a randomly selected small subset of the data, by default 150. In all test cases, the optimization of the model parameters was done with the ADAM [34] optimizer as implemented by pytorch [35] with a learning rate of 0.0005, $\beta_1 = 0.8$, $\beta_2 = 0.9$, and a weight decay of 0.0001. No scheduler was used to change the learning rate during the training process.

2.3 Application of alternative approaches

2.3.1 Non-negative matrix factorization

Non-negative matrix factorization aims to deconstruct the original data matrix $\mathbf{X}$ into multiple components:

$$\mathbf{X} = \mathbf{W}\mathbf{H}, \tag{2}$$

where $\mathbf{W} \in \mathbb{R}^{s \times n_c}$, $\mathbf{H} \in \mathbb{R}^{n_c \times n}$ and n_c denotes the number of components the data matrix has been broken into. We used the matrix $\mathbf{H}$ to cluster genes in two ways. First, we used $\mathbf{H}$ directly, where each gene is defined by a vector of dimension n_c. Alternatively, we first reduce the dimensionality of $\mathbf{H}$ before conducting clustering on the low-dimensional space. In all cases, NMF is implemented by `sklearn.decomposition.NMF` with default parameters.

2.3.2 Applying geneTrajectory

GeneTrajectory is applied with code downloaded from https://github.com/KlugerLab/GeneTrajectory with default parameters.

2.3.3 Constructing an embedding for genes using UMAP

By default, we constructed an embedding for genes with UMAP by first transposing the data matrix, $\mathbf{X}$, after normalization and extracting top 20 principal components before applying UMAP (**n_neighbour**=5).

2.4 Computing cross entropy between k-nearest neighbor graphs

We computed the cross–entropy between KNN graph constructed using the original space and the KNN graph constructed using UMAP embeddings to quantify the extent of their differences. The

fuzzy set cross entropy [14] is defined as:

$$C(\mathbf{p},\mathbf{q}) = \sum_{i,j} p_{ij} \log\left(\frac{p_{ij}}{q_{ij}}\right) + (1-p_{ij})\log\left(\frac{1-p_{ij}}{1-q_{ij}}\right)$$
$$= \sum_{i,j} p_{ij}\log(p_{ij}) + (1-p_{ij})\log(1-p_{ij}) - \left(\sum_{i,j} p_{ij}\log(q_{ij}) + (1-p_{ij})\log(1-q_{ij})\right), \quad (3)$$

where $\mathbf{p}$ and $\mathbf{q}$ are the KNN graphs constructed from the high-dimensional and low-dimensional space respectively. It is easy to see that only the second term:

$$-\left(\sum_{i,j} p_{ij}\log(q_{ij}) + (1-p_{ij})\log(1-q_{ij})\right), \quad (4)$$

is dependent on the quality of the UMAP projection, and hence we use this term to quantify the difference between the two KNN graphs. However, as both graphs are sparse, the result obtained by the direct application of Eq 4 will be dominated by unconnected nodes. To alleviate this problem, we sample 1000 positive edges from $\mathbf{p}$, and for each positive edge we sample 5 negative edges. This balanced set of positive and negative edges is used as input for Eq 4. For each pair of graphs, $\mathbf{p}$ and $\mathbf{q}$, we repeat this process 20000 times.

2.5 Datasets

The toy datasets we used are constructed in two steps. First, we define the topology of the data on a low-dimensional space. Then, via a topology preserving transformation, we lift the low-dimensional data into higher dimensions. For the "linear", "2-branch", and "4-branch" toy datasets, we define low-dimensional topology using the scattered dots illustrated in Figure S1B. Next, we use UMAP to project them from a low-dimensional space to a high-dimensional space. For the toy data with complex topology, we first sampled from the 3D shapes directly before projecting them onto high dimensions using UMAP. Raw data containing the mesh representation of these 3D shapes can be found at https://www.craftsmanspace.com/free-3d-models/topological-mesh-modeling-examples.html.

For real scRNA-seq datasets, we preprocessed them all in the same way where we first normalized each cell by its library size and then normalized each gene by z-scoring. 3000 highly variable genes were selected by scanpy [36] prior to applying ID. See below for detailed descriptions of the datasets we used:

2.5.1 Blood differentiation

The blood differentiation dataset [37, 38] was downloaded from https://github.com/HaghverdiLab/velocity_notebooks/tree/main/datasets/HSPC. This dataset contains 2430 cells with a mean UMI 17925. It has been enriched for long-term hematopoetic stem cells. Cells within this dataset have not reached final states of mature blood cells.

2.5.2 Mouse hippocampus

The mouse hippocampus dataset [39, 40] was downloaded from http://pklab.med.harvard.edu/velocyto/DentateGyrus. It contains 18213 cells with a mean UMI of 4677.

2.5.3 Mouse microglia

The mouse microgila dataset [41] was downloaded from https://cellxgene.cziscience.com/collections/4828d33d-fb26-42e7-bf36-18293b0eec85. This dataset contained 12330 sorted microglia with a mean UMI of 3682.

2.5.4 Mouse hair follicle

The mouse hair follicle dataset [30] was downloaded from https://figshare.com/articles/dataset/Processed_Seurat_objects_for_GeneTrajectory_inference_Gene_Trajectory_Inference_for_Single-cell_Data_by_Optimal_Transport_Metrics_/25243225. The dataset contains both wild type cells and cells with Wnt knock-out. It contains a total of 10328 cells with a mean UMI of 13948.

2.5.5 Human lung samples

The human lung organoid data [42] was downloaded from https://cellxgene.cziscience.com/collections/e9cf4e8d-05ed-4d95-b550-af35ca219390. It is collected from two separate human donors, dissociated and used to generate lung organoid (O1). O1 derived transitionally differentiated (TD) cells are cultured in lung organoid growth medium, PneumaCult EX-plus medium and PneumaCult-ALI maintenance medium to generate the O2, EX and the ALIEX model. This data contains 15806 cells with a mean UMI of 3833.

The human lung sample data [43] was downloaded from https://cellxgene.cziscience.com/collections/73cf6939-3caa-4105-bc57-e073ee885a28. It is obtained from two human transplant donor lungs that are considered unsuitable for transplantation. Cells are isolated from the bronchial rings, dissociated and then cultured in four different culture media. This dataset contains 10224 cells with a mean UMI 2826.

2.6 Quantifying the extent of randomness across gene sets

Suppose G_1 and G_2 are the feature genes from dataset 1 and dataset 2 respectively. We are interested in quantifying if an ID-derived subset of G_1, S_1, and an ID-derived subset of G_2, S_2 have significant overlap. This is achieved by randomly sampling $|S_1|$ genes from G_1 and $|S_2|$ genes from G_2 to obtain two new gene sets $\mathcal{S'}_1$ and $\mathcal{S'}_2$. By default, we repeated this 15000 times to obtain a distribution of $|\mathcal{S'}_1 \cap \mathcal{S'}_2|$, i.e. the expected size of the intersection between S_1 and S_2 if they had been drawn at random from G_1 and G_2.

3 Results

A major difficulty in the analysis of high-dimensional data is that the notion of distance defined by different feature sets need not agree with each other. For example, Figure 1A plots a simple four-dimensional system (gene 1 through gene 4) where dots of the same color are the same cells. On the space defined by gene 1 and gene 2, it is clear that cells lie along a 1-dimensional curve along which they progressively transform from dark blue to light blue. However, two light blue cells which are close to each other on the space defined by gene 1 and gene 2 can be very far apart in the space defined by gene 3 and gene 4 (Figure 1A, middle panel). In addition, note that the "Y" shaped trajectory jointly defined by gene 3 and gene 4 would be missed if one were only provided the space defined by gene 1 and gene 3 (Figure 1A, bottom panel), such as in the case where gene 4 is not detected.

As in this example, in the case of scRNA-seq data, some genes can jointly define a linear trajectory while others jointly define a branched trajectory. To detect these patterns, it is best to identify feature sets ahead of time to avoid constructing a space in which these contributions are indistinguishable (like the one defined by gene 1 and gene 3). We assume that information encoded in high-dimensional data can be mapped to some lower-dimensional space [44, 45, 46, 47, 48, 49, 50]. Consequently, we can "compress" the data and represent it differently. For example, our four-dimensional toy example illustrated in Figure 1A can be represented by the distance along the 1-dimensional curve in the space defined by gene 1 and gene 2. In other words, by representing this four-dimensional system with the color that we have provided for each cell, we reduce the original 4D system to a 1D coordinate. If we now initialize our system at a light blue region (one that is close to the arrows on Figure 1A, for example) and change the new coordinate by just a little bit (move from one light blue cell to another light blue cell), the cell may not travel far on the space defined by gene 1 and gene 2, but may traverse a long distance on the space defined by gene 3 and gene 4 because the next light blue cell may be located on the other branch (Figure 1A, middle panel). This illustrates how we can use the response of each gene to minor perturbation in the low-dimensional space to set them apart. With this reasoning, we hypothesized that if we first find some alternative parametrization of a high-dimensional system

