

RECOVERING TIME-VARYING NETWORKS FROM SINGLE-CELL DATA

Anonymous authors

Paper under double-blind review

ABSTRACT

Gene regulation is a dynamic process that underlies all aspects of human development, disease response, and other key biological processes. The reconstruction of temporal gene regulatory networks has conventionally relied on regression analysis, graphical models, or other types of relevance networks. With the large increase in time series single-cell data, new approaches are needed to address the unique scale and nature of this data for reconstructing such networks. Here, we develop a deep neural network, Marlene, to infer dynamic graphs from time series single-cell gene expression data. Marlene constructs directed gene networks using a self-attention mechanism where the weights evolve over time using recurrent units. By employing meta learning, the model is able to recover accurate temporal networks even for rare cell types. In addition, Marlene can identify gene interactions relevant to specific biological responses, including COVID-19 immune response, fibrosis, and aging, paving the way for potential treatments.

1 INTRODUCTION

Biological systems are dynamic, changing over time in response to various stimuli and events. To construct accurate models of biological activity during development, disease progression, treatment response, and other biological processes, it is essential to track their evolution over time (Bar-Joseph et al., 2012). Studying the *regulation* of these dynamic processes is key for understanding the underlying mechanisms that drive the response and for identifying potential interventions that can serve as cures for diseases (Silverman et al., 2020).

Much of the research in this area is focused on the reconstruction of regulatory networks (Karlebach & Shamir, 2008; Badia-I-Mompel et al., 2023). These networks comprise a subset of proteins known as transcription factors (TFs), which regulate the activity of all other genes and proteins within the cell. However, these gene regulatory networks (GRNs) are not static. Instead, both the active nodes (proteins) and the edges (genes) change over time (Gasch et al., 2000; Luscombe et al., 2004). To reconstruct such networks, researchers often integrate static data—such as the type of nodes in the network—with dynamic data, such as time series measurements of node activity (gene expression profiles). Early work in this area employed microarrays and ChIP-chip data (Lee et al., 2002; Wang et al., 2006; Blais & Dynlacht, 2005; Gilchrist et al., 2009) followed by time series next-generation RNA-seq data (Wang et al., 2009), and most recently, single-cell RNA-seq (scRNA-seq) data (Ding et al., 2022; Nguyen et al., 2021; Matsumoto et al., 2017).

Several computational methods have been proposed over the last two decades to reconstruct such dynamic GRNs (Bar-Joseph et al., 2003; Yosef et al., 2013; Schulz et al., 2012; Yan et al., 2021). Some of these methods utilized time-varying graphical models including Hidden Markov models, Markov random fields, and Dynamic Bayesian Networks (Ahmed & Xing, 2009; Song et al., 2009; Dondelinger et al., 2013; Zhu & Wang, 2015). Other approaches attempted to use regression or to learn temporal precision matrices using extensions of the graphical lasso algorithm (Wang et al., 2020; Hallac et al., 2017).

While such models successfully reconstructed some processes (Ahmed & Xing, 2009; Kim et al., 2012), they are less suitable for more recent types of data, most notably scRNA-seq time series. First, the larger size of the data presents a challenge for traditional graphical models. Also, prior methods do not directly account for the fact that multiple cells are profiled for each time point. Finally, prior

054 methods do not leverage larger models, such as neural networks, which have demonstrated signif-
 055 icant performance improvements across various learning tasks (Ching et al., 2018; Angermueller
 056 et al., 2016).

057 Very recently, a few methods have been proposed for using deep learning to recover static GRNs
 058 (Shrivastava et al., 2022; Shu et al., 2021). However, they cannot be directly used to capture dynamic
 059 GRNs (i.e., enforcing learning between time points). Two recent approaches, Dictys and CellOracle
 060 (Wang et al., 2023b; Kamimoto et al., 2023) can infer dynamic GRNs, however, these methods
 061 depend on data types like ATAC-seq, which provides direct information about TF binding sites but
 062 is less prevalent and harder to obtain.

063 Beyond the realm of biology, the inference of dynamic graphs using neural networks has garnered
 064 significant attention. This problem has found applications in diverse domains, including information
 065 retrieval, molecular graphs, and traffic forecasting (Zhu et al., 2021; Zhang et al., 2020). While there
 066 are similarities between these problems and the dynamic GRN problem, there are also significant
 067 differences that make it hard to extend these methods for time series scRNA-seq. The problem
 068 of inferring temporal graphs is usually defined by recovering a series of graph adjacency matrices
 069 $\mathbf{A}_t \in \mathbb{R}^{n \times n}$ where n is the number of nodes and each node is a k -dimensional feature vector.
 070 However, when dealing with scRNA-seq data, the problem becomes: given a gene expression matrix
 071 $\mathbf{X}_t \in \mathbb{R}^{c \times g}$, where c is the number of cells (samples) and g is the number of genes (features), we are
 072 interested in recovering gene networks $\mathbf{A}_t \in \mathbb{R}^{g \times g}$, i.e., *graphs of features* rather than nodes (cells).

073 In this paper, we present a novel deep learning framework that effectively addresses the challenges
 074 discussed above for reconstructing dynamic GRNs. Our contribution is three-fold. First, we demon-
 075 strate that existing deep learning methods for temporal graph structure learning can be adapted
 076 for scRNA-seq data analysis. To achieve this, we perform a *gene featurization* step by leveraging
 077 set-like architectures such as DeepSets or Set Transformers (Zaheer et al., 2017; Lee et al., 2019).
 078 Second, we construct dynamic graphs by applying a self-attention mechanism (Bahdanau et al.,
 079 2014) to these gene feature vectors. To model dynamics, we draw inspiration from EvolveGCN
 080 where a gated recurrent unit (GRU) evolves the weights of a graph neural network (Pareja et al.,
 081 2020). However, unlike EvolveGCN, our approach uses a GRU to evolve the weights of key and
 082 value projection matrices in the self-attention module. This allows for the construction of dynamic
 083 graphs that capture regulatory interactions over time. Lastly, GRNs are highly dependent on cell
 084 functions, hence, separate GRNs need to be learned for each cell type. A single scRNA-seq dataset
 085 may combine cells of multiple types, some of which are rare cell populations. To this end, we
 086 employ a model-agnostic meta-learning (MAML) (Finn et al., 2017) training procedure by treating
 087 each cell type as a “task” to be learned. With this approach, the model quickly adapts to tasks with
 088 few samples, enabling the reconstruction of dynamic graphs even for rare cell types.

089 We apply our **meta** learning approach for inferring temporal **gene** regulatory networks (Marlene)
 090 to three publicly available scRNA-seq datasets. The first is a time series SARS-CoV-2 mRNA
 091 vaccination dataset of human peripheral blood mononuclear cells (PBMCs) (Zhang et al., 2023).
 092 The second dataset is a human lung aging atlas from the Human Cell Atlas Project (Regev et al.,
 093 2017; Sikkema et al., 2023). The third dataset is from a study of lung fibrosis using a mouse lung
 094 injury model (Strunz et al., 2020). All three datasets incorporate several time points, thus enabling
 095 a longitudinal analysis of the relevant biological responses through the inference of dynamic, cell
 096 type-specific GRNs. As we show, our method is able to reconstruct accurate networks for these
 097 datasets, significantly improving upon prior methods proposed for this task.

098 2 METHODS

100 2.1 PROBLEM SETUP

101 Consider a gene expression matrix $\mathbf{X} \in \mathbb{R}^{c \times g}$ where c is the number of cells and g is the number
 102 of genes. In the human genome, g varies from 25,000 to 30,000, while the number of cells could
 103 be between a couple thousand to a few million. In the setting of dynamic graphs, we assume the
 104 existence of a time point for each row (cell), leading to a time series $\tilde{\mathbf{X}} := \{\mathbf{X}_1, \dots, \mathbf{X}_T\}$ with
 105 $\mathbf{X}_t \in \mathbb{R}^{c_t \times g}$. Here, the number of cells c_t may vary with t . We are interested in recovering a
 106 series of directed graphs $\tilde{\mathcal{G}} := \{\mathcal{G}_1, \dots, \mathcal{G}_T\}$ where each $\mathcal{G}_t = \{\mathcal{N}, \mathcal{E}_t\}$. The set of nodes is the
 107 set of genes, i.e., $\mathcal{N} = [g]$, and we assume this set is static over time. The dynamic edge sets

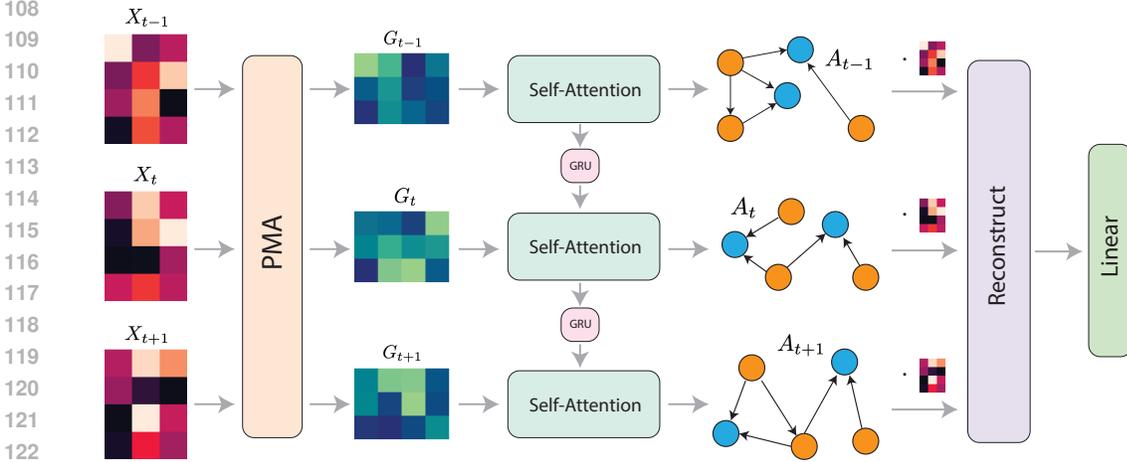


Figure 1: Overview of Marlene. Marlene takes as input gene expression data in the form of a cell-by-gene matrix. It then performs gene featurization via the pooling by multihead attention (PMA) mechanism which returns a gene feature matrix. This matrix is then inputted into a self-attention module to obtain a gene network in the form of an adjacency matrix. The weights of the self-attention module evolve from one time point to the next via a gated recurrent unit (GRU). The expression of transcription factors and the recovered graph are used to reconstruct the full gene expression vector. Finally, the reconstructed matrix is used to predict the cell type for the batch. The network is trained in a model-agnostic meta-learning fashion where each cell type is treated as a “task” to be learned, thus enabling the model to quickly adapt to cell types with low representation.

