ISAC CYTOMETRY
Journal of Quantitative Cell Science PART A

Learning Pathway Dynamics from Single-Cell Proteomic Data: A Comparative Study

Kleio-Maria Verrou,[1] Ioannis Tsamardinos,[1,2] Georgios Papoutsoglou[1*]

[1]Computer Science Department, University of Crete, Heraklion, Greece

[2]Gnosis Data Analysis PC, Heraklion, Greece

Received 18 April 2019; Revised 7 November 2019; Accepted 13 January 2020

Grant sponsor: FP7 Ideas: European Research Council, Grant number617393; CAUSALPATH

Additional Supporting Information may be found in the online version of this article.

*Correspondence to: Georgios Papoutsoglou, University of Crete–Department of Computer Science, Voutes Campus, GR-70013 Heraklion, Crete, Greece.
Email: papoutsoglou@csd.uoc.gr

Published online 26 February 2020 in Wiley Online Library (wileyonlinelibrary.com)

DOI: 10.1002/cyto.a.23976

Abstract

Single-cell platforms provide statistically large samples of snapshot observations capable of resolving intrercellular heterogeneity. Currently, there is a growing literature on algorithms that exploit this attribute in order to infer the trajectory of biological mechanisms, such as cell proliferation and differentiation. Despite the efforts, the trajectory inference methodology has not yet been used for addressing the challenging problem of learning the dynamics of protein signaling systems. In this work, we assess this prospect by testing the performance of this class of algorithms on four proteomic temporal datasets. To evaluate the learning quality, we design new general-purpose evaluation metrics that are able to quantify performance on (i) the biological meaning of the output, (ii) the consistency of the inferred trajectory, (iii) the algorithm robustness, (iv) the correlation of the learning output with the initial dataset, and (v) the roughness of the cell parameter levels though the inferred trajectory. We show that experimental time alone is insufficient to provide knowledge about the order of proteins during signal transduction. Accordingly, we show that the inferred trajectories provide richer information about the underlying dynamics. We learn that established methods tested on high-dimensional data with small sample size, slow dynamics, and complex structures (e.g. bifurcations) cannot always work in the signaling setting. Among the methods we evaluate, Scorpius and a newly introduced approach that combines Diffusion Maps and Principal Curves were found to perform adequately in recovering the progression of signal transduction although their performance on some metrics varies from one dataset to another. The novel metrics we devise highlight that it is difficult to conclude, which one method is universally applicable for the task. Arguably, there are still many challenges and open problems to resolve. © 2020 The Authors. *Cytometry Part A* published by Wiley Periodicals, Inc. on behalf of International Society for Advancement of Cytometry.

Key terms

signaling; pathway; dynamics; single-cell data; trajectory inference; evaluation metrics; mass cytometry.

A major challenge in understanding the functional properties of cells is to determine the temporal characteristics of the mechanisms that regulate them. A class of such mechanisms is signaling pathways, which are essentially biochemical reaction programs that allow cells to sense and distribute information in a timely fashion (1). Typically, environmental perturbations activate these programs by modulating specific surface receptor-molecules. Incoming information is then relayed intracellularly through a cascade of protein–protein interactions whereby the functional activation of one protein activates the next in the sequence, and so on. Depending on pathway topology and stochastic variation of endogenous proteins the transient activation signals can vary substantially (2,3). As a result, the identification of signaling dynamics becomes an intriguingly difficult task.

To address this challenge, single-cell platforms enable querying the behavior of molecular mechanisms, while taking into account cell-to-cell heterogeneity in a high-throughput manner (4). On such basis, mass cytometry was recently introduced as a

new single-cell format that captures information from about 40 molecular quantities in thousands of cells (5). Such high content data is ideal for the comprehensive analysis of entire signaling pathways (6). Ideally, to study a signaling pathway, one needs to observe its components in the context of a time course. However, a mass cytometer creates cross-sectional data precluding the follow-up of the same cell over time. An approximation to a time-course, then, is to perform repeated sampling at time intervals. Still, information about a cell cannot be reliably shared across separate measurements (7).

Key to improve the temporal resolution is to consider that when cells of the same type are exposed to the same trigger, they will respond in the same but not necessarily synchronized manner due to their inherent heterogeneity (8). In other words, each cell in a population will be executing the same cascade program at a stochastically different rate. On this basis, if we regard signaling as a continuous succession of protein activation states then, a typical mass cytometry measurement constitutes a mixture of cells whose protein levels map to one of these activation states. This means that the temporal information is implicit in the distribution of the data. To recover, therefore, the dynamics of signaling requires the unmixing of single cells by reordering them in a way that best leverages the protein activation pattern.

The problem of ordering multivariate units appears in many different scientific contexts starting from archeology more than a century ago (9). In biology, Magwene et al. were the first to adopt this concept and proposed an algorithm that is able to order temporal microarray samples and reveal differentiation patterns in bacteria and yeast (10). Then, Gupta and Bar-Joseph and Czibula et al. proposed solving this problem through a probabilistic instance of the traveling salesman problem or using reinforcement learning, respectively (11,12). More recently, the advent of single-cell transcriptomics dramatically expanded the field of application (13). Already there is a large arsenal of computational techniques that exploit the intercellular heterogeneity and arrange cells by their progression along the dynamic process of interest. To the best of our knowledge, however, these methods, referred to as trajectory inference (TI) methods, have so far been applied only to the reconstruction of the dynamics of differentiation and proliferation pathways.

In this work, we examine the potential of these developments to recovering the temporal characteristics of signaling pathways. To this end, we perform an extensive comparison of state-of-the-art TI methods using as input four mass cytometry temporal datasets from the public domain. TI performance evaluation is not a trivial task since we are missing the gold standard of the actual biological time of each cell. Therefore, we employ five different metrics, four of which presented here for the first time. Each metric evaluates a different property of TI that is; (i) its consistency with respect to prior biological knowledge; (ii) its consistency with respect to the time when the data were collected, (iii) the robustness; (iv) the association of the data with the proposed cell ordering; and (v) the roughness of the output trajectory. Our results show that, in principle, TI can be used to recover the temporal characteristics of signal transduction. Despite the fact that there can be no clear consensus as to which one method performs best there are several candidates that are fit for the task.

Materials and Methods

TI Framework

Consider a dataset that captures a population of single cells as they progress through a biological process. We denote this dataset by D and we let D consist of n vectors in $\mathbb{R}^p$, where n is the number of sampled cells and p is the number of assayed molecules. Since the main source of variation in D is the process under study, cell profiles will not be randomly scattered in $\mathbb{R}^p$; rather, they will lie on a low-dimensional noisy surface often referred to as a manifold. Effectively, a path should exist that extends across this manifold and traces cells as their molecular profiles evolve through the biological process. Based on this principle, the goal of TI methods is to learn the shape of the path and arrange cells along its course such that the position of each cell reflects its similarity to other cells.

Formally, the process of interest is considered continuous and its path can be described by a vector-valued function of the form $\vec{r}(t) = \left\langle f_1(t), f_2(t), \ldots, f_p(t) \right\rangle$, where each component function f represents the temporal evolution of measured species. Graphically, imagine that $\vec{r}(t)$ is a vector at the origin whose terminal point traces a curve in $\mathbb{R}^p$ as t varies; similarly to the trajectory of a moving particle but in p dimensions. In this way, a point in the image of $\vec{r}$ is also a point in $\mathbb{R}^p$ representing the state of the cell at a particular point in time. Let the vector $\vec{\varphi}(s) = \vec{r}(s) + \vec{\varepsilon}(s)$, where $\vec{\varepsilon}$ denotes the noise vector, be a sample of such point on $\vec{r}$ at an arbitrary time s. Then, the dataset D constitutes a set of poorly ordered samples of $\vec{\varphi}$ vectors; that is, $D = \left\{ \vec{\varphi}(s_1), \vec{\varphi}(s_2), \ldots, \vec{\varphi}(s_n) \right\}$. Sampling time parameter s is a latent random variable, often referred to as the pseudo-time, which reflects the progress of the biological process rather than the time when the data were collected.

