

SIMULATION-FREE STRUCTURE LEARNING FOR STOCHASTIC DYNAMICS

Anonymous authors

Paper under double-blind review

ABSTRACT

Modeling dynamical systems and unraveling their underlying causal relationships is central to many domains in the natural sciences. Various physical systems, such as those arising in cell biology, are inherently high-dimensional and stochastic in nature, and admit only partial, noisy state measurements. This poses a significant challenge for addressing the problems of modeling the underlying dynamics and inferring the network structure of these systems. Existing methods are typically tailored either for structure learning or modeling dynamics at the population level, but are limited in their ability to address both problems together. In this work, we address both problems simultaneously: we present STRUCTUREFLOW, a novel and principled simulation-free approach for jointly learning the structure and stochastic population dynamics of physical systems. We showcase the utility of STRUCTUREFLOW for the tasks of structure learning from interventions and dynamical (trajectory) inference of conditional population dynamics. We empirically evaluate our approach on high-dimensional synthetic systems, a set of biologically plausible simulated systems, and an experimental single-cell dataset. We show that STRUCTUREFLOW can learn the structure of underlying systems while simultaneously modeling their conditional population dynamics — a key step toward the mechanistic understanding of systems behavior.




Figure 1: Overview of STRUCTUREFLOW for joint structure learning and dynamical inference.

1 INTRODUCTION

Unraveling the underlying structure of physical systems from their stochastic dynamics, given noisy and partial measurements, is a central problem in many areas of science. Many physical systems, notably in cell and molecular biology, operate in out-of-equilibrium regimes, are high-dimensional, and are subject to intrinsic stochasticity (Elowitz et al., 2002; Schiebinger et al., 2019). This poses a significant challenge for the task of deciphering the underlying system structure and modeling the resulting dynamics. Effectively addressing this problem is a crucial step toward gaining a mechanistic understanding of systems and the ability to predict their behavior under natural and perturbational conditions. This would, in turn, provide practitioners with a map for how to control and guide a system’s response towards desirable states (Dixit et al., 2016; Norman et al., 2019). In cellular and molecular biology, such a tool would have significant implications in advancing our ability to interpret and model cell fate, development, response to disease (Rizvi et al., 2017; Binnewies et al., 2018; Gulati et al., 2020; Molè et al., 2021), predict perturbational responses in tissues and patients (Ji et al., 2021; Bunne et al., 2023; Peidli et al., 2024; Atanackovic et al., 2024), and facilitate experimental design (Zhang et al., 2023; Huang et al., 2024).

Currently, there exist two broad classes of approaches for deciphering the stochastic dynamics of physical systems from *data*. The first class, which we refer to as **dynamical inference**¹, commonly

¹This task is also commonly labeled and referred to as *trajectory inference* (Hashimoto et al., 2016; Weinreb et al., 2018; Schiebinger et al., 2019; Tong et al., 2020; Neklyudov et al., 2022). Since the concept of “trajectory”

deals with constructing an estimate of the underlying vector field from partial and noisy measurements. From a good estimate of the vector fields, qualitative properties of dynamics can be interpreted, valuable for understanding cell dynamics and cell fate (Qiu et al., 2022; MacArthur, 2022). In this work, we focus on the preceding task of learning good estimates of vector fields from partially observed population snapshot data.

The second (and in our work, concurrent) class, **structure learning**², aims to reconstruct the *directional* relationships between each of the variables of a given system. Existing methods are typically designed to tackle either task in isolation and are limited in their ability to address both problems together (Zheng et al., 2018; Huynh-Thu and Geurts, 2018; Gao et al., 2018; Brouillard et al., 2020; Atanackovic et al., 2023; Zhang, 2024). We posit that knowledge of the underlying vector field is advantageous for this task, and thus we address both problems simultaneously under a single-model.

To achieve this, we propose STRUCTUREFLOW, a principled and *simulation-free* method for joint dynamical inference and structure learning. Building on recent advances in score and flow matching ([SF]²M) (Tong et al., 2024) and entropy-regularized optimal transport (EOT) (Cuturi, 2013; Shi et al., 2023), STRUCTUREFLOW learns a probability flow ordinary differential equation (PF-ODE) from snapshot data. This allows it to model the continuous evolution of a system’s population while simultaneously inferring the underlying network structure embedded directly within the model’s parameterization.

The simulation-free framing of STRUCTUREFLOW is critical for high-dimensional systems common in cellular biology, as it avoids computationally expensive trajectory simulations during training. To best facilitate the joint learning task, we introduce a novel parameterization: an autonomous (time-independent) vector field represented by a Neural Graphical Model (NGM) (Bellot et al., 2022; Dinh and Ho, 2020) which captures the stationary system structure, while a time-dependent score function captures the evolving stochastic dynamics. We present an overview of our framework in Figure 1 and outline our core contributions below:

- We formulate joint dynamical inference and structure learning as a **multi-timepoint** Schrödinger Bridge (SB) problem.
- We introduce STRUCTUREFLOW, a novel simulation-free approach tailored for simultaneously learning the underlying network structure and conditional population dynamics of a stochastic dynamical system from noisy and partial observations.
- We construct a comprehensive empirical evaluation for the joint inference task over an assortment of systems: (i) on high-dimensional synthetic systems, (ii) an collection of biologically plausible simulated systems, and (iii) an experimental single cell dataset with genetic interventions. We use our evaluation pipeline to showcase the application of STRUCTUREFLOW for the joint structure learning and dynamical inference tasks, all while building an extensive benchmark of state-of-the-art methods.

2 BACKGROUND AND PRELIMINARIES

2.1 PROBLEM: MODELING STOCHASTIC POPULATION DYNAMICS

We consider a dynamical model in \mathbb{R}^d described by the stochastic differential equation (SDE):

$$d\mathbf{x}_t = \mathbf{v}_t(\mathbf{x}_t) dt + \boldsymbol{\sigma} d\mathbf{B}_t, \quad (1)$$

where $\mathbf{x}_t \in \mathbb{R}^d$ is the state at time t , $\mathbf{v}_t : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the drift, $d\mathbf{B}_t$ are Brownian motion increments, and $\boldsymbol{\sigma} \in \mathbb{R}^{d \times d}$ is the diffusion coefficient matrix. For simplicity, we shall assume that $\boldsymbol{\sigma} = \sigma \mathbf{I}$ for some $\sigma > 0$, although the concepts we introduce can all be generalized to the setting of anisotropic noise.