$\mathcal{E}_t = \{(u, v, w)\}_{u, v \in \mathcal{N}, w \in \mathbb{R}}$ denote directed weighted links between genes, where the source u is the gene that regulates the expression of its target v , and w is the strength of this relationship. The source nodes are called transcription factor genes (TFs).

We can alternatively characterize each graph \mathcal{G}_t by the corresponding adjacency matrix $\mathbf{A}_t \in \mathbb{R}^{g \times g}$. Denote $\tilde{\mathbf{A}} := \{\mathbf{A}_1, \dots, \mathbf{A}_T\}$. Since TFs control the expression of their target genes, the underlying GRNs should, in principle, allow the recovery of the full expression profile for a cell. In other words, $\tilde{\mathbf{X}} = f(\tilde{\mathbf{X}}^{\text{TF}}, \tilde{\mathbf{A}})$ where $\tilde{\mathbf{X}}^{\text{TF}}$ denotes the expression of all TFs. The function f is unknown as it involves intricate interactions among genes, including combinatorial effects. For instance, certain scenarios exist where TFs cooperate with each-other to activate a gene, while in other instances, the activation requires some TFs to be active and others to be inactive (Wise & Bar-Joseph, 2015).

In existing deep learning literature, f is sometimes modeled using autoencoders (Seninge et al., 2021; Shu et al., 2021; Wang et al., 2023a). However, reconstructing the full gene expression vector is challenging as the data is extremely sparse and conventional reconstruction losses, such as mean squared error, tend to emphasize overall averages. Since GRNs are dependent on cell function, we hypothesize that simplifying the problem by predicting cell types may improve the accurate recovery of GRNs. In other words, given a temporal batch of cells of the same type $\tilde{x}^{\text{TF}} \in \mathbb{R}^{T \times \text{batch size} \times |\text{TFs}|}$, we consider the classification problem $y = f(\tilde{x}^{\text{TF}}, \tilde{\mathbf{A}})$ where y is the known cell type label for the batch. Finally, the task of learning $\tilde{\mathbf{A}}$ given a batch \tilde{x} becomes

$$\arg \min_{\tilde{\mathbf{A}}} \text{CrossEntropyLoss}(y, f(\tilde{x}^{\text{TF}}, h(\tilde{x}))), \quad \tilde{\mathbf{A}} := h(\tilde{x}) \quad (1)$$

for choices of functions f and h where h uses the expression data to obtain the adjacency matrices.

2.2 ARCHITECTURE OF MARLENE

In this work, we propose a neural network architecture called Marlene that effectively learns dynamic GRNs (Figure 1). Marlene consists of three main steps. The first two steps address the choice of h in (1), while the last addresses the choice of f .

In the first step, we apply a gene featurization step by treating a batch of cells as a set of elements. The DeepSet architecture introduces a pooling operator that allows the neural network to be invariant to the order of input samples, effectively treating the input as a set (Zaheer et al., 2017). Similarly, the Set Transformer architecture is designed to process sets of data via attention-based operators that are permutation invariant (Lee et al., 2019). Specifically, the pooling by multihead attention (PMA) aggregation scheme introduced in Set Transformers outputs a matrix of k vectors $\mathbf{H} \in \mathbb{R}^{k \times g}$ for an arbitrary input set $\mathbf{X} \in \mathbb{R}^{c \times g}$. Each of the k output vectors has a specific meaning such as statistics of the input data. By applying the PMA operator to a temporal batch of cells, we obtain a gene-by-feature matrix $\mathbf{G} = \mathbf{H}^\top \in \mathbb{R}^{g \times k}$ that encodes information about the cells from each time point. PMA consists of a multihead attention block (MAB) where the input \mathbf{X} consists of key vectors, and the query is a learnable set of k vectors $\mathbf{S} \in \mathbb{R}^{g \times k}$. In this work, we use a shared PMA layer for all time points assuming that the specific key statistical properties are invariant (though their value obviously changes for different time points). Given $\tilde{x} = [x_1, \dots, x_T]$, we have

$$\tilde{\mathbf{G}} = \text{PMA}(\tilde{x})^\top := \text{MAB}(\mathbf{S}, \tilde{x})^\top := (\tilde{\mathbf{M}} + \text{rFF}(\tilde{\mathbf{M}}))^\top \quad (2)$$

where $\mathbf{M} = \mathbf{S} + \text{Multihead}(\mathbf{S}, x, x) \in \mathbb{R}^{g \times k}$ and rFF is a row-wise feedforward layer. For completeness, these operations are defined in the appendix.

Note that if cells of different types are mixed in the same batch, the statistics derived by the PMA step may not capture cell type-specific information. Consequently, in a single batch, we only include cells of one type.

In the second step, Marlene learns temporal adjacency matrices using a self-attention mechanism. To model dynamics, we draw inspiration from EvolveGCN which performs model adaptation using a GRU (Pareja et al., 2020). Unlike EvolveGCN, which uses the GRU to update the weights of a graph convolution layer, we use a GRU to evolve the key and query projection weights of the self-attention module. Since most time series we deal with contain very few time points, a GRU should suffice and not suffer from vanishing gradient problems. Similar to EvolveGCN, we apply a summarization step via top k pooling to reduce the gene feature matrix to a square matrix for the GRU (Appendix).

More precisely, we initialize self-attention weights $\tilde{\mathbf{W}}_0^Q, \tilde{\mathbf{W}}_0^K \in \mathbb{R}^{k \times k}$ and two recurrent units $\text{GRU}_Q, \text{GRU}_K$. Given the time sequence of gene feature matrices $\tilde{\mathbf{G}}$ obtained from the previous step, temporal adjacency matrices are constructed in the following recurrent fashion for all $t \in [T]$:

$$\mathbf{Z}_t = \text{TopK}(\mathbf{G}_t) \in \mathbb{R}^{k \times k} \quad (3)$$

$$\mathbf{W}_t^Q = \text{GRU}_Q(\mathbf{Z}_t, \mathbf{W}_{t-1}^Q), \quad \mathbf{W}_t^K = \text{GRU}_K(\mathbf{Z}_t, \mathbf{W}_{t-1}^K) \quad (4)$$

$$\mathbf{Q}_t = \mathbf{G}_t \mathbf{W}_t^Q, \quad \mathbf{K}_t = \mathbf{G}_t \mathbf{W}_t^K \quad (5)$$

$$\mathbf{A}_t = \text{softmax} \left(\frac{\mathbf{Q}_t \mathbf{K}_t^\top}{\sqrt{k}} \right). \quad (6)$$

Here, \mathbf{W}_t^Q and \mathbf{W}_t^K serve as hidden states for the respective GRUs. The GRUs dynamically adapt self-attention weights, influencing which TFs specific genes should attend to in subsequent time steps. Consequently, the evolution of these weights is constrained. We also restrict the columns of \mathbf{A}_t (i.e., sources) to p known TFs in the TRRUST database (Han et al., 2018) which greatly reduces the number of parameters to be learned. Therefore, in our implementation $\mathbf{A}_t \in \mathbb{R}^{g \times p}$.

Next, we perform a gene expression reconstruction step based on the expression of TFs and the inferred adjacency matrices. This is followed by any number of fully connected layers with nonlinear activation functions σ . Finally, we sum across output vectors to obtain a logit vector with the same dimension as the number of cell types in the data:

$$\tilde{y} = \text{Pool}(\text{Linear}(\dots \sigma(\text{Linear}(\tilde{x}^{\text{TF}} \tilde{\mathbf{A}}^\top))))). \quad (7)$$

Network depth can be introduced at all three levels by stacking MAB layers during gene featurization, stacking GRUs, or stacking linear layers at the end.

2.3 META LEARNING FOR RARE CELL TYPES

ScRNA-seq datasets often originate from biological samples that exhibit cellular heterogeneity, potentially containing multiple distinct cell types. Some of these cell subpopulations are rare and are

216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269

Table 1: Time series scRNA-seq datasets used in this study.

Dataset	Number of				Metadata	
	Cells	Genes ¹	TFs	Cell Types	Time Points	Sample
SARS-CoV-2	113,271	1899	556	7	d0, d2, d10, d28	PBMCs (human)
HLCA	27,953 ²	2433	674	11	Ages < 35, 35 – 50, ≥ 50	lung (human)
Fibrosis	22,758	1217	433	6	PBS, d3, d7, d10, d14, d21, d28	lung (mouse)

¹ Only showing the number of genes overlapping with the TRRUST database.
² We randomly sampled cells from 11 cell types.

represented by a small number of cells in the sample (Jindal et al., 2018). Since we are concerned with the discovery of cell type-specific temporal GRNs, learning such large graphs for these rare cell types may not be feasible and lead to overfitting. Since many interactions are shared across cell types (Chasman & Roy, 2017), we employ the model-agnostic meta-learning framework (MAML) (Finn et al., 2017). MAML is specifically designed to enable neural networks to adapt to novel tasks with limited training samples (i.e., few shot learning). By treating each cell type as a “task”, the MAML training paradigm facilitates the recovery of dynamic graphs for rare cell types. We begin by adapting model parameters through multiple optimization steps using a batch of support examples (cells). These adapted parameters are then evaluated on a separate set of query cells, followed by a meta-update.

During the adaptation step, we perform gradient descent, while for the meta-update, we employ the Adam optimizer (Kingma & Ba, 2014). During training, gradient clipping proves crucial to prevent overfitting of the MAML adaptation step to the cell type under consideration.

3 EXPERIMENTS

To validate our approach, we use three public scRNA-seq datasets (Table 1): a human SARS-CoV-2 mRNA vaccination dataset, a lung aging atlas (The Human Lung Cell Atlas—HLCA), and a mouse lung fibrosis dataset (Zhang et al., 2023; Sikkema et al., 2023; Strunz et al., 2020). To assess the quality of the inferred networks, we draw upon two databases of TF-gene regulatory interactions, which have been curated from the scientific literature—TRRUST and RegNetwork (Han et al., 2018; Liu et al., 2015). For the human genome, TRRUST contains 8427 unique validated regulatory edges, while RegNetwork contains 150,405. Note that certain edges lack a corresponding TF or gene in the expression data, so the numbers used for the analysis are smaller. We used only the genes that were present in TRRUST for all three datasets. For the mouse lung dataset we used the corresponding mouse networks for both databases. To match the number of links in these databases, we selected the top 2% of edges for all methods. For Marlene, this was done by sparsifying the self-attention matrix to retain only the top scoring edges. The significance of the overlap was carried out via Fisher’s exact test (Fisher, 1922). All p -values were corrected for multiple testing using the Benjamini-Hochberg procedure (Benjamini & Hochberg, 1995).