To solve the problem of estimating $\vec{r}$, TI methods try to order cell profiles based on their pairwise similarities. To this end, let there be a similarity function g that reflects the pairwise similarity between any two cells. Then, the goal is to bijectively map each cell to an integer between 1 to n so that the similarities remain consistent. Consistency means that given any three cells whose ordered indexes are i, j, and k, where $i < j < k$, then it must hold that $g\left(\vec{\varphi}(s_i), \vec{\varphi}(s_j)\right) \geq g\left(\vec{\varphi}(s_i), \vec{\varphi}(s_k)\right)$ and $g\left(\vec{\varphi}(s_j), \vec{\varphi}(s_k)\right) \geq g\left(\vec{\varphi}(s_i), \vec{\varphi}(s_k)\right)$. Intuitively the reordering of cells could be addressed by sorting their pairwise Euclidean distances. However, as cells run through $\mathbb{R}^p$ in complex ways this will not preserve the dynamics of the process under study. For example, the process path may consist of cells in sequential state transition represented as points along a nonlinear progression curve (e.g. a signaling process); or, in more complex cell state

transitions such as branching (e.g. hematopoietic differentiation) or cyclic ones (e.g. cell cycle). Thus, to reveal the underlying patterns in the data, it is preferable, if not necessary, to assess the geometry of the low-dimensional manifold in which they lie.

To this direction, the majority of TI algorithms integrate two independent procedures. The first procedure is to compute a compact representation of the data either by assuming the trajectory of cells on the manifold or learning one de novo. The trajectories assumed are typically unidimensional, that is, linear or cyclic; or multidimensional, that is, bi/multifurcating, tree-like, connected graphs or disconnected graphs (14). The second procedure is the TI itself where, depending on the assumed or learnt representation the best path is sought. The dependence on the compact representation allows the division of TI approaches in two classes: the curve-based and the graph-based class of methods. Both attempt a reduced dimension learning step at the beginning. Approaches typically used in this step are, principal components analysis (PCA), *t*-distributed stochastic neighbor embedding (tSNE), diffusion maps, and independent components analysis (ICA) (15–19). Then, each class tries to find the best possible ordering within its own conceptual framework.

The main difference between the two classes is that curve-based methods try to infer the trajectory of $\vec{r}$ immediately after the low-dimensional representation is created. Trajectory inference in most curve-based methods is built around principal curves (20). In essence, principal curves regress a smooth, nonparametric curve through the center of the data manifold. Because this curve resolves the variability in the temporal dimension, the projection point of a cell to it will be indicating how far this cell has progressed through the dynamic process. Commonly, the geodesic distance from the beginning of the curve to the projected point represents the pseudo-time. Then, the ordering of cells results from the ordering of the respective pseudo-times assigned to each one of them. Of particular importance here is to note that the determined ordering is directionless because the direction of time cannot be recovered without prior knowledge. For this reason, most algorithms require the user to provide a starting position, for example, if the dynamic process involves differentiation, a cell at a pluripotent stage would be chosen as a starting point. Other approaches used for finding such curve are Gaussian Processes and Differential Equations (21–23).

Graph-based methods, on the other hand, do not try to reconstruct $\vec{r}$; rather they try fit a segmented line to the data manifold (in Ref (10), this is called polygonal reconstruction). For this, they take one extra step during dimensionality reduction that tries to simplify the data representation further by use of connected graphs, for example, a minimum spanning tree or a nearest neighbor graph (18,24). The nodes in the graph represent cells or clusters of cells in two dimensions, while the edges indicate proximity based on some (dis) similarity measure between cells. To reconstruct the process dynamics graph-based methods subtly connect the nodes in the graph creating a path. This path can be either the longest path in the graph (19,25,26) or a consolidation of short paths (24,27). Commonly, the endpoints of this path are not known so in both cases the user can define a specific node as an origin of the process, similarly to the starting position of the curve in curve-based methods. Additionally, the methods of this class may also assign a pseudo-time to each cell; again by projection to the inferred path.

TI in the Context of Signaling

In literature, there exist more than 50 TI methodologies for the analysis of single-cell data all of which have been employed to infer the trajectory of dynamic gene expression mechanisms, such as cell proliferation and differentiation (14). The main difference between these processes and signaling is the timescale in which they occur. A signaling cascade typically lasts from a few minutes to an hour, while the transition between consecutive differentiation states may take hours or days to complete. This means that a single-cell gene expression dataset would be capturing most of differentiating cell types whereas a mass cytometry dataset would imprint only a snapshot of the signaling progression in each of these populations. In terms of collected single cell data, therefore, signaling dynamics are better studied in the context of a time course (28).

Each step in a signaling cascade may itself occur in variable times, hence, even if we collect data snapshots over time, the physical sampling frequency should be such that ensures the true dynamics of the system are accurately captured. Furthermore, each snapshot rarely recapitulates a homogenous set of cells from a single signaling state due to the cell-to-cell variation of response rates that signaling processes exhibit (29). As a result, the recorded experimental time is expected to capture convoluted information of the relevant dynamics. These reasons underscore the importance of decoupling signaling states from one another by TI methods.

Regarding the process trajectory, this highly depends on the strength and duration of the upstream stimuli and the overall signaling network architecture. For example, if the perturbation is spontaneous and the network includes a negative feedback loop the shape of the dynamics will be transient, resembling a positively skewed curve (30). Then, depending on the duration of the effect, the response may be prolonged. In general, the signaling response dynamics can be described as being transient (positively skewed), sustained, of short-duration or pulsatile. To depict pulsatile patterns is very difficult in mass cytometry as it requires time series data to be collected. However, the other types of dynamic responses are in principle detectable.

Finally, a typical task in mass cytometry is that of immunophenotyping whereby cells are manually assigned to cell types so that each population in the data is identified and analyzed separately. This task is better known as gating. If cells remain ungated, they would consist of a mixture of cell types each of which encodes a different signaling program. As such, the dynamics of signaling would be different between cells in the dataset and the structure of $\vec{r}$ would be multidimensional. On the other hand, if cells are gated, $\vec{r}$ becomes

unidimensional and potentially easier to reconstruct as each specific cell type, and its subpopulations are expected to share many common signaling circuitries. At this point, we should highlight that the term "linear" trajectory refers to the sequential state transition of cells in the manifold that can be represented as points along the progression line $\vec{r}$; geometrically, this could also be a curve. Moreover, the shape of the trajectory should not be confused with the signaling network architecture that is usually depicted as a graph of interacting proteins.

TI Methods Suitable for Recovering the Dynamics of Signaling Pathways

Despite the plethora of TI methods, not all of them are suitable in the signal transduction setting. For example, if the data are gated methods that output only multidimensional trajectories are not suitable. This is because the signaling program should apply the same to all cells of the same population regardless of the time the population snapshot was captured. Of course, it is possible that ungated subpopulations not following the same signaling program with the parent population may exist. However, for the most cases, the data are expected to encode the dynamics of protein activation of only one singling response. Another feature that does not apply here is the necessity of user supervision to the algorithm. Many methods request the user to provide key nodes of the process trajectory beside an estimated origin of the process. Locating such nodes in datasets of differentiating cells is fairly easy using prior knowledge on the maturity of cells. In signaling datasets, however, the population of single cells does not change. Even the identification of the starting position is problematic in the signaling datasets because the activation level of the initial receptor-molecules is rarely measured. Therefore, it is preferable to employ methods that require user supervision as little as possible.

Under these restrictions, we evaluate below 12 TI methodologies whose software is freely available. For consistency, we based our analysis in the R programming environment in which the majority of TI methods are also available as free packages. Our choice for freely available methods is because the goal of this work is not only to evaluate the ability of TI methods to solve the signaling reconstruction problem, but also to provide practical recommendations for the application of these methods to researchers in the field. We also did not qualify methods that were unable to run successfully. Furthermore, to evaluate methods with minimum possible bias, we selected only methods that can operate without any user intervention.