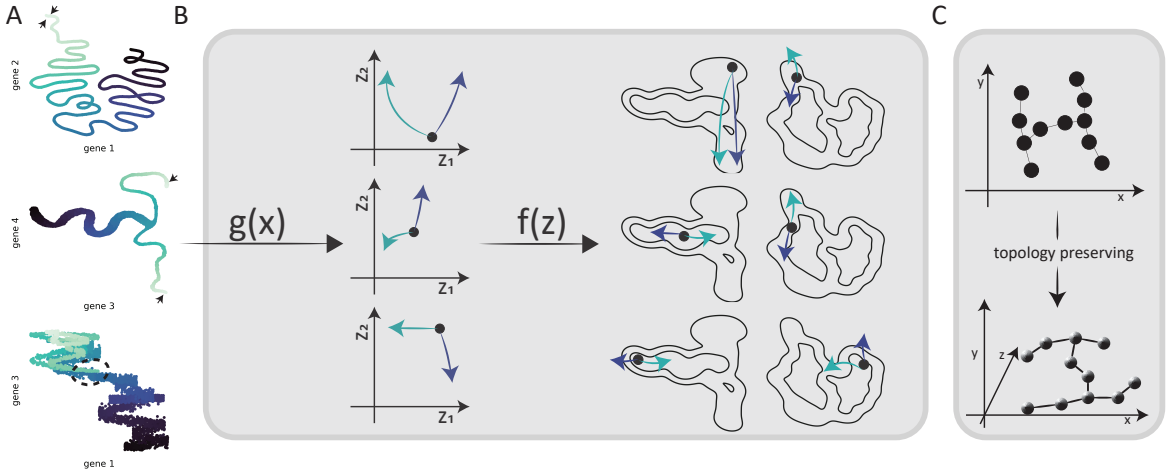

Figure 1: ID overview. A: Example 4-dimensional toy system. The space defined by gene 1 and gene 2 defined a linear trajectory; and the space defined by gene 3 and gene 4 defined a branched trajectory. Color encodes the distance traversed along the manifold in the space defined by gene 1 and gene 2, i.e., the arc length. Arrows on the top and middle panel point to the same cells. The circled region on the bottom panel illustrates a location where there is a "discontinuity" in arc length. B: In ID, high dimensional data (x) is mapped onto a low dimensional space (z). This latent representation of the cell is then perturbed in different ways and the response of the original space is used to identify different gene sets by clustering their response. C: Construction of the toy datasets. The topology of the data is first defined on a low dimensional space and then mapped onto higher dimensions via a topology–preserving transformation.

(x) that maps it to a lower-dimensional space (z), we may be able to cluster features based on how they respond to a small perturbation of the latent representation of the original data (Figure 1B).

To test this hypothesis, we generated toy datasets by first defining a topological structure on the low-dimensional space and mapping it to a higher dimension via a topology preserving transformation (Figure 1C). When this procedure is applied to a linear topological structure, we observed that the resulting "gene expression" pattern can show diverse dynamics, tracing a non-linear curve in the high-dimensional space (Figure S1A).

3.1 Identifying distinct topological structure

To solve the problem that different features can contribute to different topological structures owing to different biological processes (e.g. differentiation, cycling, aging, and so forth), we developed Identification of Distinct topological structures (ID), which first maps the high-dimensional data (x) onto a lower-dimensional space (z) and clusters genes based on how x changes with respect to changes in z (see details in Methods). To test how ID performs, we first constructed a toy dataset that contains a linear and a branched structure (Figure 2A, first panel, and Figure S1B) that follows a "similar order" (an example where cells follow a different order can be found at Figure 2B). In other words, as cell progresses along the linear structure, it also progresses somewhat smoothly along the branched structure. We first defined the linear and the branched topology on the two-dimensional space and lifted each structure to a 50-dimensional space, resulting in a gene expression matrix that contains 1000 cells (samples) and 100 genes (features) for each cell (Figure 2A, second panel). By constructing a gene embedding using ID, we observed that genes contributing to the formation of these different topological structures form two well-defined clusters (Figure 2A, third panel). Alternatively, when we attempted the naïve approach of using UMAP [14] on the transpose of the gene expression matrix, we observed that this approach failed to separate the two clusters (Figure 2A, inset in the third panel). With this simple test, we illustrated that ID at least out-performs the most naïve of approaches. We then sought to test how UMAP would perform if the two structures become more "different". To do so, instead of coming up with another topological structure, we simply changed the order of cells such that when a cell travels along the "Y"-shaped structure, it no longer traverses the linear structure (Figure 2B, in

contrast to Figure 2A). Instead, the cell "jumps around" on the linear structure. With orders of cells altered, we observed that the UMAP projection of the transpose of the gene expression matrix also produces two distinct clusters, each representing a distinct topological structure (Figure 2B).

To summarize, we observe that when the "connectivity" of the two structures is very different, the baseline method of constructing a UMAP embedding of the transpose of the data matrix is sufficient to separate features that made up the two distinct structures. When the ordering of the two structures is similar, we observe that this baseline approach fails, while ID remains capable of correctly identifying features responsible for distinct structures with almost perfect accuracy.

We next sought to investigate how separable these structures would be as a function of the number of genes defining each of them. To investigate how ID performs when the ratio between the linear and the branched component differs, we fixed the system dimension to 100 and gradually increased the size of the linear dimension from 2 to 98. When the two topological structures share a similar order, we observed that ID failed to separate features when the size of the linear dimension was too small, having an adjusted rand index (ARI) close to 0. As the size of the linear dimension got greater than 15, ID started to have a near perfect performance (Figure 2C, top panel; Figure S2, top panel). As expected, when the order defined by the two classes of features is similar, the UMAP projection of the transpose of the gene expression matrix cannot separate the two classes of features regardless of their ratio (Figure 2C, top panel; Figure S2, top panel). Similarly, when the order of the two processes is different, we observed that neither ID nor UMAP can separate the features by the topological structure they represent when the size of the linear dimension is too small. However, consistent with previous observations, both approaches begin to have near-perfect performance when the size of the linear dimension grew larger than 15 (Figure 2C, bottom panel; Figure S2, bottom panel).

Next, we compared the performance of ID against other well established methods: non-negative matrix factorization and geneTrajectory [30]. To test how each method performed, we first conducted the comparison in a situation with three distinct topological structures, each composed of the same number of genes. By changing the total dimension of the system, we observed that ID out-performed other three methods in discerning constituents of the three topological structures irrespective of the dimension of the system (Figure 2D). In this test case, NMF and geneTrajectory consistently under-performed while the performance of UMAP improved as the size of the system increases (Figure 2E, to panel). We also tested how the run time of each method scales with the number of features and observed that when the size of the system is small, ID is the least efficient method. However, as the size of the system increases, the run time of geneTrajectory exceeds that of ID, most likely a consequence of its internal construction of a gene-gene distance matrix (Figure 2E, bottom panel).

In the previous example, when we used NMF, we decomposed the gene expression matrix into three components because our toy data contained three topological structures (see methods for details regarding NMF). However, we observed that the performance of NMF appeared to depend non-linearly on the number of components that we decompose the data matrix into (Figure S3A). By decomposing the data matrix into a wider range of components and computing the resulting ARI, we observed that the performance of NMF increased when we decomposed the data matrix to 30 components but then decreased when we attempted to decompose our data further (Figure S3B). However, without knowing the feature labels ahead of time, one cannot find the optimal number of components to decompose the data matrix into with ARI. To examine whether there are alternative approaches to select the optimal number of components, we conducted a much finer sweep (Figure S4). To our surprise, when we plotted the NMF reconstruction error against the number of components, we observed that the elbow region of this curve overlaps with the ARI peak. This thus provides a potential way to select the optimal number of components. We also investigated whether similar problem is also present in ID by testing how it performs when we change the dimension of the latent space. Unlike NMF, we observed that ID can achieve a near perfect performance as long as the dimension of the latent space is sufficiently high (Figure S5). To investigate potential reason why geneTrajectory under-performed, we constructed a toy dataset that contains only a linear component. In this situation, we observed that geneTrajectory outputs a "meaningful" projection of the features (Figure 2F, top panel), where genes that are nearby on the geneTrajectory space shared similar expression pattern across cells; and genes far apart had anti-correlated patterns (Figure 2F, bottom panel). However, the geneTrajectory space can neither capture the topology of the data nor provide information to order these genes. When overlaying the expression of the genes on top of the UMAP projections of cells, we observed that the dynamics of genes can be more complicated than what geneTrajectory had envisioned (Figure 2F, bottom panel),

peaking in more than one position while still jointly defining a curve in the high dimensional space.