We compare Marlene against several popular static gene regulatory network inference methods included in the BEELINE benchmark (Pratapa et al., 2020) and beyond, such as PIDC, GENIE3, GRNBoost2, SCODE, and DeepSEM (Chan et al., 2017; Huynh-Thu et al., 2010; Moerman et al., 2019; Matsumoto et al., 2017; Shu et al., 2021), which are applied independently to each time point. DeepSEM is a deep generative model based on structural equation modeling. We also compare against time-varying graphical lasso (TVGL)¹, a method that models temporal precision matrices (Hallac et al., 2017), and to a deep neural network that utilizes the S4 module (GraphS4mer) (Tang et al., 2022; Gu et al., 2021).

During inference, we obtain multiple A_t for different batches and average them. We train Marlene with a batch size of 16 cells and also use 16 seeds in the PMA layer. For MAML, we use 5 inner steps. The model with the lowest loss is selected for GRN inference. For the meta-update, we use a decaying learning rate starting with $1e-4$, while for the inner step we use $1e-3$ for both datasets. Experiments were performed using an NVIDIA RTX 3060 and took only a few minutes per run.

¹We used the implementation of <https://github.com/fdtomasi/regain>

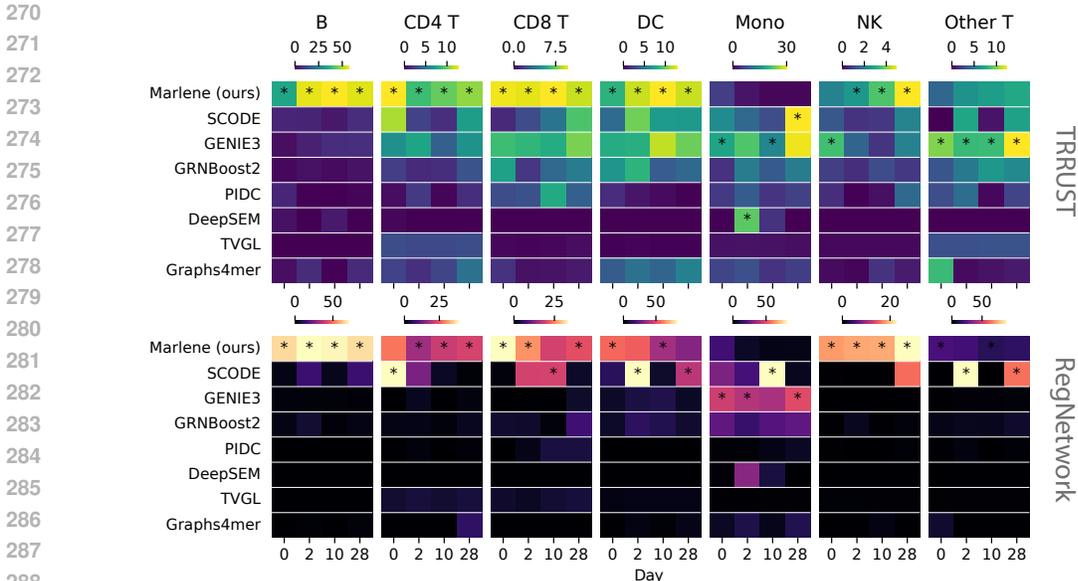


Figure 2: Overlap analysis of the SARS-CoV-2 vaccination dataset. Showing $-\log_{10}(\text{FDR})$ values from a Fisher’s exact test measuring the overlap between predicted TF-gene interactions in reconstructed networks and two TF-gene interaction databases, TRRUST (top) and RegNetwork (bottom). Cell types are shown as columns. Best performing method is starred.

3.1 CASE STUDY 1: SARS-COV-2 VACCINATION

The SARS-CoV-2 vaccination dataset consists of peripheral blood mononuclear cells (PBMCs) from six healthy donors at four time points (days 0, 2, 10, and 28) (Zhang et al., 2023). Day 0 samples were obtained before vaccination. We removed the “Other” cell type group and kept the remaining seven. These include B cells, dendritic cells (DC), monocytes (Mono), natural killer cells (NK), and various types of T cells.

3.1.1 MARLENE RECOVERS ACCURATE GENE REGULATORY NETWORKS

Analysis results using Marlene and prior methods is presented in Figure 2. As can be seen, Marlene outperformed competing methods in the dynamic GRN inference task for 5 of the cell types, yielding statistically significant results across time points. Specifically, for B cells, Marlene successfully identified more than 800 regulatory links within the RegNetwork framework at each time point ($\text{FDR} \leq 1e-67$), surpassing the performance of the second-best method, SCODE, which detected 579 links (day 2, $\text{FDR} \leq 1e-15$). Analogous findings were observed for natural killer cells, where Marlene identified over 600 RegNetwork links at each time point ($\text{FDR} \leq 1e-18$). In comparison, the second-ranking method, SCODE, showed a significant overlap for only one time point (day 28). Upon examining the TRRUST database, we observed less pronounced differences in the results. Nonetheless, Marlene obtained higher overlap for 5 out of 7 cell types followed by GENIE3, which performed well for monocytes and the “Other T” category.

3.1.2 MARLENE RECOVERS REALISTIC DYNAMIC TRANSITIONS

The analysis so far has primarily focused on individual time points. Next, we turned our attention to assessing the quality of graph transitions between consecutive time points. Specifically, we examined whether the learned graphs demonstrated smooth transitions over time. To evaluate this, we computed the intersection-over-union (IoU) score for edges between time points t and $t + 1$ (Figure 3a). Notably, our findings revealed that for most cell types, Marlene exhibited the lowest IoU score during the initial period (days $0 \rightarrow 2$), followed by higher scores during days $2 \rightarrow 10$, and $10 \rightarrow 28$. This pattern aligns with our expectations, as variations in gene expression are likely

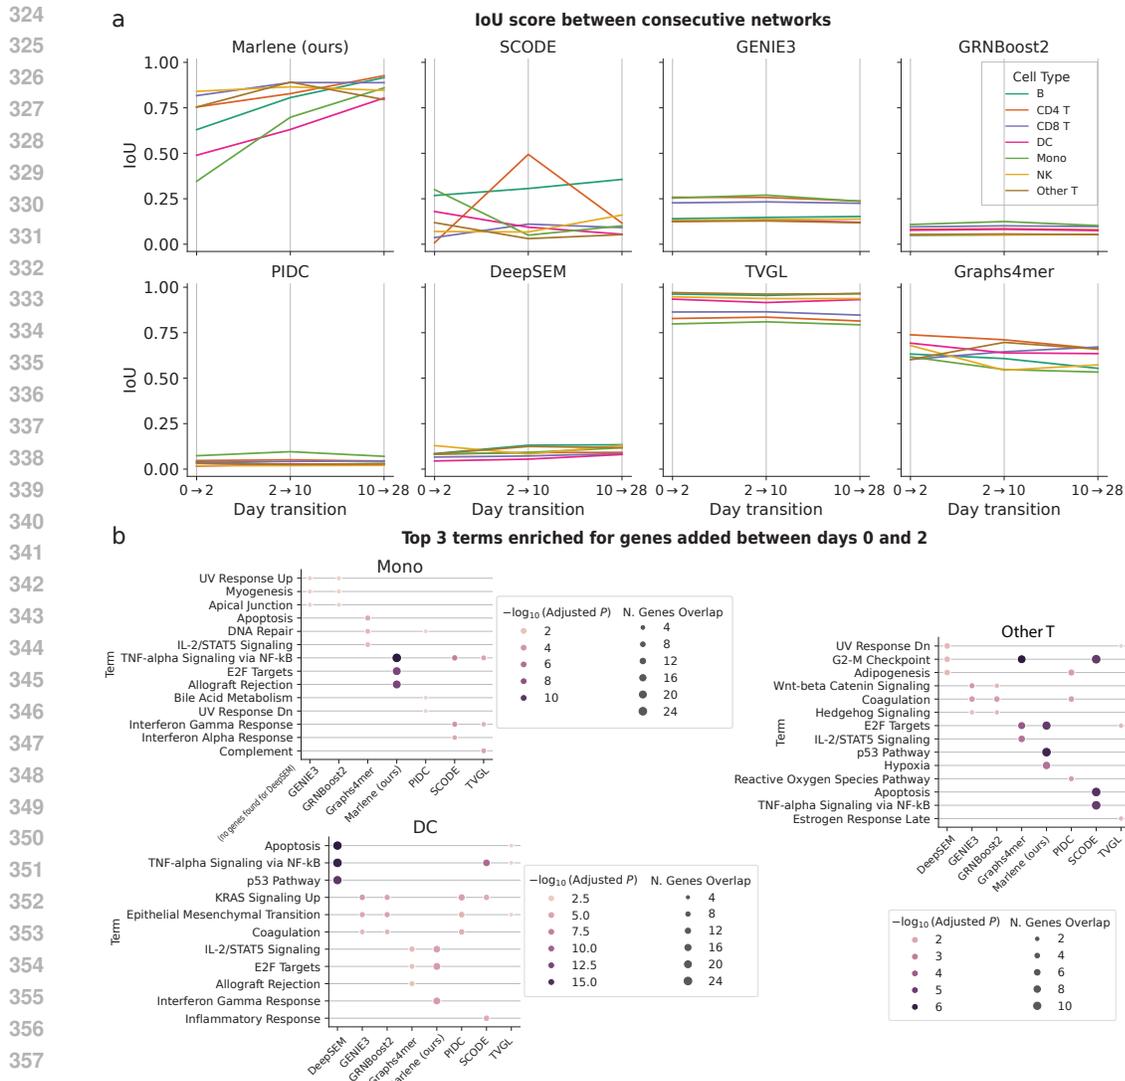


Figure 3: Temporal analysis of the predicted gene regulatory networks for the SARS-CoV-2 vaccine dataset. (a) Intersection-over-union (IoU) scores between consecutive graphs. (b) For each method, top 3 MSigDB terms enriched for genes that were regulated at day 2 but not day 0.

to be most pronounced during the early post-vaccination period (days 0 → 2). From the methods we compare against, GENIE3, GRNBoost2, SCODE, PIDC, and DeepSEM exhibited significantly lower IoU scores across all temporal transitions, likely due to their lack of dynamic modeling and the fact that they were ran independently per time point. TVGL, on the other hand, showed high IoU scores which remained close to constant over time. Finally, Graphs4mer displayed a reduction in IoU scores over time, which is unlikely given the expected immediate immune response.