Table 1 lists the final set of TI methods that we compare and Supplementary Table 1 briefly summarizes their mechanics. For a more detailed description, the reader is referred to the original references. The first two methods in Table 1, that is, Scorpius and Slingshot, were recently found to be the best performing ones in a comparative study on differentiation data (14). Because Slingshot allows the user to choose between several different dimensionality reduction methods, we assessed its performance in conjunction with PCA, tSNE, and Diffusion Maps. In total, we employ 12 different TI approaches; eight graph-based approaches (4 of which are variants of Slingshot) and four curve-based ones. All graph-based methods are able to output both unidimensional (linear) and branched trajectories, while all curve-based ones output only unidimensional (linear) trajectories. Two of the curve-based methods we employ, namely tSNE with Principal Curves (tPC) and Diffusion Maps with Principal Curves (DMPC), are regarded as baseline methods each of which combines a well-established dimensionality reduction method as a first step, that is tSNE (35) and Diffusion Maps (36), followed by a TI step using Principal Curves. tPC was introduced as an extension to another curve-based method that constructs branched trajectories (33). However, the combination of Diffusion Maps with Principal Curves (DMPC) is first reported in this work. All methods were run using the default input parameters (see Supplementary Table 2). Supplementary Table 3 provides justifications for the other methods we considered suitable in this work but were excluded from the analysis shown below.

Data

We employ data from a public mass cytometry dataset where 11 different signaling activators were applied independently on Peripheral Blood Mononuclear Cells (PBMCs) and the levels of 10 surface and 14 intracellular markers were

Table 1. List of the final set of TI methods

METHOD	TI PRINCIPAL	OUTPUT	REF.
SCORPIUS	Curve	Linear	(27)
Slingshot with PCA (with 2, 5 and 10 PCs)	Graph	Linear or branched	(31)
Slingshot with tSNE	Graph	Linear or branched	(31)
Slingshot with diffusion maps	Graph	Linear or branched	(31)
MonocleDDRTree	Graph	Tree	(17)
SLICER	Graph	Linear or branched	(32)
TSCAN	Graph	Linear or branched	(15)
DeLorean	Curve	Linear	(22)
tSNE with principal curve	Curve	Linear	(33)
Diffusion maps with principal curve	Curve	Linear	(20,34)

measured in eight sequential time points (0 min, 1 min, 5 min, 15 min, 30 min, 60 min, 120 min, and 240 min) (36). Data were already pre-processed and manually gated for cell type and the response of 14 distinct cell populations to these activators are available on http://community.cytobank.org/cytobank upon request from the authors of the original publication (37). For TI evaluation, we selected the gated data on the response of $CD4^+$ and $CD8^+$ T-cells to Orthovanadate (pVO_4) and tetradecanoylphorbol acetate (PMA)/ionomycin perturbation. The selection of T-cells is based on the fact that this population is highly abundant; hence, it provides a large pool of single cell samples. In addition, the signaling pathways that the T-cells express are generally well established. On the other hand, the selected activators are well-known to lead to the strong perturbation of several intracellular signaling pathways, resulting in the general T-cell activation with high signal-to-noise ratio. The cell parameters we employ for TI are the recorded intracellular markers (pNFkB, pp38, pStat5, pAkt, pStat1, pSHP2, pZap70, pStat3, pSlp76. pBtk, pPlcg2, pErk. pLat and pS6).

Evaluation methodology

Despite the fact that the field of TI is constantly growing, reliable metrics to evaluate the performance of developed methods are scarce (13). Moreover, in the case of learning the dynamics of signaling pathways there is, to our knowledge, no available study. For this reason, we designed four novel performance evaluation metrics and employed one more from the available literature. Those that we designed serve to (i) evaluate the biological interpretation of the inferred dynamics, (ii) assess the quality of the inferred ordering, (iii) estimate the robustness with respect to perturbation to the input data; the sensitivity of the TI methods to the specific sample and (iv) quantify the relation between the data and the inferred trajectory. The former metric is context-specific but can be applied to any kind of TI where signaling dynamics are involved. On the other hand, the second metric is a general one for when time course data are available, while the latter two are general metrics suitable for any biological context where TI can be applied. We also adopt a metric developed to quantify the smoothness of the resulting trajectories and assess the variability in consecutive protein levels (22). Supplementary Table 5 (Information Metrics) indicates the algorithmic steps for each of the proposed metrics. A brief overview is given next.

The first metric we designed acts as biological validation; we refer to this metric as the *biological consistency* metric. Its main assumption is that signaling dynamics approximate a transient profile. On such basis, the activation level of one protein is expected to reach its maximum before the level of the protein that is next in the sequence of the signaling cascade. This means that the order of peak activation pseudotimes should correspond to the order of the respective proteins in known signal transduction cascades. Thereupon, our algorithm first scales the output pseudo-time between 0 and 1 so that the TI results are comparable. To reduce the variability in the trajectories and locate the position of the peak, a smoothing spline is fitted between the pseudo-time and the levels of each protein. The positions of the peaks are located from the projection of the maximum smoothed values on the pseudo-time axis. The resulting peak positions from all proteins are, in turn, sorted and all consecutive protein pairs are recorded. We employ established signaling pathways to examine the degree that each protein pair conforms to background biological knowledge (38). To this end, we quantify the consistency of a TI method by calculating the number of times the correct pairwise sequence is recovered. Because we provide no supervision to the algorithms the peak positions will depend on the inferred starting point of the trajectory. For this reason, we determine the consistency over 100 randomly chosen starting cells and assign the maximum score to the algorithm.

The second metric assesses the quality of the inferred latent variable, which is the pseudo-time, using the physical time of data collection as prior knowledge. Particularly, we define the agreement between experimental time and pseudotime as the probability of the inferred ordering of two cells being the same as their ordering with respect to the experimental time. To calculate this quality measure, we sample 10K pairs of cells collected at different experimental timepoints and enumerate the orderings that are in agreement. Because the cell ordering with respect to the pseudotime also depends on the starting point of the inferred trajectory, we calculate the probabilities achieved over 100 randomly chosen staring cells and, as previously, we assign the maximum one to the respective algorithm. We refer to this metric as the *experimental time consistency*.

The third metric tries to quantify the robustness of TI methods by virtue of their inherent stochasticity and data subsampling. To achieve this goal, each method is run once and the inferred ordering is recorded in a vector. A uniform random sample of single cells, typically 20%, is then drawn from the input time-course dataset, and TI is repeated. To measure the robustness between the two runs, we employ the Spearman and Kendall Rank correlation. The former measures the monotonicity in the ordering between the common cells of the initial ordering vector and the subsampled cell ordering vector while, the latter, measures how identical the relative positions of the common single cells are. If the TI method is robust, both Rank correlations would approximate +1 or − 1 depending on the starting position of the curve/path. Because the ordering is directionless, this starting position is critical for the calculation of the ranks. To avoid any stringent rule such as setting a common starting position in all TI attempts, we find the maximum possible correlation between the two vectors by cyclically shifting the elements of the subsampled ordering vector. This procedure is iterated 100 times and the average maximum Rank correlations are calculated, in absolute values, for each algorithm.

The fourth metric we devised seeks to assess the correlation of the inferred ordering with the input time-course dataset. Possible highly correlation implies that the TI method exploited the most of the information embedded in the data. To evaluate this, we calculate the distance correlation between

the vector of cell orders and the data (39). Distance correlation is a general statistical measure of multivariate association that is insensitive to nonlinear relationships and can be drawn to characterize independence between vectors in arbitrary dimensions. On such basis, a distance correlation equal to zero indicates independence between the inferred ordering and the original data. These properties render distance correlation suitable in the context of quantifying the association between the pseudo-time and the initial variables of the dataset.

Finally, the fifth metric we employ evaluates how rough the trajectory of each protein is by calculating the differences in consecutive protein levels after ordering (22). Low metric values indicate smooth trajectories. To see whether the score is statistically significantly different than that of a random ordering, a one-sided *t*-test is also applied. As performance score for this metric we find the average of all individually significant scores.