Next, we investigated how ID would perform in a more complex scenario. In this test case, we sampled from the linear and branched structure, as well as six other distinct topological structures from 2D manifolds embedded in the 3D space (Figure 3A). In total, we sampled from eight different structures, and projected each onto a 50 dimensional space via a topology preserving transformation (see methods for detail), resulting in a final data matrix with a 400 dimensional feature space. When we visualize this data on a 2D plane, we observed that low dimensional projection of cells had a branched structure and displayed no characteristics of the 3D shapes we sampled from (Figure 3B). Interestingly, when we swept across the number of components for NMF, we observed that the peak position has shifted from our previous example (Figure S6A and Figure S7). It is worth noting that the run time for NMF increases as the number of components increases. Consequently, the need to sweep and search for the optimal number of components makes NMF more computationally expensive than ID (Figure S7, bottom panel). In this situation, we still observed that ID can separate the distinct classes of features well, out-performing all other methods and achieve near–perfect performance in almost all trials (Figure 3C, D and Figure S6B, top panel). To investigate potential reasons why the original embedding (Figure 3B) failed to contain characteristics of the topological structure it contains, we computed the cross entropy between the k-nearest neighbor (KNN) graph computed from the original data matrix and that from the UMAP embedding of samples. As expected, compared to a 2D UMAP projection, we observed that the 3D UMAP projection can significantly reduce the cross entropy (Figure S6B, bottom panel). There are a few reasons for this discrepancy. For example, it is possible that the data cannot be faithfully represented in the 2D space. If one tries to project a sphere onto a plane, the front and the back of the sphere, originally far away from each other in 3D, will now be forced to become neighbors. If this were our case, we would expect the UMAP projections of certain cells to appear as if they lie on top of the others. However, this does not seem to be the case for the UMAP projection of cells illustrated in Figure 3B, where the color of each dot is the order of cells defined by the linear dimension. If some overlapping event occurred, we should observe something similar to that illustrated in Figure 1A (the space defined by gene 1 and gene 3). This observation implies that while there are high-dimensional structures contained within the KNN graph constructed from the original space, which resulted in the decrease of cross entropy as embedding dimension increased from two to three, the topological information from the complex 3D shapes is masked by, in this case, that of the branched structure. This observation again underscores the importance of classifying features ahead of time before the construction of any cell-cell distance graph. When we used the eight clusters of features captured by ID (Figure 3C) individually to define distances between cells, we observed that the most suitable dimension to embed the subsetted original high-dimensional data depended on the features that we select, as expected (Figure 3E and Figure S6C). For each cluster of genes we IDed, when we embed the data onto the appropriate low-dimensional space, we observed that the embedded data now possess the same characteristics of the low-dimensional topology with which we constructed the high-dimensional data (Figure 3F, G and Figure S6D).

Together, our results showed that ID can reveal distinct topological structures that are masked within the data and illustrated the superiority of ID over existing methods.

3.2 ID delineates cellular differentiation

To see how ID can help us understand biological systems, we tested how it performed in diverse single cell RNA-seq datasets from distinct experiments. We first looked at two datasets that respectively charted the differentiation of hematopoetic stem cells [38, 37] and the development of the mouse hippocampus [39, 40]. Applying ID to these two datasets revealed two major clusters of genes on their gene embeddings, with the larger cluster of genes defining a tree-like topology, which corresponded to the annotated cell types, and a smaller cluster of genes defining a ring-like topology, which corresponds to the phases of cell cycle (Figure 4A, B). In these gene embeddings, we also noticed the presence of small clusters of genes that are typically located far away from the main clusters. We observed that genes within these clusters are typically lowly expressed and contained little information. To investigate if it is worthwhile to analyze the data with just the tree-like genes, we produced UMAP embedding of cells using all features and compared how they are different from those created with only the tree-like genes. While we observed that these embeddings are largely similar, as expected from the small size of the ring-like genes (Figure 4A, B and Figure S8A, B), we did observe that the incorporation of additional genes produced some notable and potentially biologically relevant differences. For example,

in the hematopoetic differentiation dataset, we observed that a group of LMPPs, MK progenitors and HSCs protrude between the branch of erythroid progenitor and myeloid progenitor, almost forming an additional branch (Figure S8A). Closer examination showed that these cells are in either the G2 or the S phase of the cell cycle (Figure S8A). When only tree-like genes were used to produce low-dimensional projections, this branch disappeared (Figure 4A). Similar phenomena were also observed in the hippocampus differentiation dataset. When all genes were used to produce a low-dimensional representation, NIPCs formed an individual branch of cells and OPCs also appeared to have multiple branches (Figure S8B). Again, we observed that these additional branches were driven by the cell cycle (Figure S8B). Similarly, if only tree-like genes were used, these branches disappeared (Figure 4B). Interestingly, in the new embedding, we find that the NIPC did not even turn into its own segment, but rather merged into the main branch coming from Nbl1 cells and extended to both the radial glia and the glial progenitors (Figure 4B).

To sum up, using two differentiation datasets, we illustrated that ID is capable of identifying gene sets that defined distinct, convolved biological processes. In our examples, it may be important to use these gene sets separately, as we have demonstrated that their joint usage can introduce additional branches in the differentiation tree, which will affect downstream analysis such as trajectory inference.

3.3 ID outlines cellular response to external perturbation

We further demonstrate how ID can generate biologically relevant information with an scRNA-seq dataset of microglia [41]. Microglia are macrophages that reside in the brain and play an important role in the development of the neuronal system with diverse, defined transcriptional states that are associated with different functions and morphologies. Here, we examined a scRNA-seq dataset of microglia sampled from mouse with or without whisker deprivation [41]

When all feature genes were used to construct a low-dimensional space, we arrived at a projection that is consistent with the original annotation, which is associated with both the cell cycle phase and the potential function of the cells (Figure S8C). Moreover, while many cell types were related to the cell cycle, this low-dimensional projection contained no obvious circular structure (Figure S8C).

By applying ID, we identified three major clusters of genes (Figure 4C). Interestingly, we observed that the largest cluster of genes can only be used to differentiate the neuronal-like and the interferon-responsive microglia, while a second, much smaller cluster of genes is sufficient to define the majority of the annotated clusters (Figure 4C). Finally, we observed that the third cluster of genes is associated with the cell cycle, evident from the low-dimensional projection of the cells constructed using these genes, which both formed a clear circular structure and is consistent with the cell cycle phase estimation (Figure 4C).

In the original study, it was concluded that whisker removal would lead to the expansion of an interferon-responsive microglia population capable of engulfing neurons thereby limiting the accumulation of DNA-damaged cells, which consequently generated a population of microglia expressing neuronal markers [41]. Originally, the relationship between the interferon–responsive microglia and neuronal-like microglia was initially established via RNA-velocity [39]. However, if this group of interferon–responsive microglia really start to "express" neuronal markers due to their engulfment of DNA-damaged neurons, how can RNA velocity, which is developed to predict the future transcriptome of a cell, anticipate the cellular target of microglial phagocytosis? This puzzling chain of thought led to our closer examination of the low-dimensional projection of cells constructed from the cluster 1 genes, which can only be used to differentiate these two cell types. Interestingly, we observed that this low-dimensional projection contained a "gap" region between the interferon-responsive and the neuronal-like microglia that are connected by other microglia population (Figure 4). This is consistent with the fact that the transition between the transcriptional state of a microglia and one with an engulfed neuron should be "discrete". In order to engulf other cells, neuronal-like microglia need to break down the engulfed neuron and return to a "ground-state" cell before it can transition to an interferon-responsive cell [51, 52, 53].

To summarize, we observed that this microglia dataset contained multiple biological processes which reduced the interpretability of the low-dimensional space it generates (Figure S8C). By using ID, we identified three distinct topological structures from the data, one of which consistent with the validated finding of the original study [41].

3.4 ID dissects the effect of gene knock-out

Hair follicle dermal condensates are essential for the formation of hair follicles whose differentiation is regulated by the Wnt signaling pathway [54, 55, 56]. Here, we examined a scRNA-seq dataset of dermal cell differentiation that is sampled both from wild-type mice and mice with Wnt knock-out [30]. This defect is known to result in the lack of formation of dermal condensates and consequently the lack of hair follicles [30, 57, 58].

When all feature genes were used to construct a low-dimensional space, we observed that this projection took the shape of a snail (Figure S8D). To examine the potential meaning of the space, we overlayed the existing annotation of genotype, cell type and cell cycle and observed that this UMAP is the most consistent with cell type. The lower dermis (LD), upper dermis (UD) and dermal condensates (DC) can be well separated along this low-dimensional space (Figure S8D). While this space can well separate cellular identity, we have observed that this space does worse at separating genotype, where the only discernible difference between WT and KO mice lies in the "tail" region of the snail, or cell cycle, whose circular structure became the "shell" of the snail (Figure S8D).

By applying ID, we observed two major clusters of genes (Figure 4D). Dimensionality reduction using cluster 1 genes as features revealed a clear, linear structure, where the lower dermis is more similar to the upper dermis than to the dermal condensate (Figure 4D). More surprisingly, the effect of Wnt knock-out is directly visible along this structure (Figure 4D). On this space, we observed that while the lower dermis of the WT and KO mice is transcriptionally similar, that of their upper dermis is not, suggesting that Wnt disrupted the formation of dermal condensates by altering the transcriptional state of the upper dermis. Similar to the previous examples, we have also identified a cluster of genes that defined the cell cycle. The low-dimensional projection of cells defined by these genes form a clear circular structure that is consistent with cell cycle phases (Figure 4D).

3.5 ID identifies conserved differentiation trajectory in human lungs

So far, we used ID to delineate cellular differentiation (Figure 4A, B), outline cellular response to external perturbation (Figure 4C), and dissect the effect of genetic knock-out (Figure 4D) by identifying genes that defined distinctive topological structures. Next, we sought to use ID to test the hypothesis that for a given system, there will be a conserved set of genes that are responsible for defining some conserved topological structures. In other words, we seek to find datasets that characterize the same system, apply ID to these datasets and cluster genes accordingly to investigate if the cluster assignment of genes will be conserved across datasets. As different datasets may be collected by different research groups or processed with different protocols under different conditions, they will inevitably contain many systematic differences. Some of these differences may be attributable to changes of the transcriptional state of the cell incurred by the environment, and others may simply be gene expression noise or batch effects. By looking at whether ID identified clusters of genes are conserved, we assessed the extent of randomness of gene sets that defined the same structure across experiments. Since one would expect batch effect to randomly impact measured transcripts, this analysis allows us to make some inference of the nature of different gene sets and determine if they represent actual biological differences. To this end, we analyzed two scRNA-seq datasets collected from the human lungs, hereafter referred to as the lung sample data [43], collected from two unsuitable transplant lungs, and the lung organoid data [42], collected from two donors, which were then cultured into different organoid models.