Next, we sought to assess the quality of the TF-gene regulatory links added between time points. For brevity, we focused specifically on the initial temporal transition (days 0 → 2), as this period is likely to witness a more significant biological response. For each cell type, we took note of all the genes that were regulated by some TF at day 2 but not at day 0. Using this set of genes z , we performed gene set enrichment analysis (GSEA) (Subramanian et al., 2005; Fang et al., 2023) using the molecular signatures database (MSigDB) (Liberzon et al., 2015). Through permutation tests, GSEA assigns an enrichment score (ES) to z reflecting its overrepresentation within the MSigDB gene set collection. We found that for many cell types, genes added by Marlene at day 2 greatly

378 overlapped with COVID-19 and SARS-CoV-2 related gene sets. For instance “Interferon Gamma
 379 Response”, which was identified as a SARS-CoV-2 antiviral response (Hilligan et al., 2023), was
 380 significantly enriched in dendritic cells (15 genes, FDR = $1e-6$). Similarly, “TNF-alpha Signaling
 381 via NF-kB”—a pathway involved in the immune response and inflammation (Hayden & Ghosh,
 382 2011)—was enriched in several cell types, as well as processes such as “Apoptosis” (cell death) and
 383 “p53 Pathway” (inhibits replication of infected cells) (Elmore, 2007; Harris & Levine, 2005). Other
 384 methods, while being enriched for relevant terms, showed a smaller gene overlap for these types
 385 (Figure 3b) or were not consistent across cell types (e.g., DeepSEM, SCODE).

386 Overall, these results suggest that Marlene is able to capture both known TF-gene links, but also
 387 genes that are relevant to the response being studied.

3.2 CASE STUDY 2: AGING AND SENESENCE IN THE LUNG

391 The Human Lung Cell Atlas (HLCA) is a large data integration effort by the Human Cell Atlas
 392 Project (Sikkema et al., 2023; Regev et al., 2017). This data combines scRNA-seq samples from 107
 393 individuals spanning an age range of 10 to 76 years, making it particularly attractive for studying
 394 aging and senescence (a form of aging characterized by the absence of cell division) (van Deursen,
 395 2014; SenNet Consortium, 2022).

396 We split the atlas into three age groups at 35 and 50 years old, thus forming a pseudotime series
 397 of length 3. We removed smokers from the dataset as these will likely confound the results. To
 398 accommodate the data in the GPU, we randomly selected cells from 11 cell types, including type II
 399 pneumocytes, endothelial cells, and monocytes.

401 Similar to the vaccination dataset, we begin the analysis by evaluating the set of regulatory links
 402 using the TRRUST and RegNetwork databases. For this dataset we find that Marlene and SCODE
 403 are the top two performing methods (Figure 4a). For some of the cell types, Marlene achieves signif-
 404 icant results, recovering more than 1000 RegNetwork links (classical monocytes, FDR = $1e-76$).
 405 Even for cell types with fewer cells, such as non-classical monocytes (with only 138 cells for the
 406 second age group), Marlene still recovered more than 800 known TF-gene links for each transition
 407 (FDR $\leq 1e-27$). SCODE performed well for some cell types such as CD1c-positive myeloid den-
 408 dritic cells and CD4-positive, alpha-beta T cells. For all other methods, the overlap was smaller
 409 (Appendix Figure 6a). Note that while SCODE is comparable for the static network (single time
 409 point) inference task, it does not utilize dynamic information.

410 We next examined the ability of different methods to capture the dynamics of the biological pro-
 411 cesses. For this, we looked at graph transitions. IoU scores show that only the temporal methods
 412 (Marlene, TVGL, Graphs4mer) capture the smooth temporal transition between time points, while
 413 other methods, including SCODE, achieve low IoU scores (Appendix Figure 6b). We performed
 414 GSEA using Jensen Diseases gene set to see if genes added by Marlene in these transitions were
 415 enriched for any age-related diseases (Pletscher-Frankild et al., 2015; Grissa et al., 2022). We found
 416 that Marlene added genes are enriched for several diseases such as arthritis, lung disease, and coro-
 417 nary artery disease. Other dynamic baselines were also enriched for relevant terms, but contained
 418 fewer marker genes (Figure 4b).

419 Finally, we also investigated whether the genes regulated at different age groups were enriched for
 420 senescence. Cellular senescence refers to a permanent arrest of cell division triggered by the accu-
 421 mulation of DNA damage (Suryadevara et al., 2024). The absence of cell division can detrimentally
 422 impact tissue regeneration and repair, thereby contributing to various age-related diseases. Here, we
 423 use the SenMayo gene set which contains 125 genes reported to be enriched for senescence (Saul
 424 et al., 2022). Only 81 of these genes overlapped with our data. We found that for 4 cell types, there
 425 was an increase in SenMayo gene regulation at the oldest age group (age > 50), suggesting that
 426 senescent cells accumulate with age as hypothesized (Figure 5).

3.3 CASE STUDY 3: FIBROSIS IN A MOUSE LUNG INJURY MODEL

428 Next, we evaluated whether Marlene could perform effectively across different species by analyzing
 429 a dataset from a mouse model of lung injury induced by the chemotherapeutic agent bleomycin
 430 (Strunz et al., 2020). The dataset included seven time points: one pre-treatment and six post-
 431

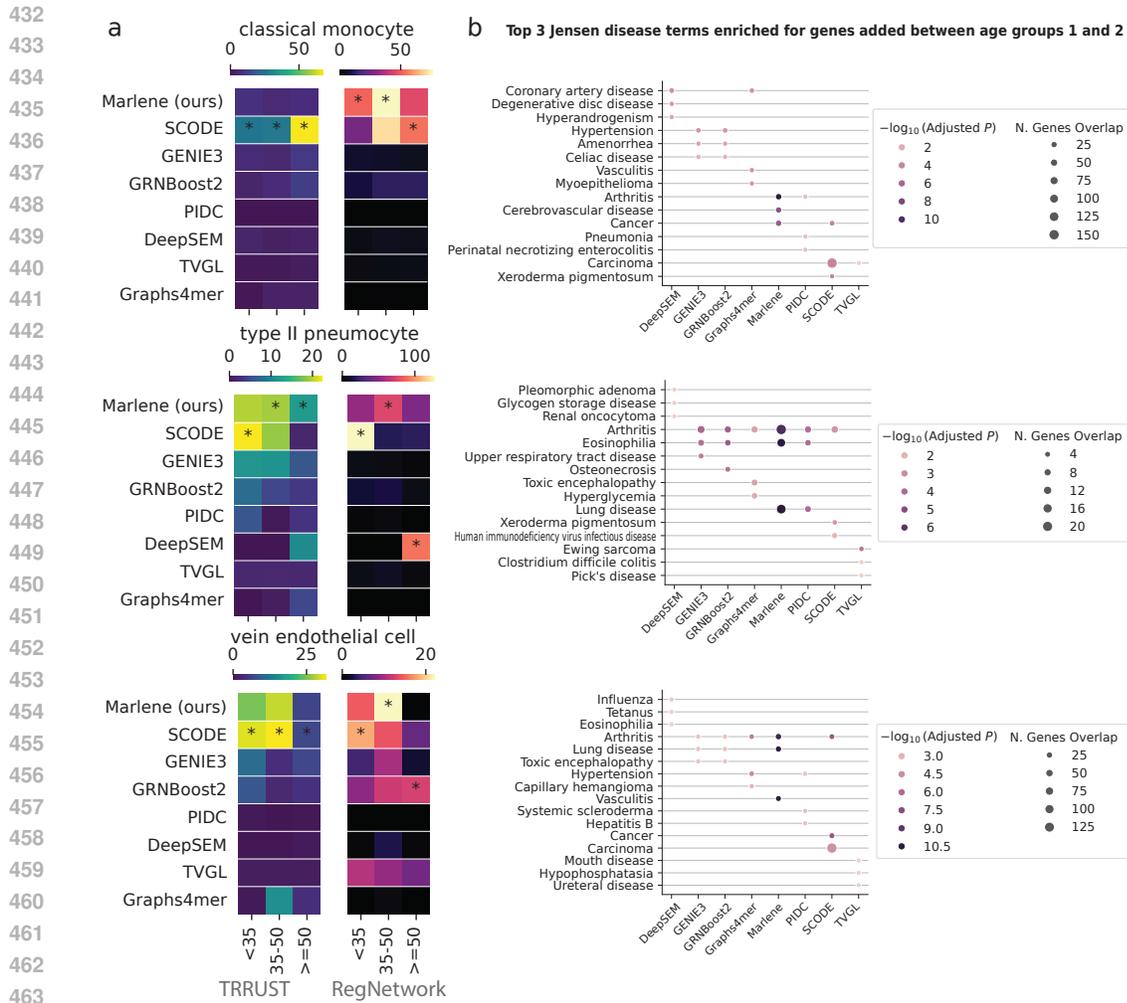


Figure 4: Results on the HLCA dataset. (a) FDR corrected p -values of Fisher exact tests reflecting the number of links that overlap with TRRUST and RegNetwork databases. (b) Top 3 Jensen Diseases terms enriched for genes added between the first and second age group.

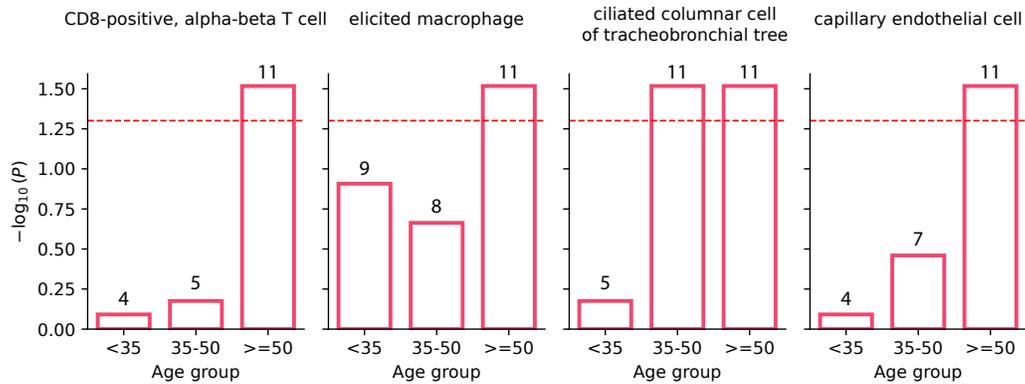


Figure 5: Enrichment for senescence using the SenMayo set. For 4 cell types, there was statistically significant enrichment for the oldest age group. We only used the top 200 regulated genes.

486 treatment intervals. After filtering out cell types with low representation and genes with low counts,
487 we retained six cell types, including B cells, T cells, and macrophages.