Results

Signaling Dynamics Reconstruction

Figure 1A shows the distribution of cells across experimental time for each protein after the activation with PMA. Similarly, the distribution of cells against pVO_4 is seen in Figure 1B. Each point refers to a cell and different colors refer to the different experimental time. The continuous curve in each subplot denotes the smoothing spline model that was fit in order to identify the positions of peak activation for each protein. We will regard these inferred dynamics as being the baseline and will refer to this peak detection approach as the *naïve* approach. It is easily seen that the peak activation occurs almost simultaneously for most of the activated proteins. This implies that the temporal resolution requires significant improvement in order to accurately depict the process dynamics.

Before applying any TI method on the data, we examined whether the heterogeneity at each time point is confounded by cell parameters such as cell cycle and cell size. Similarly, we examined whether the main effect in the data is signaling or some other cell behavior such as differentiation or proliferation. For the former case, we employed information on cell size and the DNA content for each cell and tried to establish the correlation of these two features with the phosphorylated protein abundances. Supplementary Figure 11A shows the respective statistically significant correlations ($p < 0.01$) in the case of pVO_4 data. In general, with the exception of pBtk, the effect of both cell features is relatively small on the phosphorylated proteins. On top of that, this faint association pattern changes in the case of PMA activation as seen in Supplementary Figure 11B. Therefore, the main driver of the heterogeneity should be the signaling process and not some other cell behavior. Accordingly, Supplementary Figure 12 displays the distribution of DNA content and cell size at each timepoint. As seen, no apparent association between experimental time and any of the two cell features exists. Therefore, the main driver of the observed dynamics is signaling and not some other biological process.

Next, we applied TI on the protein activation data. Figure 2 shows the result of using Slingshot with t-SNE on the dataset where $CD4^+$ T cells were perturbed with PMA. Similar figures of results for the other TI methods and datasets can be found in the Supplementary Figures 1–10. Figure 2A illustrates the distribution of cells in the low-dimensional space. To be comparable with the results of the naïve approach in Figure 1, the same color-coding was used. From the geometry of the low-dimensional embedding, it is easily observed that t-SNE manages to separate adequately the samples from each

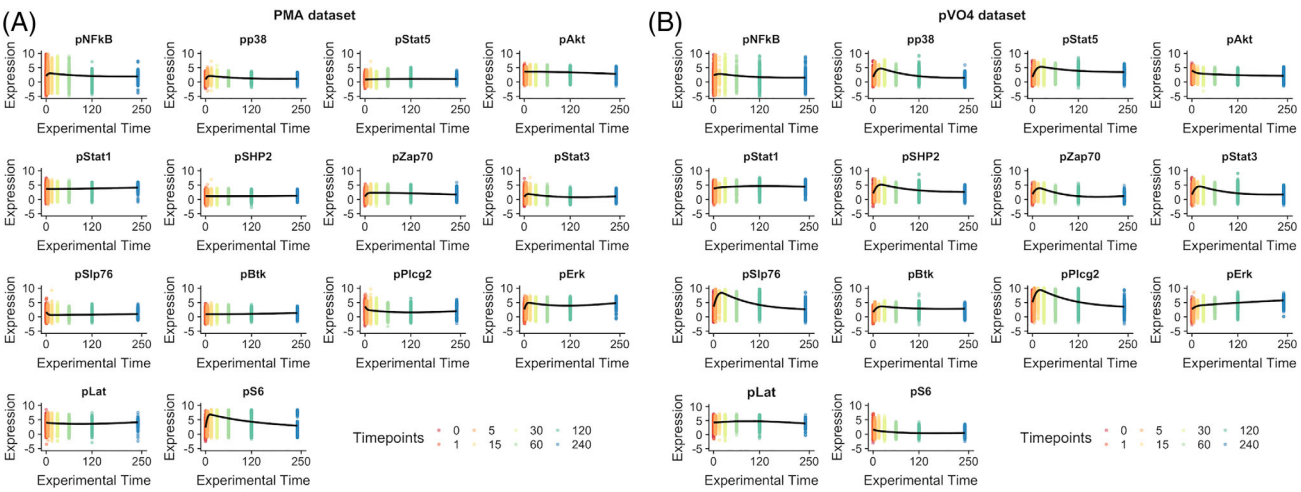

Figure 1. Baseline trajectory inference after the activation with PMA/ionomycin (A) and pVO4 (B). Each subplot in both (A) and (B) depicts the abundance of each of the 14 activated proteins against experimental time, where each point refers to a cell. Different cell colors refer to the different experimental timepoints. The time is measured in minutes. The continuous curve in each subplot denotes the smoothing spline that was fit in order to identify the positions of peak activation for each protein. The peak activation occurs almost simultaneously for most of the activated proteins implying that the temporal resolution requires significant improvement in order to accurately depict the process dynamics. [Color figure can be viewed at wileyonlinelibrary.com]

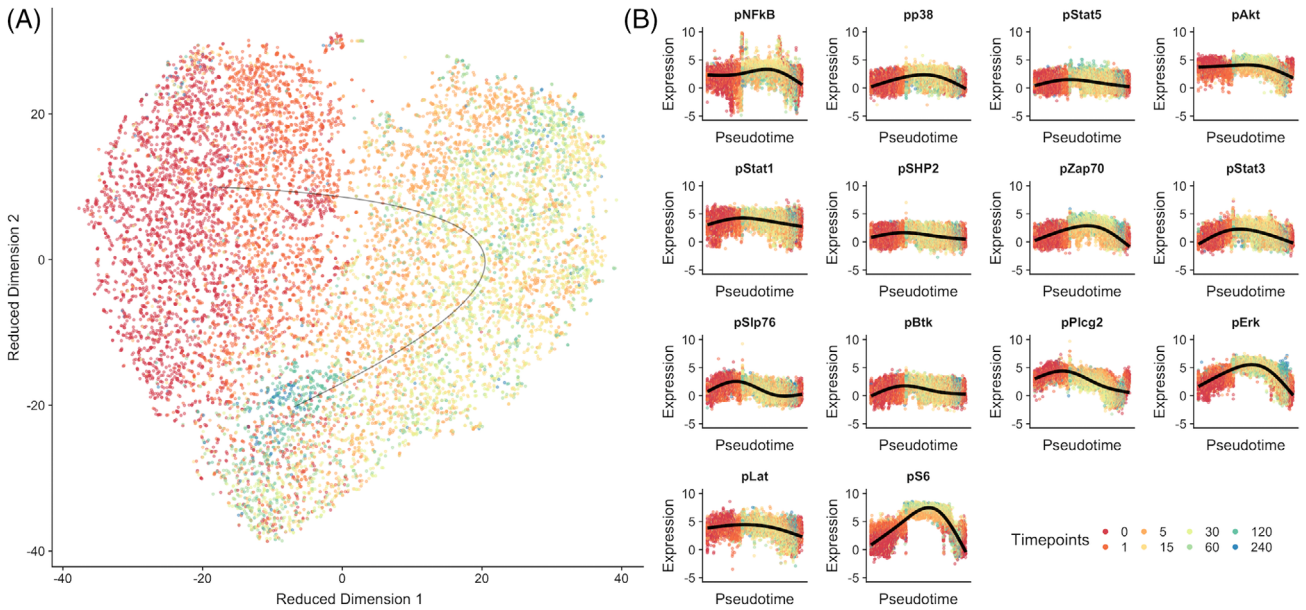

Figure 2. Result of slingshot with tSNE on the PMA dataset. (A) The distribution of cells in the first two tSNE components. Each cell is depicted as a dot. Colors match the experimental time points from which the cells where measured. (B) The abundances of the 14 proteins are plotted against the pseudo-time derived from slingshot (using a minimum spanning tree and principal curves). The continuous lines denote the fitted spline model used for peak detection. As far as the peaks in phosphorylation models are concerned, the method lead to a variety in the pseudo-time, constructing realistic trajectories. [Color figure can be viewed at wileyonlinelibrary.com]

temporal measurement. Then, if we imagine a trajectory that starts from the top left corner and follows a counter-clockwise direction, we can see that the order of experimental timepoints is 0, 240, 120, 60, 30, 15, 5, 2, and 1 min. Such trajectory illustrates a critical point in TI, that is, that the ordering is directionless. This proves the necessity (in some cases) of setting an initial position to the trajectories and the strong need of metrics that takes this into account. Figure 2B displays the protein abundances against pseudo-time after reordering cells by the final trajectory. To identify the position of the peak, a cubic spline was fit. Arguably, spline fitting is necessary in graph-based methods whose output is a connected path rather than a continuous line as in the case of curve-based methods. However, to maintain equality, we employ a spline model for peak detection in curve-based results as well. Compared to the naïve approach, the positions of the peaks between the proteins are well separated. These results imply that TI can, in principle, improve the temporal resolution of the collected data.