We first examined the lung organoid data [42]. When all features were used to construct a low-dimensional representation of cells, we observed that each organoid model occupies a distinct region of the space (Figure S9A) and the annotated cell types can be well separated within each such region (Figure S9A). Similar to our previous examples, we also observed on this space the formation of clusters of cells consequent to their cell cycle phases (Figure S9A). Interestingly, we observed that cells sampled from the EX (PneumaCult EX-plus medium) model had a Y-shaped path of differentiation. Closer examination of the metadata revealed that each branch of the "Y" is composed of cells that were harvested from the same donor (Figure S9A, inset). In summary, when all features were used to construct a low-dimensional representation of cells, the resulting space is jointly affected by cell type, cell cycle phases, organoid model, and donor from which the cells were harvested.

Applying ID to the lung organoid data, we observed three major clusters of genes (Figure 5A, inset), with the largest cluster defining a space for cells that is similar to that constructed from all features. We

then explored the low-dimensional space defined by the other clusters of genes. On the space defined by the cluster 2 genes, we observed a conserved "barrel"-shaped differentiation trajectory (Figure 5B). Despite the differences among these organoid models, we observed that the O1 and O2 organoid had an almost identical "barrel"-shaped trajectory while that for the other two models, ALIEX (cells cultured in air-liquid interface with PneumaCult-ALI medium) and EX (cells cultured in PneumaCult EX-plus medium), showed considerable differences (Figure 5B). A similar observation was also made in the space defined by the cluster 3 genes, which was observed to be associated with the cell cycle. In this space, we observed a striking similarity among the low-dimensional projection of the O1, O2 and the EX model, all possessing a ring-like structure that is consistent with the annotated cell cycle phases (Figure 5B). Interestingly, the ALIEX model stood out from the other three models. In the space defined by the cluster 2 genes, the basal cells and secretory cells of the ALIEX model clearly differed from those of the other organoid models; and in the space defined by the cluster 3 genes, the ALIEX model is the only organoid model from which we failed to observe a ring structure (Figure 5B). While some of the organoid models remained distinct from each other in both the "differentiation" space and the "cell cycle" space, the extent of their differences have become much smaller compared to that illustrated in Figure S9A. In addition, we strikingly observed in the "differentiation" and "cell cycle" spaces that the effect of donor is almost completely absent in all four organoid models (Figure S9B), suggesting that ID may have identified a conserved set of genes that define the differentiation trajectory and cell cycle phases of lung epithelial cells.

These results suggest that the lung organoid dataset contained a large amount of technical variation due to both donor and organoid model, but that by using ID, we can identify gene sets that are less influenced by this source of variation. While this could suggest that we have found a gene set that robustly defines the differentiation trajectory of lung epithelial cells, we may also have observed something else. If technical variations due to experimental covariates (such as donor) randomly act on all detected genes with some probability, then there will always be some gene sets that were "left alone", showing no obvious relationship with experimental covariates in the same way as our detected differentiation or cell cycle-related genes. In principle, to show that these genes were not left alone by technical variations, it would be desirable to study another dataset that contained similar information and show that the same gene sets can be extracted and that they characterize similar relationships between cells.

To this end, we examined the lung sample dataset [43], which is sampled from two discarded transplant lungs and subsequently cultured in four different types of media. Similar to the lung organoid data, the lung sample data also contained a large amount of technical variation, evident from the low-dimensional space constructed from all feature genes (Figure S10A). By using ID, we again identified multiple gene sets, with cluster 1 genes capturing most of the technical variations (Figure 5C). Strikingly, when we examined the space defined by the cluster 2 genes from the lung sample data, we again observed a "barrel"-shaped differentiation trajectory and it is not only conserved across cells cultured in different media but also conserved across cells sampled from different donors (Figure 5D and Figure S10B).

We next assessed the extent of randomness of these gene sets and directly investigated their very nature. For example, if the genes that defined the "barrel"-shaped differentiation trajectory are the same across the two datasets, then it would suggest that they are the defining properties of lung epithelial cell differentiation and not simply the "survivors" of the technical variations introduced by experimental covariates. To quantify the extent of the randomness of these gene sets, we tested, for a given gene set, if the observed number of intersecting genes would be different than if we were to sample genes randomly (see method for more detail). We first conducted this test for the cluster 1 and cluster 2 genes, which respectively captured the variation induced by experimental covariates and cellular differentiation. Interestingly, we observed that the number of genes that are conserved across the two datasets to be significantly smaller than the null model for the cluster 1 genes (Figure 5E) and the converse with cluster 2 genes (Figure 5F). With the cluster 3 genes, we also observed that the number of intersecting gene to be much larger than the null model (Figure S10C). When we examined the low-dimensional space defined by the cluster 3 genes, we observed a ring-like structure similar to that observed in the lung organoid data (Figure 5B and Figure S10D).

To sum up, we first used ID to independently identify gene sets that defined cellular differentiation, cell cycle progression and captured the most amount of technical variation in two human lung datasets. In both datasets, we observed that the variation introduced by experimental covariates to be reduced

if not eliminated in both the differentiation and cell cycle space. To elucidate if the differentiation and cell cycle related genes that we discovered are biologically meaningful or simply "survivors" of technical variation, we investigated if their identities are conserved across datasets. By using our custom permutation test, we observed that both the differentiation and cell cycle gene set to be highly conserved and that the gene set that captured the most amount of technical variation to be highly different. This observation suggests that ID can be used to identify key genes that define differentiation/cell cycle and more importantly, the expression of these key regulators is independent of factors such as donor or cell culture type.

4 Conclusion

Here, we introduced ID, a simple, efficient algorithm to uncover hidden topological structures from data through the construction of an alternative parametrization of the system. We demonstrated the utility of this approach through extensive benchmarks on toy datasets and against other existing methods. When the state of the cell is jointly influenced by extrinsic (experimental covariates) and intrinsic (cell cycle phases) factors, we showed that ID was capable of identifying gene sets associated with distinct topological structures to enable the better understanding of differentiation and the cellular response to both environmental and genetic perturbations.

Through our investigation, we observed that distinct gene sets can provide conflicting measures of cell-cell distance. For example, cells of the same cell type may not to be in the same phase of the cell cycle. As a result, the joint usage of cell cycle genes and cell type markers to define cell state can confound cell type annotation and produce misleading low-dimensional visualization of the cells (Figure 4C, D and Figure S8C, D). This is likely the reason why geneTrajectory under-performed in our test cases. As the toy data we constructed are composed of multiple gene sets, each providing a distance measure that conflicts with that from the other sets, the cell-cell graph geneTrajecory constructs as its first step is automatically an inadequate description of the true, multi-faceted cell-cell relationship.

In addition to identifying features that encode different topological structures, we found that ID can also be a powerful tool for integrative analysis. By applying ID on two independently collected scRNA-seq data of the human lung, we observed in each a set of genes that is sensitive to experimental covariates, a set of genes that defines a "barrel"-shaped differentiation trajectory of lung epithelial cells, and a final set of genes that charts cell cycle progression. We further elucidate that the differentiation and cell cycle related genes are highly conserved across donor and other experimental covariates via permutation test. In the future, it will be worthwhile to employ the same strategy to investigate if there are topological structures conserved across related datasets, and if the same structure is defined by the same set of genes. By looking at genes that robustly define the progression along a certain topological structure, we are more likely to focus on the key regulator of the process of interest.

ID constructs an alternative parametrization of the system by a variational autoencoder (VAE) [33] and identifies related features by looking for those that respond similarly to a small perturbation in the latent space. We believe ID is not the only way to identify these different gene sets. For example, the VAE component of ID can be replaced with other approaches with similar functionality, such as a normalizing flow across dimension [59], but they may incur some additional computational cost. Simpler approaches may also be used to achieve the same goal. In our analysis, we showed that UMAP of the transpose of the data matrix is sufficient when the order defined by the different structure is vastly different (Figure 2). However, in more complex scenarios as shown in Figure 3, UMAP of the transpose of the data matrix becomes unable to separate features that encode distinct structures. We also showed that NMF performed well when the data is broken down into an optimal number of components. However, NMF becomes more inefficient as the number of component increases (Figure S4), and finding such optimal number of component is itself highly inefficient.

A critical assumption we have made in our analysis is that each gene only participates in a single structure. While this assumption should hold under most circumstances, it would be exciting to adapt ID for more complex datasets, such as those that contain both "normal" and "diseased" cells. In such cases, genes that are supposed to be part of a "circle" may become part of a "branch" due to dysregulated gene expression.