Given $\mathbf{x}_0 \sim p_0$, where p_0 is a density over \mathbb{R}^d , the dynamics of equation 1 gives rise to a family of *marginals* $(p_t)_{t \in [0, 1]}$, where p_t denotes the marginal probability distribution of the random variable \mathbf{x}_t at time t , which are characterized by the accompanying *Fokker-Planck equation* (FPE):

$$\partial_t p_t = -\nabla \cdot (p_t(\mathbf{x}_t) \mathbf{v}_t(\mathbf{x}_t)) + \frac{1}{2} \nabla \cdot (\sigma^2 \nabla p_t(\mathbf{x}_t)). \quad (2)$$

The solution to the FPE at time t , $p_t(\mathbf{x})$, gives the probability density of the population at time t , starting from $p_0(\mathbf{x})$. Importantly, equation 2 can be reformulated as an equivalent *probability flow* ordinary differential equation (ODE) (Song et al., 2021; Maoutsa et al., 2020):

$$\partial_t p_t(\mathbf{x}) = -\nabla \cdot (p_t(\mathbf{x}) \mathbf{u}_t(\mathbf{x})), \quad \mathbf{u}_t(\mathbf{x}) = \mathbf{v}_t(\mathbf{x}) - \frac{1}{2} \sigma^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}). \quad (3)$$

has been frequently used in the literature to refer to several non-equivalent things, in this work, we prefer to name this task *dynamical inference*.

²Also commonly referred to in literature as a *network inference*, *causal discovery*, and *system identification*.

In the above, \mathbf{u}_t is the *probability flow field* and is related to the drift of the SDE equation 1 up to the addition of a term involving the score, $\mathbf{s}_t(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$. As will be apparent in Section 2.2, the probability flow formulation of the FPE equation 2, and hence the SDE equation 1, opens up an avenue for addressing dynamics inference without the need for costly simulation-based methods.

Inference problem setup. We consider a setting where empirical snapshot observations are available at multiple timepoints t_0, \dots, t_{T-1} under different conditions $c \in \mathcal{C}$, yielding a set of empirical marginal distributions $\{\hat{p}_{t_i}^{(c)}\}_0^{T-1}$ over $\mathbf{x}_{t_i} \in \mathbb{R}^d$ for each condition c . Each marginal comprises N_i i.i.d. samples assumed to arise from some unknown SDE of the form equation 1. We assume the noise level σ is known, while the vector field \mathbf{v}_t is unknown³. From here, our objective is twofold: **(1; dynamical inference)** approximate the (*conditional*) vector field \mathbf{v}_t of the underlying stochastic dynamics, and **(2; structure learning)** recover a directed graph $\mathbf{A} \in \mathbb{R}^{d \times d}$ that captures dependencies among d variables and represents the underlying data generative process of the system. We formulate this joint inference task through the lens of Schrödinger Bridges.

We highlight orthogonal fields of work, such as Kipf et al. (2018) and Frishman and Ronceray (2020) which jointly learn interactions and dynamics, but require fully observed trajectories, making them inapplicable to single cell time-series experiments (which produce population snapshots only). Further, neural relational inference (Kipf et al., 2018) assumes deterministic trajectories, whereas we consider the setting of trajectories adhering from stochastic dynamics.

2.2 SIMULATION-FREE SCHRÖDINGER BRIDGES VIA SCORE AND FLOW MATCHING

The Schrödinger Bridge Problem (SBP) is concerned with finding the most likely stochastic evolution transporting a source density q_0 to a target density q_1 , given a reference process that encodes prior knowledge of the dynamics. In its dynamical form, the SBP is typically formulated in terms of *laws of stochastic processes*, i.e. probability measures on the path space $\mathcal{C}([0, 1], \mathbb{R}^d)$ that describe the distribution of entire sample trajectories. Writing \mathbb{P} to be the law of a process transporting q_0 to q_1 and \mathbb{Q} to be the reference measure, we seek a stochastic process \mathbb{P}^* satisfying:

$$\mathbb{P}^* = \arg \min_{\mathbb{P}: p_0=q_0, p_1=q_1} \text{KL}(\mathbb{P} \parallel \mathbb{Q}) \quad (4)$$

for marginals p_t of \mathbb{P} , and $\text{KL}(\mathbb{P} \parallel \mathbb{Q}) = \int d\mathbb{P} \log(\frac{d\mathbb{P}}{d\mathbb{Q}})$ is the Kullback-Leibler divergence. When the reference process \mathbb{Q} is the Brownian motion, a key result from Schrödinger Bridge theory (Föllmer, 1988) is that the solution of equation 4 takes the form of a mixture of Brownian bridges \mathbb{Q}_{xy} with respect to a *coupling* π of the distributions (q_0, q_1) :

$$\mathbb{P}^* = \int \mathbb{Q}_{xy} d\pi(x, y). \quad (5)$$

Furthermore, the SBP coupling π amounts to the solution of an entropic optimal transport problem (Léonard, 2014):

$$\pi^* = \arg \min_{\pi \in \Pi(q_0, q_1)} \frac{1}{2} \mathbb{E}_{\pi} \|x - y\|_2^2 + \varepsilon \text{KL}(\pi | q_0 \otimes q_1) \quad (6)$$

with regularization parameter $\varepsilon = \sigma^2 > 0$. In practice when (q_0, q_1) are discrete distributions, this can be solved extremely efficiently using the Sinkhorn algorithm (Cuturi, 2013). While the coupling π is easy to compute from samples, finding a solution to the dynamical SBP in continuous time is desirable for modelling continuous dynamics. With a notable exception being the case of Gaussian measures (Bunne et al., 2022) where analytical expressions are available, the dynamical SBP equation 4 does not admit a straightforward solution for $\mathbf{v}_{\text{SB}}(t, \mathbf{x})$, defined as the drift of the SDE, $d\mathbf{x}_t = \mathbf{v}_{\text{SB}}(t, \mathbf{x}_t) dt + \sigma d\mathbf{B}_t$, which underlies the Schrödinger Bridge process \mathbb{P} .

