

# Learning cell-specific networks from dynamical single cell data

Stephen Y Zhang <sup>\*1</sup> and Michael P H Stumpf<sup>1</sup>

<sup>1</sup>Mathematics and Statistics, University of Melbourne, Australia

## Abstract

Cell dynamics and biological function are governed by changing patterns of gene expression. Intricate gene interaction networks orchestrate these changes. Inferring these interactions from data is a notoriously difficult inverse problem. The majority of existing network inference methods work at the population level. They construct static representations of gene regulatory networks, and they do not naturally allow us to infer differences in gene regulation across heterogeneous cell populations. Here we build upon recent dynamical inference methods that model single cell dynamics using Markov processes, which leads to an information-theoretic approach, locaTE, which employs the localised transfer entropy to infer cell-specific, causal gene regulatory networks. LocaTE uses high-resolution estimates of dynamics and geometry of the cellular gene expression manifold to inform inference of regulatory interactions. We find that this approach is superior to static inference methods, often by a significant margin. We demonstrate that factor analysis can give detailed insights into the inferred cell-specific GRNs. In application to three experimental datasets, we demonstrate superior performance and additional insights compared to standard static GRN inference methods. For example, we recover key transcription factors and regulatory interactions driving mouse primitive endoderm formation, pancreatic development, and haematopoiesis. For both simulated and experimental data, we find that locaTE provides a powerful, efficient and scalable network inference method that allows us to distil cell-specific networks from single cell data.

## 1 Introduction

Cell identity is controlled by the dynamics of gene regulation networks. Identifying and characterising gene regulatory interactions is a central aim of cell and systems biology. A frequently cited advantage of single cell assays is that population heterogeneity is captured, allowing for the detection and study of rare cellular phenotypes that would otherwise be unobservable from bulk studies [1]. This also allows us to detect variability within groups of cells which would conventionally be regarded as being of the same “type”. The popularity of single cell data has spurred rapid advances in data analysis and modelling methods such as trajectory inference [2], dimensionality reduction [3], and network inference [4, 5].

In contrast to bulk assays where only population-averaged expression can be measured, single cell data is more abundant, often vastly so. Modern datasets allow us to overcome the “large- $p$  small- $N$  problem” [6], in which the number of predictions outnumber the number of samples [7]. This problem has been plaguing gene regulatory network inference because the number of predictions grows quadratically with the number of genes considered. In this vein, a collection of methods have been developed for network inference from single cell data, some of which build on methods originally intended for bulk expression data, and others which are specifically tailored to single cell data. The recent surveys in [8, 9, 5] provide an overview of the current methodology. The overwhelming majority of network inference methods aim to reconstruct a single *static* network [5, 8] that describes the set of possible interactions occurring within an observed population. Given the biological importance of cellular heterogeneity, a natural expectation is that variations in transcriptional state may correspond to variations in (cell state dependent) regulatory interactions which cannot be represented as static networks [7], and indeed has been observed empirically in previous studies [10, 11].

---

\*stephenz at student.unimelb.edu.au

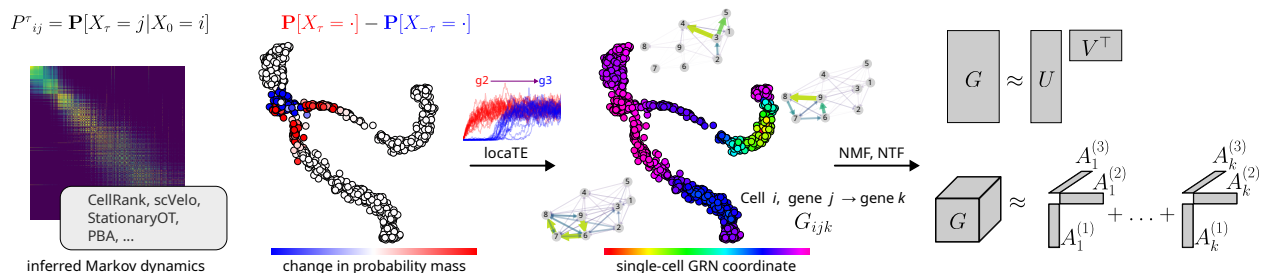


Figure 1: **Cell-specific network inference from estimated dynamics.** LocaTE takes as input a transition matrix  $P$  that encodes inferred cellular dynamics as a Markov chain on the cell state manifold. By considering the coupling  $(X_\tau, X_{-\tau})$ , locaTE produces an estimate of transfer entropy for each cell  $i$  and each pair of genes  $(j, k)$ .

The clear need for methods that can learn cell-specific networks [7, 12], is now leading to new classes of network inference methods [13, 14]. These methods seek to infer multiple networks (e.g. one per cell type, or even one per cell) from single-cell resolved data, thus allowing us to learn about variation in regulatory relationships over time and between conditions. Such methods generally rely on neighbourhood or cluster information to construct networks that are cell, time, or context-specific. Existing methods for cell-specific network inference have focused on inferring undirected networks. If from gene expression measurements we hope to infer *causal* or *directed* interactions, then information about underlying dynamics in a cell population is essential [15]. To this end, many methods for static network inference take advantage of pseudotemporal orderings to inform network inference [16, 17, 18, 19, 20]. The one-dimensional nature of pseudotime necessarily imposes a total ordering over cells, resulting in a loss of finer structure of the cell state manifold. Recent work on trajectory inference from a Markov process viewpoint [21, 22, 23, 24] depart from this framework and are able to model more complex dynamics on the set of observed cell states directly, and are free from assumptions on the trajectory topology (and do not require coarse graining of the data).

Here we adopt the transfer entropy (TE) [25] as an information-theoretic measure of causality. Transfer entropy, like other information-theoretical measures, is a model-free framework which has been used widely in contexts such as neuroscience where abundant time-series data are available. In a Gaussian setting it reduces to the linear Granger causality [26]. Transfer entropy has been used to predict static networks using a pseudotemporal ordering [16, 17]. We demonstrate that the transfer entropy can be adapted to infer cell-specific networks without imposing a prescribed ordering of cells.

Inference of cell-specific networks will help us to address problems in developmental biology, where it is often the transient dynamics that are the most interesting since these are likely to correspond to major epigenetic changes [7]. Our approach is tailored to such problems since we make no restricting simplifying assumptions: the cell state space is treated as a continuum, and *a priori* no partitioning or clustering is required to group or classify cells.

## 2 Method

**Local transfer entropy of expression dynamics** We represent the transcriptional state of a cell as a vector of expression levels  $X = (X^{(1)}, \dots, X^{(N)})$ . Transfer entropy does not make assumptions on the specific form of a model, but for the sake of illustration let us consider a model of cellular dynamics where the expression level of each transcript  $X^{(i)}$  is governed by a stochastic differential equation (SDE) of the form

$$dX_t^{(i)} = f^{(i)}(X_t^{(1)}, \dots, X_t^{(N)})dt + \sigma^{(i)}(X_t^{(1)}, \dots, X_t^{(N)})dB_t. \quad (1)$$

Note that this framework includes the chemical Langevin equation (CLE) [27], a biophysically derived model of single cell expression dynamics. Under the Itô interpretation [28, 29], for a small time interval  $\tau$  we may write approximately,

$$X_{t+\tau}^{(i)} = X_t^{(i)} + f^{(i)}(X_t^{(1)}, \dots, X_t^{(N)})\tau + \sigma^{(i)}(X_t^{(1)}, \dots, X_t^{(N)})\sqrt{\tau}Z^{(i)}, \quad (2)$$

where the  $Z^{(k)}$  are independent and identically distributed Gaussian random variables. It is straightforward to condition on  $\{X_t^{(i)} = x\}$ ,

$$X_{t+\tau}^{(i)} | \{X_t^{(i)} = x\} = x + f^{(i)}(\{X_t^{(k)}\}_{k \neq i}, x)\tau + \sigma^{(i)}(\{X_t^{(k)}\}_{k \neq i}, x)\sqrt{\tau}Z^{(i)}. \quad (3)$$

For any other gene  $X^{(j)}, j \neq i$ , following [30, Section 4.2], the *transfer entropy* with time lag  $\tau$  is defined to be

$$\text{TE}_\tau(X^{(j)} \rightarrow X^{(i)}) := \mathbf{E}_{X_t^{(i)}} \left[ \mathbf{I}(X_{t+\tau}^{(i)}, X_t^{(j)} | X_t^{(i)}) \right], \quad (4)$$

where

$$\mathbf{I}(X, Y | Z) := \mathbf{E}_{X, Y} \left[ \log \rho_{XY|Z} - \log(\rho_{X|Z} \otimes \rho_{Y|Z}) \right] \quad (5)$$

is the mutual information of two random variables,  $X, Y$ , conditional on a third variable,  $Z$ . Under the discrete approximation, the transfer entropy will be positive whenever  $f^{(i)}$  or  $\sigma^{(i)}$  depend on  $X^{(j)}$ , since then one has  $\mathbf{I}(X_{t+\tau}^{(i)}, X_t^{(j)} | X_t^{(i)}) > 0$ . In the biophysically relevant setting of the CLE, the functions  $f$  and  $\sigma$  share the same functional dependence on  $\{X^{(k)}\}_k$  when  $X^{(i)} = x$ , and so for infinitesimal time-scales  $\tau$  a non-zero transfer entropy implies regulation, i.e.  $\partial_j f^{(i)} > 0, \partial_j \sigma^{(i)} > 0$ . In the practically relevant setting where the time-scale  $\tau$  is short, the data processing inequality [31] guarantees that direct interactions will dominate indirect interactions, i.e. for a feed-forward loop  $X^{(i)} \rightarrow X^{(j)} \rightarrow X^{(k)}$ , one has that  $I_{ik} < \min\{I_{ij}, I_{jk}\}$  [32].

The concept of transfer entropy has previously been used for inference of static networks where pseudotemporally ordered expression data are available [16, 17]. However, in these methods the expectations in (5), (4) are taken with respect to the set of lagged tuples of cells along a total ordering. This introduces a strong dependence on the input pseudotemporal ordering; it also results in loss of resolution: any variation in regulatory activity over the expression space is integrated out to yield a static network. This can be problematic since biological trajectories may exhibit bifurcating, cyclic or converging trajectories, potentially requiring multiple runs of GRN inference on different subsets of a dataset.

We propose to take, for each observed cell  $\hat{X}$ , the expectation (5) with respect to the evolution described by the transition kernel  $P^\tau$  of (1) starting from some neighbourhood  $\mathcal{N}(\hat{X})$  of the cell. For short time scales,  $\tau$ , and an appropriately sized neighbourhood,  $\mathcal{N}(\hat{X})$ , (i.e. large enough to take expectations, but not so large as to lose local detail), the resulting dynamics are therefore local. This is best understood as a conditional transfer entropy,

$$\text{TE}_\tau \left[ X^{(j)} \rightarrow X^{(i)} \mid X_0 \in \mathcal{N}(\hat{X}) \right] := \mathbf{E}_{X_0 \in \mathcal{N}(\hat{X})} \mathbf{I}(X_\tau^{(i)}, X_0^{(j)} | X_0^{(i)}).$$

This enables us to infer regulatory interactions that are specific to the local cell context. Some information-theoretic static inference methods [4, 16] condition on a third gene in order to remove indirect interactions. Instead of conditioning separately on each potential indirect interaction, our framework conditions on a local region of the expression space. This can be interpreted as conditioning on some (potentially non-linear) function of the overall gene expression profile that is determined by the construction of the cellular manifold (e.g. in the case of PCA, this would be an affine combination).

Finally, since we model dynamics as a general Markov process, no assumptions on the ordering of cells or topology of trajectories is made or necessary.

**Inference in practice** Typically we do not have information about the drift and diffusion terms [28] in (1); dynamical information must instead be *inferred* from observed expression data. This *trajectory inference* problem has been widely studied [22, 33, 2]. Since we are interested in cell-specific networks, we restrict our focus to a class of methods that model cellular dynamics as a Markov process involving drift and diffusion over a manifold of cell states [21, 23, 34, 24]. Let  $\mathcal{X} = \{x_i\}_{i=1}^N$  be a sample from a population of cells, where  $x_i \in \mathbb{R}^M$  is the vector of mRNA expression values for cell  $i$ ;  $\mathcal{M}$  denotes a graph constructed from  $\mathcal{X}$  that can be thought of as approximating some underlying manifold of cell states [35]. We assume that we have access to a transition matrix  $P \in \mathbb{R}^{N \times N}$  that is supported on  $\mathcal{M}$  and that adequately describes the (“true”) dynamics of (1). In practice, as we do in Section 3, we could use a range of recently reported methods to

estimate  $P$ . Then  $\mathcal{M}$  equipped with  $P$  encodes a discrete Markov chain that approximates the true cellular dynamics. In what follows, we will write  $\llbracket n \rrbracket = \{1, \dots, n\}$  for  $n \in \mathbb{N}$ .

For any cell,  $x \in \mathcal{M}$ , we can construct (using e.g. a uniform distribution, truncated Gaussian, etc.) a probability distribution supported on its neighbourhood,  $\pi_0^x$ . We can then consider a Markov process  $X_t, t \geq 0$  that starts from  $X_0 \sim \pi_0^x$  and evolves under  $P$ . For some fixed  $\tau$  the resulting *coupling* is

$$(X_0, X_\tau) \sim \text{diag}(\pi_0^x)P^\tau =: \gamma_\tau^x. \quad (6)$$

For a pair of genes,  $(j, k)$ , we quantify the local transfer entropy for the causal relationship  $j \rightarrow k$  from knowledge of the coupling  $\gamma_t^x$  and the gene expression states,

$$\text{TE} \left[ X_0^{(j)} \rightarrow X_\tau^{(k)} \right] = \mathbf{E}_{(X_\tau, X_0) \sim \gamma_\tau^x} \left[ \mathbf{I}(X_0^{(j)}, X_\tau^{(k)} | X_0^{(k)}) \right], \quad (7)$$

Repeating this computation for all cells,  $x_i \in \mathcal{M}$ , and gene pairs,  $j, k \in \llbracket M \rrbracket$ , we arrive at a tensor,  $\hat{G} \in \mathbb{R}^{N \times M \times M}$ , such that,  $\hat{G}_{ijk}$ , is the transfer entropy score of  $j \rightarrow k$  in the neighbourhood of cell,  $x_i$ . Since the TE score matrix,  $\hat{G}_i$ , for each cell has been learned from a local neighbourhood, it is potentially noisy. Furthermore, it is impractical (due to limited data and computational resources) to condition on potential indirectly interacting genes [16]. Thus, TE scores may contain noise and spurious signals resulting from indirect interactions. To deal with this, we first filter interaction scores using context likelihood of relatedness (CLR) [36, 4] and then solve a manifold-regularised optimisation problem to both de-noise the interaction signals, and to attenuate noise and indirect interactions [16, 13].

**Context likelihood of relatedness.** Given raw TE scores  $\hat{G}$ , we use the context likelihood of relatedness [37] to filter scored interactions. The CLR algorithm has been shown to be an efficient mechanism for filtering out indirect interactions in the context of mutual information scores [37] and other information theoretical measures [4]. For a pair of genes  $(i, j)$  and given an interaction score matrix,  $A$ , we compute  $z_i$  (resp.  $z_j$ ) to be the  $z$ -score of  $A_{ij}$  with respect to  $A_i$ . (resp.  $A_j$ ). Then we define the *weighted* CLR score for  $(i, j)$  to be

$$\tilde{A}_{ij} = \frac{A_{ij}}{2} \sqrt{\max(0, z_i)^2 + \max(0, z_j)^2}. \quad (8)$$

Applying CLR filtering along the first axis of  $\hat{G}$ , we obtain the filtered tensor,  $\tilde{G}$ . We modify the original approach of [37] and weigh the CLR score by the initial MI value  $\hat{G}_{ij}$ . This is important since CLR was originally designed to filter interactions in static networks. In cell-specific networks, very few edges may be “active” in any given cellular context; then entire rows or columns of  $\hat{G}$  may consist only of noise. Computing  $z$ -scores along those rows or columns would put both noise and signal on the same scale.

**Smoothing and denoising** The tensor  $\tilde{G}$  contains a (noisy) matrix of filtered interaction scores, one for each cell. For notational convenience we write  $\tilde{G}$  as  $\tilde{G} \in \mathbb{R}^{N \times M^2}$ , i.e. each row is a length- $M^2$  unfolded score matrix. We propose to solve the optimisation problem  $G = \arg \min_X \mathcal{L}(X; L, \tilde{G})$  where

$$\mathcal{L}(X; L, \tilde{G}) = \frac{1}{2} \sum_{i=1}^N w_i \|X_{i \cdot} - \tilde{G}_{i \cdot}\|_2^2 + \frac{\lambda_1}{2} \text{tr}(X^\top LX) + \lambda_2 \sum_{i=1}^N w_i \|X_{i \cdot}\|_1, \quad (9)$$

$\lambda_1, \lambda_2$  are hyper-parameters,  $w$  is a vector of cell weights, and  $L$  is the symmetric graph Laplacian on  $\mathcal{M}$ . The term associated with  $\lambda_1$  is a manifold regularisation term corresponding to an expectation that regulatory relationships should vary smoothly with changes in cell state, i.e.  $\text{tr}(X^\top LX)$  is large for rapidly fluctuating  $X$ . The term associated to  $\lambda_2$  is a Lasso [6] term that encourages  $X$  to be sparse, reflecting our knowledge of biological networks [15, 13]. Together, the objective (9) encourages both parsimony and sharing of information along the cell state manifold. The problem is a case of L1-L2-regularised least squares, and can be solved efficiently using the scheme described in the appendix.

**Using a backward transition matrix** Since above we consider couplings of the process at times  $(0, \tau)$ , TE scores may reflect the underlying dynamics up to a shift in time. To remedy this, we consider couplings symmetric in time, i.e. between  $(-\tau, \tau)$ . While time-reversal of a diffusion-drift process away from equilibrium is generally not a well-posed problem, in practice, suitable backward operators have been constructed by transposing the transition matrix [23, 38] (see Appendix A.2 for further discussions). Given an approximate backward operator,  $Q$ , a time-symmetric coupling,  $\gamma_\tau^x = (Q^\top)^\tau \text{diag}(\pi_0^x) P^\tau$ , can be constructed.