In this work, we raised the issue that the relationship between cells is more complicated than their pairwise distances on the gene expression space, addressed it using ID, and exemplified its broad

usefulness across different contexts. We envision that ID can be useful in a variety of ways. First, we demonstrate through multiple examples that careful feature selection is required to avoid the generation of misleading data visualization. Second, our result underscores the importance of decoding the "driving force" of the formation of each structure, and dissecting the interactions/dependence between different structures (how cell cycle interacts with differentiation). One useful way to understand the system is to reconstruct the system. By identifying genes associated with each structure,ID allows one to reduce the daunting task of characterizing the dynamics of a high-dimensional system into a much simpler problem of characterizing the dynamics of a small number of weakly coupled low-dimensional systems. This suggests that unlike the existing paradigm of parametrizing dynamics on a single low-dimensional space [60, 61, 62], the correct approach to model the dynamics of a cell is to simultaneously evolve multiple subspaces. This same strategy can also be adopted to characterize multi-modal data. With this general framework, we get one step closer to the making of a computational model that can further our understanding of the system, predict how the system will respond to external perturbation, and eventually guide the design of some strategy for control.

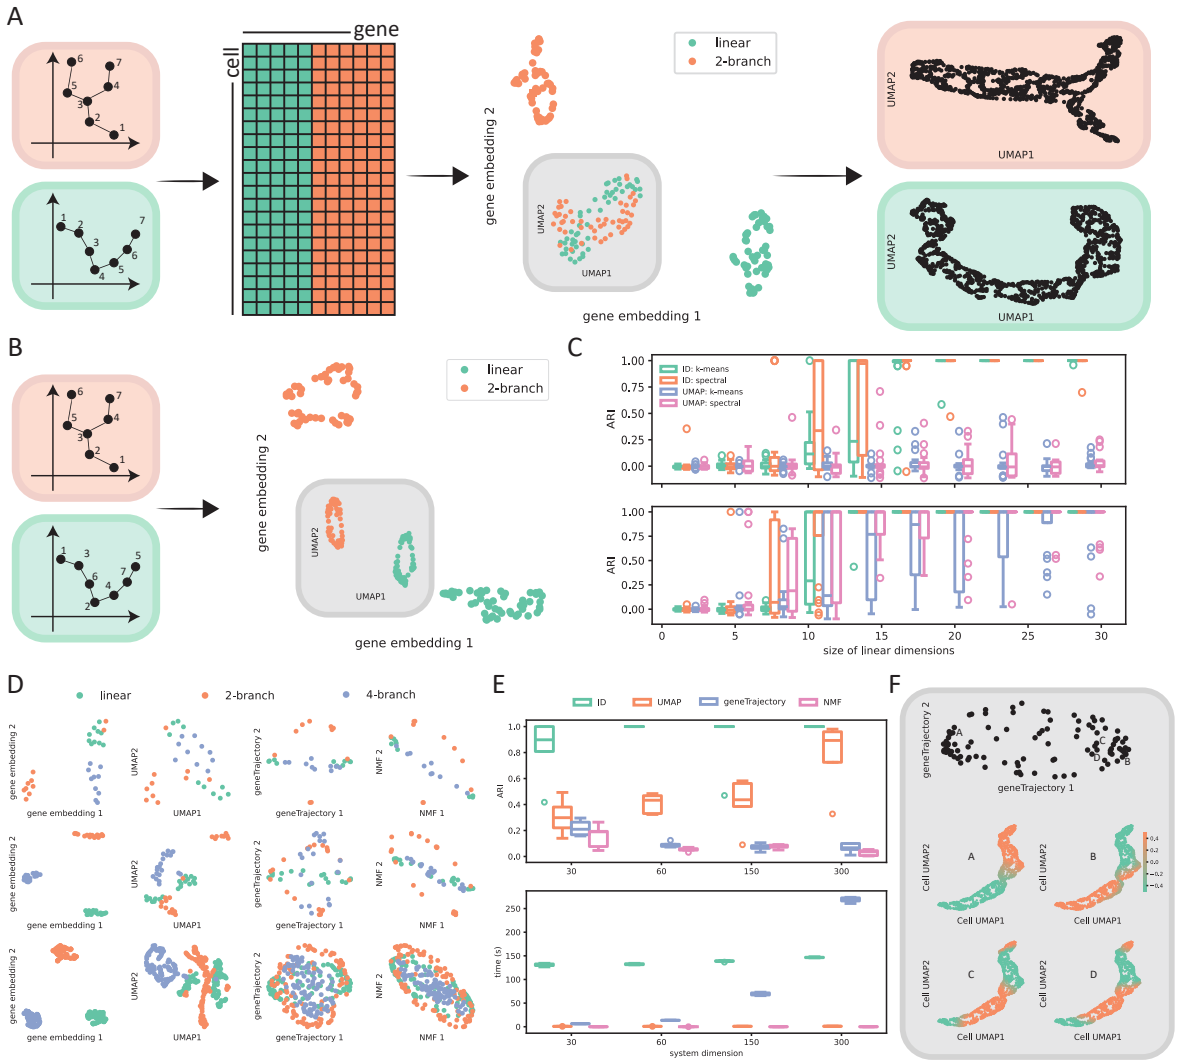

Figure 2: Applying ID on simple toy datasets. A: A flow chart that describes how we generated the data and how we used ID to classify genes. We first defined two topological structures on a low-dimensional space and mapped them both to a high-dimensional space before concatenating them together (panel 1 and 2). We then used ID to construct a gene embedding (panel 3, inset: gene embedding constructed with UMAP). Using the two clusters of genes as defined by ID, we observed that the UMAP of cells contains both a linear and a branched sub-component. B: An example where the "order" defined by the two topological structure differs greatly. In this case, both ID and UMAP can construct gene embeddings that allow the perfect classification of genes. C: Comparing ID and UMAP in classifying genes. Top: examples when the "order" is similar. Bottom: examples when the "order" is different. 20 experiments were conducted for each condition. D: Example embeddings of genes produced by different algorithms in toy datasets that contain three different structures. The toy dataset used in first, second and third row contained 30, 60, 300 genes respectively. E: ARI and run time for each method as a function of total system dimension. Each trial was repeated five times. F: Top: Embedding produced by geneTrajectory from a toy dataset with only "linear" genes. Bottom: Expression of labeled genes overlayed on the UMAP of cells.

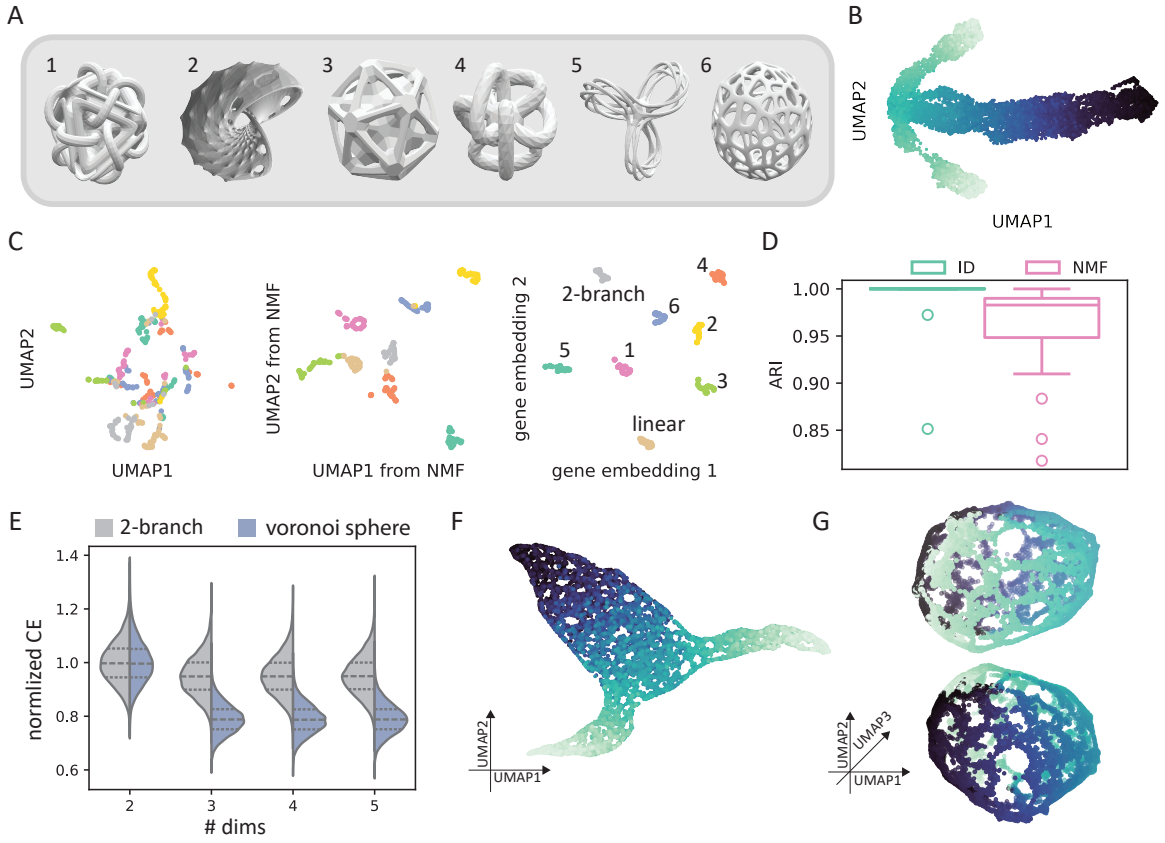

Figure 3: Applying ID on complex toy dataset. A: Manifolds from which we sampled to construct the toy dataset, in addtion to the linear and branched structures previously described. 1, Octahedron; 2, Clifford torus; 3, Dodecahedron; 4, Genus-6 torus; 5, Genus-3 surface; 6, Voronoi sphere. B: UMAP visualization of samples constructed using all features; only the branched structure is evident. C: Embedding for genes constructed by UMAP, NMF and ID; ID clusters are well–defined and correspond to the structures present in the data. D: ARI computed using k-means clustering on the embedding for genes constructed from NMF or ID from 20 individual trials. E: Cross–entropy between the KNN graph constructed from the original space and the KNN graph constructed from the UMAP projection of samples with different embedding dimension. F: 2D UMAP projection of cells using only the grey genes from Figure 3C. G: 3D UMAP projection of cells using only the dark blue genes from Figure 3C.