488
489 In this analysis, Marlene outperformed competing methods in four of the six cell types when bench-
490 marked against the RegNetwork database, specifically in alveolar epithelial cells, dendritic cells,
491 endothelial cells, and macrophages. For T cells, TVGL showed slightly better performance. When
492 evaluated against the TRRUST database, SCODE performed well in four cell types, while Marlene
493 surpassed it in the remaining two. The differing results between two databases may reflect their
494 incomplete coverage, highlighting the need for further refinement.

495 Finally, all static baselines, including SCODE, showed low IoU scores across time points, indi-
496 cating their inability to capture temporal evolution. In contrast, Marlene, showed increasing IoU
497 scores over time, suggesting ongoing lung regeneration following treatment which slowly stabilizes.
498 Figures illustrating these findings are provided in the appendix.

500 4 DISCUSSION

501 Gene regulation is a dynamic process that underlies all biological systems. Understanding which
502 TFs regulate which genes, and when this regulation occurs, provides insights into these dynamic
503 processes which can lead to better treatment options. For instance, understanding what TF-gene
504 links are disrupted could help researchers discover drugs targets for specific TF-gene connections.
505

506 To improve on current methods for reconstructing time varying regulatory networks, we use the ex-
507 pressive capabilities of deep neural networks to model the dynamic regulation of genes. Specifically,
508 we focused on inferring dynamic networks from scRNA-seq data.

509 Our proposed method, Marlene, constructs dynamic graphs from time series data. Marlene begins
510 with a set pooling operator based on PMA to extract gene features. These gene features are then
511 used to construct dynamic graphs via a self-attention mechanism. The weights of the self-attention
512 block are updated through the use of GRUs. Additionally, by employing MAML, we help Marlene
513 uncover graphs even for rare cell types. However, Marlene optimizes the prediction of cell type label
514 rather than gene expression. As such, Marlene is not currently equipped to determine the impact of
515 perturbations including gene knockouts or overexpression experiments. Exploring the integration of
516 causal inference capabilities into Marlene represents a promising direction for future research.

517 We demonstrated the effectiveness of Marlene in recovering dynamic GRNs using three datasets: a
518 SARS-CoV-2 vaccination dataset, a lung aging atlas, and a mouse dataset of fibrosis. In all three
519 datasets, Marlene successfully identified many validated TF-gene links from the TRRUST and Reg-
520 Network databases across various cell types. It also accurately modeled the temporal dynamics of
521 these connections. Some prior methods ignored the temporal aspect, leading to little similarity be-
522 tween consecutive networks. Other methods integrated all time points together, leading to very sim-
523 ilar networks for each time point. In contrast, Marlene accurately recovered the variation dynamics,
524 which is often characterized by strong rewiring following treatment that later stabilizes. In addition,
525 Marlene identified many relevant edges. For instance, in the lung aging data, several dynamic edges
526 were enriched for age-related diseases, such as arthritis. Meanwhile, in the SARS-CoV-2 data, these
527 dynamic links were enriched for immune response processes. Prior methods captured some known
528 edges, however, the overall results were less significant. By providing better models to explain dis-
529 ease and vaccine response, researchers can zoom in on the specific mechanisms targeted which in
530 turn can lead to better treatments. Code will be made publicly available on publication.

531 5 LIMITATIONS

532
533 While successful, Marlene has a few limitations. The datasets we used in this study, while typical for
534 scRNA-seq time series, consisted of only a few time points. For longer sequences, the GRU opera-
535 tion may suffer from vanishing gradient problems (Pascanu et al., 2012). In such scenarios, the S4
536 module may be preferred as it has been shown to model long sequences better than traditional GRUs
537 (Gu et al., 2021). In addition, using a large number of genes for training, results in quadratic growth
538 in memory consumption due to the need to store adjacency matrices. This led us to restrict the set of
539 genes for each of the two studies. A more efficient implementation or alternative approaches such
as FlashAttention (Dao et al., 2022) can lead to better ability to utilize all genes profiled.

REFERENCES

- 540
541
542 Amr Ahmed and Eric P. Xing. Recovering time-varying networks of dependencies in social and
543 biological studies. *Proceedings of the National Academy of Sciences*, 106(29):11878–11883,
544 July 2009.
- 545 Christof Angermueller, Tanel Pärnamaa, Leopold Parts, and Oliver Stegle. Deep learning for com-
546 putational biology. *Mol. Syst. Biol.*, 12(7):878, July 2016.
- 547
548 Pau Badia-I-Mompel, Lorna Wessels, Sophia Müller-Dott, Rémi Trimbour, Ricardo O Ramirez Flo-
549 res, Ricard Argelaguet, and Julio Saez-Rodriguez. Gene regulatory network inference in the era
550 of single-cell multi-omics. *Nat. Rev. Genet.*, 24(11):739–754, November 2023.
- 551 Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly
552 learning to align and translate. *arXiv [cs.CL]*, September 2014.
- 553
554 Ziv Bar-Joseph, Georg K. Gerber, Tong Ihn Lee, Nicola J. Rinaldi, Jane Y. Yoo, François Robert,
555 D. Benjamin Gordon, Ernest Fraenkel, Tommi S. Jaakkola, Richard A. Young, and David K. Gif-
556 ford. Computational discovery of gene modules and regulatory networks. *Nature Biotechnology*,
557 21(11):1337–1342, November 2003.
- 558 Ziv Bar-Joseph, Anthony Gitter, and Itamar Simon. Studying and modelling dynamic biological
559 processes using time-series gene expression data. *Nat. Rev. Genet.*, 13(8):552–564, July 2012.
- 560
561 Y Benjamini and Y Hochberg. Controlling the false discovery rate: a practical and powerful ap-
562 proach to multiple testing. *Journal of the royal statistical society series b-methodological*, 57(1):
563 289–300, 1995.
- 564 Alexandre Blais and Brian David Dynlacht. Constructing transcriptional regulatory networks. *Genes
565 & Development*, 19(13):1499–1511, July 2005.
- 566
567 Cătălina Cangea, Petar Veličković, Nikola Jovanović, Thomas Kipf, and Pietro Liò. Towards sparse
568 hierarchical graph classifiers. *arXiv [stat.ML]*, November 2018.
- 569
570 Thalia E Chan, Michael P H Stumpf, and Ann C Babbie. Gene regulatory network inference from
571 single-cell data using multivariate information measures. *Cell Syst.*, 5(3):251–267.e3, September
572 2017.
- 573 Deborah Chasman and Sushmita Roy. Inference of cell type specific regulatory networks on mam-
574 malian lineages. *Curr. Opin. Syst. Biol.*, 2:130–139, April 2017.
- 575
576 Travers Ching, Daniel S Himmelstein, Brett K Beaulieu-Jones, Alexandr A Kalinin, Brian T Do,
577 Gregory P Way, Enrico Ferrero, Paul-Michael Agapow, Michael Zietz, Michael M Hoffman, Wei
578 Xie, Gail L Rosen, Benjamin J Lengerich, Johnny Israeli, Jack Lanchantin, Stephen Woloszynek,
579 Anne E Carpenter, Avanti Shrikumar, Jinbo Xu, Evan M Cofer, Christopher A Lavender, Srinivas
580 C Turaga, Amr M Alexandari, Zhiyong Lu, David J Harris, Dave DeCaprio, Yanjun Qi,
581 Anshul Kundaje, Yifan Peng, Laura K Wiley, Marwin H S Segler, Simina M Boca, S Joshua
582 Swamidass, Austin Huang, Anthony Gitter, and Casey S Greene. Opportunities and obstacles for
583 deep learning in biology and medicine. *J. R. Soc. Interface*, 15(141), April 2018.
- 584 Tri Dao, Daniel Y Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and
585 memory-efficient exact attention with IO-awareness. *arXiv [cs.LG]*, May 2022.
- 586
587 Jun Ding, Nadav Sharon, and Ziv Bar-Joseph. Temporal modelling using single-cell transcriptomics.
588 *Nature Reviews Genetics*, 23(6):355–368, June 2022.
- 589 Frank Dondelinger, Sophie Lèbre, and Dirk Husmeier. Non-homogeneous dynamic Bayesian net-
590 works with Bayesian regularization for inferring gene regulatory networks with gradually time-
591 varying structure. *Machine Learning*, 90(2):191–230, February 2013.
- 592
593 Susan Elmore. Apoptosis: a review of programmed cell death. *Toxicol. Pathol.*, 35(4):495–516,
June 2007.

- 594 Zhuoqing Fang, Xinyuan Liu, and Gary Peltz. GSEApY: a comprehensive package for performing
595 gene set enrichment analysis in python. *Bioinformatics*, 39(1):btac757, January 2023.
596
- 597 Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation
598 of deep networks. *arXiv [cs.LG]*, pp. 1126–1135, March 2017.
- 599 R A Fisher. On the interpretation of χ^2 from contingency tables, and the calculation of P. *J. R. Stat.*
600 *Soc.*, 85(1):87, January 1922.
601
- 602 A P Gasch, P T Spellman, C M Kao, O Carmel-Harel, M B Eisen, G Storz, D Botstein, and P O
603 Brown. Genomic expression programs in the response of yeast cells to environmental changes.
604 *Mol. Biol. Cell*, 11(12):4241–4257, December 2000.
- 605 Daniel A. Gilchrist, David C. Fargo, and Karen Adelman. Using ChIP-chip and ChIP-seq to study
606 the regulation of gene expression: Genome-wide localization studies reveal widespread regulation
607 of transcription elongation. *Methods*, 48(4):398–408, August 2009.
608
- 609 Dhouha Grissa, Alexander Junge, Tudor I Oprea, and Lars Juhl Jensen. Diseases 2.0: a weekly
610 updated database of disease-gene associations from text mining and data integration. *Database*
611 *(Oxford)*, 2022:baac019, March 2022.
- 612 Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured
613 state spaces. *arXiv [cs.LG]*, October 2021.
614
- 615 David Hallac, Youngsuk Park, Stephen Boyd, and Jure Leskovec. Network inference via the time-
616 varying graphical lasso. *KDD*, 2017:205–213, August 2017.
- 617 Heonjong Han, Jae-Won Cho, Sangyoung Lee, Ayoung Yun, Hyojin Kim, Dasom Bae, Sunmo Yang,
618 Chan Yeong Kim, Muyoung Lee, Eunbeen Kim, Sungcho Lee, Byunghee Kang, Dabin Jeong,
619 Yaeji Kim, Hyeon-Nae Jeon, Haein Jung, Sunhwee Nam, Michael Chung, Jong-Hoon Kim, and
620 Insuk Lee. TRRUST v2: an expanded reference database of human and mouse transcriptional
621 regulatory interactions. *Nucleic Acids Res.*, 46(D1):D380–D386, January 2018.
622
- 623 Sandra L Harris and Arnold J Levine. The p53 pathway: positive and negative feedback loops.
624 *Oncogene*, 24(17):2899–2908, April 2005.
- 625 Matthew S Hayden and Sankar Ghosh. NF-kb in immunobiology. *Cell Res.*, 21(2):223–244, Febru-
626 ary 2011.
627
- 628 Kerry L Hilligan, Sivaranjani Namasivayam, Chad S Clancy, Paul J Baker, Samuel I Old, Victo-
629 ria Peluf, Eduardo P Amaral, Sandra D Oland, Danielle O’Mard, Julie Laux, Melanie Cohen,
630 Nicole L Garza, Bernard A P Lafont, Reed F Johnson, Carl G Feng, Dragana Jankovic, Olivier
631 Lamiable, Katrin D Mayer-Barber, and Alan Sher. Bacterial-induced or passively administered
632 interferon gamma conditions the lung for early control of SARS-CoV-2. *Nat. Commun.*, 14(1):
633 8229, December 2023.
- 634 Vân Anh Huynh-Thu, Alexandre Irrthum, Louis Wehenkel, and Pierre Geurts. Inferring regulatory
635 networks from expression data using tree-based methods. *PLoS One*, 5(9):e12776, September
636 2010.
- 637 Aashi Jindal, Prashant Gupta, Jayadeva, and Debarka Sengupta. Discovery of rare cells from volu-
638 minous single cell expression data. *Nat. Commun.*, 9(1):4719, November 2018.
639
- 640 Kenji Kamimoto, Blerta Stringa, Christy M Hoffmann, Kunal Jindal, Lilianna Solnica-Krezel, and
641 Samantha A Morris. Dissecting cell identity via network inference and in silico gene perturbation.
642 *Nature*, 614(7949):742–751, February 2023.
- 643 Guy Karlebach and Ron Shamir. Modelling and analysis of gene regulatory networks. *Nat. Rev.*
644 *Mol. Cell Biol.*, 9(10):770–780, October 2008.
645
- 646 Man-Sun Kim, Jeong-Rae Kim, Dongsan Kim, Arthur D. Lander, and Kwang-Hyun Cho. Spa-
647 tiotemporal network motif reveals the biological traits of developmental gene regulatory networks
in *Drosophila melanogaster*. *BMC Systems Biology*, 6(1):31, May 2012.