**Extracting a static network** Cell-specific networks can be aggregated across a cell population to generate a static network. Averaging over individual cell-specific networks can be understood as applying the so-called “tower” property of conditional expectation. Here we observe that inferred cell-specific networks become less informative near the endpoints of trajectories, where temporal information is lost. Similarly, averaging may result in signals from small subpopulations of cells to be diluted. Thus, some reweighting strategies (or heuristics) will be useful to infer appropriate static networks.

**Computational considerations** The problem of cell-specific network inference is at least  $\mathcal{O}(NM^2)$  for  $N$  cells and  $M$  genes, since each cell-gene-gene triplet must be considered. For each triplet  $(i, j, k)$  locaTE calculates the joint distributions of  $(X_\tau^{(k)}, X_0^{(j)}, X_0^{(k)})$  under the coupling  $\gamma^{(i)}$ . While in the worst-case (if  $\gamma$  were dense) this would have complexity  $\mathcal{O}(N^2)$ , in all our applications we consider sparse  $\gamma$  where transitions only occur between nearby cells, in which the complexity is  $\mathcal{O}(|\gamma^{(i)}|) \ll N^2$ . The CLR filtering step is also  $\mathcal{O}(NM^2)$ , and the regression step primarily involves iterative matrix products of sparse  $N \times N$  against  $N \times M^2$ . We implemented locaTE using the Julia programming language [39] to take advantage of its speed for numerical computation.

We note that with exception of the regression step, which is global, all computational steps are easily parallelised across cells. The calculation of TE scores can alternatively be parallelised across gene pairs. This allows locaTE to take advantage of large scale parallelisation when available, potentially drastically reducing the compute time required. We find empirically that a large portion of computational burden arises from the construction of the joint distribution,  $(X_\tau^{(k)}, X_0^{(j)}, X_0^{(k)})$ . Since this is effectively an accumulation operation, this can be significantly accelerated using GPU computation. We observe speedups of 100 – 200 $\times$  using a NVIDIA V100 card<sup>1</sup>. Other steps in the locaTE pipeline (regression, factor analysis) amount to linear algebra operations and are also amenable to GPU acceleration. For datasets in excess of  $\sim 10^4$  cells, the computational cost could be greatly reduced by first grouping sampled cells into metacells, using e.g. the approach in [41].

## 3 Results

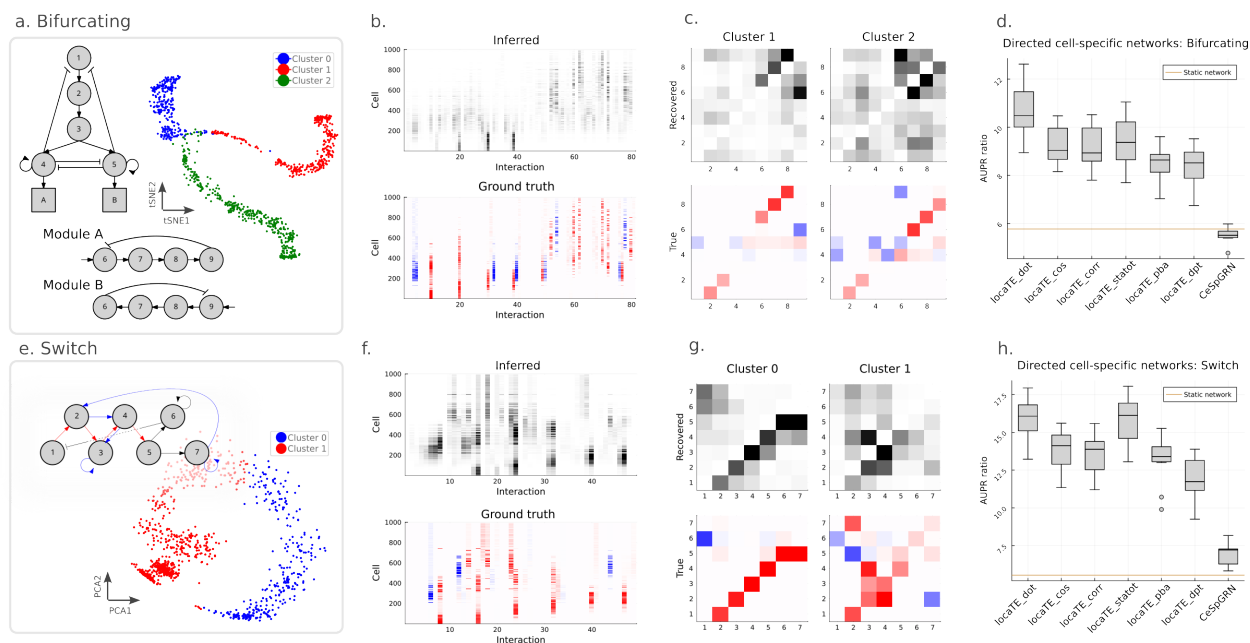
We first discuss two simulation examples where the true regulatory interactions are known by design. To illustrate the advantages of cell-specific networks, we consider examples where the network is rewired over time. Results for a wider range of “simple” developmental trajectories are provided in Appendix B.2.

### 3.1 Simulated data: branch-dependent logic

**Overview** We start by considering a scenario where a static network cannot capture the mechanisms underlying gene regulation. In the network in Figure 2(a) a bifurcation driven by a toggle-switch feeds into one of two “gene expression modules”. Modules A and B involve the same genes  $\{g6, g7, g8, g9\}$ , but with very different interactions among the genes: the flow of information in each module is the mirror image of the other.

A static network (in the graphical modelling sense) is insufficient to describe this system, since the directionality of some interactions depends on cellular context. Context dependence can alternatively be understood as arising from the presence of higher-order interactions. Consider gene 7, which may either be

<sup>1</sup>TE calculation for each cell in the Hayashi et al. mESC-Endoderm example [40] took 5ms (GPU, NVIDIA V100) vs. 847ms (CPU single core, Intel Xeon Gold 6242 @ 2.8 GHz)



**Figure 2: Simulated networks.** (a) Bifurcating network: 5-gene bifurcating subnetwork feeding into two modules A, B involving the same 4 genes  $g_6, g_7, g_8, g_9$ . 4-gene modules A, B each involve different regulatory interactions involving the same species; tSNE embedding of simulated cell profiles, coloured by inferred diffusion pseudotime cluster. (b) LocaTE output as a  $N \times M^2$  matrix in order of increasing pseudotime (top), shown alongside the ground truth Jacobians (bottom). (c) Interactions inferred by locaTE averaged by cluster (top) shown against ground truth interactions averaged by cluster. (d) AUPRC ratio for locaTE and other methods across 10 sampled simulation datasets, in the setting of cell-specific inference. We show as a horizontal line the AUPRC ratio achieved by perfect knowledge of the static network. (e) Switch network: Blue (resp. red) interactions are active when  $(g_6 \wedge (\neg g_5))$  is FALSE (resp. TRUE). Notably, gene 2 interacts both directly and indirectly with gene 4, depending on the system state. (f) LocaTE output as a  $N \times M^2$  matrix in order of increasing pseudotime (top), shown alongside the ground truth Jacobians (bottom). (g) Interactions inferred by locaTE averaged by cluster (top) shown against ground truth interactions averaged by cluster. (h) AUPRC ratio for locaTE and other methods across 10 sampled simulation datasets, in the setting of cell-specific inference. We show as a horizontal line the AUPRC ratio achieved by perfect knowledge of the static network.

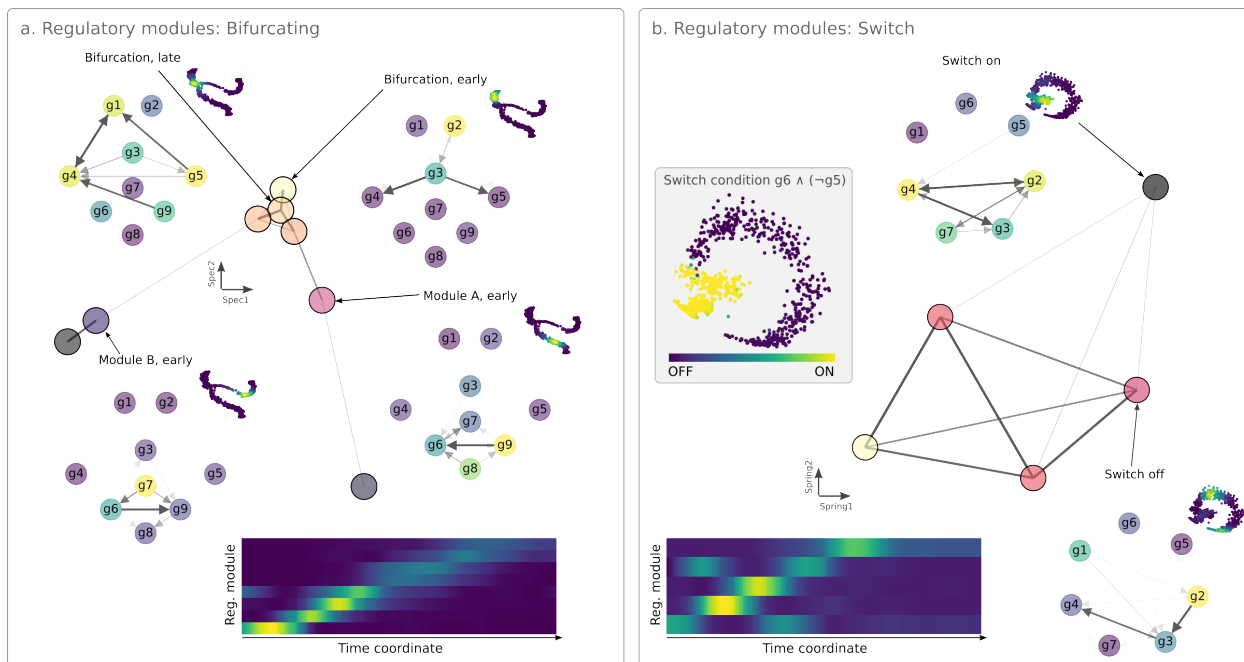


Figure 3: **Regulatory modules using locaTE-NMF.** (a) Top: Regulatory modules for the simulated bifurcating network found by locaTE-NMF with  $k = 8$ , shown as a graph spectral layout. Selected modules are shown (top 10% of edges), along with the module activity on tSNE coordinates in inset. Bottom: smoothed module activities against simulated time. (b) Top: Regulatory modules for the simulated switch network found by locaTE-NMF with  $k = 5$ , shown as a graph spring layout. Selected modules are shown (top 20% of edges), along with the module activity on PCA coordinates in inset. Bottom: smoothed module activities against simulated time.

activated by gene 6 (in module A) or by gene 8 (in module B). This could be understood conceptually as a *hyperedge*, since the corresponding Boolean rule would be

$$g7 \leftarrow (g6 \wedge (g4 \wedge \neg g5)) \vee (g8 \wedge (\neg g4 \wedge g5)).$$

Inference of higher-order interactions is known to be an even more challenging statistical problem than inference for pairwise interactions [15, 4]. Without trying to address this problem in generality, cell-specific networks allow us to disentangle interactions for a certain subclass of higher-order interactions: those which can be understood to be locally first-order, conditional on some latent variable that is encoded in the cellular context. Since existing cell-specific network inference methods [13, 14] infer undirected edges only, such methods are unable to distinguish the direction of information flow in this example and thus would predict identical networks for either branch.

**Simulation** We simulate the network in Figure 2(a) using the BoolODE software package [5]. This uses a chemical Langevin equation scheme for simulating biologically plausible expression dynamics:  $dX_t = f(X_t)dt + g(X_t)dB_t$  where  $f, g$  have specific forms dictated by the reaction network. From 1,000 independent realisations of the system’s trajectories, we generate a set of 1,000 sampled cells by sampling a cell state from each trajectory at a time chosen uniformly at random. As a measure of ground truth interactions for each cell  $x_i$ , we compute the corresponding Jacobians  $J_{ijk} = \partial_j f_k(x_i)$ . In Appendix B.1 we document the Boolean rules used to implement this network.

**Inferring dynamics** As input our method requires a cell-state transition matrix  $P$ . We construct an assumed “ground truth” matrix  $P$  from the velocity of sampled cells in the expression space (“RNA velocity”) as follows. For each observed cell,  $x_i$ , we calculate the corresponding velocity vector,  $v_i = f(x_i)$ . From this,

we can construct transition probabilities to a neighbourhood of  $x_i$  (in gene expression space),

$$P_{ij} = \frac{\exp\left(\frac{\langle x_j - x_i, v_i \rangle}{\sigma}\right)}{\sum_{k \in \mathcal{N}[i]} \exp\left(\frac{\langle x_k - x_i, v_i \rangle}{\sigma}\right)} \mathbf{1}\{j \in \mathcal{N}[i]\} \quad (10)$$

where  $\mathbf{1}$  denotes the 0-1 indicator function. We set  $\mathcal{N}$  to be the  $k$ -NN neighbourhood with  $k = 25$ , and chose the bandwidth  $\sigma$  following the median heuristic of [23]. We also consider other approaches for constructing  $P$ : as in [38], the inner product can be replaced with cosine similarity or Pearson correlation between  $x_j - x_i$  and  $v_i$  in order to estimate transition probabilities depending only on the orientation of  $v_i$ .

In addition to velocity-based methods, we consider a transition matrix based on diffusion pseudotime [42], which corresponds to a random walk biased in the direction of increasing pseudotime. We also consider the methods StationaryOT [34] and PBA [21], which were shown to have good performance for inferring stochastic dynamics without depending on pseudotime or velocity estimates.

**Inferring causal interactions** We apply the method described in Section 2 with a backward matrix obtained via the transpose (see Section A.2), for a range of parameter values  $(\tau, \lambda_1, \lambda_2)$ . To construct the initial neighbourhood distribution  $\pi_0^x$  we use quadratically regularised optimal transport [43], which produces a local neighbourhood density similar in effect to a truncated Gaussian. In Figure 2(b) we show the individual cell-specific networks obtained for  $\tau = 3, \lambda_1 = 25, \lambda_2 = 10^{-3}$  represented as vectors against the pseudotemporal ordering, and in Figure 2(c) we show the averaged networks over the cell clusters corresponding to the two branches. We observe overall that the de-noised result  $G$  resembles the ground truth. The averaged networks for module A and module B are directed and reflect the mirrored connectivities described in Figure 2(a). Averaging networks is convenient for summarising the inference output for cell-specific networks, but may obscure the cell-specific nature of the inference result. To illustrate this, in Figure S7(b) we show the interaction scores for the two edges  $g_6 \rightarrow g_7, g_7 \rightarrow g_6$ , which are specific to modules A and B, respectively. We observe that the flow of information between genes 6 and 7 is reversed on opposite branches, demonstrating that locaTE can resolve the cell-specific directionalities of edges.

**Comparison to other methods** There is a lack of methods for inferring cell-specific GRNs. CeSpGRN [13] is perhaps the method most similar to locaTE in terms of the problem it seeks to address: inferring a  $N \times G \times G$  array of cell-specific gene-gene interaction scores. It is limited to undirected interactions due to use of the (Gaussian) Graphical Lasso. For a larger set of comparisons we also consider methods for static network inference. We consider three information-theoretic methods: TENET [17] and Scribe [16], both of which infer static directed networks from pseudotime-ordered single-cell data, as well as PIDC [4] which infers an undirected network using partial information decomposition. We additionally consider SCODE [18] and (for the experimental datasets) GRISLI [20], which are ODE-based models for directed inference.

**Cell-specific edge detection** For each triple  $(i, j, k)$  with  $i \in \llbracket N \rrbracket$  and  $j, k \in \llbracket M \rrbracket$ , we treat the problem of detecting an edge  $j \rightarrow k$  in cell  $x_i$  as a binary classification problem with some threshold  $q$ . To construct the set of true positives, we consider the matrix  $\Pi J$ , where  $\Pi$  is a neighbourhood transition matrix such that  $\Pi_{i \cdot} = \pi_0^{x_i}$ . The motivation for considering this instead of simply using the raw Jacobians,  $J$ , is that cell-specific interactions are necessarily inferred using neighbourhood information. Strict cell-wise comparison to a ground truth would be overly stringent and sensitive to small perturbations in the expression space. By computing  $\Pi J$ , the ground truth signal is smoothed over the same neighbourhood that was used for inference, allowing for more robust assessment of classifiers. We classify all potential edges for which  $(\Pi J)_{ijk} > 0.5$  to be “true” edges.

In Figure S7(c) we show the Precision-Recall (PR) curves calculated for locaTE, as well as raw transfer entropy (TE), transfer entropy followed by CLR filtering (TE+CLR), and the method CeSpGRN [13] which infers cell-specific undirected networks. We find that the CLR filtering step improves upon raw TE scores in terms of prediction performance, and that locaTE improves substantially upon both TE and TE+CLR. This demonstrates that both the CLR filtering and the regularised regression step are essential for inference: this can be understood as a combined effect of (1) removal of indirect interactions using the CLR algorithm, and (2) sharing of information along the cell state manifold and attenuation of noise using the smooth and



sparse regression. We also compute a ground-truth static network by aggregating the cell-specific ground truth networks and calculate a static network AUPRC, corresponding to the best-case performance for a static method, i.e. if all cells were estimated to have the same static network.

To investigate the performance of locaTE further we use the BoolODE package to generate 10 datasets using different kinetic parameters for each simulation [5, Methods]. Thus, variation across datasets results both from the random sampling of cells, as well as differences in the simulated biophysics. For each dataset, we infer cell-specific networks using locaTE using several methods for trajectory inference. We consider three approaches using velocity information: dot-product, cosine similarity, and Pearson correlation. We also consider StationaryOT and PBA, as well as a transition kernel based on diffusion pseudotime [42]. In Figure 2(d) we summarise performance in terms of AUPRC ratio, that is the AUPRC for inferred cell-specific networks relative to that of a random predictor.