Numerous works aim to build approximations to \mathbf{v}_{SB} and hence solutions to the dynamical SBP (Chen et al., 2022; Bortoli et al., 2021; Shi et al., 2023; Tong et al., 2024). In particular, Tong et al. (2024) proposes to use score matching (Hyvärinen, 2005) and flow matching (Lipman et al., 2023a) ([SF]²M) as a natural methodology for building the approximation to \mathbf{v}_{SB} in order to solve for \mathbb{P} . Note $\mathbf{v}_{\text{SB}}(t, \mathbf{x})$ simply corresponds to the drift \mathbf{v}_t defined by the SBP. The key observation is to leverage the reciprocal process characterization of \mathbb{P} (Léonard, 2014) as a mixture of Brownian bridges, and the fact that Brownian bridges conditioned on endpoints $(\mathbf{x}_0, \mathbf{x}_1) =: \mathbf{z}$ admit closed form expressions for their

³Throughout we *assume* this data-generation model holds for every condition c .

probability flow $\mathbf{v}_t^\circ(\mathbf{x}|\mathbf{z})$ and score $\nabla \log p_t(\mathbf{x}|\mathbf{z})$:

$$\mathbf{v}_t^\circ(\mathbf{x}|\mathbf{z}) = \frac{1-2t}{t(1-t)}(\mathbf{x} - (t\mathbf{x}_1 + (1-t)\mathbf{x}_0)) + (\mathbf{x}_1 - \mathbf{x}_0); \quad \nabla \log p_t(\mathbf{x}|\mathbf{z}) = \frac{t\mathbf{x}_1 + (1-t)\mathbf{x}_0 - \mathbf{x}}{\sigma^2 t(1-t)}. \quad (7)$$

Together, these provide the probability flow characterisation of the Brownian bridge. With this in hand, and leveraging the fact that the dynamical SBP can be constructed as a mixture of Brownian bridges (equation 5), Tong et al. (2024) propose to construct a neural approximation to the SB flow field $\mathbf{v}^\theta(t, \mathbf{x}) \approx \mathbf{v}(t, \mathbf{x}) := \mathbf{v}_t^\circ(t, \mathbf{x}) - \frac{\sigma^2}{2} \nabla \log p_t(\mathbf{x})$. The flow field $\mathbf{v}^\theta(t, \mathbf{x})$ is trained together with a score field $\mathbf{s}^\theta(t, \mathbf{x})$ using the conditional score and flow matching loss

$$\mathcal{L}_{[\text{SF}]^2\text{M}}(\theta) = \mathbb{E}_{t, \mathbf{z}, \mathbf{x}} [\|\mathbf{v}_t^\theta(\mathbf{x}) - \mathbf{v}_t^\circ(\mathbf{x}|\mathbf{z})\|^2 + \lambda(t)\|\mathbf{s}_t^\theta(\mathbf{x}) - \nabla \log p_t(\mathbf{x}|\mathbf{z})\|^2]. \quad (8)$$

In the above, $\lambda(t) > 0$ weighs the score as a function of time t , and $\mathbf{z} := (\mathbf{x}_0, \mathbf{x}_1) \sim \pi(\mathbf{x}_0, \mathbf{x}_1)$ where π is the entropic OT coupling in equation 5 obtained via the Sinkhorn algorithm. Following Proposition 3.4 of Tong et al. (2024), minimising $\mathcal{L}_{[\text{SF}]^2\text{M}}(\theta)$ then approximates the solution to the SBP in equation 4 under mild conditions.

In the above, $\lambda(t) > 0$ weighs the score as a function of time t , and $\mathbf{z} := (\mathbf{x}_0, \mathbf{x}_1) \sim \pi(\mathbf{x}_0, \mathbf{x}_1)$ where π is the entropic OT coupling in equation 5 obtained via the Sinkhorn algorithm. Following Proposition 3.4 of Tong et al. (2024), minimising $\mathcal{L}_{[\text{SF}]^2\text{M}}(\theta)$ then approximates the solution to the SBP in equation 4 under mild conditions. Additionally, Chen et al. (2023); Theodoropoulos et al. (2025); Rohbeck et al. (2025) model a multi-marginal SBP, where successive time-marginal couplings are not computed independently, but rather globally across all timepoints.

3 STRUCTUREFLOW: JOINT INFERENCE OF STRUCTURE AND DYNAMICS

Our objective is to jointly infer the network structure and population dynamics of the underlying stochastic system. STRUCTUREFLOW achieves this in a novel simulation-free manner. In this section we introduce the central components of our framework, namely: (i) improved parameterization tailored for the joint task, (ii) modeling conditional stochastic population dynamics from interventional data, and (iii) simulation-free training for learning conditional stochastic population dynamics.

3.1 STRUCTURAL VECTOR FIELD PARAMETERIZATION

We consider a time-independent (or *autonomous*) vector field, which models the underlying structure of a system, and a time-dependent score function accounting for stochasticity. We assume that the underlying system structures (graphs) are stationary between adjacent marginals $(\hat{p}_{t_i}^{(c)}, \hat{p}_{t_{i+1}}^{(c)})$. We posit that this assumption yields an easier *joint* inference problem, allowing us to parameterize a *single* structure within the autonomous vector field.

Autonomous vector field and time-dependent score parameterization. From equation 1 we seek to recover the autonomous vector field \mathbf{v} , while the objective in (Tong et al., 2024) targets the SB flow \mathbf{v}_t° . The quantities of the ground truth process equation 1 are related via $\mathbf{v}_t^\circ(\mathbf{x}_t) = \mathbf{v}(\mathbf{x}_t) - \frac{\sigma^2}{2} \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$, where $\mathbf{v}_t^\circ, \nabla \log p_t$ are the probability flow field and the score of the Brownian bridge pinned at $(t_i, \mathbf{x}_i), (t_{i+1}, \mathbf{x}_{i+1})$ respectively. We therefore make the parametrization choice $\hat{\mathbf{v}}_t(\mathbf{x}_t) = \mathbf{v}^\theta(\mathbf{x}_t) - \frac{\sigma^2}{2} \mathbf{s}_t^\phi(\mathbf{x}_t)$. Since time-independent vector fields, \mathbf{v} , are a function of state, and not spurious temporal correlations, they are better suited for modeling underlying structures which are stationary over time, as we assume. This leads to a more principled modeling approach.