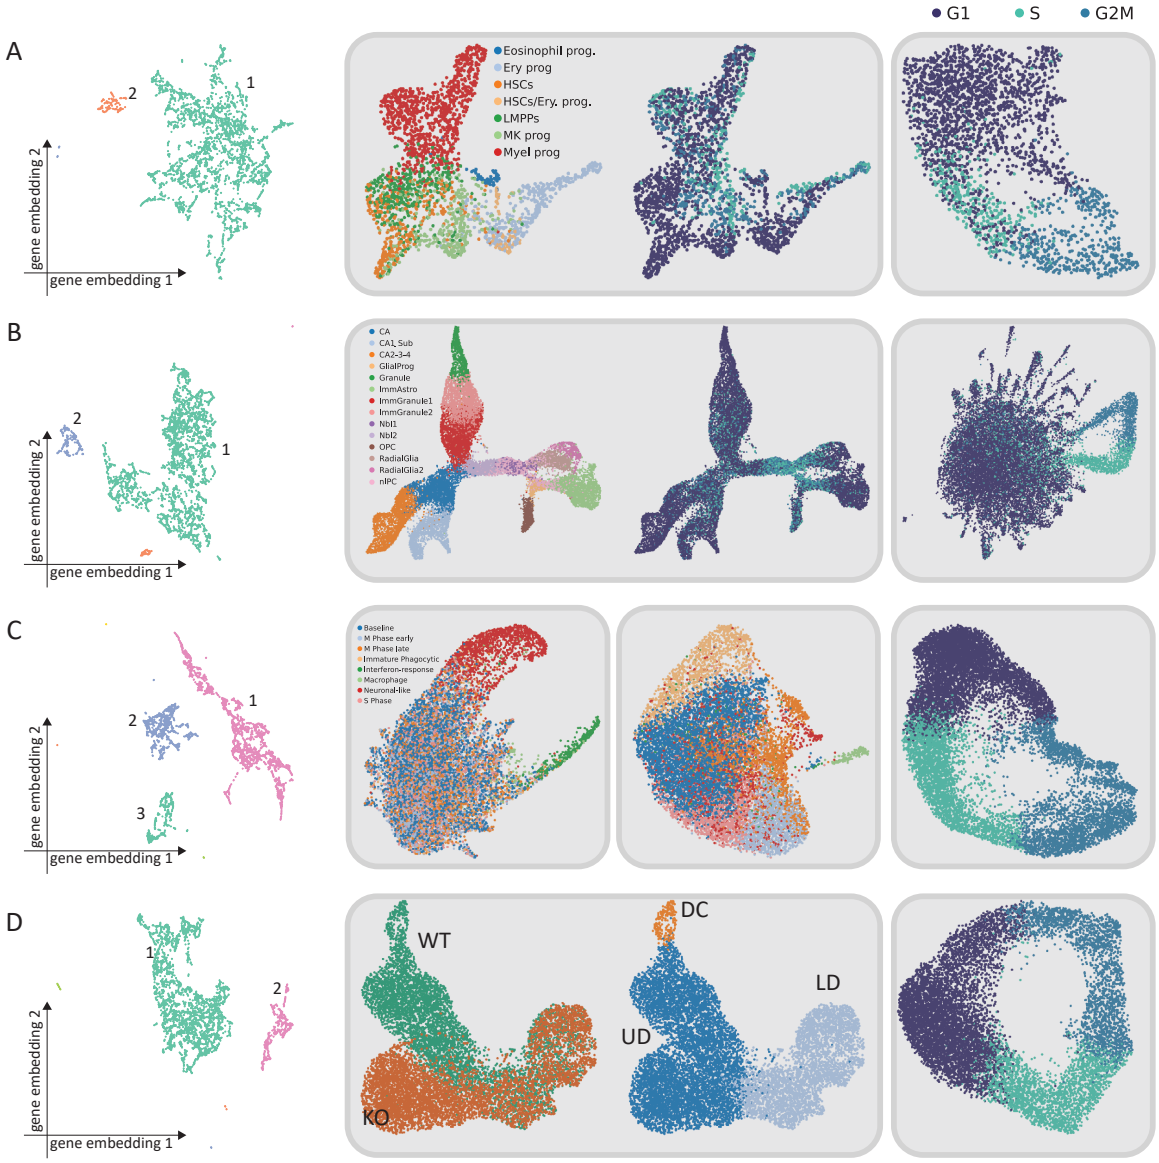

Figure 4: Applying ID to scRNA-seq datasets. ID is applied to (A) hematopoetic differentiation [37, 38], (B) hippocampus development [39, 40], (C) microglia [41], and (D) hair follicle [30] scRNA-seq datasets. The resulting gene embedding for each dataset is plotted on the first column. In each row, each box contains the UMAP projection of cells constructed using its associated gene clusters (e.g., UMAP in the first box is constructed from the cluster 1 genes.)

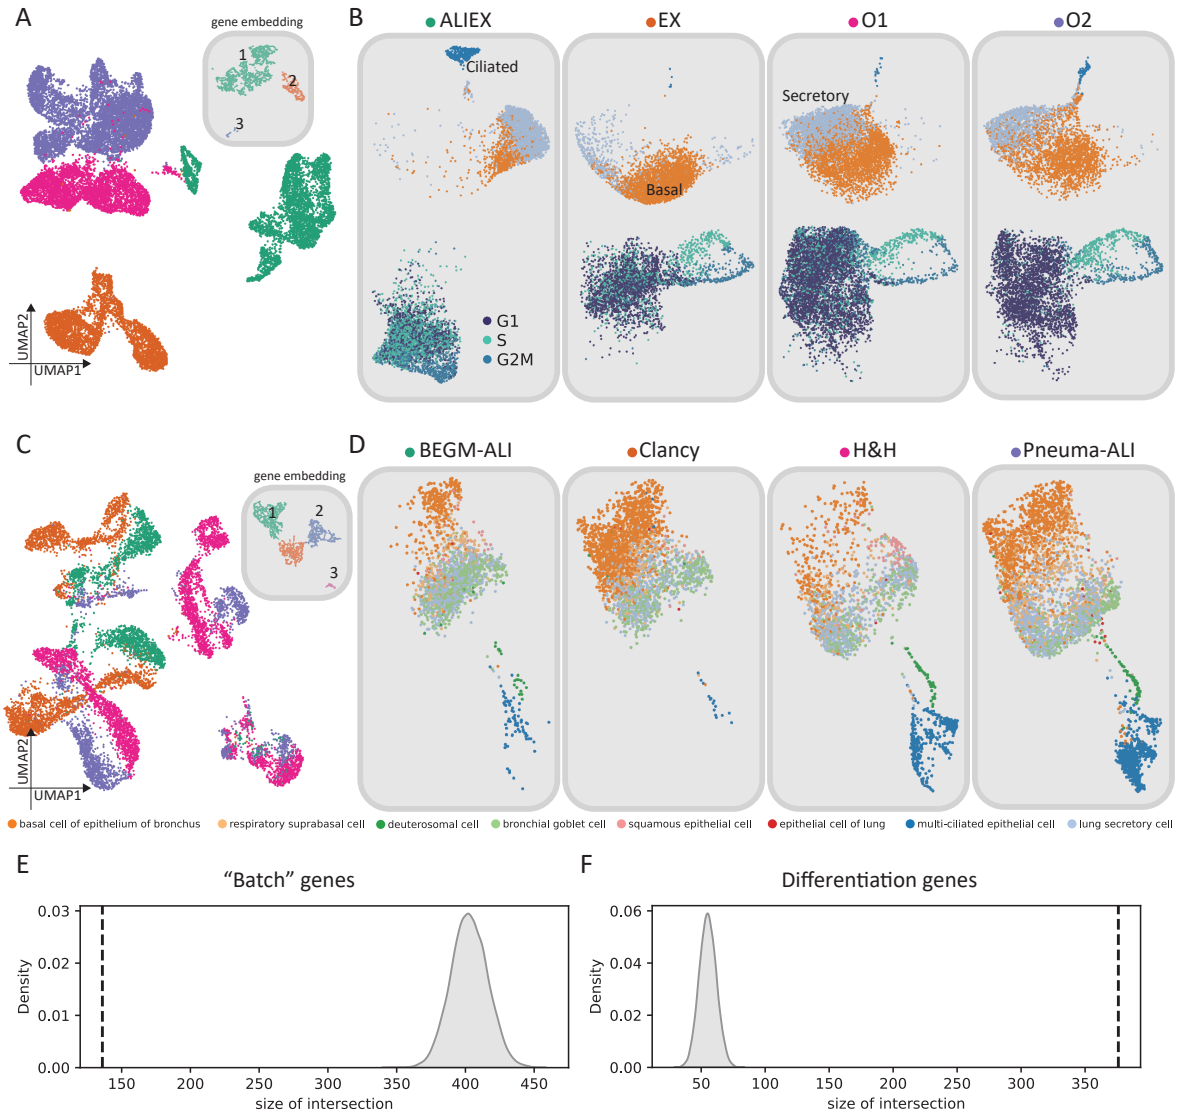

Figure 5: Applying ID to scRNA-seq datasets collected from the human lungs. ID is applied to the lung organoid dataset [42] (A, B) and the lung sample dataset [43] (C, D). A: UMAP projection of cells in the lung organoid data constructed from cluster 1 genes (gene embedding as inset). B: UMAP projection of cells constructed from cluster 2 (top) and cluster 3 (bottom) genes respectively. Each column contains cells in the same organoid model. Colors denote cell type annotation. C: UMAP projection of cells in the lung sample data constructed from cluster 1 genes (gene embedding as inset). D: UMAP projection of cells constructed from cluster 2. Each column contains cells cultured in the same media. Colors denote cell type annotation. E: The observed number of intersect between the cluster 1 genes from both datasets (dashed line) and the null distribution constructed by random sampling. F: The observed number of intersect between the cluster 2 genes from both datasets (dashed line) and the null distribution constructed by random sampling. In panel B and D, the UMAP projection of cells are plotted on the same coordinate system.