These results show that the trajectory reconstructions used as input to locaTE can affect the inference result: the dot-product velocity kernel outperforms the cosine and correlation velocity kernels. This is not unexpected, since cosine and correlation kernels depend only on the orientation of the velocity vector and not on its magnitude – thus losing information. StationaryOT performs similarly to cosine and correlation kernels, and the PBA and the pseudotime kernel perform worst. However, we find that all locaTE-based methods outperform the undirected method CeSpGRN, indicating that the availability of even potentially poor dynamical information can lead to improved inference of directed interactions. All methods perform at least as well as the static network baseline, illustrating the value of cell-specific network inference methods in general.

**Static networks as aggregates** As most methods in the literature deal with the static inference problem, we investigate the performance of locaTE by averaging local networks across cells to form a static network (see Figure S7(e) for representative examples). We show AUPRC ratios in Figure S7(d) for networks inferred by locaTE, as well as other selected static inference methods. In this setting we find that locaTE performs similarly well across trajectory inference methods, and that CeSpGRN performs marginally better. This can be understood, at least in parts, because the aggregated networks for modules A and B together in Figure 2(a) would be indistinguishable from the corresponding undirected network (as we show in the next example, this is not always the case). Competing static methods all perform worse, with PIDC performing best, and SCODE performing worst, broadly in keeping with the conclusions of [5].

**Downstream factor analysis** Instead of using the de-noising regression described previously, we can also apply techniques from factor analysis to the tensor of TE values. Methods such as non-negative matrix factorisation (NMF) and non-negative tensor factorisation (NTF) seek a simplified, low-rank description of the data. In the context of cell-specific GRNs, we are faced with an  $N \times G \times G$  array. By applying NMF (or NTF) to the flattened  $N \times G^2$  matrix (resp.  $N \times G \times G$  tensor), we effectively seek groups or *modules* of interactions that are active together in some region of the cell state space; here the regulatory state of a given cell is represented as a linear combination of some inferred archetypal regulation patterns. We derive a numerical scheme for finding matrix and tensor low-rank representations of the TE score matrix that incorporates the same smoothing and sparsity priors as (9). This can be used as a drop-in replacement for the (full-rank) problem (9). We describe these in detail in Appendix A.4 (NMF) and Appendix A.5 (NTF).

In Figure 3(a) we illustrate the utility of NMF and NTF for finding simplified representations of regulatory interactions inferred using locaTE for the bifurcating example. Applying locaTE-NMF with  $k = 8$  factors, we find regulatory modules whose activities are localised in the cell state manifold. Following a similar approach to [44], an undirected graph was constructed from the locaTE-NMF gene regulatory modules with edge affinities calculated from the Wasserstein distance between the normalised module activities. We present a spectral layout of this graph, which exhibits the bifurcating structure that reflects the underlying process. Examining the top edges of selected modules shows that key portions of the regulatory dynamics are recovered. Two distinct phases of the bifurcation are found: the first corresponding activation of genes 4 and 5 by gene 3, and the second corresponding to deactivation of gene 1 by genes 4 and 5. Along each of the branches, locaTE-NMF uncovers elements of the different regulatory logics that involve the same genes.

Tensor factorisation differs from matrix factorisation in that the learned atoms have the additional constraint of being separable. Because of this added structure, tensor decompositions are often found to have

better properties (identifiability of factors, robustness to noise [45]). We apply locaTE-NTF to find 16 rank-one factors of the scGRN tensor, and we apply  $k$ -means clustering to group the rank-one atoms into 8 clusters based on their activities. We find that the resulting clusters largely resemble the ones found by NMF. Unlike NMF however, the clusters can be decomposed into smaller rank-one components, and appear less noisy than the atoms found by NMF owing to the additional constraint. In Figure S8 we show how clusters 1 and 3 from the NTF representation split into rank-one components, demonstrating that the rank constraint leads to finer grained representations of regulatory states.

### 3.2 Simulated data: switch-like behaviour

We next consider a scenario where differential regulation occurs along a linear trajectory instead of a bifurcation. This inference problem is more subtle than a bifurcating process given the lack of a clear branching or clustering structure upon which regulatory interactions depend. We consider a model network depicted in Figure 2(e): this network either contains a feed-forward loop  $g2 \rightarrow g3 \rightarrow g4$ , or  $g2 \rightarrow g4 \rightarrow g3$ , depending on the position of a cell along the trajectory. The interaction  $g2 \rightarrow g4$  is first *indirect* and then becomes *direct*: the nature of this differential regulation is such that static GRN inference methods would observe a superposition of these interactions, i.e.  $g2 \rightarrow g3 \rightarrow g4$  as well as  $g2 \rightarrow g4$ . As an example, this could potentially lead to the  $g2 \rightarrow g4$  interaction being erroneously identified as an indirect interaction and thus being pruned, if filtering heuristics such as CLR or the data-processing inequality were used [37, 17]. Cell-specific networks may be able to disentangle these two effects.

We implement this network in BoolODE and simulate 1,000 cells. As before, we apply locaTE using a transition kernel derived from the simulated velocity information and the same parameters as in the bifurcating dataset. We find from Figure 2(f, g) that the inferred interactions accurately capture the true directed interactions. In Figure 2(h) we demonstrate that locaTE significantly outperforms competing methods in terms of AUPRC, for both cell-specific and static networks. locaTE substantially outperforms CeSpGRN in both settings for this example, in contrast to the previous example, where CeSpGRN performed well for static inference.

We next apply locaTE-NMF to find  $k = 5$  regulatory modules following the same procedure as for the bifurcating example. In Figure 3(b) we show a force-directed layout of the regulatory module graph, upon which we label modules that correspond to the regulatory logic when the switch condition  $g6 \wedge (\neg g5)$  is respectively on and off. In contrast to the results shown in Figure 2(g), which made use of a separate clustering step, this demonstrates that locaTE-NMF can identify distinct regulatory states in an unsupervised manner independent of clustering.

Together, our simulated network analyses demonstrate that locaTE is able to use dynamical information in the form of transition matrices to improve inference of directed interactions, often substantially so. Unsurprisingly, better estimates of dynamics lead to better results for network inference. Perhaps what is more remarkable is that, even after aggregating cell-specific networks produced by locaTE, the resulting static networks generally outperform competing static inference methods (see Figures S7(d, e), S9(d, e)).

### 3.3 mESC differentiation dataset using StationaryOT

We first consider the dataset of Hayashi et al. (2018) [40], consisting of 456 cells and 100 TFs. This dataset is a time-series of mouse embryonic stem cells differentiating through to primitive endoderm, sampled at 0, 12, 24, 48, 72h. Cells in this differentiation process follow a linear trajectory, as can be seen in Figure 4(a). To infer trajectories, we apply StationaryOT [34] using quadratically-regularised optimal transport to ensure a sparse transition kernel (see Section B.3 for details). Using this transition matrix, we infer cell-specific networks with locaTE, resulting in a  $456 \times 100 \times 100$  array of interaction scores.

Since no ground truth is available for cell-specific networks in this experimental dataset, we are only able to validate the output of locaTE against a static benchmark. In [19] the authors constructed a reference network from known regulatory interactions in the embryonic stem cell atlas from pluripotency evidence (ESCAPE) database [46], which consists of 19 regulators and 100 targets. Starting from the array of cell-specific interaction scores, we obtain a static  $100 \times 100$  GRN by averaging interaction scores across cells and compute precision-recall curves relative to the interactions in our benchmark network (Figure 4(c)). Cells from the last 10% of pseudotime were excluded since informative dynamics are lost at the end of the observed

trajectory. For comparison, we run CeSpGRN, which infers cell-specific undirected networks, and a collection of static inference methods. We compare a subnetwork of the inferred networks to the ESCAPE reference. In the directed setting, we find that locaTE outperforms other methods in terms of AUPRC, and performs well in terms of EP (0.496, compared to a best EP of 0.553 for TENET and baseline of 0.343). For undirected inference locaTE is out-performed by SCODE (AUPRCs of 0.453 and 0.483 respectively). Together with the observation that SCODE performs poorly for directed inference, this suggests that, despite being a method for directed inference, SCODE may be prone to detecting interactions but with the wrong directionality; PIDC [4] is the second best performing method in terms of AUPRC (0.456).

Since locaTE infers cell-specific networks, we can ask questions at finer resolution than possible at the population level. To demonstrate the power of this type of analysis, we use locaTE-NMF with  $k = 8$  factors to find regulatory modules that are active at different stages along the trajectory (Figure 4(d)). Hierarchical clustering of the inferred factor activities produces three groups, corresponding to “early”, “mid”, and “late” trajectory stages. For each cluster we compute a representative network by taking weighted averages of NMF factors, and the top 0.25% of edges for each stage are shown in Figure 4(e). It is clear that the nature of the inferred regulatory interactions changes drastically between stages. Included among the top inferred interactions are the canonical ones previously reported to be associated with maintenance of pluripotency and embryonic development, including *Nanog-Zfp42* [47], *Foxd3-Nanog* [48] and *Sall4-Klf2* [49] which are active early in the trajectory. Top inferred interactions for the late trajectory include *Epas1 (Hif2a)-Pou5f1 (Oct4)*, which is known to play a role in embryonic development [50]; and *Epas1 (Hif2a)-Bhlhe40 (Dec1)*, which plays a role in suppression of *Deptor* [51], a stemness factor known to be downregulated during ES differentiation [52].

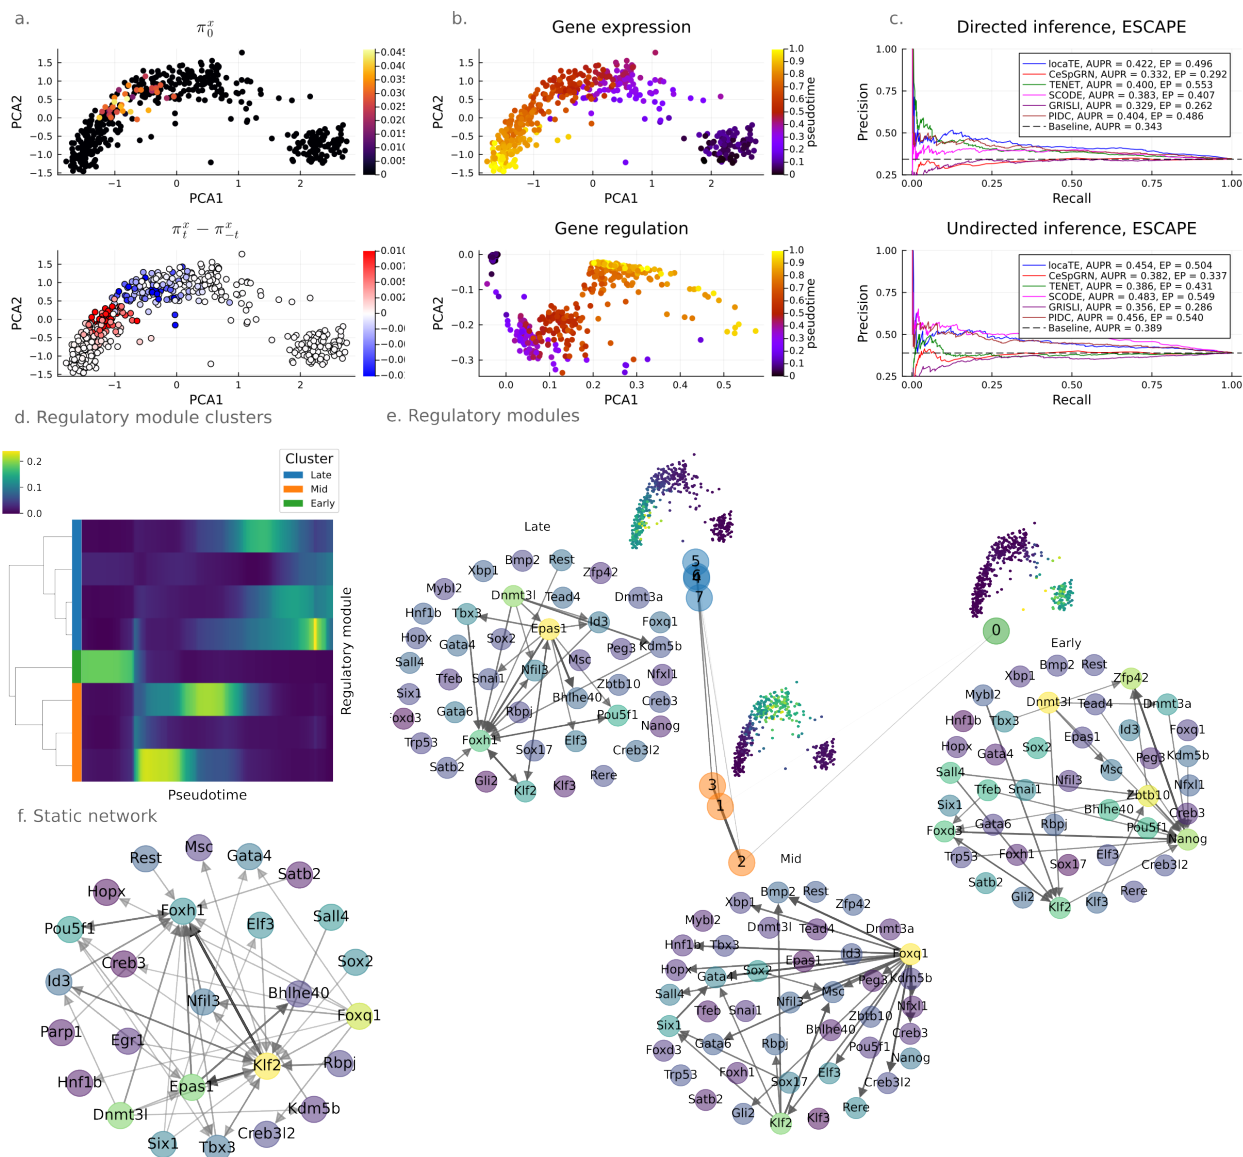
For each stage, we rank potential regulators by their out-edge eigenvector centrality and show the 25 top-ranked genes for each stage (see Figure S13(a)). We find that many of the top genes are known to be associated with stem cell development, including *Nanog*, *Pou5f1 (Oct4)*, *Sox2* [48], *Gata6* [53], *Zfp42* [47]. These gene lists are input to Metascape [54] to find corresponding Gene Ontology (GO) annotations (Figure S13). Among the gene set annotations found for the early trajectory, we find that enriched pathways with the highest confidence are those associated with pluripotency and embryonic morphogenesis. Annotations for the mid and late trajectory largely correspond to cell fate commitment, endoderm formation, and organ development. These annotations reflect the developmental progression of embryonic stem cell differentiation to primitive endoderm.

### 3.4 Pancreas development dataset with RNA velocity

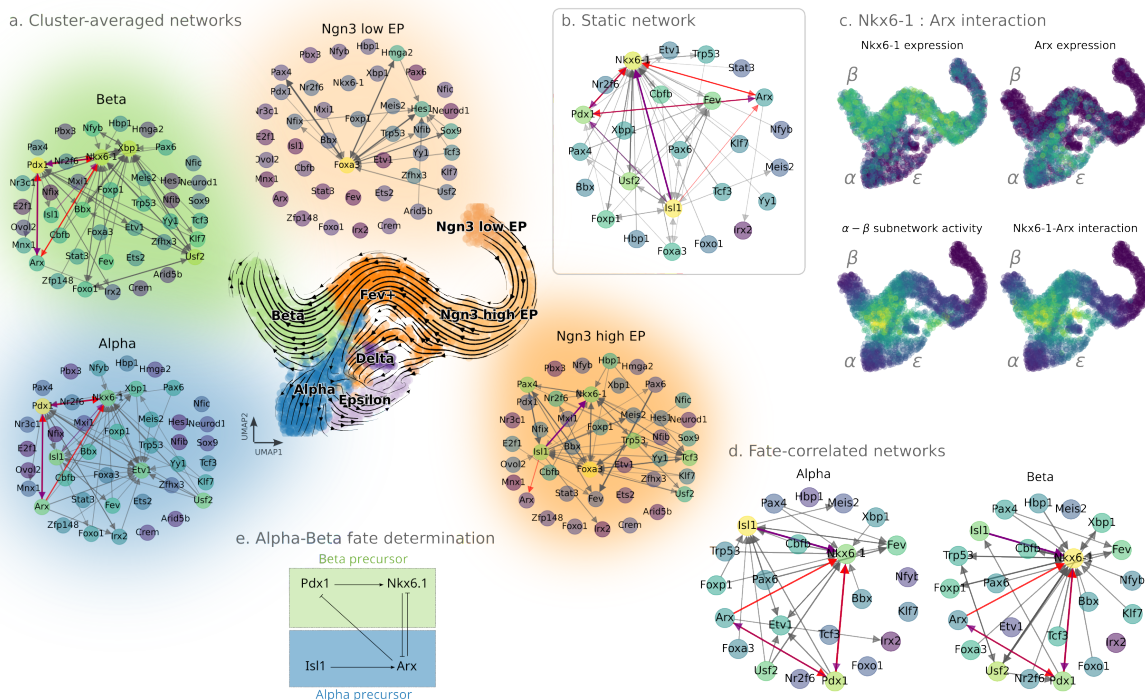
Next we consider a murine pancreatic development dataset of Bastidas-Ponce et al. [55]; here aspects of the cellular dynamics are captured in the form of RNA velocity estimates. The dataset consists of 2,531 cells, and we filter for 82 transcription factors following the preprocessing of [20]. We apply the scVelo [56] package to produce RNA velocity estimates, and CellRank [23] with a cosine kernel to produce a transition matrix encoding the inferred dynamics on the cell set.

We apply locaTE with the RNA-velocity derived transition matrix, and obtain a  $2,531 \times 82 \times 82$  array of cell-specific interaction scores. In Figure 5(a) we show a UMAP embedding of the full dataset with streamlines visualising RNA-velocity dynamics. To illustrate the heterogeneity in the networks inferred by locaTE, we also colour cells by their corresponding cell-specific GRN using the first diffusion component in “scGRN-space” (see Figure S14(b)). This reveals a clear gradient along the differentiation trajectory, suggesting that the networks output by locaTE capture a progression in regulatory interactions over developmental time. The highest-scoring interactions involve the genes *Fos* and *Egr1*, which are well known to play a role in stress response [58]. *Fos* expression has been documented to be an dissociation-induced artifact in scRNA-seq [59], and we therefore conservatively filter out interactions involving *Fos* or *Egr1* in subsequent analyses.