Modeling dynamic structural dependencies via neural graphical vector fields. We consider a neural structural model following the definition of Bellot et al. (2022), where we assume there exists functions $\mathbf{v}_1, \dots, \mathbf{v}_d$ such that for $j = 1, \dots, d$, and $\mathbf{v}_j : \mathcal{X}^d \rightarrow \mathbb{R}$, where \mathcal{X}^d is a bounded subset of \mathbb{R}^d . We can rewrite equation 1 under the neural dynamic structural model formalism as

$$dx_j(t) = \mathbf{v}_j(\mathbf{x}(t)) dt + \boldsymbol{\sigma}_j dB_j(t), \quad \mathbf{x}(0) \sim p_0, \quad (9)$$

where $dx_j(t)$ denotes the structural dependency for the instantaneous rate of change of the j^{th} variable dependent on all state observations \mathbf{x} at time t .⁴ We use an NGM (Bellot et al., 2022) to parameterize the dynamic structural relationships defined in equation 9.

NGMs are a class of neural graphical model parameterizations designed to represent complex variable dependencies in a computationally efficient manner (Shrivastava and Chajewska, 2023).

⁴We note that we are using $\mathbf{x}(t)$ interchangeably with \mathbf{x}_t .

Bellot et al. (2022) extend this formulation for the structure learning of continuous-time dynamical systems, motivating their application in our framework. The NGM lets us parameterize the structural relationships of a system with d variables directly within the *autonomous* vector field $v(\mathbf{x}_t)$. We use an NGM to approximate the *autonomous* component of the ground truth process described in equation 1 ($v(\mathbf{x}_t)$ instead of $\mathbf{v}_t(\mathbf{x}_t)$). The NGM parameterization of the *autonomous* vector field is defined as:

$$\mathbf{v}^{\theta_j}(\mathbf{x}_t) = \psi(\dots \psi(\psi(\mathbf{x}_t \theta_j^A) \theta_j^1) \dots) \theta_j^K, \quad j = 1, \dots, d, \quad (10)$$

where $\theta_j^A \in \mathbb{R}^{d \times h}$ denotes a weight matrix (or graph layer) representing the dynamic dependencies between $dx_j(t)$ and $\mathbf{x}(t)$, $\theta_j^k \in \mathbb{R}^{h \times h}$ for $k = 1, \dots, K-1$ are the weight matrices of each corresponding hidden layer, $\theta_j^K \in \mathbb{R}^{h \times 1}$ is the final layer’s weight matrix yielding a scalar output, and $\psi(\cdot)$ is an activation function. We assume variable dependencies defined by θ_j^A are sparse and enforce sparsity on θ_j^A using Group Lasso regularization during training. We include further details regarding the NGM and its optimization for dynamic structure learning in Appendix A.1, as well as principled theoretical justification for recovering the underlying network structure in Appendix A.3.

3.2 MODELING CONDITIONAL POPULATION DYNAMICS

STRUCTUREFLOW explicitly learns conditional stochastic dynamics (trajectories) using both interventional and observational data and simultaneously infers the underlying system structure. We model interventions as ideal *knockouts*, i.e. where for a given intervention, a variable is entirely removed from the system, and the intervention does not directly affect other variables of the system. For an intervention c , we define a binary mask $M^{(c)} \in \{0, 1\}^{d \times d}$ as:

$$M_{ji}^{(c)} = \begin{cases} 0 & \text{if } i = c, j \neq c \\ 1 & \text{otherwise} \end{cases} \quad (11)$$

Where i is the possible source of influence and j is the target. The mask is structured to represent the severed outgoing influences of the variable c . In the observational setting, no mask is applied. Under condition c , the vector field’s parameters become $\theta_j^{A(c)}$, obtained by an element-wise product with $M^{(c)}$: $\theta_j^{A(c)} = M^{(c)} \odot \theta_j^A$. The drift is then evaluated as:

$$\mathbf{v}_c^{\theta_j}(\mathbf{x}_t | c) = \psi(\dots \psi(\psi(\mathbf{x}_t \theta_j^{A(c)}) \theta_j^1) \dots) \theta_j^K. \quad (12)$$

We can then recover our underlying matrix estimate $A \in \mathbb{R}^{d \times d}$ as $A_{ji} = \|(\theta_j^A)_{i,:}\|_2$. Similarly, we condition the score $\mathbf{s}_t^\phi(\mathbf{x}_t)$ on c , i.e., $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t | \mathbf{k}^{(c)})$ where $\mathbf{k}^{(c)} \in \{0, 1\}^d$ is a conditional input vector with $k_i^{(c)} = 1$ if i is perturbed under c , otherwise $k_i^{(c)} = 0$.

3.3 SIMULATION-FREE END-TO-END TRAINING

STRUCTUREFLOW builds on the [SF]²M framework introduced by Tong et al. (2024) to learn a neural approximation of v without simulation. We use Entropic Optimal Transport (EOT) to pair points drawn from temporally adjacent snapshots $\hat{p}_{t_i}^{(c)}$ and $\hat{p}_{t_{i+1}}^{(c)}$. We use the Sinkhorn algorithm (Cuturi, 2013) to estimate the probabilistic EOT couplings between pairs of distributions **only a single time before training starts, and do not need to be recomputed**. From each coupling, we draw paired samples (x_0, x_1) across $[t_i, t_{i+1}]$ and compute the corresponding target drift of the probability flow \mathbf{v}_t^ϕ and target score $\mathbf{s}_t = \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t | c)$. We parametrize the drift of the probability flow as:

$$\hat{\mathbf{v}}_{t,c}(\mathbf{x}_t; \Theta | M^{(c)}, \mathbf{k}^{(c)}) = \mathbf{v}_c^\theta(\mathbf{x}_t | M^{(c)}) - \frac{\sigma^2}{2} \mathbf{s}_t^\phi(\mathbf{x}_t | \mathbf{k}^{(c)}). \quad (13)$$