- 648 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv [cs.LG]*,
649 December 2014.
- 650
- 651 Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosior, Seungjin Choi, and Yee Whye Teh. Set trans-
652 former: A framework for attention-based permutation-invariant neural networks. In Kamalika
653 Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on*
654 *Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 3744–3753.
655 PMLR, 2019.
- 656 Tong Ihn Lee, Nicola J. Rinaldi, François Robert, Duncan T. Odom, Ziv Bar-Joseph, Georg K.
657 Gerber, Nancy M. Hannett, Christopher T. Harbison, Craig M. Thompson, Itamar Simon, Julia
658 Zeitlinger, Ezra G. Jennings, Heather L. Murray, D. Benjamin Gordon, Bing Ren, John J. Wyrick,
659 Jean-Bosco Tagne, Thomas L. Volkert, Ernest Fraenkel, David K. Gifford, and Richard A. Young.
660 Transcriptional Regulatory Networks in *Saccharomyces cerevisiae*. *Science*, 298(5594):799–804,
661 October 2002.
- 662 Arthur Liberzon, Chet Birger, Helga Thorvaldsdóttir, Mahmoud Ghandi, Jill P Mesirov, and Pablo
663 Tamayo. The molecular signatures database (MSigDB) hallmark gene set collection. *Cell Syst.*,
664 1(6):417–425, December 2015.
- 665
- 666 Zhi-Ping Liu, Canglin Wu, Hongyu Miao, and Hulin Wu. RegNetwork: an integrated database
667 of transcriptional and post-transcriptional regulatory networks in human and mouse. *Database*
668 (*Oxford*), 2015:bav095, September 2015.
- 669
- 670 Nicholas M Luscombe, M Madan Babu, Haiyuan Yu, Michael Snyder, Sarah A Teichmann, and
671 Mark Gerstein. Genomic analysis of regulatory network dynamics reveals large topological
672 changes. *Nature*, 431(7006):308–312, September 2004.
- 673 Hiroataka Matsumoto, Hisanori Kiryu, Chikara Furusawa, Minoru S H Ko, Shigeru B H Ko, Norio
674 Gouda, Tetsutaro Hayashi, and Itoshi Nikaido. SCODE: an efficient regulatory network inference
675 algorithm from single-cell RNA-seq during differentiation. *Bioinformatics*, 33(15):2314–2321,
676 August 2017.
- 677
- 678 Thomas Moerman, Sara Aibar Santos, Carmen Bravo González-Blas, Jaak Simm, Yves Moreau,
679 Jan Aerts, and Stein Aerts. GRNBoost2 and arboreto: efficient and scalable inference of gene
680 regulatory networks. *Bioinformatics*, 35(12):2159–2161, June 2019.
- 681 Hung Nguyen, Duc Tran, Bang Tran, Bahadir Pehlivan, and Tin Nguyen. A comprehensive survey
682 of regulatory network inference methods using single cell RNA sequencing data. *Briefings in*
683 *Bioinformatics*, 22(3):bbaa190, May 2021.
- 684
- 685 Aldo Pareja, Giacomo Domeniconi, Jie Chen, Tengfei Ma, Toyotaro Suzumura, Hiroki Kaneza-
686 shi, Tim Kaler, Tao Schardl, and Charles Leiserson. EvolveGCN: Evolving graph convolutional
687 networks for dynamic graphs. *Proc. Conf. AAAI Artif. Intell.*, 34(04):5363–5370, April 2020.
- 688 Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural
689 networks. *arXiv [cs.LG]*, November 2012.
- 690
- 691 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
692 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Ed-
693 ward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,
694 Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance
695 deep learning library. *arXiv [cs.LG]*, December 2019.
- 696
- 697 Sune Pletscher-Frankild, Albert Pallejà, Kalliopi Tsafou, Janos X Binder, and Lars Juhl Jensen.
698 DISEASES: text mining and data integration of disease-gene associations. *Methods*, 74:83–89,
699 March 2015.
- 700 Aditya Pratapa, Amogh P Jalihal, Jeffrey N Law, Aditya Bharadwaj, and T M Murali. Benchmark-
701 ing algorithms for gene regulatory network inference from single-cell transcriptomic data. *Nat.*
Methods, 17(2):147–154, February 2020.

- 702 Aviv Regev, Sarah A Teichmann, Eric S Lander, Ido Amit, Christophe Benoist, Ewan Birney,
703 Bernd Bodenmiller, Peter Campbell, Piero Carninci, Menna Clatworthy, Hans Clevers, Bart De-
704 plancke, Ian Dunham, James Eberwine, Roland Eils, Wolfgang Enard, Andrew Farmer, Lars
705 Fugger, Berthold Göttgens, Nir Hacohen, Muzlifah Haniffa, Martin Hemberg, Seung Kim, Paul
706 Klenerman, Arnold Kriegstein, Ed Lein, Sten Linnarsson, Emma Lundberg, Joakim Lunde-
707 berg, Partha Majumder, John C Marioni, Miriam Merad, Musa Mhlanga, Martijn Nawijn, Mihai
708 Netea, Garry Nolan, Dana Pe'er, Anthony Phillipakis, Chris P Ponting, Stephen Quake, Wolf
709 Reik, Orit Rozenblatt-Rosen, Joshua Sanes, Rahul Satija, Ton N Schumacher, Alex Shalek, Ehud
710 Shapiro, Padmanee Sharma, Jay W Shin, Oliver Stegle, Michael Stratton, Michael J T Stubb-
711 ington, Fabian J Theis, Matthias Uhlen, Alexander van Oudenaarden, Allon Wagner, Fiona Watt,
712 Jonathan Weissman, Barbara Wold, Ramnik Xavier, Nir Yosef, and Human Cell Atlas Meeting
713 Participants. The human cell atlas. *Elife*, 6, December 2017.
- 714 Dominik Saul, Robyn Laura Kosinsky, Elizabeth J Atkinson, Madison L Doolittle, Xu Zhang,
715 Nathan K LeBrasseur, Robert J Pignolo, Paul D Robbins, Laura J Niedernhofer, Yuji Ikeno, Di-
716 ana Jurk, João F Passos, Latonya J Hickson, Ailing Xue, David G Monroe, Tamara Tchkonja,
717 James L Kirkland, Joshua N Farr, and Sundeep Khosla. A new gene set identifies senescent cells
718 and predicts senescence-associated pathways across tissues. *Nat. Commun.*, 13(1):4827, August
719 2022.
- 720 Marcel H. Schulz, William E. Devanny, Anthony Gitter, Shan Zhong, Jason Ernst, and Ziv Bar-
721 Joseph. DREM 2.0: Improved reconstruction of dynamic regulatory networks from time-series
722 expression data. *BMC Systems Biology*, 6(1):104, August 2012.
- 723 Lucas Seninge, Ioannis Anastopoulos, Hongxu Ding, and Joshua Stuart. VEGA is an interpretable
724 generative model for inferring biological network activity in single-cell transcriptomics. *Nat.*
725 *Commun.*, 12(1):5684, September 2021.
- 727 SenNet Consortium. NIH SenNet consortium to map senescent cells throughout the human lifespan
728 to understand physiological health. *Nat Aging*, 2(12):1090–1100, December 2022.
- 729 Harsh Shrivastava, Xiuwei Zhang, Le Song, and Srinivas Aluru. GRNUlar: A deep learning frame-
730 work for recovering single-cell gene regulatory networks. *J. Comput. Biol.*, 29(1):27–44, January
731 2022.
- 732 Hantao Shu, Jingtian Zhou, Qiuyu Lian, Han Li, Dan Zhao, Jianyang Zeng, and Jianzhu Ma. Mod-
733 eling gene regulatory networks using neural network architectures. *Nat. Comput. Sci.*, 1(7):491–
734 501, July 2021.
- 735 Lisa Sikkema, Ciro Ramírez-Suástegui, Daniel C Strobl, Tessa E Gillett, Luke Zappia, Elo Madis-
736 soon, Nikolay S Markov, Laure-Emmanuelle Zaragosi, Yuge Ji, Meshal Ansari, Marie-Jeanne
737 Arguel, Leonie Apperloo, Martin Banchemo, Christophe Bécavin, Marijn Berg, Evgeny Chichel-
738 nitskiy, Mei-I Chung, Antoine Collin, Aurore C A Gay, Janine Gote-Schniering, Baharak
739 Hooshir Kashani, Kemal Inecik, Manu Jain, Theodore S Kapellos, Tessa M Kole, Sylvie Leroy,
740 Christoph H Mayr, Amanda J Oliver, Michael von Papen, Lance Peter, Chase J Taylor, Thomas
741 Walzthoeni, Chuan Xu, Linh T Bui, Carlo De Donno, Leander Dony, Alen Faiz, Minzhe Guo,
742 Austin J Gutierrez, Lukas Heumos, Ni Huang, Ignacio L Ibarra, Nathan D Jackson, Preetish
743 Katur Lakshminarasimha Murthy, Mohammad Lotfollahi, Tracy Tabib, Carlos Talavera-López,
744 Kyle J Travaglini, Anna Wilbrey-Clark, Kaylee B Worlock, Masahiro Yoshida, Lung Biological
745 Network Consortium, Maarten van den Berge, Yohan Bossé, Tushar J Desai, Oliver Eickelberg,
746 Naftali Kaminski, Mark A Krasnow, Robert Lafyatis, Marko Z Nikolic, Joseph E Powell, Jayaraj
747 Rajagopal, Mauricio Rojas, Orit Rozenblatt-Rosen, Max A Seibold, Dean Sheppard, Douglas P
748 Shepherd, Don D Sin, Wim Timens, Alexander M Tsankov, Jeffrey Whitsett, Yan Xu, Nicholas E
749 Banovich, Pascal Barbry, Thu Elizabeth Duong, Christine S Falk, Kerstin B Meyer, Jonathan A
750 Kropski, Dana Pe'er, Herbert B Schiller, Purushothama Rao Tata, Joachim L Schultze, Sara A
751 Teichmann, Alexander V Misharin, Martijn C Nawijn, Malte D Luecken, and Fabian J Theis. An
752 integrated cell atlas of the lung in health and disease. *Nat. Med.*, 29(6):1563–1577, June 2023.
- 753 Edwin K Silverman, Harald H H W Schmidt, Eleni Anastasiadou, Lucia Altucci, Marco Angelini,
754 Lina Badimon, Jean-Luc Balligand, Giuditta Benincasa, Giovambattista Capasso, Federica Conte,
755 Antonella Di Costanzo, Lorenzo Farina, Giulia Fiscon, Laurent Gatto, Michele Gentili, Joseph