Figure 3 describes the result from applying the DMPC method on the $CD4^+$, pVO_4 dataset. As before, Figure 3A displays the distribution of cells in the low-dimensional space created during dimensionality reduction, the diffusion maps in this case. According to the geometry of the manifold one optimal trajectory should start from the top-left corner and follow, again, a counter-clockwise direction. This time, however, the experimental timepoints appear to be in their physical ordering. Indeed, based on the resulting ordering of cells, the inferred trajectory appears to have recovered the dynamics of the signaling pathway realistically without the need of shifting or reversing the final ordering (Fig. 3B).

Our results also indicate that some algorithms may output multidimensional trajectories. Supplementary Figure 14 depicts such an example from when we apply Slingshot in conjunction with PCA where the five first principal components are used in the dimensionality reduction step. In the main, if the underlying trajectory is multidimensional, one would expect each inferred branch to include cells from all experimental timepoints since signaling is a continuous process. However, we systematically find that some temporal measurements are exclusively detected in one of the inferred branches. In the analysis that follows, therefore, we excluded those algorithms whose inference was a multidimensional trajectory.

Performance Evaluation Results

Figure 4 depicts the comparative performance evaluation of the TI methods according to the different metrics. The exact values are indicated in Supplementary Table 4. As background pathway for calculating the biological consistency, we employed the one that was described in the original manuscript of the data (36). In general, there is no clear consensus as to which method performs the best in all datasets. However, in all cases, we examine the majority of methods that did not predict a multidimensional trajectory outperform the naïve approach in terms of the biological consistency. Furthermore, all methods have scored high in experimental time consistency with respect to a random reconstruction. Taken together these results indicate that TI is indeed improving the temporal resolution of the data in a way that is both biologically and experimentally consistent.

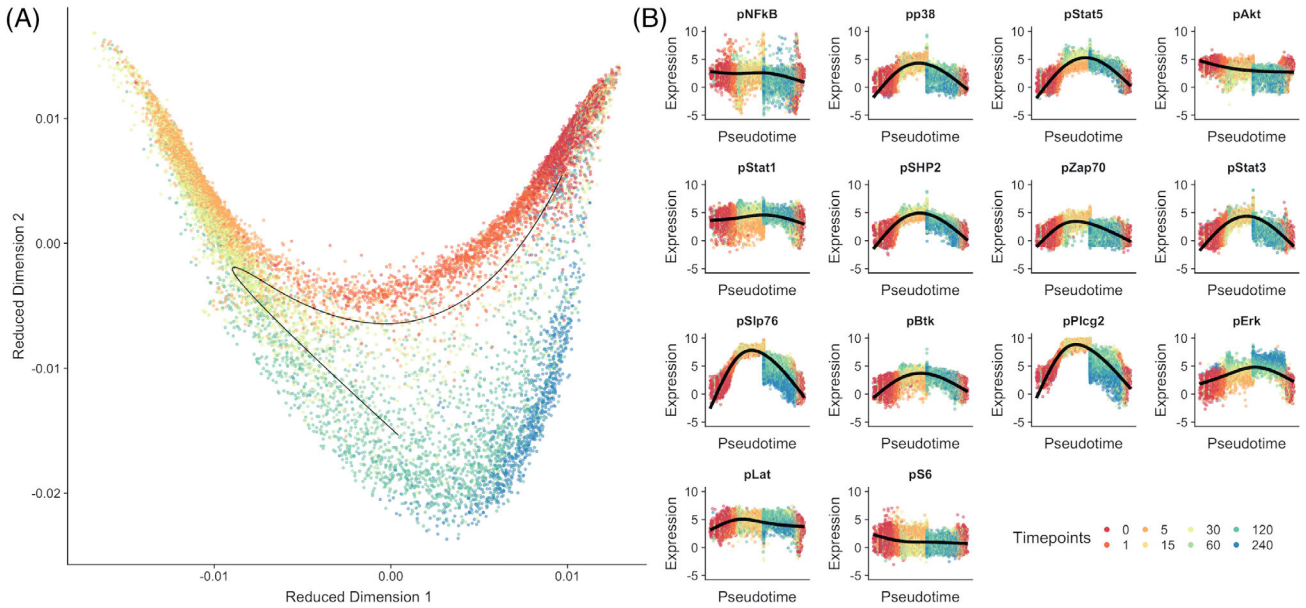

Figure 3. Result of diffusion maps with principal curves (DMPC) on the pVO4 dataset. As shown in Figure 2, (A) denotes the distribution of cells in the low-dimensional space generated by the diffusion maps algorithm. Similarly, (B) illustrates the abundances of the 14 proteins against the inferred pseudo-time along with their peak detection model fit. [Color figure can be viewed at wileyonlinelibrary.com]

In more detail, Figure 4A summarizes the scores on the $CD4^+$ T cells dataset when these were perturbed by PMA. Clearly, all methods that predict a sensible trajectory outperform the random model in all metrics with the exception of tPC for the roughness, which shows a comparable score. Then, in terms of biological consistency, the methods DMPC, Slingshot with tSNE and tPC perform equally well. In terms of time consistency, DMPC scores highest followed closely by Slingshot with tSNE. Similarly, DMPC achieved the lowest roughness value and high-distance correlation. The highest distance correlation was scored by TSCAN. On the other hand, SCORPIUS and Slingshot with tSNE are the most robust methods in terms of both Kendal and Spearman Rank correlation. Taken together, one may conclude that the DMPC has an overall better performance followed by Slingshot with tSNE. However, in terms of robustness its score is average.

Along the same lines, Figure 4B describes the respective TI scores on $CD4^+$ T cells when perturbed by pVO_4. As before, all methods outperform the random model with the exception of tPC and SCORPIUS that scored worse in biological consistency. Regardless, it is difficult to conclude which method is the best. For example, in terms of biological consistency, distance correlation and roughness, Slingshot with two-dimensional PCA scored highest. However, its robustness scores are of the lowest ones. On the other hand, Slingshot with t-SNE performs best in terms of experimental time consistency and robustness (Spearman), while it is second best in terms of roughness and biological consistency. In contrast, it is has among the lowest scores in distance correlation. These results indicate that selecting the best scoring method is not trivial. In addition, the performance of each method according to the same metrics appears to change between different datasets. Figure 4C,D illustrates the performance of TI algorithms on the datasets, where $CD8^+$ T cells were collected. Particularly, Figure 4C illustrates result for when PMA was used. As in the case of $CD4^+$, T cells for the same perturbation, all TI methods scored better than the random or the naïve model with the exception of SCORPIUS. Then, the most biologically consistent method was DeLorean followed by MonocleDDRTree, while the most consistent one with the experimental time measurements was DMPC. DMPC scores also highest in distance correlation followed, again by MonocleDDRTree. MonocleDDRTree scores also high in roughness but is not as robust as tPC and SCORPIUS were the most robust algorithms. Overall, we could say that the trajectory inferred by DMPC maintains its relation to the input data using most of the information they contain in a fairly robust manner. However, it performs marginally when detecting biologically relevant protein sequences. On the other hand, MonocleDDRTree performs comparably to DPMC for most metrics. On top of that, it has the second best consistency with the underlying biology. Finally, Figure 4D depicts the performance of TI methods when pVO_4 was applied. Surprisingly, TSCAN shows an overall better performance being among the highest scoring methods in all performance measures.