for conditional probability flow parametrization $\hat{\mathbf{v}}_{t,c}$ under condition c . We optimize model parameters $\Theta = (\theta, \phi)$ by regressing to the targets \mathbf{v}_t^ϕ and \mathbf{s}_t via the STRUCTUREFLOW loss

$$\mathcal{L}_{\text{SF}}(\Theta) = \sum_c \mathbb{E} \left[(1 - \alpha) \|\hat{\mathbf{v}}_{t,c}(\mathbf{x}_t; \Theta | M^{(c)}, \mathbf{k}^{(c)}) - \mathbf{v}_t^\phi\|^2 + \alpha \|\mathbf{s}_{t,c}^\phi(\mathbf{x}_t | \mathbf{k}^{(c)}) - \mathbf{s}_t\|^2 \right], \quad (14)$$

where the expectation is over $t \sim \mathcal{U}(0, 1)$ and $\mathbf{x}_t \sim p_t(\mathbf{x})$, and $0 \leq \alpha \leq 1$ weighs the vector field loss and score loss. During training, we apply Group Lasso on θ_j^A to encourage the model to learn sparse dependencies. We include detailed pseudo code for training STRUCTUREFLOW in Algorithm 1.

We note that we choose to solve the EOT problem to pair successive time marginals ($\hat{p}_{t_i}^{(c)}$ and $\hat{p}_{t_{i+1}}^{(c)}$) independently of the other marginals $\hat{p}_t | t \notin \{i, i+1\}$, not globally across all time marginals like is done in Chen et al. (2023); Theodoropoulos et al. (2025); Rohbeck et al. (2025). STRUCTUREFLOW is flexible however, in that one would simply swap the coupling algorithm to solve the global multi-

marginal problem instead with no additional changes to training. This could be used for other means as well, such as incorporating biological priors.

4 RELATED WORK

Structure learning for continuous dynamics. There exist several works for structure learning of continuous dynamics. Notably, Bellot et al. (2022) introduce a continuous-time framework for inferring structure of the underlying dynamical system from time-series data. Wang et al. (2023) extend this approach to the stochastic and Bayesian setting. However, these approaches rely in computationally expensive neural ODE solvers during training. Recent lines of work have introduced structure learning methods for continuous dynamics specifically tailored for biological systems and the inference of gene regulatory networks (GRNs) (Huynh-Thu et al., 2010; Huynh-Thu and Geurts, 2018; Gao et al., 2018; Tokumasu et al., 2023). However, all these methods assume a selection of fixed ODE parameterizations, linear dynamics, or perturbation-specific recovery, limiting their ability to model stochastic trajectories and complex network interactions. Symbolic regression-based approaches off a complimentary line of work (Brunton et al., 2016; Sun et al., 2023; Dakhmouche et al., 2025). Dakhmouche et al. (2025) also solve for continuous population dynamics (dynamical/trajectory inference) and provide an interpretable model via the learned symbolic equations, but do not explicitly recover a network structure. Brunton et al. (2016); Sun et al. (2023) also use symbolic regression to recover closed-form equations, but assume fully observed trajectories.

Joint inference of structure and dynamics. While both dynamical inference and structure learning are central to understanding stochastic systems, many existing methods treat these tasks in a decoupled fashion (Qiu et al., 2022; Sha et al., 2024). These procedures are centered around learning a dense vector field from the observed data, subsequently extracting structural information post-hoc. Moreover, they do not incorporate any prior on sparsity or explicitly model structure during training. Tong et al. (2024) provide a preliminary investigation into the use of the NGM architecture in conjunction with [SF]²M for jointly inferring structure and dynamics, but do so in a limited setting, i.e. don't consider interventions and assume access to a significant quantity of time-points. We demonstrate later (Section 5 Figure 3) that the trivial application of the NGM architecture under the [SF]²M setup does not achieve competitive performance. Lin et al. (2025) propose a neural ODE-based approach for the joint inference task, but do not model stochastic dynamics, only consider two time-points, and require simulation during training (akin to Bellot et al. (2022)). STRUCTUREFLOW addresses these limitations for the joint task.

5 EXPERIMENTS

We evaluate STRUCTUREFLOW on both synthetic and real-world datasets to assess its ability to recover underlying network structure and infer stochastic population dynamics. Our experiments focus on three key capabilities: structure learning of dynamical systems from population data, dynamical inference across left-out time-points, and prediction of responses to unseen interventions (knockouts). We also include a scaling study to highlight our method's computational efficiency. We provide further details for these experiments in Appendix A, B, and E.

5.1 EXPERIMENTS ON SYNTHETIC HIGH-DIMENSIONAL SYSTEMS

We evaluate STRUCTUREFLOW for the structure learning task across varying system dimensionality using randomly generated Erdos-Renyi graphs (Erdős and Rényi, 1959) and derive synthetic data via linear SDE simulation (see Appendix A.6 for details). We vary the dimensionality of the system d from 10 to 500 variables and the sparsity (proportion of non-zero edges) from 5% to 40% of the underlying graph which defines the data-generative process of the dynamical system. We consider a simulation-based approach, NGM-NeuralODE (NGM-NODE) (Bellot et al., 2022), and Reference Fitting (RF) (Zhang, 2024),

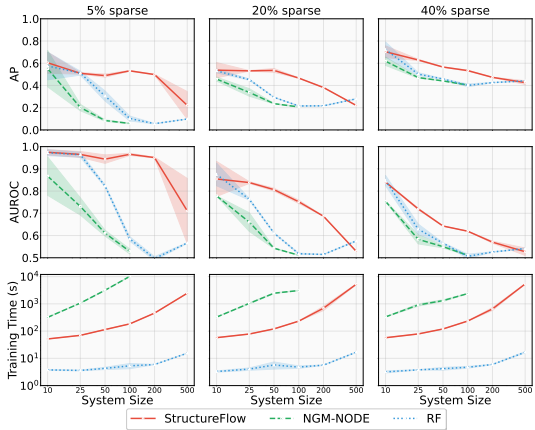


Figure 2: **STRUCTUREFLOW yields improved structure learning performance when scaled to high-dimensional systems.** We compare with NGM-NODE and RF on synthetic linear systems with varying dimensionality (d) and system (graph) sparsity levels (5%, 20%, 40%).