As a first summary of the inference output, we produce an  $80 \times 80$  static GRN by averaging (as before we exclude the last 10% of pseudotime). Figure 5(b) shows the top 1% of edges from this network. We observe that the genes with highest out-edge eigenvector centrality include *Nkx6-1*, known to be a master regulator of pancreatic beta cell development [57, 60], as well as *Isl1* and *Pdx1*, known determinants of alpha and beta cell fate [61]. In [57] the authors report evidence for a 4-gene network (henceforth referred to as the  $\alpha/\beta$  network) involving *Nkx6-1*, *Pdx1*, *Isl1* and *Arx*, that drives alpha and beta fate determination (Figure 5(e)). We find that all 4 genes have high out-edge centrality in the inferred static network, and furthermore we find



**Figure 4: ES-endoderm development dataset.** (a) PCA embedding of 456 cells along an ES-endoderm differentiation trajectory, coloured by (top) example neighbourhood density (bottom) change in density along evolution of Markov process. (b) PCA embedding of (top) expression profiles (bottom) cell-specific interaction networks, coloured by diffusion pseudotime. (c) Precision-recall curves against the ESCAPE reference for static inference, both directed and undirected. (d) Regulatory module activities found using locaTE-NMF, clustered by activity along the trajectory. (e) Graph layout of regulatory modules found by locaTE-NMF colored by cluster together with inferred networks (top 0.25% of all edges) for early, mid and late module clusters. Network nodes (TFs) are coloured by out-edge eigenvector centrality. Overall activity for each cluster of modules is shown (inset). (f) Static network produced by averaging locaTE scGRNs, excluding final 10% of pseudotime.



**Figure 5: Pancreatic development dataset.** (a) UMAP embedding of murine pancreatic development dataset [55], with smoothed RNA velocity vector field shown. Cells are coloured according to their cluster annotation. Averaged scGRNs are shown for selected clusters corresponding to Ngn3<sup>-</sup>, Ngn3<sup>+</sup> endocrine progenitor, alpha and beta cells respectively. Top 1% of edges are shown, after filtering for edges involving stress response gene Fos and Egr1. Direct (indirect) edges belonging to the  $\alpha/\beta$  subnetwork are shown in red (purple). Nodes are coloured by out-edge eigenvector centrality. (b) Static GRN constructed by averaging over first 90% of pseudotime. Top 1% of edges are shown, after filtering for stress response genes. (c) Nkx6.1-Arx expression and interaction scores in UMAP coordinates, and average activity of the 4-gene  $\alpha/\beta$  subnetwork of [57]; (d) Putative driver interactions associated with beta-cell fate. (e) Alpha-beta fate determination network of [57].

that the  $\alpha/\beta$  network is recapitulated in the top-1% edges inferred network, as highlighted by the coloured edges in Figure 5(a, b).

This dataset contains annotations for cell types. To investigate differences in regulatory state between subpopulations, we produce cluster-averaged networks for each celltype. In Figure 5(a) we show the top 1% edges of cluster-averaged, filtered networks for the Ngn3<sup>-</sup> EP (endocrine progenitor), Ngn3<sup>+</sup> EP, beta and alpha clusters (see Appendix S14 for remaining clusters). We note that the inferred network for Ngn3<sup>-</sup> EP cells includes Sox9-Hes1. Sox9 is known to play a key role in maintaining the pancreatic progenitor cell state by regulating Hes1 [62], an inhibitor of Neurog3 (Ngn3) [61]. This interaction is no longer present in the Ngn3<sup>+</sup> network, further evidence for its role in maintaining the Ngn3<sup>-</sup> cell state. This interaction also does not appear in the static GRN in Figure 5(b), illustrating the need for cell-specific networks in capturing localised signals, which would be static networks or population averaging. Out-edge eigenvector centrality of the Ngn3<sup>-</sup> EP cluster identifies additional potential regulators including Foxa3, known to be implicated in endodermal organ formation [63]; and Hmga2, a chromatin-associated protein expressed during development that does not directly regulate downstream genes, but appears to have widespread indirect effects via chromatin interactions [64].

Clusters corresponding to alpha and beta committed cells yield networks that are distinct from the progenitor clusters. In particular, we find that the  $\alpha/\beta$  network is prominent in both subpopulations. This implies that these genes play key roles in the fate determination mechanism. To investigate this at a finer scale we visualise expression levels and interaction scores for the Nkx6.1-Arx interaction in Figure 5(c). We observe a fate decision boundary from expression levels that reflects the competition between Nkx6.1 and Arx corresponding to beta and alpha fate. The interaction scores for both the Nkx6.1-Arx edge and the entire  $\alpha/\beta$  subnetwork are also strongest near this boundary. Examining again the out-edge eigenvector centrality, we identify Xbp1 as a top regulator specific to the alpha and beta subpopulations, recently found to be critical to maintenance of beta-cell identity [65].

Modelling of cellular dynamics as Markov processes naturally allows estimation of fate commitment in terms of absorption probabilities [21, 23, 34] and so examining correlation of gene interactions with fate commitment can be used to construct a putative “lineage driver” network that is correlated to a cell fate of interest. In Figure 5(d) we construct such a network for the alpha and beta cell lineages. As we expect, the  $\alpha/\beta$  network is strongly correlated to beta cell fate. We find that Isl1 and Nkx-6.1 have higher out-edge eigenvector centrality towards the alpha and beta lineages respectively, reflecting their roles in the known network.

### 3.5 Haematopoiesis with extensions of RNA velocity

One advantage of our Markov framework is that it is agnostic to how the dynamics are inferred. Currently, inference, interpretation and downstream use of RNA velocity estimates is a contentious topic [66, 67] and remains an active research area. Marot-Lassauzaie et al. [68] propose  $\kappa$ -velo, an alternative method for estimating RNA velocity and provide an example application to a haematopoiesis dataset where the more established method scVelo [56] fails to recover expected trajectories. We apply locaTE to this dataset using inferred dynamics computed by  $\kappa$ -velo. A bifurcation in the trajectories between myeloid and erythroid lineages is a prominent feature in this data set (see Figure 6(a)). Erythroid-myeloid cell fate determination is governed by a genetic switch comprised of mutually antagonistic Gata1 and Spi1 (known also as PU.1) [69] (see illustration in Figure 6(f)). In Figure 6(b) we show the locaTE interaction score for the Gata1-Spi1 interaction alone along with the smoothed RNA velocity vector field. We find that the peak activity corresponds to a “dead zone” in the velocity field, while in the surrounding regions downstream, the velocity fields are coherent and indicative of lineage commitment. This is of practical interest [68] as one interpretation of this could be that it corresponds to cells which have high fate plasticity and are still in the “decision-making” process.

We apply locaTE-NMF to find  $k = 16$  regulatory modules, and in Figure 6(d) we show a layout of the graph constructed from modules in the same manner as Figure 3 (see Figure S16(a) for all modules). From this we can see two clear branches corresponding to erythroid and myeloid fate. We find a module (see Figure 6(d), labelled Myel-Ery) which appears concentrated near the bifurcation point on the UMAP coordinates. This module is centered around activity of Spi1 and Gata1, as well as Mef2c, which is known to be involved in lymphoid-myeloid fate determination [70]. In this module locaTE-NMF detects the key Gata1-

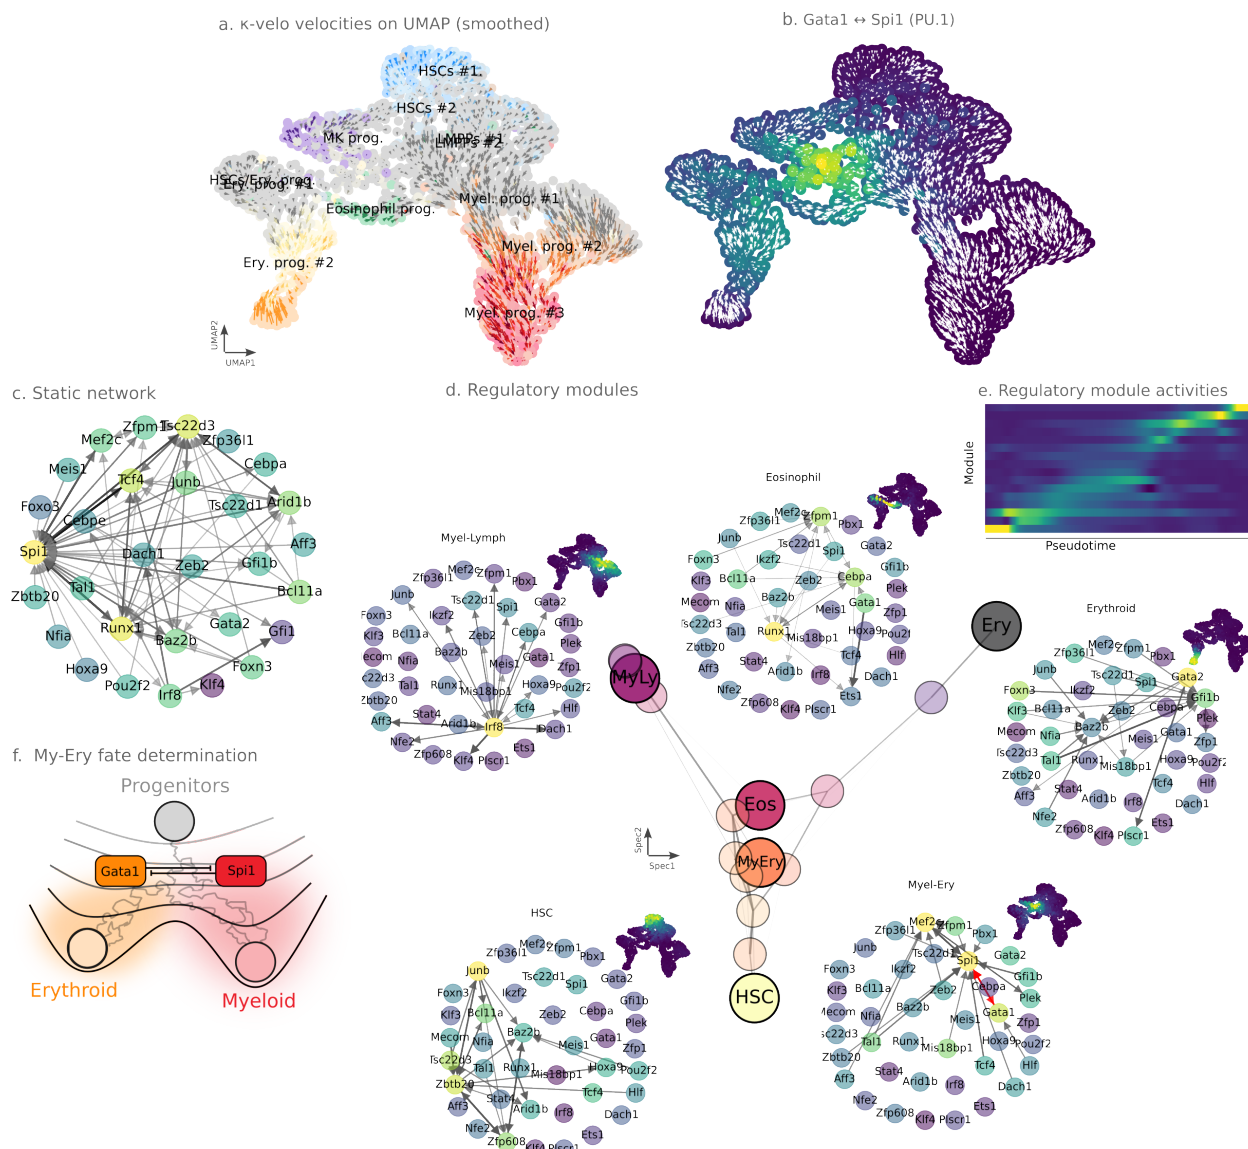


Figure 6: **Early-stage haematopoiesis.** (a) Smoothed velocities found by  $\kappa$ -velo shown on UMAP coordinates, colored by cluster annotation. (b) Average locaTE score for Gata1-Spi1 interactions shown on UMAP coordinates along with smoothed velocities. (c) Static network (top 1% of all edges) produced by averaging locaTE scGRNs along the first 90% of pseudotime. (d) Graph layout of regulatory modules found by locaTE-NMF colored by proximity (measured by mean first passage time) from the HSC module, shown alongside representative networks of selected modules (top 0.25% of all edges) corresponding to HSC, Myeloid-Erythroid decision boundary, Erythroid, Eosinophil and Myeloid-Lymphoid. (e) Smoothed regulatory module activities shown against pseudotime. (f) Myeloid-Erythroid fate determination schematic.

Spi1 (PU.1) interaction (edge highlighted in red). In Figure 6(d, e) we display other regulatory modules found corresponding to various degrees of fate commitment, including haematopoietic stem cells (HSCs). These networks uncover known hallmarks of haematopoietic development: the myeloid-lymphoid progenitor (labelled Myel-Lymph) module prominently features Irf8, understood to play a key role in the development of myeloid and lymphoid lineages [71, 72] (note that the lymphoid lineage is not captured in this dataset). The erythroid module is centered on a different set of genes such as Gfi1b [73]. The HSC module features Junb prominently, which has been implicated in homeostasis in the long-term HSC population [74]. The HSC module contains interactions involving Baz2b, a master regulator which has been found to reprogram lineage-committed cells to a multipotent state [75]. In the static network produced by aggregating locaTE scGRNs (see Figure 6(c)) we find that interactions such as Gata1-Spi1 are obscured due to the population averaging. Cell-specific methods such as locaTE that are required to learn interpretable representations of regulatory states in an unsupervised fashion without the need for clustering.

## 4 Conclusion

Biology is dynamic and a static representation of a network can never do justice to what we observe phenotypically. To discern molecular causes requires experimental and theoretical advances to progress in lock-step. The power of single cell data, and the insights we can gain from such data, crucially rely on flexible statistical analyses that allow us to dissect the differences between gene regulatory networks (and programs) and how they differ between cells.

Gaining deeper mechanistic insights into single-cell data requires a framework to reconstruct cell-specific, causal networks. Our method, locaTE, relies on computing the transfer entropy (TE) [25, 30] for pairs of genes to measure causal relationships. The approach uses approximations of the underlying dynamics but makes no extraneous assumptions on linearity of interactions, the distribution of gene expression counts, or the topology of the trajectory.

By using neighbourhoods in gene expression space we can infer cell-specific networks without going to the extreme  $N = 1$  case. How to pool information across cells most efficiently will ultimately decide on the success of such methods and the insights they provide. Temporal ordering helps enormously, but we are still, by and large, not able to collect meaningful single cell transcriptomic time-course data. Pseudo-temporal ordering or some of the many versions of RNA velocity entering the literature can provide useful information, as can, related multi-omic data, such as scATAC-seq.

Directions for future work include investigating approaches using which local information can help dealing with technical zeros (dropout) in single cell data. We investigated the application of data-imputation methods for the pancreatic dataset; however, we found that imputation worsened performance for all inference methods considered. This effect has been reported previously to be due to spurious signals being amplified [76] between non-interacting genes.

The problem of inferring cellular dynamics, upon which application of methods such as locaTE crucially depend, is far from solved. In particular, accurately modelling and quantifying RNA velocity remains challenging in practice [66, 38]. Metabolic labelling assays may also provide a more accurate source of dynamical information [24]. Lastly, a further route for improving locaTE and network inference in general is to use relevant prior information to guide network inference.

## Acknowledgements

SYZ would like to thank Geoffrey Schiebinger, Léo Diaz and Wilder Scott for insightful discussions. The first author acknowledges support from the Australian Government Research Training Program and an Elizabeth and Vernon Puzey Scholarship. MPHS gratefully acknowledges funding from the Australian Research Council through a Laureate Fellowship (FL220100005), and the University of Melbourne through DRM funding.



## A Method

### A.1 Code availability

An open-source Julia implementation of locaTE is available at <https://github.com/zsteve/locaTE.jl>.

### A.2 Constructing a backward operator

It is straightforward to write:

$$\begin{aligned} Q_{ij} &= \mathbf{P}[X_{-t} = x_j | X_0 = x_i] = \frac{\mathbf{P}[X_0 = x_i | X_{-t} = x_j] \mathbf{P}[X_{-t} = x_j]}{\mathbf{P}[X_0 = x_i]} \\ &= \frac{(P^t)_{ji} \mathbf{P}[X_{-t} = x_j]}{\mathbf{P}[X_0 = x_i]} \end{aligned}$$

If  $P$  encoded a reversible Markov chain and we chose  $\mathbf{P}[X_{-t} = \cdot]$  to be the stationary distribution, this would give us time-reversal at equilibrium. Since we are interested in the behaviour of the process away from equilibrium, we must prescribe  $\mathbf{P}[X_{-t} = \cdot]$  away from equilibrium. In practice, we find that taking  $\mathbf{P}[X_{-t} = \cdot] = \text{Unif}$  works well – this is equivalent to simply taking the transpose of  $P$  and rescaling, as done in [23]. We refer the reader to the analysis presented in Section 4.3 of [38] for a more detailed discussion on the transpose as a backward operator and its interpretation in the continuous limit.