Taken together, one could say that DPMC is a TI method that performs well under PMA perturbation. Then, in all cases, we tested SCORPIUS is among the most robust and consistent with respect to the experimental time methods. However, in the majority of cases, it fails to detect the underlying biology. On the other hand, Slingshot with tSNE or two-dimensional PCA appear to perform well in the $CD4^+$ T cell datasets. This is however an exception because all other dimensionality reduction variants of Slingshot have failed to

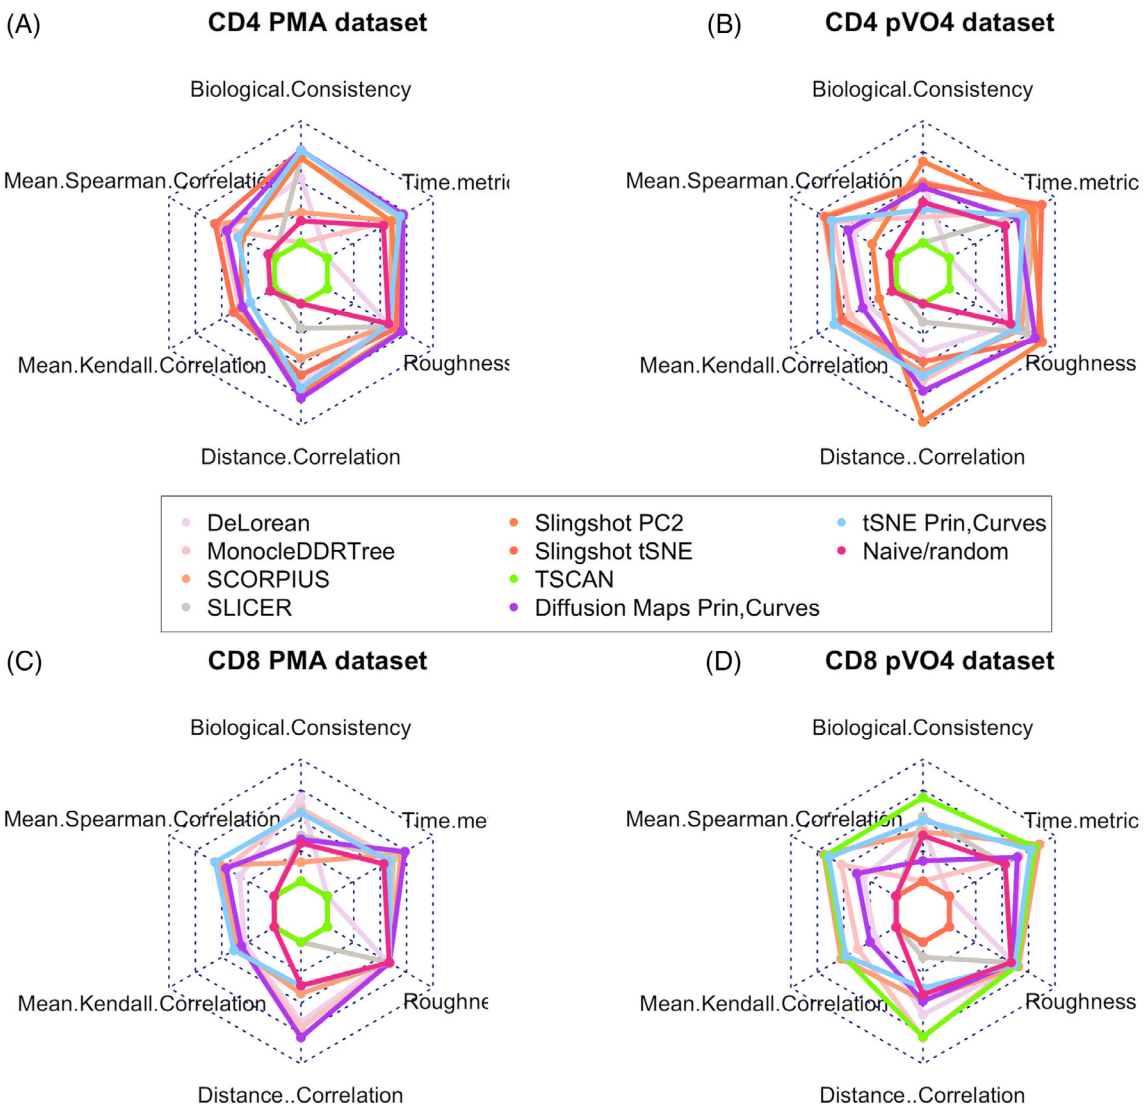

Figure 4. Performance evaluation summary. Each subfigure illustrates a spider plot where the TI scores on each metric are indicated for each method. Each metric is appended to the corner of the plot, while the performance of each TI method is described with a different color. The random (baseline) approach is indicated as a seventh method. For the biological consistency metric, the random approach refers to the naive approach described in the text where the true biological time is used. Because in roughness smaller values are better, the reported values are in reverse, that is, 2^{-} roughness. (A) The scores of each TI method on the $CD4^{+}$ PMA dataset. (B) The scores of each TI method on the $CD4^{+}$ pVO4 dataset (C) The scores of each TI method on the $CD8^{+}$ PMA dataset. (D) The scores of each TI method on the $CD8^{+}$ pVO4 dataset. [Color figure can be viewed at wileyonlinelibrary.com]

detect a sensible trajectory in all other datasets. For this reason, their performance was not evaluated. Of note is that in all results of Figure 4 there appears to be a correlation between the biological consistency scores and distance correlation and between experimental time consistency and robustness scores.

Discussion

TI Methods Can in Principle Be Used for the Reconstruction of Signaling Dynamics

The focus in this work was to explore the possibility of using TI methods in order to describe the dynamics of signaling pathways from single-cell data. The reason we employ temporal data and not just one cross-section as when modeling the progression of differentiation is because differentiation is a relatively slow process and a typical cross-section should capture cells from all its different stages. In contrast, the timescale of signaling is a lot shorter and one single-cell measurement is not guaranteed to capture the transition of cells between all states of an activated pathway. On the other hand, in the case of temporal data, there is no direct mapping between cells measured at sequential intervals. This means that to capture the underlying dynamics correctly the sampling frequency must be such that the abundance distributions of the proteins measured will remain unimodal over

time. Unfortunately, in mass cytometry the sampling frequency is such that the time snapshots rarely represent a homogenous set of cells from one single signaling state. For example, in the data, we employed the protein distributions often display bimodal or other forms of broad distributions (Supplementary Fig. 13). Therefore, the first goal in this work was to examine whether TI methods increase the temporal resolution of the data, thus setting the ground for the recovery of the underlying process progression.

To address this hypothesis we devised a novel metric that quantifies whether the inferred dynamics are consistent with known signaling pathways (biological consistency metric). Validation of TI methods with known biology is common among studies in the field. However, it is less complicated for these studies because they refer to differentiation where tracking known phenotypic markers along progression lines or comparing against a known trajectory is straightforward (15,24). At the same time, however, this makes such validation schemes study-specific and cannot be generally applied for the evaluation of the TI methods under any given circumstance. Therefore, the validation methodology we proposed employs the cell ordering for setting the order of the measured variables (Supplementary Table 5). The latter ordering directly reflects the causality of the underlying biology; if a protein comes before another in a dynamic process then its activation will happen earlier in time. Here, one may object that this metric cannot be applied in the case of a pulsatile pattern. Still, even in the case of such dynamic profile, the time of maximum of the first pulse will be denoting the time when the respective protein was activated. On the other hand, this philosophy can be directly transferred to the differentiation domain. Although, there, the dynamics are multimodal and the second ordering problem becomes more complicated. We leave such an extension as future work.

Returning to our hypothesis, we have shown that the main contributor to the system dynamics is signaling rather than confounding factors such differentiation or proliferation. On such basis, our results indicated that almost all methods that predict sensible trajectories are able to score higher than the naïve model in all datasets. This shows that TI methods are able to leverage important biological information on the process of signal transduction by enhancing the temporal resolution of single cell signaling data.