Figure 3: **STRUCTUREFLOW is consistently a top performing structure learning method across simulated biological systems.** Here, we use interventional (with *knockouts*) and observational (no *knockouts*) data, and report average precision (AP) and area under the ROC curve (AUROC) scores.

Table 1: **STRUCTUREFLOW outperforms baseline methods for dynamical inference on simulated biological systems.** Shown is a comparison of dynamical inference methods for learning conditional population dynamics across synthetic biological systems. We include $[SF]^2M$ for comparison, but note that $[SF]^2M$ does not infer underlying network structure (i.e. cannot address joint task). Colors indicate **1st best performer** and **2nd best performer** for models that perform the joint inference task.

	TF		CY		LL		HSC (Curated)		BF		SW	
	$W_2 \downarrow$	MMD1	$W_2 \downarrow$	MMD1	$W_2 \downarrow$	MMD1	$W_2 \downarrow$	MMD1	$W_2 \downarrow$	MMD1	$W_2 \downarrow$	MMD1
$[SF]^2M$	0.761 ± 0.014	0.134 ± 0.005	0.517 ± 0.015	0.153 ± 0.002	0.847 ± 0.053	0.190 ± 0.010	0.664 ± 0.011	0.083 ± 0.003	0.660 ± 0.007	0.140 ± 0.004	0.572 ± 0.008	0.215 ± 0.005
RF	1.376 ± 0.090	0.191 ± 0.000	2.086 ± 0.000	0.326 ± 0.000	2.120 ± 0.007	0.296 ± 0.000	0.950 ± 0.000	0.066 ± 0.000	1.495 ± 0.000	0.242 ± 0.000	1.371 ± 0.000	0.290 ± 0.000
OTVelo	1.874 ± 0.262	0.582 ± 0.039	1.972 ± 0.144	0.622 ± 0.034	2.878 ± 0.361	0.647 ± 0.015	1.942 ± 0.006	0.603 ± 0.002	1.847 ± 0.277	0.591 ± 0.042	1.940 ± 0.144	0.577 ± 0.022
TIGON	1.409 ± 0.294	0.063 ± 0.022	1.224 ± 0.053	0.044 ± 0.003	1.840 ± 0.293	0.066 ± 0.039	0.666 ± 0.363	0.012 ± 0.014	1.394 ± 0.294	0.066 ± 0.021	0.957 ± 0.190	0.017 ± 0.009
STRUCTUREFLOW	0.789 ± 0.012	0.135 ± 0.004	0.576 ± 0.012	0.140 ± 0.004	0.842 ± 0.031	0.196 ± 0.008	0.683 ± 0.011	0.068 ± 0.003	0.694 ± 0.017	0.142 ± 0.007	0.610 ± 0.010	0.220 ± 0.006

as baselines, since they are capable of the joint structure learning and trajectory inference task. We report results in Figure 2 showing area under the receiver operating characteristic curve (AUROC), average precision (AP), and training time (seconds). We observe that STRUCTUREFLOW exhibits favorable scaling performance as dimensionality of the system increases. While the baseline methods are competitive on small systems, we see that STRUCTUREFLOW consistently maintains strong performance in high-dimensional settings while exhibiting improved computational efficiency in terms of training time compared to NGM-NODE, especially for large d . We remark that RF is the most computationally efficient in this regard as it makes assumptions of linearity of the underlying system. Due to this, RF suffers from poor expressivity and exhibits poor performance on the second element of the joint task – dynamical (trajectory) inference. We show this in the following sections.

5.2 EXPERIMENTS ON SIMULATED BIOLOGICAL SYSTEMS

We consider six simulated biological systems: trifurcating (TF), cyclical (CY), long linear (LL), swirling (SW), bifurcating (BF), and a curated system mimics hematopoietic stem cell (HSC) differentiation. All datasets are simulated via BoolODE (Pratapa et al., 2020) under both observational and interventional (knockout) conditions across multiple time-points (see details in Appendix A). See Appendix E for extended synthetic system results and ablations. In this section, we empirically evaluate STRUCTUREFLOW for the joint structure learning and dynamics inference tasks.

STRUCTUREFLOW effectively infers network structure of simulated biological systems. To evaluate how well STRUCTUREFLOW recovers the underlying network structure, we extract the inferred graph from the first layer of the parameterized vector field (defined in equation 10) which interprets the learned weight matrix as a proxy for variable-variable (gene-gene) interactions. We use AP and AUROC to evaluate how closely the STRUCTUREFLOW inferred graphs recapitulate the ground truth structure used to simulate the respective synthetic biological systems.⁵ We compare STRUCTUREFLOW to a variety of baseline methods (some of which are tailored for structure learning in biological systems): RF Zhang (2024), OTVelo Zhao et al. (2024), TIGON Sha et al. (2024), dynGENIE3 Huynh-Thu and Geurts (2018), SINCERITIES Gao et al. (2018), SCODE

⁵In this work, we do not threshold the learned graphs/structures. Instead, we use AP/AUROC metrics to evaluate structure learning performance across all possible thresholds. We leave extensions for learning thresholded structures for future work.