References

[1] C.H. Waddington. *The Strategy of the Genes*. Routledge, 0 edition, April 2014.

[2] Caleb Weinreb, Alejo Rodriguez-Fraticelli, Fernando D. Camargo, and Allon M. Klein. Lineage tracing on transcriptional landscapes links state to fate during differentiation. *Science*, 367(6479):eaaw3381, February 2020.

[3] Simon L. Freedman, Bingxian Xu, Sidhartha Goyal, and Madhav Mani. A dynamical systems treatment of transcriptomic trajectories in hematopoiesis. *Development*, 150(11):dev201280, June 2023.

[4] J. M. W. Slack. *From Egg to Embryo: Regional Specification in Early Development*. Cambridge University Press, 2 edition, May 1991.

[5] Ravi Allada and Brian Y. Chung. Circadian Organization of Behavior and Physiology in *Drosophila*. *Annual Review of Physiology*, 72(1):605–624, March 2010.

[6] Ray Zhang, Nicholas F. Lahens, Heather I. Ballance, Michael E. Hughes, and John B. Hogenesch. A circadian gene expression atlas in mammals: Implications for biology and medicine. *Proceedings of the National Academy of Sciences*, 111(45):16219–16224, November 2014.

[7] John S. O'Neill and Akhilesh B. Reddy. Circadian clocks in human red blood cells. *Nature*, 469(7331):498–503, January 2011.

[8] Bingxian Xu, Dae-Sung Hwangbo, Sumit Saurabh, Clark Rosensweig, Ravi Allada, William L. Kath, and Rosemary Braun. Temperature-driven coordination of circadian transcriptional regulation. *PLOS Computational Biology*, 20(4):e1012029, April 2024.

[9] Chunghun Lim and Ravi Allada. Emerging roles for post-transcriptional regulation in circadian clocks. *Nature Neuroscience*, 16(11):1544–1550, November 2013.

[10] Zhilin Qu, W. Robb MacLellan, and James N. Weiss. Dynamics of the Cell Cycle: Checkpoints, Sizers, and Timers. *Biophysical Journal*, 85(6):3600–3611, December 2003.

[11] Timothy S. Gardner, Milos Dolnik, and James J. Collins. A theory for controlling cell cycle dynamics using a reversibly binding inhibitor. *Proceedings of the National Academy of Sciences*, 95(24):14190–14195, November 1998.

[12] Andrea Riba, Attila Oravecz, Matej Durik, Sara Jiménez, Violaine Alunni, Marie Cerciat, Matthieu Jung, Céline Keime, William M. Keyes, and Nacho Molina. Cell cycle gene regulation dynamics revealed by RNA velocity and deep-learning. *Nature Communications*, 13(1):2865, May 2022.

[13] Wayne Stallaert, Katarzyna M. Kedziora, Colin D. Taylor, Tarek M. Zikry, Jolene S. Ranek, Holly K. Sobon, Sovanny R. Taylor, Catherine L. Young, Jeanette G. Cook, and Jeremy E. Purvis. The structure of the human cell cycle. *Cell Systems*, 13(3):230–240.e3, March 2022.

[14] Leland McInnes, John Healy, and James Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction, September 2020. arXiv:1802.03426 [stat].

[15] Laleh Haghverdi, Maren Büttner, F Alexander Wolf, Florian Buettner, and Fabian J Theis. Diffusion pseudotime robustly reconstructs lineage branching. *Nature Methods*, 13(10):845–848, October 2016.

[16] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.

[17] Xiaojie Qiu, Qi Mao, Ying Tang, Li Wang, Raghav Chawla, Hannah A Pliner, and Cole Trapnell. Reversed graph embedding resolves complex single-cell trajectories. *Nature Methods*, 14(10):979–982, October 2017.

[18] Jiarui Ding and Aviv Regev. Deep generative model embedding of single-cell RNA-Seq profiles on hyperspheres and hyperbolic spaces. *Nature Communications*, 12(1):2554, May 2021.

[19] Jiarui Ding, Anne Condon, and Sohrab P. Shah. Interpretable dimensionality reduction of single cell transcriptome data with deep generative models. *Nature Communications*, 9(1):2002, May 2018.

[20] Daniel Lee and H. Sebastian Seung. Algorithms for Non-negative Matrix Factorization. In T. Leen, T. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems*, volume 13. MIT Press, 2000.

[21] Daniel D. Lee and H. Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, October 1999.

[22] Pentti Paatero, Unto Tapper, Pasi Aalto, and Markku Kulmala. Matrix factorization methods for analysing diffusion battery data. *Journal of Aerosol Science*, 22:S273–S276, 1991.

[23] Pentti Paatero and Unto Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2):111–126, June 1994.

[24] P Anttila, P Paatero, U Tapper, and O Jarvinen. Source identification of bulk wet deposition in Finland by positive matrix factorization. *Atmospheric Environment*, 29(14):1705–1718, 1995.

[25] Hanna Mendes Levitin, Jinzhou Yuan, Yim Ling Cheng, Francisco Jr Ruiz, Erin C Bush, Jeffrey N Bruce, Peter Canoll, Antonio Iavarone, Anna Lasorella, David M Blei, and Peter A Sims. *De novo* gene signature identification from single-cell RNA-seq with hierarchical Poisson factorization. *Molecular Systems Biology*, 15(2):e8557, February 2019.

[26] Karin Pelka, Matan Hofree, Jonathan H. Chen, Siranush Sarkizova, Joshua D. Pirl, Vjola Jorgji, Alborz Bejnood, Danielle Dionne, William H. Ge, Katherine H. Xu, Sherry X. Chao, Daniel R. Zollinger, David J. Lieb, Jason W. Reeves, Christopher A. Fuhrman, Margaret L. Hoang, Toni Delorey, Lan T. Nguyen, Julia Waldman, Max Klapholz, Isaac Wakiro, Ofir Cohen, Julian Albers, Christopher S. Smillie, Michael S. Cuoco, Jingyi Wu, Mei-ju Su, Jason Yeung, Brinda Vijaykumar, Angela M. Magnuson, Natasha Asinovski, Tabea Moll, Max N. Goder-Reiser, Anise S. Applebaum, Lauren K. Brais, Laura K. DelloStritto, Sarah L. Denning, Susannah T. Phillips, Emma K. Hill, Julia K. Meehan, Dennie T. Frederick, Tatyana Sharova, Abhay Kanodia, Ellen Z. Todres, Judit Jané-Valbuena, Moshe Biton, Benjamin Izar, Conner D. Lambden, Thomas E. Clancy, Ronald Bleday, Nelya Melnitchouk, Jennifer Irani, Hiroko Kunitake, David L. Berger, Amitabh Srivastava, Jason L. Hornick, Shuji Ogino, Asaf Rotem, Sébastien Vigneau, Bruce E. Johnson, Ryan B. Corcoran, Arlene H. Sharpe, Vijay K. Kuchroo, Kimmie Ng, Marios Giannakis, Linda T. Nieman, Genevieve M. Boland, Andrew J. Aguirre, Ana C. Anderson, Orit Rozenblatt-Rosen, Aviv Regev, and Nir Hacohen. Spatially organized multicellular immune hubs in human colorectal cancer. *Cell*, 184(18):4734–4752.e20, September 2021.

[27] Russell Z. Kunes, Thomas Walle, Max Land, Tal Nawy, and Dana Pe'er. Supervised discovery of interpretable gene programs from single-cell data. *Nature Biotechnology*, 42(7):1084–1095, July 2024.

[28] Sarah Chen, Aviv Regev, Anne Condon, and Jiarui Ding. CellUntangler: separating distinct biological signals in single-cell data with deep generative models, January 2025.

[29] Jonathan Karin, Yonathan Bornfeld, and Mor Nitzan. scPrisma infers, filters and enhances topological signals in single-cell data using spectral template matching. *Nature Biotechnology*, 41(11):1645–1654, November 2023.

[30] Rihao Qu, Xiuyuan Cheng, Esen Sefik, Jay S. Stanley Iii, Boris Landa, Francesco Strino, Sarah Platt, James Garritano, Ian D. Odell, Ronald Coifman, Richard A. Flavell, Peggy Myung, and Yuval Kluger. Gene trajectory inference for single-cell data by optimal transport metrics. *Nature Biotechnology*, 43(2):258–268, February 2025.