Robustness of TI Methods

Despite the positive first outcome, the results also indicate that there can be no clear consensus as to which method is the best performing one. This is commonplace in the field because almost every study wherein a new TI method is proposed also describes a new metric (13,14). Many of these metrics are case-study-specific; hence, they cannot be used for the general evaluation of TI methods. For example, TI methods are independently evaluated on whether they are robust to: (i) changes in variables (40,41), (ii) subsampling (15,17,18,33,42) and (iii) their inherent stochasticity (24,43,44). Robustness to changes in the variables we employ for TI cannot be applied in mass cytometry data because the number of measured quantities is low compared to that of single-cell RNAseq data. Therefore, the evaluation methodology we propose here quantifies robustness in subsampling and inherent stochasticity based on two important principles. First, it is agnostic of any prior cell labeling or known trajectory and second; it ignores any user supervision even that of setting the starting point of the trajectory. Both of these allow our robustness evaluation methodology to be generally applicable in any TI method performance evaluation.

From our results, we learnt that the performance of the same method may vary considerably between different datasets. This is because the signal-to-noise ratio can vary substantially between single cell measurements and TI methods are very sensitive to noise. Along the same lines, there is also the artificial noise introduced during preprocessing of the data that has been shown to disturb the quality of multivariate analyses (45). Then, in terms of robustness, SCORPIUS was the overall best performing method followed by tPC in three datasets ($CD4^+$ under pVO_4 and CD8+ under PMA and pVO_4) and Slingshot with tSNE in two datasets ($CD4^+$ under PMA and pVO_4). One reason for the solid performance of SCORPIUS is that it is a curve-based method built to operate without any prior information about the dynamic process of interest (27). In addition, its dimensionality reduction mechanism is based on correlation distance between any two cells, a fact that makes its inferences quite resilient to noise.

Quality of the Inferred Trajectories and Association with the Original Data

In terms of quality of the inferred orderings, we devised a metric that evaluates the consistency between the inferred and the experimental cell orderings (experimental time consistency). That is, the quality of the increase in the temporal resolution that each TI method achieves. This measure is similar to the consistency measure described by Canoodt et al (27). Our results strongly indicate that TI methods manage to learn, to a significant extent, the order of cells over the physical time they were collected only by looking at the protein abundance levels. The main limitation of this metric is that we do not know how well the inferred pseudo-time reflects the true underlying biological time. For example, let us assume that the score of an algorithm is 0.8. This means that about 20% of the cells, it reorders were not positioned as expected from their time of measurement. However, this does not mean that they have been misplaced. In fact, this is what the inference concerns; to identify those cells whose experimental time is not in line with their biological time and reposition them accordingly. Therefore, in contrast to the biological consistency metric, since we cannot control for the true biological time the value of this metric should be carefully considered when drawing conclusions about the true underlying biology.

Another limitation of the temporal consistency performance measure is that it can be applied only when the inferred trajectory is unidimensional. In general, our analysis is liberal in that it includes several TI methods that predict multidimensional process trajectories even though this is highly unlikely when gated data are available. Our results

provide evidence of this hypothesis as the algorithms predicting a multidimensional trajectory fail to reflect the fact that signaling is a continuous process. For this reason, we have excluded some algorithms from our comparative analysis. Regardless, in the case of a sensible multidimensional trajectory, this metric would have to be adjusted for each branched path. We also leave such extension as future work.

Another metric we employ for quantifying the quality of the output trajectories is the roughness. In theory, the less variant is the outcome, the better the dynamics of the process are described. For the $CD4^+$ T cell datasets Slingshot and DMPC showed that they reliably output the smoothest trajectories. Likewise, SCORPIUS and DMPC for the $CD8^+$ T cell datasets. Even so, there is no guarantee that less variable trajectories will possess any information about the process imprinted on the multivariate data or describe just a pattern of noise. In fact, a close inspection of the roughness scores in all datasets concludes that the variability of the outcome of some methods can be comparable to that of a random ordering. It is critical therefore to examine how associated the generated dynamics are to the process that the original data have documented.

Typically, this can be achieved by comparing against a known process progression (27,44) or features derived from the original data such as clusters (46). Instead, what we proposed here is to employ the distance correlation, a statistical measure that directly assesses the independence between the ordering vector and the input data without the requirement of background knowledge or performing any other data analysis (e.g. clustering). Then, as expected, our distance correlation results show that the random ordering and the original data are indeed independent. In addition, most methods score relatively high indicating that they retain most of the information embodied in the data. Based on these scores DMPC performs best in two datasets, while TSCAN and Slingshot with two-dimensional PCA in the other two.

There are two interesting observations here. One is that whenever a method scores best in distance correlation, it also scores high in biological consistency. In fact the correspondence between the two metrics appears in the scores of several TI methods. Simply, this shows that the more information an algorithm retrieves from the data (i.e. higher distance), the better its prediction will be with respect to the underlying biology. Likewise, our second observation is that in many cases a method's robustness score relates positively to its experimental time consistency score. This is intuitive in the sense that for an algorithm to provide reproducible results the consistency of its inferred trajectories with the known physical time of data collection should be high. Taken together, these observations highlight the merits of the proposed metrics in validating and interpreting the result of TI algorithms.

Conclusions

The purpose of this work was to explore the possibility of using TI to model the dynamics of signaling. This first investigation clearly shows that the task is, in principle, possible. Although no particular method performed best in all datasets, we can safely suggest Scorpius and DMPC as good candidates. What we learnt from our analysis is that good performance depends on many different aspects, to name a few, biological perturbation, data quality, dimensionality reduction method selection and method stochasticity. Preliminary and anecdotal testing, which we leave as future work, also showed that filtering out proteins with minimal abundance variability improves the quality of TI. We also learnt that good-quality algorithms (e.g. Slingshot) that have been tested on high-dimensional data with small sample size, slow dynamics, and complex structures (e.g. bifurcations) do not work out-of-the-box for this setting. Fortunately, the modular framework of TI allows the combination of virtually any dimensionality algorithm method with any trajectory inference algorithm creating a large space of candidate methods for solving this problem. For example, we show here for the first time a new TI method that combines Diffusion Maps with Principal Curves (DMPC) and performs more than adequately across the proposed metrics. Regarding these metrics, our goal was to provide the community with multiple performance quality measures that are generally applicable for the evaluation of TI methods. Our findings indicate that the metrics we developed not only assess the algorithmic performance from many different perspectives but also provide intuitive interpretations of the respective results. Of course, there are still some aspects that require improvement in order to address more complicated process structures such as those with bimodal dynamics, branched trajectories, etc. Arguably there are still many challenges and open problems to be solved. Scripts to reproduce the preceding results are available from GitHub (https://github.com/mensxmachina/SignalingDynamicsReconstruction).

Acknowledgments

The research leading to these results has received fundingfrom the European Research Council under the EuropeanUnion's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement no. 617393; CAUSALPATH – NextGeneration Causal Analysis project. Funding for open access charge: ERC.

Literature Cited

1. Campbell N, Reece J. Biology. 7th ed. San Francisco: Benjamin Cummings, 2004.
2. Kholodenko BN, Hancock JF, Kolch W. Signalling ballet in space and time. Nat Rev Mol Cell Biol 2010;11:414–426.
3. Sigal A, Milo R, Cohen A, Geva-Zatorsky N, Klein Y, Liron Y, Rosenfeld N, Danon T, Perzov N, Alon U. Variability and memory of protein levels in human cells. Nature 2006;444:643–646.
4. de Vargas RL, Claassen M. Computational and experimental single cell biology techniques for the definition of cell type heterogeneity, interplay and intracellular dynamics. Curr Opin Biotechnol 2015;34:9–15.
5. Ornatsky O, Bandura D, Baranov V, Nitz M, Winnik M. A, Tanner S. highly multiparametric analysis by mass cytometry. J Immunol. Methods 2010;361:1–20.
6. Spitzer MH, Nolan GP. Mass Cytometry: Single cells, many features. Cell 2016;165: 780–791.
7. Finck R, Simonds EF, Jager A, Krishnaswamy S, Sachs K, Fantl W, Pe'er D, Nolan GP, Bendall SC. Normalization of mass cytometry data with bead standards. Cytom. Part A 2013;83:483–494.
8. Alberts B, Johnson A, Lewis J, Raff M, Roberts K, Walter P. Molecular biology of the cell. 4th ed. New York: Garland Science, 2002.
9. Liiv I. Seriation and matrix reordering methods: An historical overview. Stat Anal Data Min 2010;3:70–91.
10. Magwene PM, Lizardi P, Kim J. Reconstructing the temporal ordering of biological samples using microarray data. Bioinformatics 2003;19:842–850.
11. Gupta A, Bar-Joseph Z. Extracting dynamics from static cancer expression data. IEEE/ACM Trans Comput Biol Bioinforma 2008;5:172–182.