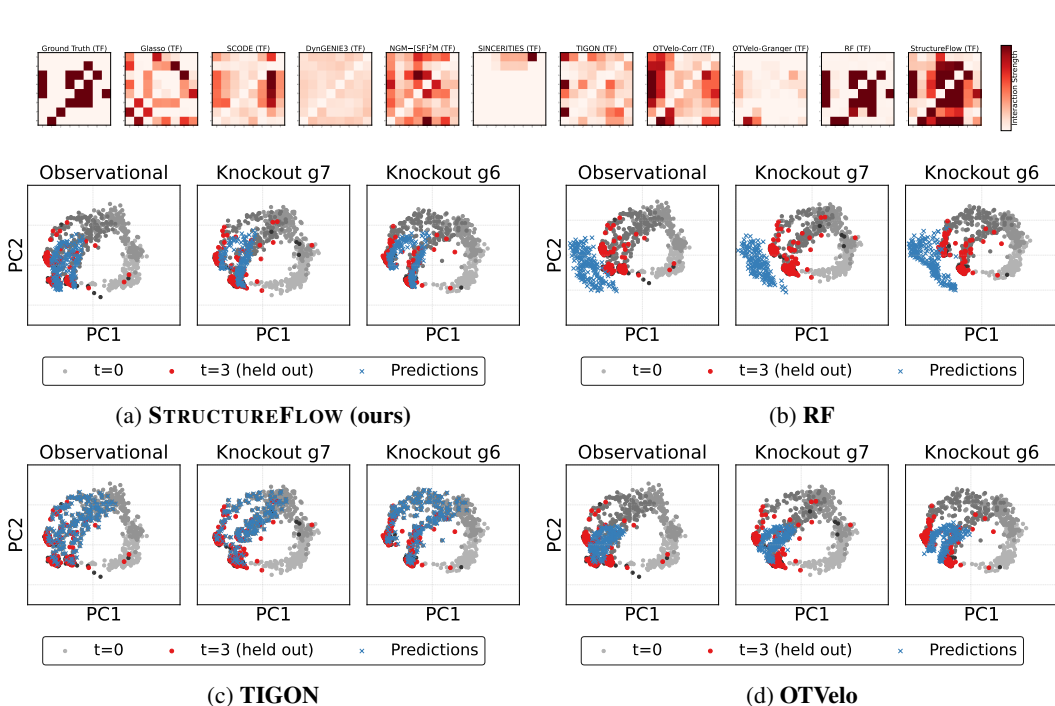


Figure 4: **Summary of inferred structure and visualization of model predicted dynamics.** (top) connectivity matrices (heatmaps) showing ground truth and inferred graphs for the **TF** (trifurcating) system. The x/y-axis labels correspond to system variables and shading indicates edge interaction strength. (bottom) visual comparison of dynamical inference using four different methods: STRUCTUREFLOW, RF, TIGON, and OTVelo, respectively. We show 2D plots of left-out timepoint prediction for the observational (wild-type) setting and for two *seen* interventions (knockouts) selected for their diverse trifurcating paths.

Matsumoto et al. (2017), graphical lasso (Glasso), and NGM-[SF]²M Tong et al. (2024). All models are trained across five random seeds. We show the quantitative results for this experiment and evaluation in Figure 3 and visualizations of the inferred structures in Figure 4 (top). We observe that STRUCTUREFLOW exhibits higher consistency as a top performing method for inferring the underlying structure across all systems and both metrics. In comparison, certain baseline methods show favorable performance in specific systems and on one metric at a time, but struggle to consistently perform across the board.

STRUCTUREFLOW effectively infers dynamics across left-out time-points in simulated biological systems. To evaluate dynamical inference performance, we consider a leave-one-out cross-validation experiment across discrete timepoints, where we iteratively exclude each timepoint t_k for $k \in \{1, \dots, T - 2\}$ during model training, then simulate via ODE from the ground truth data at t_{k-1} to t_k . We use the Wasserstein-2 distance (W_2) and Maximum Mean Discrepancy (MMD) to evaluate the distributional distances between the predicted and ground truth distributions at t_k , evaluating STRUCTUREFLOW’s ability to infer trajectory dynamics.

We consider RF, OTVelo, and TIGON as baselines since they are designed to address both inference task. We additionally include a comparison to [SF]²M as a baseline method for dynamical inference, but we note [SF]²M is not capable of jointly inferring the network structure. We show quantitative results for this experiment in Table 1 and visualization of the predictions in Figure 4a and Figure 4b. We observe that STRUCTUREFLOW outperforms both baselines RF and OTVelo, and a top performer relative to TIGON, while consistently out performing TIGON on the W_2 . We remark that although TIGON yields strong performance on the dynamical inference task, it does not yield competitive results on the concurrent structure learning task (as shown in Figure 3).

STRUCTUREFLOW infers structure and dynamics in the unbalanced setting. We show a use case of STRUCTUREFLOW in the unbalanced setting. To extend

Table 2: **Unbalanced STRUCTUREFLOW**

	AUROC \uparrow	AP \uparrow	W_2 \downarrow	MMD \downarrow
TIGON	0.540 \pm 0.014	0.427 \pm 0.047	2.355 \pm 0.591	0.038 \pm 0.041
STRUCTUREFLOW	0.735 \pm 0.005	0.358 \pm 0.006	1.672 \pm 0.532	0.029 \pm 0.031
U-STRUCTUREFLOW	0.838 \pm 0.010	0.509 \pm 0.025	1.670 \pm 0.438	0.027 \pm 0.033

STRUCTUREFLOW to the unbalanced setting (U-STRUCTUREFLOW), we replace the balanced Sinkhorn algorithm in our EOT procedure with the unbalanced Sinkhorn divergence and incorporate a growth rate field to model proliferation-induced density shifts. We evaluate our method on a growth-based bifurcation system. Following the setup of Zhang et al. (2025), we use the underlying gene regulatory network and boolean logic of the standard BoolODE bifurcation dataset, but we impose lineage specific proliferation rates where one of the attractors has a high proliferation rate and the other has a near-zero growth rate. We observe unbalanced-STRUCTUREFLOW yields the best performance for the joint task (Table 2). We discuss further details in Appendix B.3.

5.3 EXPERIMENTS ON REAL DATA (RENGE) SYSTEM

We consider an interventional time-series dataset of human induced pluripotent stem cells (iPSC) introduced by (Tokumasu et al., 2023). Interventions are conducted via clustered regularly interspaced short palindromic repeats (CRISPR) (knockouts) on 23 transcription factors (TFs). Cell states are then sequenced via single-cell RNA-seq across 4 time bins for each CRISPR intervention. We process the dataset following the procedure of Zhang (2024) and select the 8 interventions with greatest observed change in population-level gene expression (see Appendix A.8).

STRUCTUREFLOW recovers underlying structure of gene regulatory networks competitively to baselines.