### A.3 (Smoothness, sparsity)-regularised regression

The problem (9) is convex and a straightforward numerical algorithm for its solution can be derived using the alternating direction method of multipliers (ADMM) [77]. With auxiliary variables  $W, Z$  with the same dimensions as  $X$ , the corresponding ADMM scheme is

$$X^{(k+1)} = (\alpha + \lambda_1 L + \rho I)^{-1} (\alpha \hat{G} + \rho(Z^{(k)} - W^{(k)})) \quad (11)$$

$$Z^{(k+1)} = \arg \min_Z \lambda_2 \|Z\|_1 + \frac{\rho}{2} \|X^{(k+1)} - Z^{(k)} + W^{(k)}\|_2^2 = \text{prox}_{\rho^{-1}\lambda_2\|\cdot\|_1}(X^{(k+1)} + W^{(k)}) \quad (12)$$

$$W^{(k+1)} = W^{(k)} + X^{(k+1)} - Z^{(k+1)} \quad (13)$$

where  $\rho > 0$  is the ADMM relaxation parameter (we take this to be 0.05),  $\alpha = \text{diag}(w)$  and

$$\text{prox}_{\lambda\|\cdot\|_1}(x) = \text{sgn}(x)(|x| - \lambda)_+,$$

interpreted elementwise.

### A.4 Regularised non-negative matrix factorisation

Now we consider the same problem setting as (9), but this time we also require that the  $N \times M^2$  matrix of recovered interactions  $G$  be low rank. For a target rank  $k$ , we seek “tall” matrices  $U \in \mathbb{R}^{N \times k}$ ,  $V \in \mathbb{R}^{M^2 \times k}$  such that  $G \approx UV^\top$ . That is, for a cell  $i$ , its scGRN is given by a linear combination of the columns of  $V$  with coefficients  $U_{i1}, \dots, U_{ik}$ :

$$G_i = (UV^\top)_i = \sum_{\ell=1}^k U_{i\ell} V_{\cdot\ell}.$$

We consider a general loss function of the form

$$\begin{aligned} \mathcal{L}(U, V; L, G) &= \frac{1}{2} \|UV^\top - G\|_2^2 + \frac{\alpha}{2} \text{tr}(VU^\top LUV^\top) - \beta \langle H, UV^\top \rangle \\ &+ \frac{\lambda_1}{2} \text{tr}(U^\top K_1 U) + \mu_1 \|U\|_1 \\ &+ \frac{\lambda_2}{2} \text{tr}(V^\top K_2 V) + \mu_2 \|V\|_1. \end{aligned}$$

We explain the loss function term by term. The first term encourages a fit to the matrix  $G$  of raw TE values. We impose a penalty on the smoothness of the low-rank reconstruction (w.r.t. the graph Laplacian  $L$ ) with weight  $\alpha$ . We add an (optional) affine term on the reconstruction, where the entries of the matrix  $H$  can be thought of as a prior on which interactions may occur. Positive (negative) entries of  $H$  would correspond to prior on edge existence (non-existence). Priors could be constructed from gene-gene correlations, prior knowledge, or additional assay data (e.g. scATAC or ChIP data). In addition, we can impose Tikhonov and sparsity penalties on the coefficients and atoms.

One can derive multiplicative update equations for the gradient descent, following the approach of [78]. For positivity-preserving updates, we decompose any Laplacian matrices into their positive and negative parts:

$$\begin{aligned} L &= D - A, \\ K_1 &= D_1 - A_1, \\ K_2 &= D_2 - A_2. \end{aligned}$$

$$\begin{aligned} U &\leftarrow \frac{U \odot (GV + \alpha AUV^\top V + \lambda_1 A_1 U + \beta HV)}{UV^\top V + \alpha DUV^\top V + \lambda_1 D_1 U + \mu_1} \\ V &\leftarrow \frac{V \odot (G^\top U + \alpha VU^\top AU + \lambda_2 A_2 V + \beta H^\top U)}{VU^\top U + \alpha VU^\top DU + \lambda_2 D_2 V + \mu_2}. \end{aligned}$$

## A.5 Regularised non-negative tensor factorisation

Let  $G$  have dimensions  $N \times M \times M$ , but for notational convenience we do our derivations in a general setting. Write for convenience  $X = S \times_{i=1}^d A^{(i)}$ , where the multilinear product is defined as

$$(A \times_k B)_{i_1, \dots, i_{k-1}, j, i_{k+1}, \dots, i_d} = \sum_{i_k=1}^{n_k} A_{i_1, \dots, i_k, \dots, i_d} B_{j, i_k}.$$

In order to derive a multiplicative update scheme, all matrices need to be decomposed into their positive and negative parts. We write  $X_{(k)}$  for the matricization of  $X$  along its  $k$ -th mode. We want to derive multiplicative update scheme for the problem

$$\begin{aligned} \mathcal{L}(S, \{A^{(i)}\}_{i=1}^3) &= \frac{1}{2} \|X - G\|_2^2 + \frac{\alpha}{2} \text{tr}(X_{(1)}^\top LX_{(1)}) - \beta \langle H, X \rangle \\ &\quad + \sum_{i=1}^3 \frac{\lambda_i}{2} \text{tr}((A^{(i)})^\top L^{(i)} A^{(i)}) \\ &\quad + \sum_{i=1}^3 \mu_i \|A^{(i)}\|_1 \end{aligned}$$

The final two terms are straightforward to differentiate. For the first term, note that

$$\partial_{A^{(k)}} \frac{1}{2} \|X - G\|_2^2 = \partial_{A^{(k)}} \frac{1}{2} \|X_{(k)} - G_{(k)}\|_2^2$$

Now employ the identity  $X_{(k)} = [S \times_{i=1}^d A^{(i)}]_{(k)} = A^{(k)} [S \times_{i \neq k} A^{(i)}]_{(k)}$ . Let  $B^{(k)} = [S \times_{i \neq k} A^{(i)}]_{(k)}$  for convenience, then one has

$$\partial_{A^{(k)}} \frac{1}{2} \|X - G\|_2^2 = A^{(k)} B^{(k)} (B^{(k)})^\top - G_{(k)} (B^{(k)})^\top.$$

The second term is more involved. First consider the case  $k = 1$ , so  $X_{(1)} = A^{(1)} B^{(1)}$ .

$$\begin{aligned} \partial_{A^{(1)}} \frac{1}{2} \langle X_{(1)}, LX_{(1)} \rangle &= \partial_{A^{(1)}} \frac{1}{2} \text{tr} \left( (B^{(1)})^\top (A^{(1)})^\top L A^{(1)} B^{(1)} \right) \\ &= L A^{(1)} B^{(1)} (B^{(1)})^\top. \end{aligned}$$

For  $k \neq 1$ , note that if we “re-fold”  $LX_{(1)}$  along mode 1, we get

$$S \times_1 (LA^{(1)}) \times_{j>1} A^{(j)}.$$

So,

$$\begin{aligned} \frac{1}{2} \langle X_{(1)}, LX_{(1)} \rangle &= \frac{1}{2} \langle S \times_{i=1}^d A^{(i)}, S \times_1 (LA^{(1)}) \times_{j>1} A^{(j)} \rangle \\ &= \frac{1}{2} \langle A^{(k)} B^{(k)}, A^{(k)} \left[ S \times_1 (LA^{(1)}) \times_{j>1, j \neq k} A^{(j)} \right]_{(k)} \rangle \\ &= \frac{1}{2} \langle A^{(k)} B^{(k)}, A^{(k)} \tilde{B}^{(k)} \rangle, \end{aligned}$$

where  $\tilde{B}^{(k)} = [S \times_1 (LA^{(1)}) \times_{j>1, j \neq k} A^{(j)}]_{(k)}$ . Then,

$$\partial_{A^{(k)}} \frac{1}{2} \langle A^{(k)} B^{(k)}, A^{(k)} \tilde{B}^{(k)} \rangle = \frac{1}{2} A^{(k)} \left( \tilde{B}^{(k)} (B^{(k)})^\top + B^{(k)} (\tilde{B}^{(k)})^\top \right).$$

For the third term,

$$\partial_{A^{(k)}} \langle H, X \rangle = H_{(k)} (B^{(k)})^\top.$$

Combining the gradient expressions with the decomposition of Laplacians into positive and negative parts as done for NMF, we arrive at the update scheme

$$\begin{aligned} A^{(1)} &\leftarrow \frac{A^{(1)} \odot (G_{(1)} (B^{(1)})^\top + \lambda_1 W_1 A^{(1)} + \alpha W A^{(1)} B^{(1)} (B^{(1)})^\top + \beta H_{(1)} (B^{(1)})^\top)}{A^{(1)} B^{(1)} (B^{(1)})^\top + \lambda_1 D_1 A^{(1)} + \alpha D A^{(1)} B^{(1)} (B^{(1)})^\top + \mu_1} \\ A^{(i)} &\leftarrow \frac{A^{(i)} \odot (G_{(i)} (B^{(i)})^\top + \lambda_i W_i A^{(i)} + \alpha A^{(i)} C_- + \beta H_{(i)} (B^{(i)})^\top)}{A^{(i)} B^{(i)} (B^{(i)})^\top + \lambda_i D_i A^{(i)} + \alpha A^{(i)} C_+ + \mu_i}, \quad i > 1, \end{aligned}$$

where for each  $i > 1$  we define for convenience

$$\begin{aligned} C_+ &= (\tilde{B}_+^{(i)} (B^{(i)})^\top + B^{(i)} \tilde{B}_+^{(i)})/2 \\ C_- &= (\tilde{B}_-^{(i)} (B^{(i)})^\top + B^{(i)} \tilde{B}_-^{(i)})/2. \end{aligned}$$

## B Results details

### B.1 Simulated data

For each boolean network considered, simulated datasets were generated using the BoolODE package [5], modified to record additionally the velocity vector of simulated cells as well as the Jacobian matrix of expression state. We provide below the boolean rules that we used for simulating the two examples in the main text; all other simulated trajectories were produced using rules provided as part of the BoolODE package.

#### Boolean rules: bifurcating

$$\begin{aligned} g1 &\leftarrow \neg(g4 \vee g5) \\ g2 &\leftarrow g1 \\ g3 &\leftarrow g2 \\ g4 &\leftarrow ((g3 \vee g4) \wedge (\neg g5)) \\ g5 &\leftarrow ((g5 \vee g3) \wedge (\neg g4)) \\ g6 &\leftarrow ((g4 \wedge (\neg(g5 \vee g9))) \vee ((g7 \vee g6) \wedge (g5 \wedge (\neg g4)))) \\ g7 &\leftarrow ((g6 \wedge (g4 \wedge (\neg g5))) \vee (g8 \wedge (g5 \wedge (\neg g4)))) \\ g8 &\leftarrow ((g7 \wedge (g4 \wedge (\neg g5))) \vee (g9 \wedge (g5 \wedge (\neg g4)))) \\ g9 &\leftarrow (((g8 \vee g9) \wedge (g4 \wedge (\neg g5))) \vee (g5 \wedge (\neg(g4 \vee g6)))) \end{aligned}$$

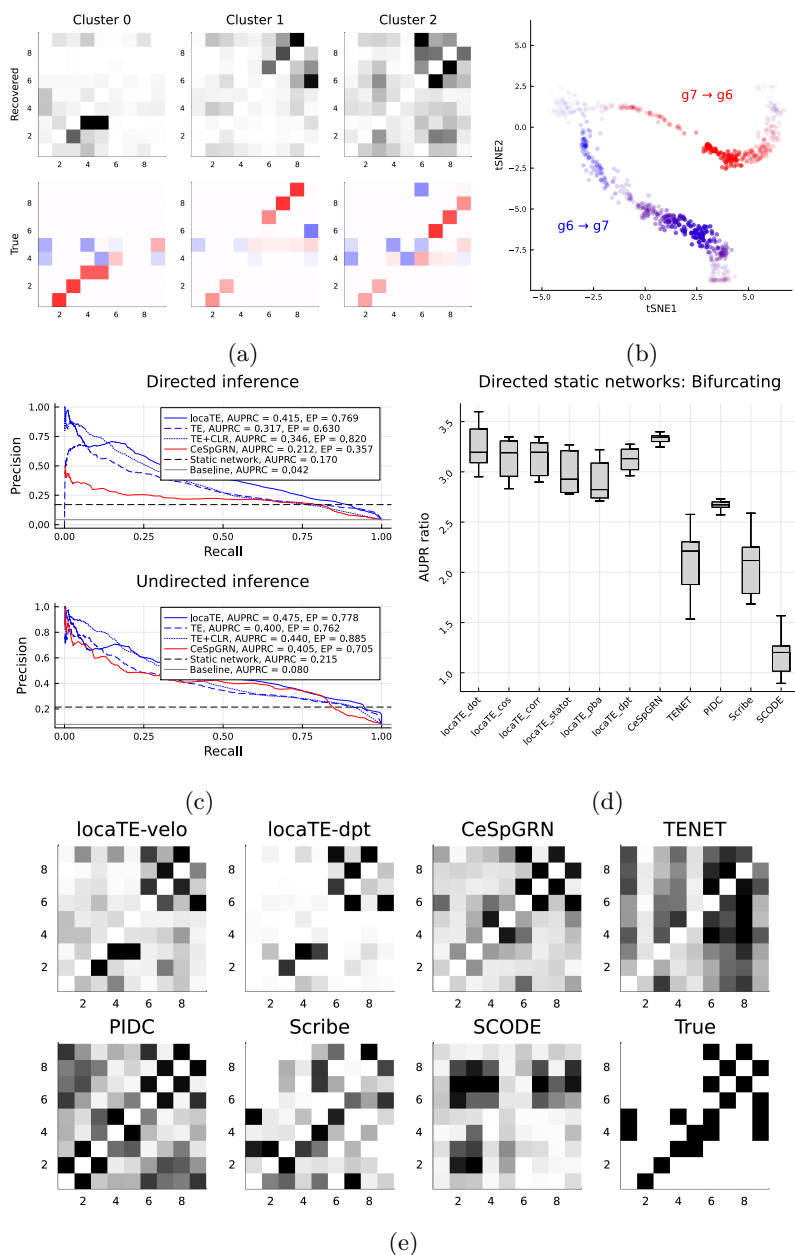


Figure S7: **Supplementary results for simulated bifurcating network.** (a) Interactions inferred by locaTE averaged by cluster (top) shown against ground truth interactions averaged by cluster. (b) Interaction scores for  $g7 \rightarrow g6$  and  $g6 \rightarrow g7$  shown for each cell on the tSNE embedding. (c) Precision-Recall (PR) curves for cell-specific edge detection from  $\hat{G}$  (raw TE values),  $\tilde{G}$  (TE + CLR values), or  $G$  (locaTE). Area under PR curve (AUPRC) and early precision at 10% recall (EP) is shown. (d) AUPRC ratio for locaTE and other methods across 10 sampled simulation datasets, in the static setting. (e) Best inferred static networks.

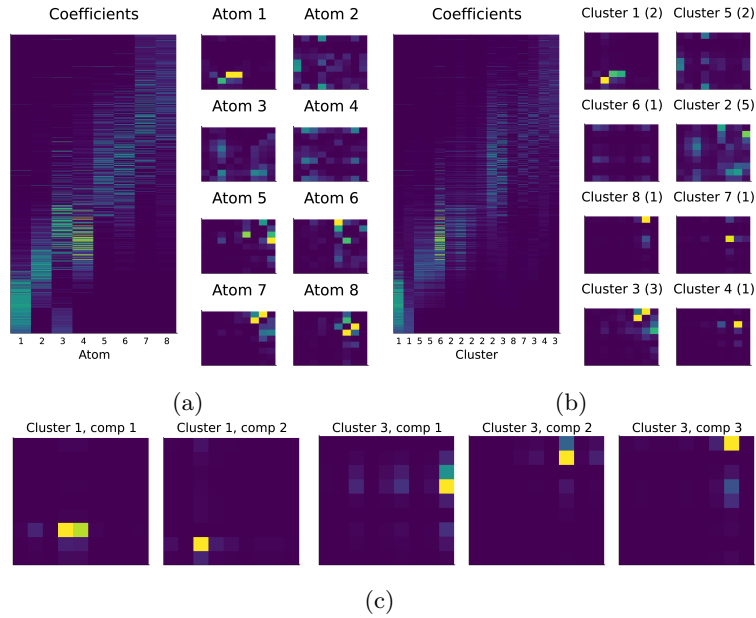


Figure S8: **Factor analysis extracts groups of differentially regulated genes.** Factor analysis results for scGRNs using (a) locaTE-NMF with  $k = 8$  factors and (b) locaTE-NTF with  $k = 16$  factors; the number of rank-one factors contributing to each cluster for NTF are shown in parentheses. (c) Rank-one atoms found by locaTE-NTF corresponding to clusters 1 and 3.