12. Czibula G, Bocicor IM, Czibula I-G. Temporal ordering of cancer microarray data through a reinforcement learning based approach. PLoS One 2013;8:e60883.
13. Cannoodt R, Saelens W, Saeys Y. Computational methods for trajectory inference from single-cell transcriptomics. Eur J Immunol 2016;46:2496–2506.
14. Saelens W, Cannoodt R, Todorov H, Saeys Y. A comparison of single-cell trajectory inference methods. Nat Biotechnol 2019;37:547–554.
15. Ji Z, Ji H. TSCAN: Pseudo-time reconstruction and evaluation in single-cell RNA-seq analysis. Nucleic Acids Res 2016;44:e117–e117.
16. Chu L-F, Leng N, Zhang J, Hou Z, Mamott D, Vereide DT, Choi J, Kendziorski C, Stewart R, Thomson JA. Single-cell RNA-seq reveals novel regulators of human embryonic stem cell differentiation to definitive endoderm. Genome Biol 2016; 17:173.
17. Qiu X, Mao Q, Tang Y, Wang L, Chawla R, Pliner HA, Trapnell C. Reversed graph embedding resolves complex single-cell trajectories. Nat Methods 2017;14:979.
18. Trapnell C, Cacchiarelli D, Grimsby J, Pokharel P, Li S, Morse M, Lennon NJ, Livak KJ, Mikkelsen TS, Rinn JL. The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells. Nat Biotechnol 2014;32: 381–386.
19. Setty M, Tadmor MD, Reich-Zeliger S, Angel O, Salame TM, Kathail P, Choi K, Bendall S, Friedman N, Pe'er D. Wishbone identifies bifurcating developmental trajectories from single-cell data. Nat Biotechnol 2016;34:1–14.
20. Hastie T, Stuetzle W. Principal curves. J Am Stat Assoc 1989;84:502–516.
21. Zwiessele M, Lawrence ND. Topslam: Waddington landscape recovery for single cell experiments. BioRxiv 2017;57778. https://doi.org/10.1101/057778.
22. Reid JE, Wernisch L. Pseudotime estimation: deconfounding single cell time series. Bioinformatics 2016;32:2973–2980.
23. Matsumoto H, Kiryu H, Furusawa C, Ko MSH, Ko SBH, Gouda N, Hayashi T, Nikaido I. SCODE: An efficient regulatory network inference algorithm from single-cell RNA-Seq during differentiation. Bioinformatics 2017;33:2314–2321.
24. Bendall SC, Davis KL, Amir ED, Tadmor MD, Simonds EF, Chen TJ, Shenfeld DK, Nolan GP, Pe'er D. Single-cell trajectory detection uncovers progression and regulatory coordination in human B cell development. Cell 2014;157:714–725.
25. Chen J, Schlitzer A, Chakarov S, Ginhoux F, Poidinger M. Mpath maps multi-branching single-cell trajectories revealing progenitor cell progression during development. Nat Commun 2016;7:11988.
26. Yotsukura S, Nomura S, Aburatani H, Tsuda K. CellTree: An R/bioconductor package to infer the hierarchical structure of cell populations from single-cell RNA-seq data. BMC Bioinform 2016;17:363.
27. Cannoodt R., Saelens W., Sichien D., Tavernier S., Janssens S., Guilliams M., Lambrecht B.N., De Preter K. and Saeys Y. (2016). Scorpius improves trajectory inference and identifies novel modules in dendritic cell development. bioRxiv, 079509.
28. Claassen M. Shooting movies of signaling network dynamics with multiparametric cytometry. In: Fienberg HG, Nolan GP, editors. High-Dimensional Single Cell Anal. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013; p. 177–189.
29. Spencer SL, Gaudet S, Albeck JG, Burke JM, Sorger PK. Non-genetic origins of cell-to-cell variability in TRAIL-induced apoptosis. Nature 2009;459:428–432.
30. Purvis JE, Lahav G. Encoding and decoding cellular information through signaling dynamics. Cell 2013;152:945–956.
31. Street K, Risso D, Fletcher RB, Das D, Ngai J, Yosef N, Purdom E, Dudoit S. Slingshot: cell lineage and pseudotime inference for single-cell transcriptomics. BMC Genomics 2018;19(1). https://doi.org/10.1186/s12864-018-4772-0.
32. Welch JD, Hartemink AJ, Prins JF. SLICER: Inferring branched, nonlinear cellular trajectories from single cell RNA-seq data. Genome Biol 2016;17:106.
33. Marco E, Karp RL, Guo G, Robson P, Hart AH, Trippa L, Yuan G-C. Bifurcation analysis of single-cell gene expression data reveals epigenetic landscape. Proc Natl Acad Sci 2014;111:E5643–E5650.
34. Haghverdi L, Buettner F, Theis FJ. Diffusion maps for high-dimensional single-cell analysis of differentiation data. Bioinformatics 2015;31:2989–2998.
35. Van Der Maaten LJP, Hinton GE. Visualizing high-dimensional data using t-sne. J Mach Learn Res 2008;9:2579–2605.
36. Bodenmiller B, Zunder ER, Finck R, Chen TJ, Savig ES, Bruggner RV, Simonds EF, Bendall SC, Sachs K, Krutzik PO. Multiplexed mass cytometry profiling of cellular states perturbed by small-molecule regulators. Nat Biotechnol 2012;30:858.
37. Kotecha N, Krutzik PO, Irish JM. Web-based analysis and publication of flow cytometry experiments. Curr Protoc Cytom 2010;53:10.17.1–10.17.24.
38. Kanehisa M, Goto S. KEGG: Kyoto encyclopedia of genes and genomes. Nucleic Acids Res 2000;28(1):27–30.
39. Székely GJ, Rizzo ML, Bakirov NK. Measuring and testing dependence by correlation of distances. Ann Stat 2007;35:2769–2794.
40. Campbell K, Ponting CP, Webber C. Laplacian eigenmaps and principal curves for high resolution pseudotemporal ordering of single-cell RNA-seq profiles. bioRxiv 2015;27219. https://doi.org/10.1101/027219.
41. Campbell K, Yau C. Ouija: Incorporating prior knowledge in single-cell trajectory learning using Bayesian nonlinear factor analysis. Biorxiv 2016;2016:60442.
42. Sharma M, Li H, Sengupta D, Prabhakar S, Jayadeva J. FORKS: Finding orderings robustly using K-means and Steiner trees. bioRxiv 2017;132811. https://doi.org/10.1101/132811.
43. Gut G, Tadmor MD, Pe'er D, Pelkmans L, Liberali P. Trajectories of cell-cycle progression from fixed cell populations. Nat Methods 2015;12:951.
44. Liu Z, Lou H, Xie K, Wang H, Chen N, Aparicio OM, Zhang MQ, Jiang R, Chen T. Reconstructing cell cycle pseudo time-series via single-cell transcriptome data. Nat Commun 2017;8:22.
45. Papoutsoglou G, Lagani V, Schmidt A, Tsirlis K, Cabrero D, Tegnér J, Tsamardinos I. Challenges in the multivariate analysis of mass cytometry data: the effect of randomization. Cytometry Part A 2019;95(11):1178–1190. https://doi.org/10.1002/cyto.a.23908.
46. van Dijk D, Sharma R, Nainys J, Yim K, Kathail P, Carr AJ, Burdziak C, Moon KR, Chaffer CL, Pattabiraman D, et al. Recovering gene interactions from single-cell data using data diffusion. Cell 2018;174:716–729.e27.