- 756 Loscalzo, Cinzia Marchese, Claudio Napoli, Paola Paci, Manuela Petti, John Quackenbush, Paolo
757 Tieri, Davide Viggiano, Gemma Vilahur, Kimberly Glass, and Jan Baumbach. Molecular net-
758 works in network medicine: Development and applications. *Wiley Interdiscip. Rev. Syst. Biol.*
759 *Med.*, 12(6):e1489, November 2020.
- 760 Le Song, Mladen Kolar, and Eric Xing. Time-Varying Dynamic Bayesian Networks. In *Advances*
761 *in Neural Information Processing Systems*, volume 22. Curran Associates, Inc., 2009.
- 762 Maximilian Strunz, Lukas M Simon, Meshal Ansari, Jaymin J Kathiriya, Ilias Angelidis,
763 Christoph H Mayr, George Tsidiridis, Marius Lange, Laura F Mattner, Min Yee, Paulina Ogar,
764 Arunima Sengupta, Igor Kukhtevich, Robert Schneider, Zhongming Zhao, Carola Voss, To-
765 bias Stoeger, Jens H L Neumann, Anne Hilgendorff, Jürgen Behr, Michael O’Reilly, Mareike
766 Lehmann, Gerald Burgstaller, Melanie Königshoff, Harold A Chapman, Fabian J Theis, and Her-
767 bert B Schiller. Alveolar regeneration through a Krt8+ transitional stem cell state that persists in
768 human lung fibrosis. *Nat. Commun.*, 11(1):3559, July 2020.
- 769 Aravind Subramanian, Pablo Tamayo, Vamsi K Mootha, Sayan Mukherjee, Benjamin L Ebert,
770 Michael A Gillette, Amanda Paulovich, Scott L Pomeroy, Todd R Golub, Eric S Lander, and Jill P
771 Mesirov. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-
772 wide expression profiles. *Proc. Natl. Acad. Sci. U. S. A.*, 102(43):15545–15550, October 2005.
- 773 Vidyani Suryadevara, Adam D Hudgins, Adarsh Rajesh, Alberto Pappalardo, Alla Karpova, Amit K
774 Dey, Ann Hertz, Anthony Agudelo, Azucena Rocha, Bikem Soygur, Birgit Schilling, Chase M
775 Carver, Cristina Aguayo-Mazzucato, Darren J Baker, David A Bernlohr, Diana Jurk, Dilyana B
776 Mangarova, Ellen M Quardokus, Elizabeth Ann L Enninga, Elizabeth L Schmidt, Feng Chen,
777 Francesca E Duncan, Francesco Cambuli, Gagandeep Kaur, George A Kuchel, Gung Lee, Heike E
778 Daldrup-Link, Helene Martini, Hemali Phatnani, Iman M Al-Naggar, Irfan Rahman, Jia Nie,
779 João F Passos, Jonathan C Silverstein, Judith Campisi, Julia Wang, Kanako Iwasaki, Karina Bar-
780 bosa, Kay Metis, Kerem Nernekli, Laura J Niedernhofer, Li Ding, Lichao Wang, Lisa C Adams,
781 Liu Ruiyang, Madison L Doolittle, Marcos G Teneche, Marissa J Schafer, Ming Xu, Mohammad-
782 javad Hajipour, Mozhgan Boroumand, Nathan Basisty, Nicholas Sloan, Nikolai Slavov, Olena
783 Kuksenko, Paul Robson, Paul T Gomez, Periklis Vasilikos, Peter D Adams, Priscila Carapeto,
784 Quan Zhu, Ramalakshmi Ramasamy, Rolando Perez-Lorenzo, Rong Fan, Runze Dong, Ruth R
785 Montgomery, Sadiya Shaikh, Sanja Vickovic, Shanshan Yin, Shoukai Kang, Sonja Suvakov, Sun-
786 deep Khosla, Vesna D Garovic, Vilas Menon, Yanxin Xu, Yizhe Song, Yousin Suh, Zhixun Dou,
787 and Nicola Neretti. SenNet recommendations for detecting senescent cells in different tissues.
788 *Nat. Rev. Mol. Cell Biol.*, pp. 1–23, June 2024.
- 789 Siyi Tang, Jared A Dunmon, Liangqiong Qu, Khaled K Saab, Tina Baykaner, Christopher Lee-
790 Messer, and Daniel L Rubin. Modeling multivariate biosignals with graph neural networks and
791 structured state space models. *arXiv [cs.LG]*, November 2022.
- 792 Jan M van Deursen. The role of senescent cells in ageing. *Nature*, 509(7501):439–446, May 2014.
- 793 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
794 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Informa-*
795 *tion Processing Systems*, 30, 2017.
- 796 Huihui Wang, Yongqing Wu, Ruiling Fang, Jian Sa, Zhi Li, Hongyan Cao, and Yuehua Cui. Time-
797 Varying Gene Network Analysis of Human Prefrontal Cortex Development. *Frontiers in Genetics*,
798 11, 2020.
- 799 Jiacheng Wang, Yaojia Chen, and Quan Zou. Inferring gene regulatory network from single-cell
800 transcriptomes with graph autoencoder model. *PLoS Genet.*, 19(9):e1010942, September 2023a.
- 801 Lingfei Wang, Nikolaos Trasanidis, Ting Wu, Guanlan Dong, Michael Hu, Daniel E Bauer, and
802 Luca Pinello. Dictys: dynamic gene regulatory network dissects developmental continuum with
803 single-cell multiomics. *Nat. Methods*, 20(9):1368–1378, September 2023b.
- 804 Yong Wang, Trupti Joshi, Xiang-Sun Zhang, Dong Xu, and Luonan Chen. Inferring gene regulatory
805 networks from multiple microarray datasets. *Bioinformatics*, 22(19):2413–2420, October 2006.

- 810 Zhong Wang, Mark Gerstein, and Michael Snyder. RNA-Seq: A revolutionary tool for transcrip-
811 tomics. *Nature Reviews Genetics*, 10(1):57–63, January 2009.
- 812
- 813 Aaron Wise and Ziv Bar-Joseph. cDREM: inferring dynamic combinatorial gene regulation. *J.*
814 *Comput. Biol.*, 22(4):324–333, April 2015.
- 815 Min Yan, Jing Hu, Huating Yuan, Liwen Xu, Gaoming Liao, Zedong Jiang, Jiali Zhu, Bo Pang,
816 Yanyan Ping, Yunpeng Zhang, Yun Xiao, and Xia Li. Dynamic regulatory networks of T cell
817 trajectory dissect transcriptional control of T cell state transition. *Molecular Therapy - Nucleic*
818 *Acids*, 26:1115–1129, December 2021.
- 819 Nir Yosef, Alex K. Shalek, Jellert T. Gaublomme, Hulin Jin, Youjin Lee, Amit Awasthi, Chuan
820 Wu, Katarzyna Karwacz, Sheng Xiao, Marsela Jorgolli, David Gennert, Rahul Satija, Arvind
821 Shakya, Diana Y. Lu, John J. Trombetta, Meenu R. Pillai, Peter J. Ratcliffe, Mathew L. Coleman,
822 Mark Bix, Dean Tantin, Hongkun Park, Vijay K. Kuchroo, and Aviv Regev. Dynamic regulatory
823 network controlling TH17 cell differentiation. *Nature*, 496(7446):461–468, April 2013.
- 824
- 825 Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and
826 Alexander J Smola. Deep sets. *Advances in Neural Information Processing Systems*, 30, 2017.
- 827
- 828 Bingjie Zhang, Rabi Upadhyay, Yuhan Hao, Marie I Samanovic, Ramin S Herati, John D Blair, Jor-
829 dan Axelrad, Mark J Mulligan, Dan R Littman, and Rahul Satija. Multimodal single-cell datasets
830 characterize antigen-specific CD8+ T cells across SARS-CoV-2 vaccination and infection. *Nat.*
831 *Immunol.*, 24(10):1725–1734, October 2023.
- 832 Qi Zhang, Jianlong Chang, Gaofeng Meng, Shiming Xiang, and Chunhong Pan. Spatio-temporal
833 graph structure learning for traffic forecasting. *Proc. Conf. AAAI Artif. Intell.*, 34(01):1177–1185,
834 April 2020.
- 835 Shijia Zhu and Yadong Wang. Hidden Markov induced Dynamic Bayesian Network for recovering
836 time evolving gene regulatory networks. *Scientific Reports*, 5(1):17841, December 2015.
- 837
- 838 Yanqiao Zhu, Weizhi Xu, Jinghao Zhang, Yuanqi Du, Jieyu Zhang, Qiang Liu, Carl Yang, and Shu
839 Wu. A survey on graph structure learning: Progress and opportunities. *arXiv [cs.LG]*, March
840 2021.
- 841
- 842
- 843
- 844
- 845
- 846
- 847
- 848
- 849
- 850
- 851
- 852
- 853
- 854
- 855
- 856
- 857
- 858
- 859
- 860
- 861
- 862
- 863

864 A APPENDIX

865 A.1 SET TRANSFORMER OPERATIONS

866 We redefine the Multihead and rFF operations from Set Transformers (Lee et al., 2019) to the ones
867 used for Marlene here.