### Boolean rules: switch

$$\begin{aligned}
 g1 &\leftarrow (\neg g6) \\
 g2 &\leftarrow ((g1 \wedge (\neg(g6 \wedge (\neg g5)))) \vee (g7 \wedge ((g6 \wedge (\neg g5)))))) \\
 g3 &\leftarrow ((g2 \wedge (\neg(g6 \wedge (\neg g5)))) \vee ((g4 \vee g3) \wedge ((g6 \wedge (\neg g5)))))) \\
 g4 &\leftarrow ((g3 \wedge (\neg(g6 \wedge (\neg g5)))) \vee (g2 \wedge ((g6 \wedge (\neg g5)))))) \\
 g5 &\leftarrow ((g4 \wedge (\neg(g6 \wedge (\neg g5)))))) \\
 g6 &\leftarrow ((g5 \vee g6)) \\
 g7 &\leftarrow (g5 \vee (g7 \wedge (g6 \wedge (\neg g5)) \wedge (\neg g2)))
 \end{aligned}$$

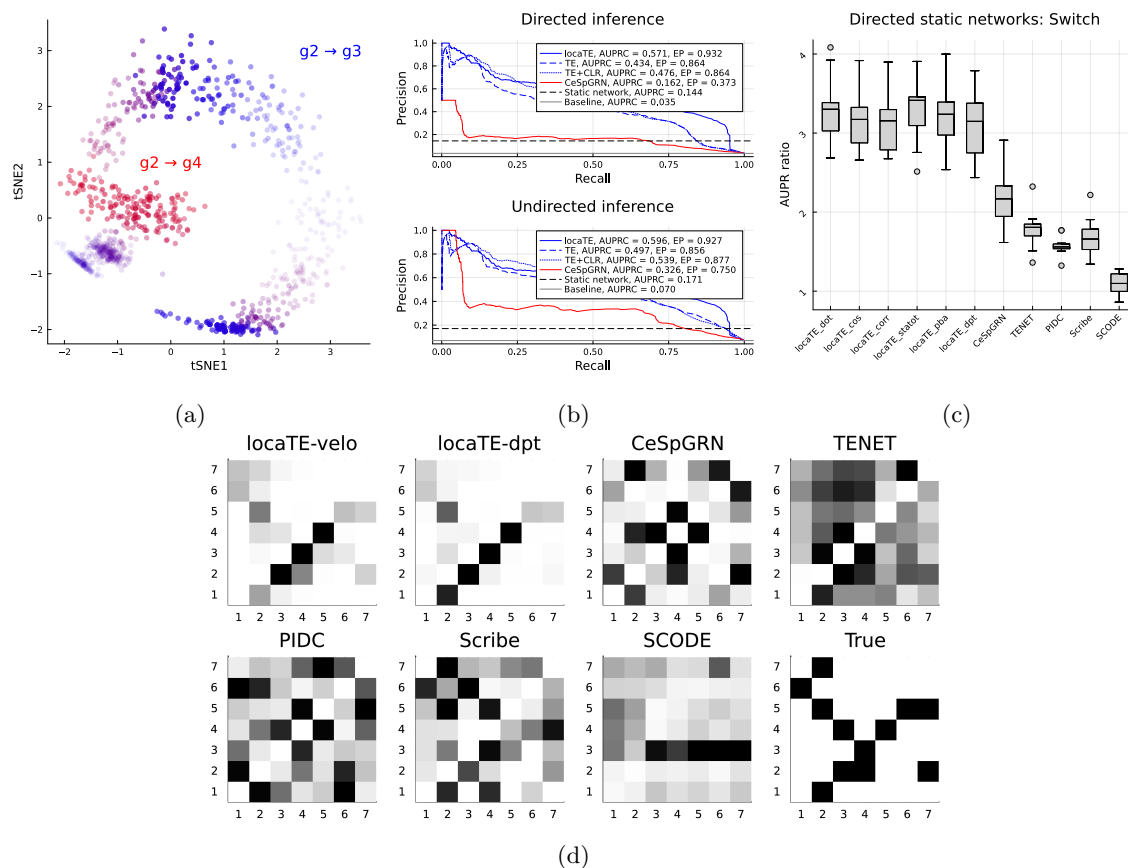


Figure S9: **Supplementary results for simulated switch network.** (a) Interaction scores for  $g2 \rightarrow g3$  and  $g2 \rightarrow g4$  for each cell on the tSNE embedding. (b) Precision-Recall (PR) curves for cell-specific edge detection from  $\hat{G}$  (raw TE values),  $\tilde{G}$  (TE + CLR values), or  $G$  (locaTE). Area under PR curve (AUPRC) and early precision at 10% recall (EP) is shown. (c) AUPRC ratio for locaTE and other methods across 10 sampled simulation datasets, in the static setting. (d) Best inferred static networks.

Dimensionality reduction was performed using PCA and a  $k$ -NN constructed with  $k = 25$ . Diffusion pseudotime was then assigned with the root cell taken to be the simulated cell with the earliest simulation time. Various transition kernels  $P$  were then constructed as described below.

**Velocity kernel** As described in the main text, the velocity kernel was constructed from evaluations of the drift term of the chemical Langevin SDE using the `cellrank.tl.kernels.{ DotProductScheme, CorrelationScheme, CosineScheme }` classes within CellRank [23], with the previously computed  $k$ -NN graph and bandwidth  $\sigma$  estimated using the median heuristic.

**Diffusion pseudotime kernel** The diffusion pseudotime (DPT) kernel was constructed from the previously estimated diffusion pseudotime ordering using the `cellrank.tl.kernels.PseudotimeKernel` class within CellRank.

**Optimal transport kernel** Cells in the top 92.5%-97.5% of pseudotime were assigned a negative flux rate  $R_i$  so that the net negative flux was  $-50$ . Remaining cells were assigned a positive flux rate to satisfy the zero net flux requirement,  $\sum_i R_i = 0$ . StationaryOT [34] with quadratically regularised optimal transport was applied with  $\Delta t = 1.0, \varepsilon = 0.05$ . The cost matrix was taken to be the matrix of pairwise graph distances constructed from the  $k$ -NN graph. To prevent short-circuiting resulting from spurious edges, edges were ranked by their random-walk betweenness centrality and the top 5% of edges were pruned.

**Population balance analysis (PBA) kernel** Population balance analysis [21] was applied using the previously constructed  $k$ -NN graph, the same flux rates  $R_i$  as used in the optimal transport kernel, and diffusivity  $D = 1.0$ .

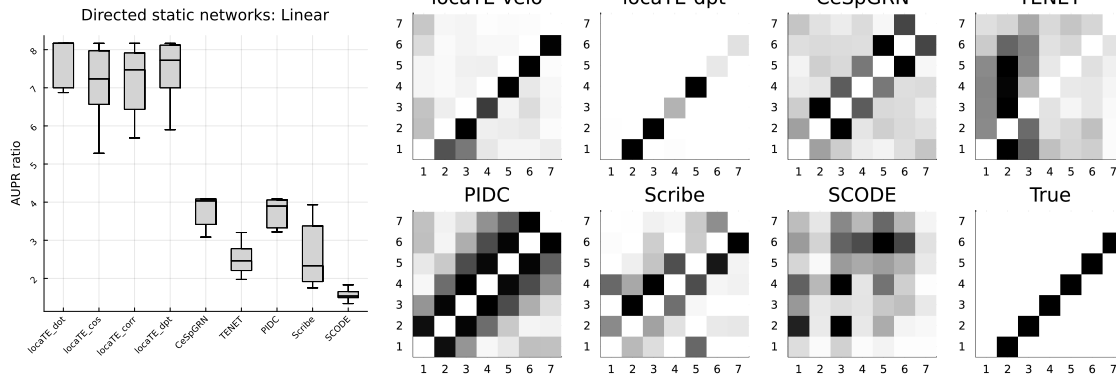
**Inferring scGRNs** LocaTE was applied using the simulated expression data and inferred transition kernels  $P$ . Expression values were discretised using the Bayesian blocks algorithm as implemented in [4]. The chemical Langevin simulation is a continuous approximation to biophysically realistic discrete expression dynamics, and as such produces expression values near zero. Expression values  $< 10^{-0.5}$  were set to zero to prevent many bins from being created near zero. A neighbourhood kernel was constructed using quadratically regularised optimal transport, and the backward kernel was constructed from  $P$  as described in Section A.2. A grid-search was performed over  $\tau \in \{1, \dots, 5\}$ ,  $\lambda_1 \in \{1, 2.5, 5, 10, 25\}$ ,  $\lambda_2 \in \{0, 0.001, 0.005, 0.01, 0.025\}$ . For comparison, CeSpGRN [13] was applied with a grid search performed over  $|D(i)| \in \{5, 25, 50, 100, 250\}$ ,  $\sigma \in \{0.05, 0.1, 0.25, 0.5, 1.0\}$ , and  $\lambda \in \{0.005, 0.01, 0.025, 0.05, 0.1\}$ .

**Evaluating scGRNs** For each simulated dataset, ground truth scGRN networks were constructed by recording the Jacobian for each sampled cell state. Inferred scGRNs were evaluated against this ground truth by constructing Precision-Recall curves using the approach described in the main text (Section 3.1). Static GRNs were constructed by averaging scGRNs over the cell population.

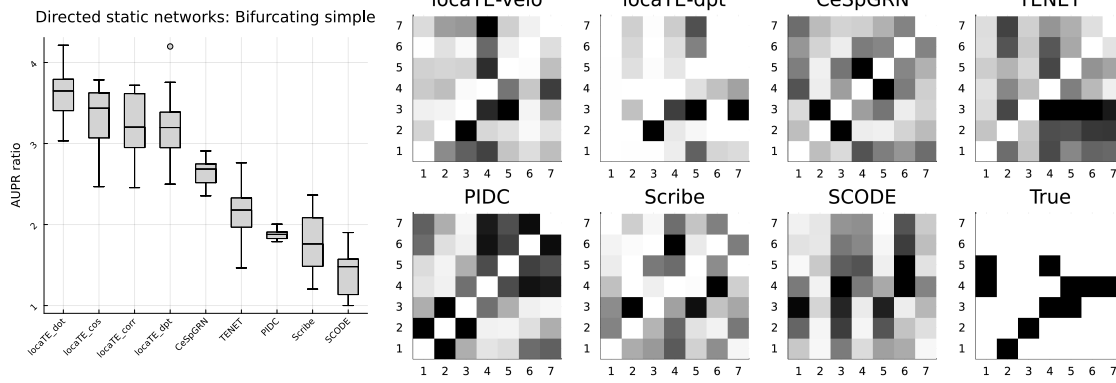
For comparison, PIDC [4], TENET [17], Scribe [16] and SCODE [18] were also applied to infer static networks. For TENET, we report the best results as measured by AUPRC ratio for a range of history lengths  $\in \{1, \dots, 8\}$ .

## B.2 Results: simple trajectories

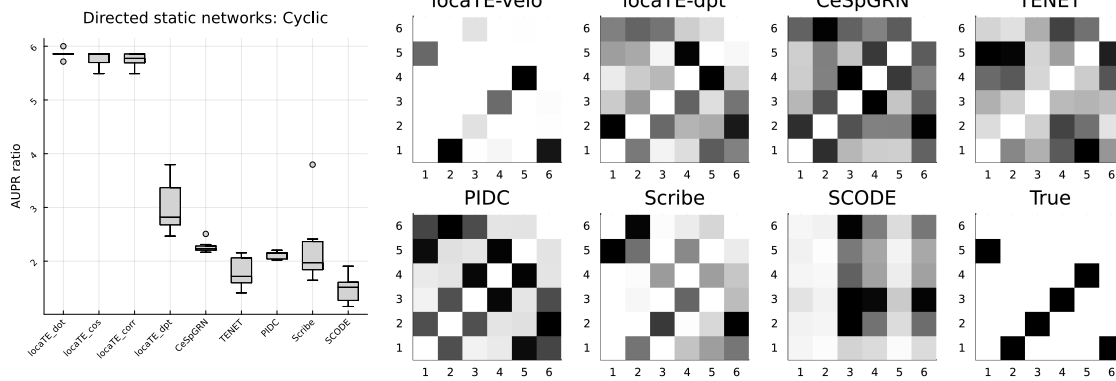
We applied locaTE with several choices of kernel to the simulated trajectory examples from [5]. For comparison, we also applied CeSpGRN, TENET, PIDC, Scribe and SCODE. A grid search was performed for locaTE, CeSpGRN and TENET, and default parameters were used for the remaining methods. Performance for each method was assessed in terms of the AUPRC ratio for directed edge inference. We summarise performance with boxplots below, and show representative examples of the recovered static networks.



(a) Linear network, 1000 cells.




(b) Simple bifurcating network, 1000 cells.

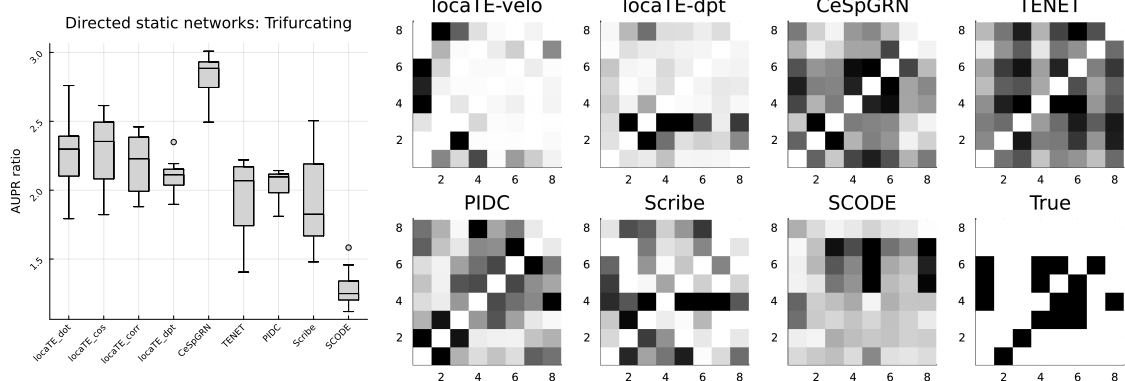


(c) Cycling network, 1000 cells.

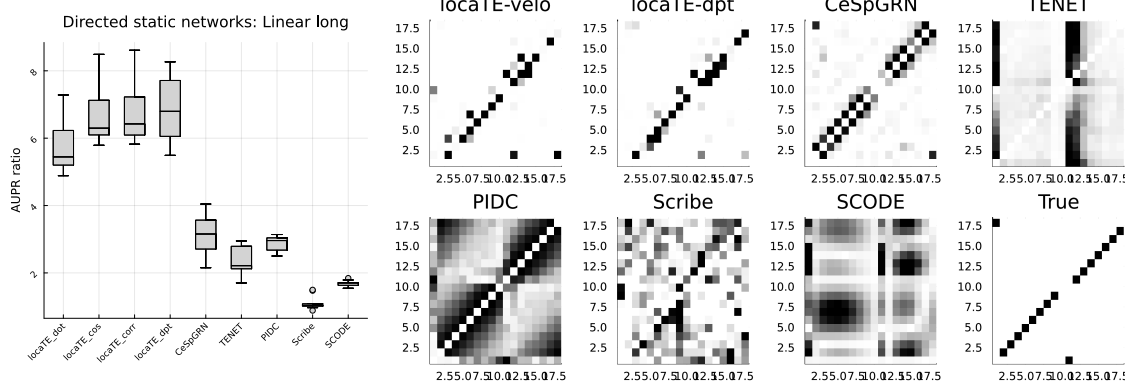




(a) Bifurcating cycling network, 1000 cells.



(b) Trifurcating network, 1000 cells.



(c) Linear long network, 1000 cells.

### B.3 mESC dataset

Log-transformed expression values for 100 TFs and pseudotime estimates were fetched from the SCODE example “data1” [18]. Cells were embedded into the top 25 principal components and a  $k$ -NN graph was constructed with  $k = 25$ . Using this  $k$ -NN graph, a cost matrix  $C$  of squared pairwise shortest-path distances was calculated using the Floyd-Warshall algorithm. StationaryOT [34] was then applied using this cost matrix, treating cells in the final 10% of pseudotime as sink cells. The flux rate  $R_i$  for sink cells was taken to be uniform such that  $\sum_{i \in \text{sinks}} R_i = -50$ , and similarly for non-sink cells so that  $\sum_{i \notin \text{sinks}} R_i = 50$ . Quadratically regularised optimal transport was applied with  $\varepsilon = \bar{C}$ , i.e. the mean value of  $C$ . The resulting transition matrix was used as the forward transition matrix for locaTE. The ESCAPE reference network of [19] was retrieved from <https://github.com/gitter-lab/SINGE-supplemental>.

LocaTE was applied to estimate TE using the transition kernel derived from StationaryOT with  $\tau = 1$ . Expression values were discretised using the Bayesian blocks algorithm as implemented in [4], with artifactually small values ( $< 10^{-2}$ ) set to zero in order to prevent many bins being created near zero. Raw TE values were first filtered using weighted CLR, and denoising regression was applied with  $\lambda_1 = 10, \lambda_2 = 10^{-3}$  using the normalised  $k$ -NN graph Laplacian.

To apply locaTE-NMF, the array of raw TE values was scaled to have unit 0.9-quantile. The target rank was set to  $k = 8$ ,  $L$  was taken to be the normalised  $k$ -NN graph Laplacian, and  $K_1 = K_2 = I$ . We chose  $\alpha = 10, \beta = 0, \lambda_1 = \lambda_2 = \mu_1 = \mu_2 = 1$ .

For comparison, PIDC [4], GRISLI [20], TENET [17], SCODE [18] and CeSpGRN [13] were applied using the same log-transformed expression values and pseudotime estimates. For TENET a history length of 1 was used, following the usage guidelines in its documentation. We found that TENET tended to perform worse with longer history lengths. For SCODE, the rank parameter was set to  $D = 4$  which is the recommended value. For CeSpGRN, we took a neighbourhood size  $|N(i)| = 25$ , bandwidth  $\sigma = 5$ , and sparsity level  $\lambda = 0.1$  in keeping with the suggested parameter ranges in [13, Section 3.1]. All other parameters were taken to be their default values.

## B.4 Pancreatic dataset

Expression data was processed and RNA velocity was estimated following the CellRank tutorial available at [https://cellrank.readthedocs.io/en/stable/cellrank\\_basics.html](https://cellrank.readthedocs.io/en/stable/cellrank_basics.html). A transition kernel derived from RNA velocity was produced using `cr.tl.kernels.VelocityKernel` using the cosine similarity scheme. Pseudotime estimates for visualisation were calculated using the function `scvelo.tl.latent_time`. Intersecting the genes that passed filtering with the list of TFs from [20] yielded a subset of 82 TFs. LocaTE was applied using log-transformed expression values of the TF subset and the transition kernel derived from RNA velocity, with  $\tau = 1$ . All other parameters were chosen to be the same as for the mESC dataset.

# C Additional figures

## C.1 mESC dataset

Additional figures for the mESC dataset.