[31] Cole Trapnell, Davide Cacchiarelli, Jonna Grimsby, Prapti Pokharel, Shuqiang Li, Michael Morse, Niall J Lennon, Kenneth J Livak, Tarjei S Mikkelsen, and John L Rinn. The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. *Nature Biotechnology*, 32(4):381–386, April 2014.

[32] Junyue Cao, Malte Spielmann, Xiaojie Qiu, Xingfan Huang, Daniel M. Ibrahim, Andrew J. Hill, Fan Zhang, Stefan Mundlos, Lena Christiansen, Frank J. Steemers, Cole Trapnell, and Jay Shendure. The single-cell transcriptional landscape of mammalian organogenesis. *Nature*, 566(7745):496–502, February 2019.

[33] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes, 2013. Version Number: 11.

[34] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, January 2017. arXiv:1412.6980 [cs].

[35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library, December 2019. arXiv:1912.01703 [cs, stat].

[36] F. Alexander Wolf, Philipp Angerer, and Fabian J. Theis. SCANPY: large-scale single-cell gene expression data analysis. *Genome Biology*, 19(1):15, December 2018.

[37] Valérie Marot-Lassauzaie, Brigitte Joanne Bouman, Fearghal Declan Donaghy, Yasmin Demerdash, Marieke Alida Gertruda Essers, and Laleh Haghverdi. Towards reliable quantification of cell state velocities. *PLOS Computational Biology*, 18(9):e1010031, September 2022.

[38] Yasmin Demerdash, Brigitte Bouman, Laleh Haghverdi, and Marieke Essers. 3071 – UNBIASED, LONGITUDINAL RESPONSE-PSEUDOTIME ANALYSIS OF THE INFLAMMATORY RESPONSE OF HSPCS AT THE SINGLE CELL LEVEL. *Experimental Hematology*, 111:S80, 2022.

[39] Gioele La Manno, Ruslan Soldatov, Amit Zeisel, Emelie Braun, Hannah Hochgerner, Viktor Petukhov, Katja Lidschreiber, Maria E. Kastriti, Peter Lönnerberg, Alessandro Furlan, Jean Fan, Lars E. Borm, Zehua Liu, David Van Bruggen, Jimin Guo, Xiaoling He, Roger Barker, Erik Sundström, Gonçalo Castelo-Branco, Patrick Cramer, Igor Adameyko, Sten Linnarsson, and Peter V. Kharchenko. RNA velocity of single cells. *Nature*, 560(7719):494–498, August 2018.

[40] Hannah Hochgerner, Amit Zeisel, Peter Lönnerberg, and Sten Linnarsson. Conserved properties of dentate gyrus neurogenesis across postnatal development revealed by single-cell RNA sequencing. *Nature Neuroscience*, 21(2):290–299, February 2018.

[41] Caroline C. Escoubas, Leah C. Dorman, Phi T. Nguyen, Christian Lagares-Linares, Haruna Nakajo, Sarah R. Anderson, Jerika J. Barron, Sarah D. Wade, Beatriz Cuevas, Ilia D. Vainchtein, Nicholas J. Silva, Ricardo Guajardo, Yinghong Xiao, Peter V. Lidsky, Ellen Y. Wang, Brianna M. Rivera, Sunrae E. Taloma, Dong Kyu Kim, Elizaveta Kaminskaya, Hiromi Nakao-Inoue, Bjoern Schwer, Thomas D. Arnold, Ari B. Molofsky, Carlo Condello, Raul Andino, Tomasz J. Nowakowski, and Anna V. Molofsky. Type-I-interferon-responsive microglia shape cortical development and behavior. *Cell*, 187(8):1936–1954.e24, April 2024.

[42] Jieun Kim, Eun-Young Eo, Bokyong Kim, Heetak Lee, Jihoon Kim, Bon-Kyoung Koo, Hyung-Jun Kim, Sukki Cho, Jinho Kim, and Young-Jae Cho. Transcriptomic Analysis of Air–Liquid Interface Culture in Human Lung Organoids Reveals Regulators of Epithelial Differentiation. *Cells*, 13(23):1991, December 2024.

[43] Elisa Redman, Morgane Fierville, Amélie Cavard, Magali Plaisant, Marie-Jeanne Arguel, Sandra Ruiz Garcia, Eamon M. McAndrew, Cédric Girard-Riboulleau, Kevin Lebrigand, Virginie Magnone, Gilles Ponzio, Delphine Gras, Pascal Chanez, Sophie Abelanet, Pascal Barbry, Brice

Marcet, and Laure-Emmanuelle Zaragosi. Cell Culture Differentiation and Proliferation Conditions Influence the *In Vitro* Regeneration of the Human Airway Epithelium. *American Journal of Respiratory Cell and Molecular Biology*, 71(3):267–281, September 2024.

[44] Jiawei Huang, Jie Sheng, and Daifeng Wang. Manifold learning analysis suggests strategies to align single-cell multimodal data of neuronal electrophysiology and transcriptomics. *Communications Biology*, 4(1):1308, November 2021.

[45] Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, February 2016.

[46] Juan A. Gallego, Matthew G. Perich, Lee E. Miller, and Sara A. Solla. Neural Manifolds for the Control of Movement. *Neuron*, 94(5):978–984, June 2017.

[47] Vasyl Alba, James E Carthew, Richard W Carthew, and Madhav Mani. Global constraints within the developmental program of the Drosophila wing. *eLife*, 10:e66750, June 2021.

[48] Giulio Bondanelli, Thomas Deneux, Brice Bathellier, and Srdjan Ostojic. Network dynamics underlying OFF responses in the auditory cortex. *eLife*, 10:e53151, March 2021.

[49] Hansem Sohn, Devika Narain, Nicolas Meirhaeghe, and Mehrdad Jazayeri. Bayesian Computation through Cortical Latent Dynamics. *Neuron*, 103(5):934–947.e5, September 2019.

[50] Peiran Gao, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya Ganguli. A theory of multineuronal dimensionality, dynamics and measurement, November 2017.

[51] Ainhoa Plaza-Zabala, Virginia Sierra-Torre, and Amanda Sierra. Autophagy and Microglia: Novel Partners in Neurodegeneration and Aging. *International Journal of Molecular Sciences*, 18(3):598, March 2017.

[52] Fushan Shi, Yang Yang, Mohammed Kouadir, Yongyao Fu, Lifeng Yang, Xiangmei Zhou, Xiaomin Yin, and Deming Zhao. Inhibition of phagocytosis and lysosomal acidification suppresses neurotoxic prion peptide-induced NALP3 inflammasome activation in BV2 microglia. *Journal of Neuroimmunology*, 260(1-2):121–125, July 2013.

[53] Harini Iyer, Kimberle Shen, Ana M. Meireles, and William S. Talbot. A lysosomal regulatory circuit essential for the development and function of microglia. *Science Advances*, 8(35):eabp8321, September 2022.

[54] Peggy Myung, Thomas Andl, and Radhika Atit. The origins of skin diversity: lessons from dermal fibroblasts. *Development*, 149(23):dev200298, December 2022.

[55] Rihao Qu, Khusali Gupta, Danni Dong, Yiqun Jiang, Boris Landa, Charles Saez, Gwendolyn Strickland, Jonathan Levinsohn, Pei-lun Weng, M. Mark Taketo, Yuval Kluger, and Peggy Myung. Decomposing a deterministic path to mesenchymal niche formation by two intersecting morphogen gradients. *Developmental Cell*, 57(8):1053–1067.e5, April 2022.

[56] Khusali Gupta, Jonathan Levinsohn, George Linderman, Demeng Chen, Thomas Yang Sun, Danni Dong, M. Mark Taketo, Marcus Bosenberg, Yuval Kluger, Keith Choate, and Peggy Myung. Single-Cell Analysis Reveals a Hair Follicle Dermal Niche Molecular Differentiation Trajectory that Begins Prior to Morphogenesis. *Developmental Cell*, 48(1):17–31.e6, January 2019.

[57] Jiang Fu and Wei Hsu. Epidermal Wnt Controls Hair Follicle Induction by Orchestrating Dynamic Signaling Crosstalk between the Epidermis and Dermis. *Journal of Investigative Dermatology*, 133(4):890–898, April 2013.

[58] Beate M. Lichtenberger, Maria Mastrogiannaki, and Fiona M. Watt. Epidermal -catenin activation remodels the dermis via paracrine signalling to distinct fibroblast lineages. *Nature Communications*, 7(1):10537, February 2016.

[59] Edmond Cunningham, Renos Zabounidis, Abhinav Agrawal, Madalina Fiterau, and Daniel Sheldon. Normalizing Flows Across Dimensions, 2020. Version Number: 1.

[60] Bingxian Xu and Rosemary Braun. Variational inference of single cell time series, August 2024.

[61] Alexander Tong, Jessie Huang, Guy Wolf, David van Dijk, and Smita Krishnaswamy. TrajectoryNet: A Dynamic Optimal Transport Network for Modeling Cellular Dynamics. 2020. Publisher: arXiv Version Number: 2.

[62] Grace Hui Ting Yeo, Sachit D. Saksena, and David K. Gifford. Generative modeling of single-cell time series with PRESCIENT enables prediction of cell trajectories with interventions. *Nature Communications*, 12(1):3222, May 2021.