868 First, we define the **Attention** operation. Let $Q \in \mathbb{R}^{k \times g}$ be the query matrix of k elements and g
869 dimensions. The Attention operation used for MAB is

$$870 \text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^\top}{\sqrt{g}} \right) V \quad (8)$$

871 where the key and value matrices are $K, V \in \mathbb{R}^{c \times g}$. Next, the Multihead attention operation with h
872 heads (Vaswani et al., 2017) is given by

$$873 \text{Multihead}(Q, K, V) = \text{concat}(O_1, \dots, O_h)W^O \quad (9)$$

874 where $O_j = \text{Attention}(QW_j^Q, KW_j^K, VW_j^V)$ for weight matrices $W_j^Q, W_j^K, W_j^V \in \mathbb{R}^{g \times g/h}$ and
875 $W^O \in \mathbb{R}^{g \times g}$ (these matrices are not to be confused with self-attention weights used consequently
876 for Marlene). In our implementation, k is the number of seeds or output vectors used for the PMA
877 layer. This is a hyperparameter that corresponds to the number of “statistic” vectors we expect to
878 learn from data. Finally, rFF is a feedforward layer such as a linear layer.

882 A.2 EVOLVEGCN OPERATIONS

883 Here, we introduce the GRU and topK pooling operations used in the second step of Marlene.

884 The topK pooling operation is needed to summarize nodes into k representative ones (Cangea et al.,
885 2018; Pareja et al., 2020). Here k is the same as the number of seeds used for PMA. Given an input
886 $\mathbf{G} \in \mathbb{R}^{g \times k}$ and a learnable vector q , the TopK operation performs the following steps:

$$887 \rho = \frac{\mathbf{G}q}{\|q\|}$$

$$888 i = \text{Top-k-indices}(\rho)$$

$$889 \mathbf{Z} = [\mathbf{G} \odot \tanh(\rho)]_i.$$

890 At time step t , given a pooled matrix \mathbf{Z}_t and hidden state \mathbf{W}_{t-1} (i.e., self-attention weights \mathbf{W}_{t-1}^Q
891 or \mathbf{W}_{t-1}^K), the standard GRU operation is:

$$892 r_t = \sigma(M_{ir}\mathbf{Z}_t + b_{ir} + M_{hr}\mathbf{W}_{t-1} + b_{hr})$$

$$893 z_t = \sigma(M_{iz}\mathbf{Z}_t + b_{iz} + M_{hz}\mathbf{W}_{t-1} + b_{hz})$$

$$894 n_t = \tanh(M_{in}\mathbf{Z}_t + b_{in} + r_t \odot (M_{hn}\mathbf{W}_{t-1} + b_{hn}))$$

$$895 \mathbf{W}_t = (1 - z_t) \odot n_t + z_t \odot \mathbf{W}_{t-1}$$

896 where σ is the sigmoid function and \odot is the Hadamard product. See also Paszke et al. (2019).

A.3 SUPPLEMENTARY FIGURE FOR THE HLCA DATASET

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

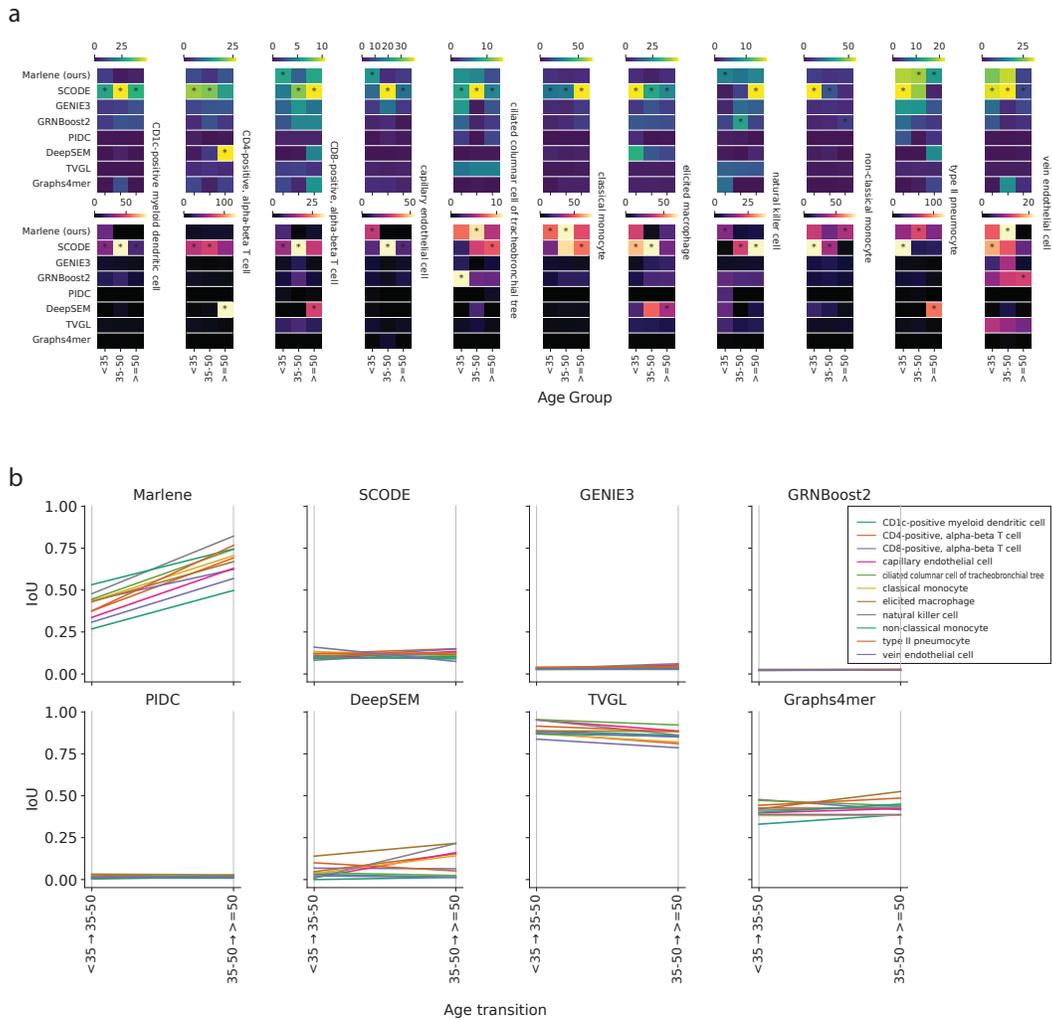


Figure 6: (a) FDR corrected p -values of Fisher exact tests reflecting the number of links that overlap with the two TF-gene databases. (b) IoU scores across time reflecting the overlap between consecutive graphs.

A.4 ANALYSIS OF THE FIBROSIS DATASET

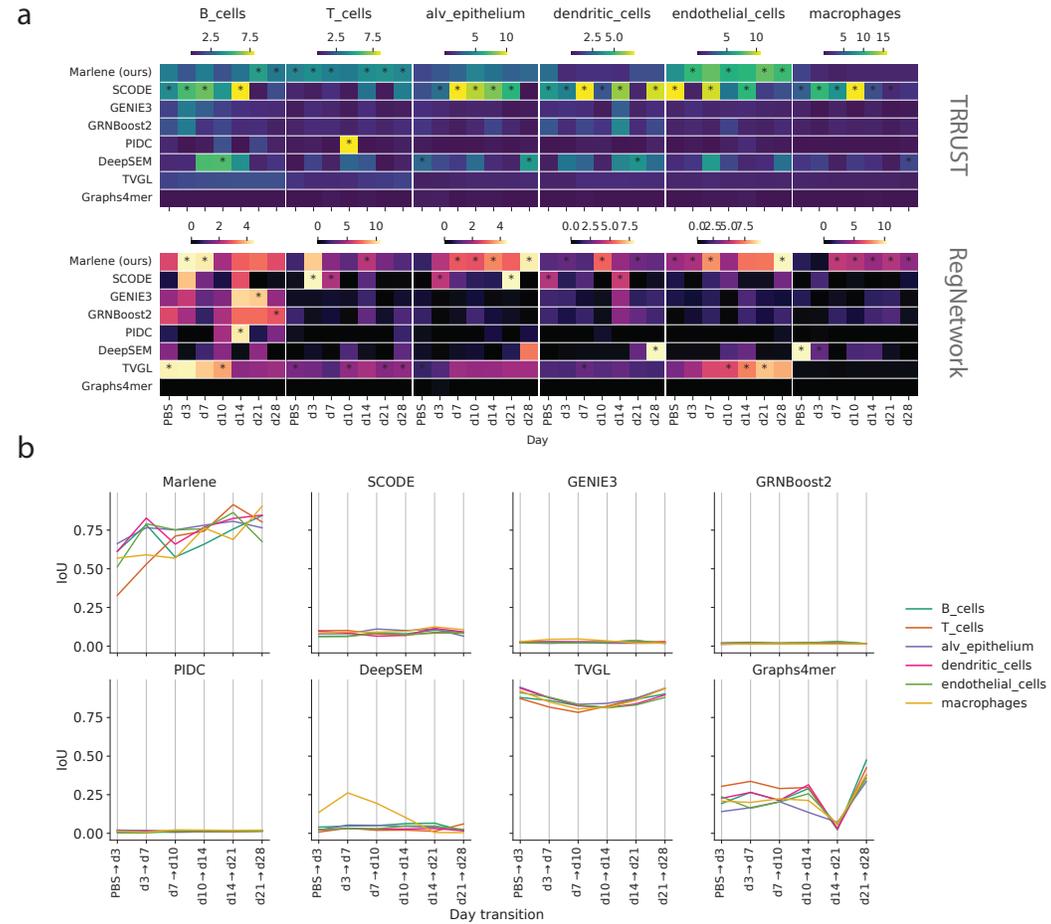


Figure 7: (a) FDR corrected p -values of Fisher exact tests reflecting the number of links that overlap with the two mouse databases. (b) IoU scores across time reflecting the overlap between consecutive graphs.