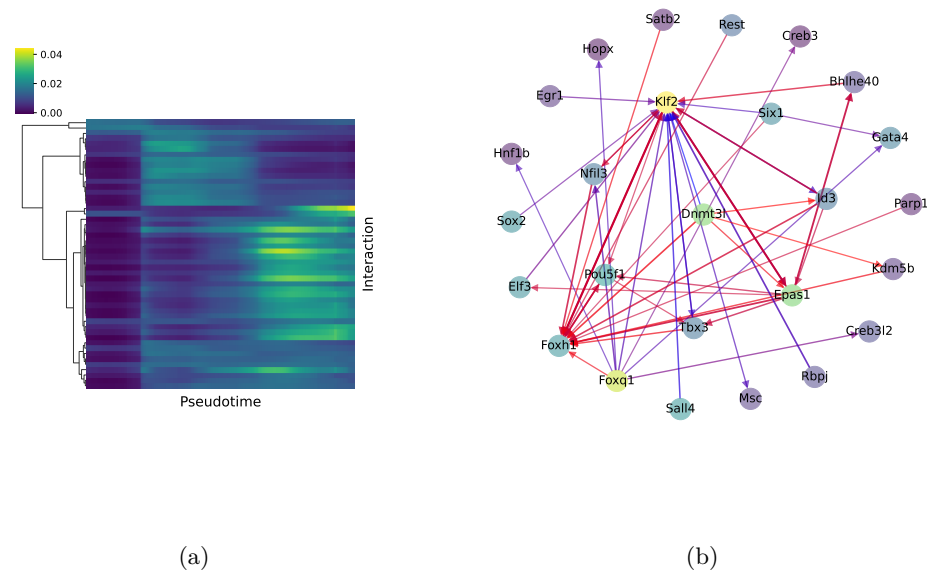


Figure S12: Smoothed plot of top 1% of inferred interactions (from static network) **(a)** clustered and shown against pseudotime **(b)** coloured by the time of peak activity, where blue corresponds to earliest, and red corresponds to latest.

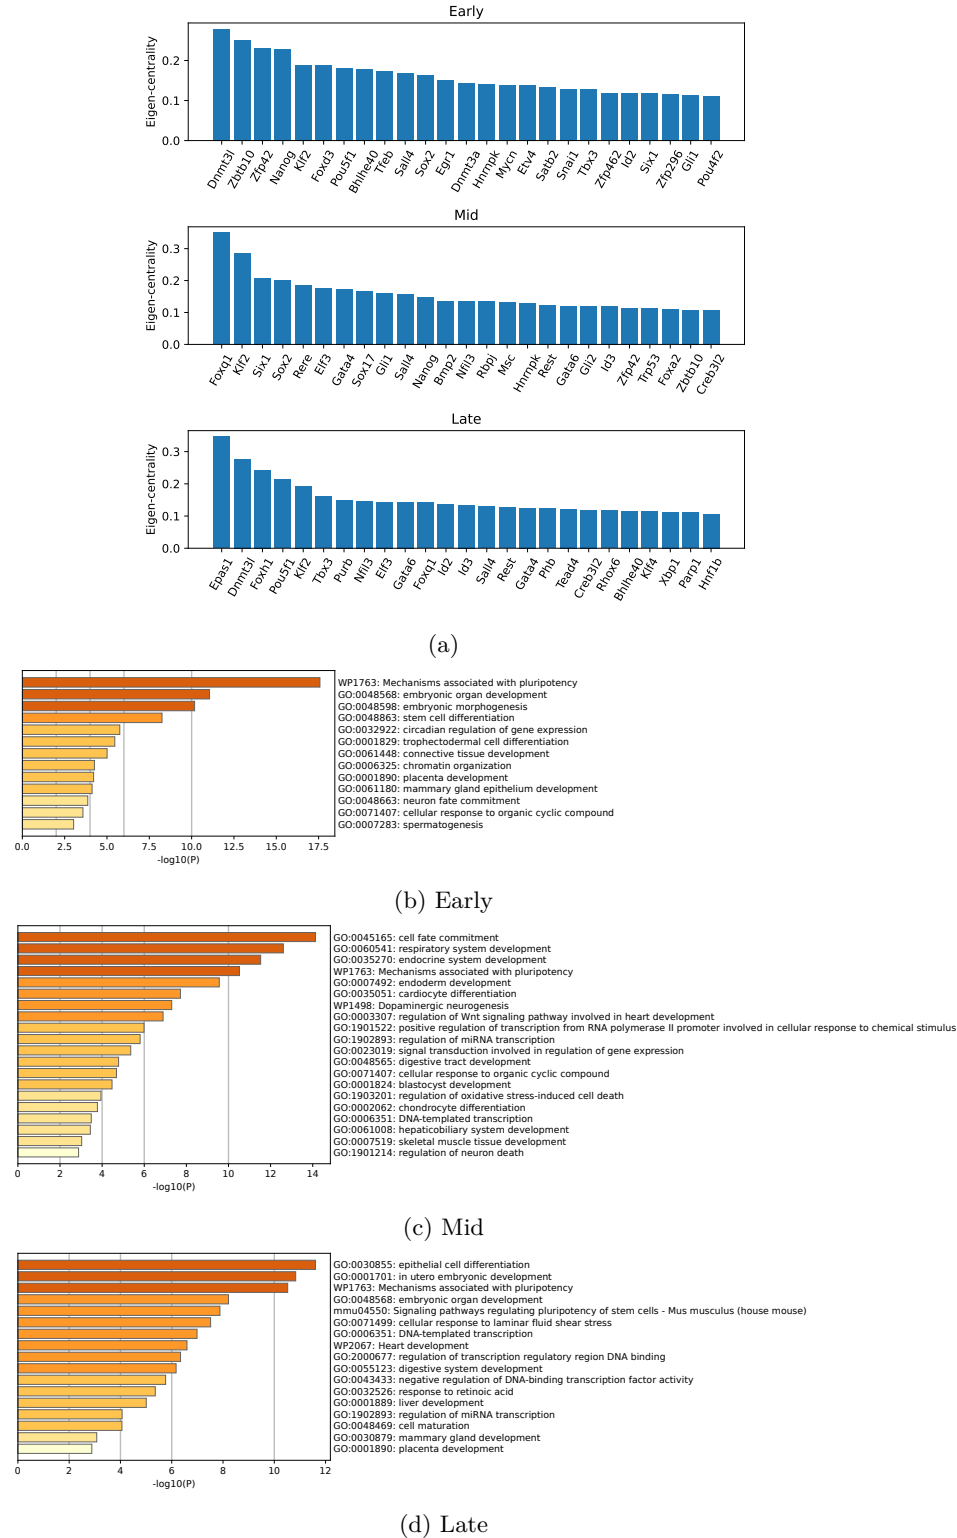


Figure S13: (a) Top 25 TFs for each stage by outgoing eigenvector centrality (b-d) GO terms found by Metascape for top 25 early/mid/late TFs.

## C.2 Pancreatic development dataset

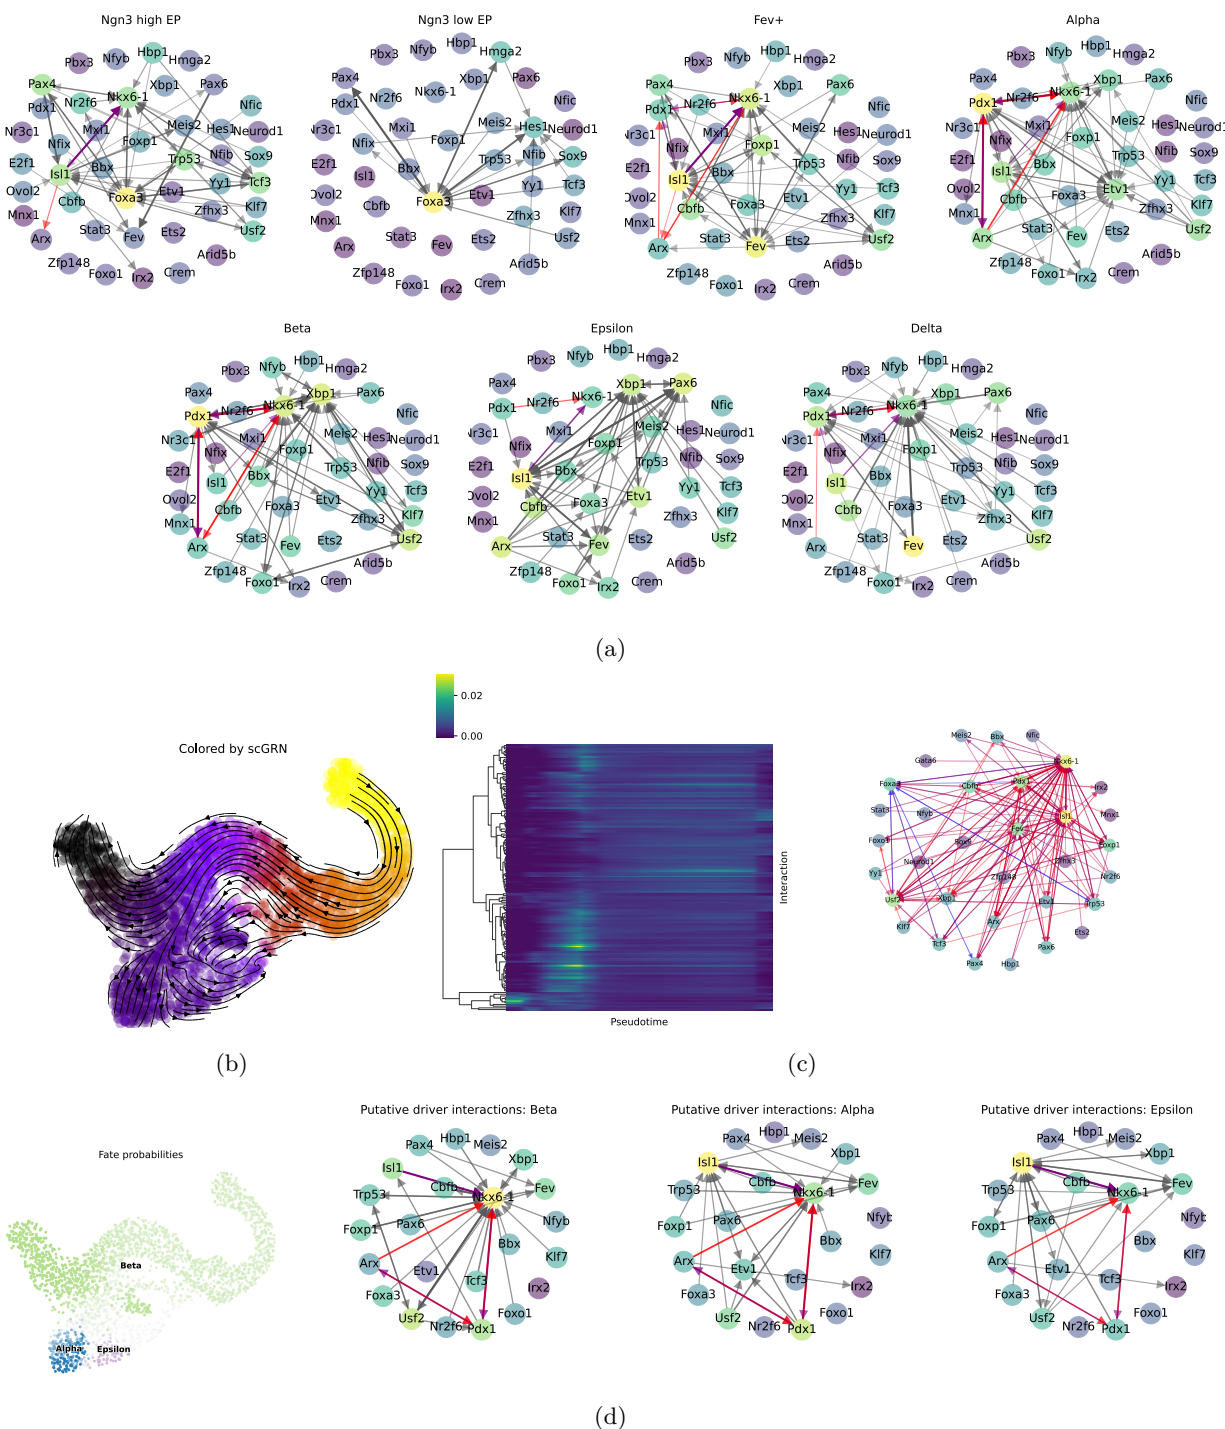





Figure S14: (a) Per-cluster networks (top 1% of interactions) (b) First diffusion component of cell-specific GRNs found by locaTE overlaid on UMAP coordinates (c) Smoothed plot of top 2.5% of interactions clustered and shown against pseudotime, and corresponding static network with edges coloured by time of peak activity (d) Putative driver networks for alpha, beta and epsilon fates.




(a)




(b) Ngn3- EP cluster



(c) Ngn3+ EP cluster



(d) Beta cluster



(e) Alpha cluster

Figure S15: (a) Top 25 TFs for each stage by outgoing eigenvector centrality (b-e) GO terms found by Metascape for top 25 TFs from Ngn3- EP, Ngn3+ EP, Beta, Alpha clusters.

### C.3 Haematopoiesis dataset

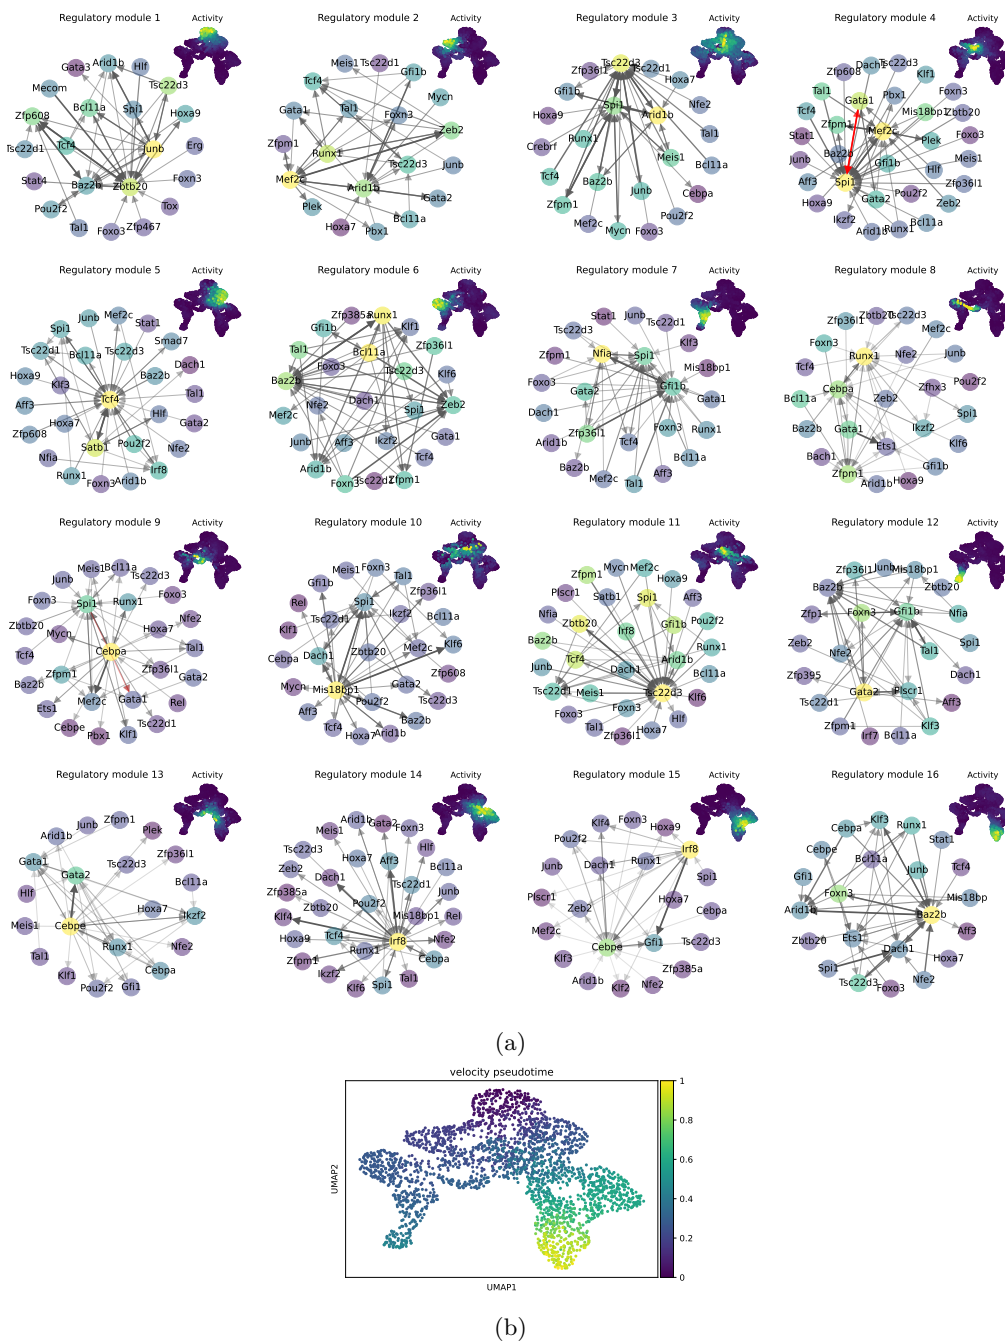





Figure S16: (a) Haematopoiesis regulatory modules (top 0.5% of edges) and module activities (inset) on UMAP coordinates. (b) Velocity pseudotime shown on UMAP coordinates highlights insufficiency of a 1D ordering.




(a)




(b) Myeloid-Erythroid module




(c) Myeloid-Lymphoid module



(d) Erythroid module



(e) Eosinophil module



(f) HSC module

Figure S17: (a) Top 25 TFs for each regulatory module ranked by out-edge eigenvector centrality. (b) Corresponding GO terms found by Metascape.



## References

- [1] Florian Buettner, Kedar N Natarajan, F Paolo Casale, Valentina Proserpio, Antonio Scialdone, Fabian J Theis, Sarah A Teichmann, John C Marioni, and Oliver Stegle. Computational analysis of cell-to-cell heterogeneity in single-cell rna-sequencing data reveals hidden subpopulations of cells. *Nature biotechnology*, 33(2):155–160, 2015.
- [2] Sophie Tritschler, Maren Büttner, David S Fischer, Marius Lange, Volker Bergen, Heiko Lickert, and Fabian J Theis. Concepts and limitations for learning developmental trajectories from single cell genomics. *Development*, 146(12):dev170506, 2019.
- [3] Etienne Becht, Leland McInnes, John Healy, Charles-Antoine Dutertre, Immanuel WH Kwok, Lai Guan Ng, Florent Ginhoux, and Evan W Newell. Dimensionality reduction for visualizing single-cell data using umap. *Nature biotechnology*, 37(1):38–44, 2019.
- [4] Thalia E Chan, Michael PH Stumpf, and Ann C Babbie. Gene regulatory network inference from single-cell data using multivariate information measures. *Cell systems*, 5(3):251–267, 2017.
- [5] Aditya Pratapa, Amogh P Jalihal, Jeffrey N Law, Aditya Bharadwaj, and TM Murali. Benchmarking algorithms for gene regulatory network inference from single-cell transcriptomic data. *Nature methods*, 17(2):147–154, 2020.
- [6] Bradley Efron and Trevor Hastie. *Computer Age Statistical Inference, Student Edition: Algorithms, Evidence, and Data Science*, volume 6. Cambridge University Press, 2021.
- [7] Michael PH Stumpf. Inferring better gene regulation networks from single-cell data. *Current Opinion in Systems Biology*, 27:100342, 2021.
- [8] Kyle Akers and TM Murali. Gene regulatory network inference in single-cell biology. *Current Opinion in Systems Biology*, 26:87–97, 2021.
- [9] Shuonan Chen and Jessica C Mar. Evaluating methods of inferring gene regulatory networks highlights their lack of performance for single cell gene expression data. *BMC bioinformatics*, 19(1):232–221, 2018.
- [10] Patrick S Stumpf, Rosanna CG Smith, Michael Lenz, Andreas Schuppert, Franz-Josef Müller, Ann Babbie, Thalia E Chan, Michael PH Stumpf, Colin P Please, Sam D Howison, et al. Stem cell differentiation as a non-markov stochastic process. *Cell Systems*, 5(3):268–282, 2017.
- [11] Elias Ventre, Ulysse Herbach, Thibault Espinasse, Gérard Benoit, and Olivier Gandrillon. One model fits all: combining inference and simulation of gene regulatory networks. *bioRxiv*, 2022.
- [12] Robrecht Cannoodt, Wouter Saelens, Louise Deconinck, and Yvan Saeys. Spearheading future omics analyses using dyngen, a multi-modal simulator of single cells. *Nature Communications*, 12(1):1–9, 2021.
- [13] Ziqi Zhang, Jongseok Han, Le Song, and Xiuwei Zhang. Inferring cell-specific gene regulatory networks from single cell gene expression data. *bioRxiv*, 2022.
- [14] Xuran Wang, David Choi, and Kathryn Roeder. Constructing local cell-specific networks from single-cell data. *Proceedings of the National Academy of Sciences*, 118(51), 2021.
- [15] Sophie Lèbre, Jennifer Becq, Frédéric Devaux, Michael P H Stumpf, and Gaëlle Lelandais. Statistical inference of the time-varying structure of gene-regulation networks. *BMC Systems Biology*, 4(1):130, 2010.
- [16] Xiaojie Qiu, Arman Rahimzamani, Li Wang, Bingcheng Ren, Qi Mao, Timothy Durham, José L McFaline-Figueroa, Lauren Saunders, Cole Trapnell, and Sreeram Kannan. Inferring causal gene regulatory networks from coupled single-cell expression dynamics using scribe. *Cell systems*, 10(3):265–274, 2020.

- [17] Junil Kim, Simon T. Jakobsen, Kedar N Natarajan, and Kyoung-Jae Won. Tenet: gene network reconstruction using transfer entropy reveals key regulatory factors from single cell transcriptomic data. *Nucleic acids research*, 49(1):e1–e1, 2021.
- [18] Hirotaka Matsumoto, Hisanori Kiryu, Chikara Furusawa, Minoru SH Ko, Shigeru BH Ko, Norio Gouda, Tetsutaro Hayashi, and Itoshi Nikaido. Scode: an efficient regulatory network inference algorithm from single-cell rna-seq during differentiation. *Bioinformatics*, 33(15):2314–2321, 2017.
- [19] Atul Deshpande, Li-Fang Chu, Ron Stewart, and Anthony Gitter. Network inference with granger causality ensembles on single-cell transcriptomics. *Cell reports*, 38(6):110333, 2022.
- [20] Pierre-Cyril Aubin-Frankowski and Jean-Philippe Vert. Gene regulation inference from single-cell rna-seq data with linear differential equations and velocity inference. *Bioinformatics*, 36(18):4774–4780, 2020.
- [21] Caleb Weinreb, Samuel Wolock, Betsabeh K Tusi, Merav Socolovsky, and Allon M Klein. Fundamental limits on dynamic inference from single-cell snapshots. *Proceedings of the National Academy of Sciences*, 115(10):E2467–E2476, 2018.
- [22] Geoffrey Schiebinger. Reconstructing developmental landscapes and trajectories from single-cell data. *Current Opinion in Systems Biology*, 27:100351, 2021.
- [23] Marius Lange, Volker Bergen, Michal Klein, Manu Setty, Bernhard Reuter, Mostafa Bakhti, Heiko Lickert, Meshal Ansari, Janine Schniering, Herbert B Schiller, et al. Cellrank for directed single-cell fate mapping. *Nature methods*, pages 1–12, 2022.
- [24] Xiaojie Qiu, Yan Zhang, Jorge D Martin-Rufino, Chen Weng, Shayan Hosseinzadeh, Dian Yang, Angela N Pogson, Marco Y Hein, Kyung Hoi Joseph Min, Li Wang, et al. Mapping transcriptomic vector fields of single cells. *Cell*, 185(4):690–711, 2022.
- [25] Thomas Schreiber. Measuring information transfer. *Physical review letters*, 85(2):461, 2000.
- [26] Lionel Barnett, Adam B Barrett, and Anil K Seth. Granger causality and transfer entropy are equivalent for gaussian variables. *Physical review letters*, 103(23):238701, 2009.
- [27] Daniel T Gillespie. The chemical langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, 2000.
- [28] Crispin Gardiner. *Stochastic Methods: A Handbook For The Natural And Social Sciences*. Springer, 2009.
- [29] MA Coomer, L Ham, and MPH Stumpf. Noise distorts the epigenetic landscape and shapes cell-fate decisions. *Cell Syst*, 13(1):83–102, 2022.
- [30] Terry Bossomaier, Lionel Barnett, Michael Harré, and Joseph T Lizier. Transfer entropy. In *An introduction to transfer entropy*, pages 65–95. Springer, 2016.
- [31] Thomas M. Cover and Joy A. Thomas. *Elements Of Information Theory*. J. Wiley, 2005.
- [32] Adam A Margolin, Ilya Nemenman, Katia Basso, Chris Wiggins, Gustavo Stolovitzky, Riccardo Dalla Favera, and Andrea Califano. Aracne: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. In *BMC bioinformatics*, volume 7, pages 1–15. BioMed Central, 2006.
- [33] Louise Deconinck, Robrecht Cannoodt, Wouter Saelens, Bart Deplancke, and Yvan Saeys. Recent advances in trajectory inference from single-cell omics data. *Current Opinion in Systems Biology*, 27:100344, 2021.
- [34] Stephen Zhang, Anton Afanassiev, Laura Greenstreet, Tetsuya Matsumoto, and Geoffrey Schiebinger. Optimal transport analysis reveals trajectories in steady-state systems. *PLoS computational biology*, 17(12):e1009466, 2021.

- [35] Kevin R Moon, Jay S Stanley III, Daniel Burkhardt, David van Dijk, Guy Wolf, and Smita Krishnaswamy. Manifold learning-based methods for analyzing single-cell rna-sequencing data. *Current Opinion in Systems Biology*, 7:36–46, 2018.
- [36] Aviv Madar, Alex Greenfield, Eric Vanden-Eijnden, and Richard Bonneau. Dream3: network inference using dynamic context likelihood of relatedness and the inferelator. *PloS one*, 5(3):e9803, 2010.
- [37] Jeremiah J Faith, Boris Hayete, Joshua T Thaden, Ilaria Mogno, Jamey Wierzbowski, Guillaume Cottarel, Simon Kasif, James J Collins, and Timothy S Gardner. Large-scale mapping and validation of escherichia coli transcriptional regulation from a compendium of expression profiles. *PLoS biology*, 5(1):e8, 2007.
- [38] Tiejun Li, Jifan Shi, Yichong Wu, and Peijie Zhou. On the mathematics of rna velocity i: theoretical analysis. *bioRxiv*, 2020.
- [39] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.
- [40] Tetsutaro Hayashi, Haruka Ozaki, Yohei Sasagawa, Mana Umeda, Hiroki Danno, and Itoshi Nikaido. Single-cell full-length total rna sequencing uncovers dynamics of recursive splicing and enhancer rnas. *Nature communications*, 9(1):1–16, 2018.
- [41] Yael Baran, Akhiad Bercovich, Arnau Sebe-Pedros, Yaniv Lubling, Amir Giladi, Elad Chomsky, Zohar Meir, Michael Hoichman, Aviezer Lifshitz, and Amos Tanay. Metacell: analysis of single-cell rna-seq data using k-nn graph partitions. *Genome biology*, 20(1):1–19, 2019.
- [42] Manu Setty, Vaidotas Kisieliovas, Jacob Levine, Adam Gayoso, Linas Mazutis, and Dana Pe’Er. Characterization of cell fate probabilities in single-cell data with palantir. *Nature biotechnology*, 37(4):451–460, 2019.
- [43] T Matsumoto, S Zhang, and G Schiebinger. Beyond knn: Adaptive, sparse neighborhood graphs via optimal transport. Submitted, 2022.
- [44] Rihao Qu, Xiuyuan Cheng, Peggy Myung, Esen Sefik, Richard Flavell, and Yuval Kluger. Gene trajectory inference for single-cell data by optimal transport metrics. *bioRxiv*, pages 2022–07, 2022.
- [45] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- [46] Huilei Xu, Caroline Baroukh, Ruth Dannenfelser, Edward Y Chen, Christopher M Tan, Yan Kou, Yujin E Kim, Ihor R Lemischka, and Avi Ma’ayan. Escape: database for integrating high-content published data collected from human and mouse embryonic stem cells. *Database*, 2013, 2013.
- [47] Wenjing Shi, Hui Wang, Guangjin Pan, Yijie Geng, Yunqian Guo, and Duanqing Pei. Regulation of the pluripotency marker rex-1 by nanog and sox2. *Journal of biological chemistry*, 281(33):23319–23325, 2006.
- [48] Guangjin Pan and James A Thomson. Nanog and transcriptional networks in embryonic stem cell pluripotency. *Cell research*, 17(1):42–49, 2007.
- [49] Chin Yan Lim, Wai-Leong Tam, Jinqiu Zhang, Haw Siang Ang, Hui Jia, Leonard Lipovich, Huck-Hui Ng, Chia-Lin Wei, Wing Kin Sung, Paul Robson, et al. Sall4 regulates distinct transcription circuitries in different blastocyst-derived stem cell lineages. *Cell stem cell*, 3(5):543–554, 2008.
- [50] Kelly L Covello, James Kehler, Hongwei Yu, John D Gordan, Andrew M Arsham, Cheng-Jun Hu, Patricia A Labosky, M Celeste Simon, and Brian Keith. Hif-2 $\alpha$  regulates oct-4: effects of hypoxia on stem cell function, embryonic development, and tumor growth. *Genes & development*, 20(5):557–570, 2006.

- [51] Hong Doan, Alexander Parsons, Shruthi Devkumar, Jogitha Selvarajah, Francesc Miralles, and Veronica A Carroll. Hif-mediated suppression of deptor confers resistance to mtor kinase inhibition in renal cancer. *Isience*, 21:509–520, 2019.
- [52] Pooja Agrawal, Joseph Reynolds, Shereen Chew, Deepak A Lamba, and Robert E Hughes. Deptor is a stemness factor that regulates pluripotency of embryonic stem cells. *Journal of Biological Chemistry*, 289(46):31818–31826, 2014.
- [53] Edward E Morrissey, Zhihua Tang, Kirsten Sigrist, Min Min Lu, Fang Jiang, Hon S Ip, and Michael S Parmacek. Gata6 regulates hnf4 and is required for differentiation of visceral endoderm in the mouse embryo. *Genes & development*, 12(22):3579–3590, 1998.
- [54] Yingyao Zhou, Bin Zhou, Lars Pache, Max Chang, Alireza Hadj Khodabakhshi, Olga Tanaseichuk, Christopher Benner, and Sumit K Chanda. Metascape provides a biologist-oriented resource for the analysis of systems-level datasets. *Nature communications*, 10(1):1–10, 2019.
- [55] Aimée Bastidas-Ponce, Sophie Tritschler, Leander Dony, Katharina Scheibner, Marta Tarquis-Medina, Ciro Salinno, Silvia Schirge, Ingo Burtscher, Anika Böttcher, Fabian J Theis, et al. Comprehensive single cell mrna profiling reveals a detailed roadmap for pancreatic endocrinogenesis. *Development*, 146(12):dev173849, 2019.
- [56] Volker Bergen, Marius Lange, Stefan Peidli, F Alexander Wolf, and Fabian J Theis. Generalizing rna velocity to transient cell states through dynamical modeling. *Nature biotechnology*, 38(12):1408–1414, 2020.
- [57] Ashleigh E Schaffer, Brandon L Taylor, Jacqueline R Benthuyssen, Jingxuan Liu, Fabrizio Thorel, Weiping Yuan, Yang Jiao, Klaus H Kaestner, Pedro L Herrera, Mark A Magnuson, et al. Nkx6. 1 controls a gene regulatory network required for establishing and maintaining pancreatic beta cell identity. *PLoS genetics*, 9(1):e1003274, 2013.
- [58] EC Baechler, FM Batliwalla, George Karypis, PM Gaffney, K Moser, WA Ortmann, KJ Espe, S Balasubramanian, KM Hughes, JP Chan, et al. Expression levels for many genes in human peripheral blood cells are highly sensitive to ex vivo incubation. *Genes & Immunity*, 5(5):347–353, 2004.
- [59] Léo Machado, Frédéric Relaix, and Philippos Mourikis. Stress relief: emerging methods to mitigate dissociation-induced artefacts. *Trends in Cell Biology*, 31(11):888–897, 2021.
- [60] Brandon L Taylor, Fen-Fen Liu, and Maike Sander. Nkx6. 1 is essential for maintaining the functional state of pancreatic beta cells. *Cell reports*, 4(6):1262–1275, 2013.
- [61] H Efsun Arda, Cecil M Benitez, and Seung K Kim. Gene regulatory networks governing pancreas development. *Developmental cell*, 25(1):5–13, 2013.
- [62] Philip A Seymour, Kristine K Freude, Man N Tran, Erin E Mayes, Jan Jensen, Ralf Kist, Gerd Scherer, and Maike Sander. Sox9 is required for maintenance of the pancreatic progenitor cell pool. *Proceedings of the National Academy of Sciences*, 104(6):1865–1870, 2007.
- [63] Richard I Sherwood, Tzong-Yang Albert Chen, and Douglas A Melton. Transcriptional dynamics of endodermal organ formation. *Developmental dynamics: an official publication of the American Association of Anatomists*, 238(1):29–42, 2009.
- [64] Shin-Heng Chiou, Madeleine Dorsch, Eva Kusch, Santiago Naranjo, Margaret M Kozak, Albert C Koong, Monte M Winslow, and Barbara M Grüner. Hmga2 is dispensable for pancreatic cancer development, metastasis, and therapy resistance. *Scientific reports*, 8(1):1–10, 2018.
- [65] Kailun Lee, Jeng Yie Chan, Cassandra Liang, Chi Kin Ip, Yan-Chuan Shi, Herbert Herzog, William E Hughes, Mohammed Bensellam, Viviane Delghingaro-Augusto, Mark E Koina, et al. Xbp1 maintains beta cell identity, represses beta-to-alpha cell transdifferentiation and protects against diabetic beta cell failure during metabolic stress in mice. *Diabetologia*, 65(6):984–996, 2022.

- [66] Gennady Gorin, Meichen Fang, Tara Chari, and Lior Pachter. Rna velocity unraveled. *bioRxiv*, 2022.
- [67] Shijie C Zheng, Genevieve Stein-O’Brien, Leandros Boukas, Loyal A Goff, and Kasper D Hansen. Pumping the brakes on rna velocity—understanding and interpreting rna velocity estimates. *bioRxiv*, pages 2022–06, 2022.
- [68] 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, 2022.
- [69] Thomas Graf and Tariq Enver. Forcing cells to change lineages. *Nature*, 462(7273):587–594, 2009.
- [70] Rachel M Gerstein. Deciding the decider: Mef2c in hematopoiesis. *Nature Immunology*, 10(3):235–236, 2009.
- [71] Hongsheng Wang and Herbert C Morse. Irf8 regulates myeloid and b lymphoid lineage diversification. *Immunologic research*, 43:109–117, 2009.
- [72] Tomohiko Tamura, Daisuke Kurotaki, and Shin-ichi Koizumi. Regulation of myelopoiesis by the transcription factor irf8. *International journal of hematology*, 101:342–351, 2015.
- [73] Voahangy Randrianarison-Huetz, Benoit Laurent, Valérie Bardet, Gerard C Blobbe, François Huetz, and Dominique Duménil. Gfi-1b controls human erythroid and megakaryocytic differentiation by regulating tgf- $\beta$  signaling at the bipotent erythro-megakaryocytic progenitor stage. *Blood, The Journal of the American Society of Hematology*, 115(14):2784–2795, 2010.
- [74] Emmanuelle Passegué, Erwin F Wagner, and Irving L Weissman. Junb deficiency leads to a myeloproliferative disorder arising from hematopoietic stem cells. *Cell*, 119(3):431–443, 2004.
- [75] Karthik Arumugam, William Shin, Valentina Schiavone, Lukas Vlahos, Xiaochuan Tu, Davide Carnevali, Jordan Kesner, Evan O Paull, Neus Romo, Prem Subramaniam, et al. The master regulator protein baz2b can reprogram human hematopoietic lineage-committed progenitors into a multipotent state. *Cell reports*, 33(10):108474, 2020.
- [76] Lam-Ha Ly and Martin Vingron. Effect of imputation on gene network reconstruction from single-cell rna-seq data. *Patterns*, 3(2):100414, 2022.
- [77] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.
- [78] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.