We evaluate STRUCTUREFLOW on its ability to recover the 103×103 directed network of inferred TF-TF interactions from the processed real data. We remark that a central challenge in evaluating structure learning performance on real biological systems (datasets) is that there rarely exists a definitively known ground truth network. In this setting, we have access to an approximation of the ground truth network for 18 genes (TFs), which are determined using a CHIP-seq reference (see E.9 for details). Models are trained using all 103 genes, but evaluation is done over 18×103 ground truth networks. We show results in Figure 5 and observe that STRUCTUREFLOW yields competitive performance to counterpart baselines. While there is considerable variance across random seeds, the best-performing instances of STRUCTUREFLOW consistently outperform the top results from baseline methods. We include additional results in Appendix E.




Figure 5: **STRUCTUREFLOW is capable of achieving better performance compared to baselines for the structure learning tasks in the real-data system.** We use AP/AUROC ratio denote structure recovery performance w.r.t. a random predictor.

STRUCTUREFLOW effectively models cell trajectories and improves prediction of left-out interventions (knockouts) in real-data setting.

We report results for dynamical inference over left out time-points and left-out knockouts on the real biological system in Table 3 and Table 4, respectively. In the left-out intervention (knockout) setting, we withhold all time marginals $t_k, k = [0, \dots, T - 1]$ for a given interventional condition during training, then evaluate performance on predicting the final timepoint. Similar to the simulated biological systems, we use W_2 and MMD to evaluate performance. OTVelo and TIGON do not model interventions, and thus are not included in the left-out intervention evaluation. For left-out-time point dynamical inference (Table 3), we observe that STRUCTUREFLOW outperforms all baselines on W_2 and yields competitive results on MMD. We observe the STRUCTUREFLOW is the best performing joint method on the left-out intervention task. We remark that although STRUCTUREFLOW shows improved generalization performance on the left-out intervention task, implying some plausibly improved learning of underlying mechanisms, we make no claim that STRUCTUREFLOW truly models the causal dependencies of the underlying data generative process. We include an equivalent experiment for the simulated biological systems in Appendix E Table 7, but note these systems present an easier generalization setting relative to the real-data system. This is an interesting result that we leave to be explored further in future work.

6 CONCLUSION

In this work, we introduced STRUCTUREFLOW, a novel simulation-free method tailored for jointly learning the underlying network structure and the conditional population dynamics of high-dimensional stochastic systems. Our work highlights the benefits of addressing the joint inference task with a single model trained via score-and-flow matching and tailored for the problem — i.e. improved performance on both tasks. We demonstrate this through a suite of empirical experiments

Table 3: **STRUCTUREFLOW shows competitive performance for dynamical inference of left-out marginals on the real (Renge) biological dataset.** We show a comparison of dynamical inference methods for learning conditional population dynamics across left-out timepoints $k = 1, 2$ on the Renge dataset. We report results over withheld k as well as the average of the two. Colors indicate **1st best performer** and **2nd best performer** for models that perform the joint inference task. Standard deviation in the last column is calculated over the averaged scores.

	Timepoint 1		Timepoint 2		Average	
	$W_2\downarrow$	MMD \downarrow	$W_2\downarrow$	MMD \downarrow	$W_2\downarrow$	MMD \downarrow
[SF] ² M	5.673 \pm 0.034	0.023 \pm 6.9e-5	5.841 \pm 0.030	0.015 \pm 1.0e-4	5.757 \pm 0.032	0.019 \pm 8.6e-5
RF	10.187 \pm 0.003	0.273 \pm 2.6e-4	9.248 \pm 0.008	0.221 \pm 5.9e-4	9.718 \pm 0.005	0.247 \pm 4.2e-4
OTVelo	10.959 \pm 0.000	0.711 \pm 0.000	10.844 \pm 0.000	0.706 \pm 0.000	10.902 \pm 0.063	0.708 \pm 3.3e-3
TIGON	6.434 \pm 0.041	0.004 \pm 1.1e-3	6.401 \pm 0.062	0.003 \pm 9.2e-4	6.417 \pm 0.016	0.003 \pm 1.2e-3
STRUCTUREFLOW	5.589 \pm 0.026	0.024 \pm 4.6e-4	5.679 \pm 0.029	0.017 \pm 1.9e-4	5.634 \pm 0.027	0.020 \pm 3.3e-4

Table 4: **STRUCTUREFLOW shows competitive performance for dynamical inference of left-out interventions on the real (Renge) biological dataset.** We show a comparison of dynamical inference methods for learning conditional population dynamics on the Renge dataset for left-out interventions. We report results for 3 left-out gene knockout conditions, as well as the average performance across conditions. **Bold** indicates the best performing method that can perform the joint task.

	NANOG		POU5F1		PRDM14		Average	
	$W_2\downarrow$	MMD \downarrow	$W_2\downarrow$	MMD \downarrow	$W_2\downarrow$	MMD \downarrow	$W_2\downarrow$	MMD \downarrow
[SF] ² M	5.726 \pm 0.015	0.011 \pm 9.4e-5	5.930 \pm 0.024	0.011 \pm 1.8e-4	6.002 \pm 0.043	0.017 \pm 1.3e-4	5.886 \pm 0.027	0.013 \pm 1.3e-4
RF	9.860 \pm 0.009	0.238 \pm 3.2e-4	10.218 \pm 0.003	0.260 \pm 2.9e-4	9.884 \pm 0.002	0.256 \pm 3.0e-4	9.988 \pm 0.005	0.251 \pm 3.1e-4
STRUCTUREFLOW	5.517 \pm 0.041	0.013 \pm 3.5e-5	5.717 \pm 0.021	0.012 \pm 2.5e-4	5.816 \pm 0.023	0.019 \pm 3.7e-4	5.683 \pm 0.028	0.015 \pm 2.2e-4

and show that STRUCTUREFLOW exhibits improved performance on the joint inference task compared to existing approaches, which either tackle the individual tasks in isolation and/or are limited in their ability to address both problems. Moreover, we showcase the application of our method on a challenging real system for recovering network structure and simultaneously predicting the system response of left-out (*unseen*) interventions.

Limitations & future work. In this work, we focused on the problem of joint structure learning and dynamical inference for stochastic population dynamics. With this come non-standard challenges, e.g. evaluating structure learning performance requires knowledge of the ground truth directed graphs which define the data generative process of the dynamical system. To address this, we considered a suite of synthetic and biologically meaningful simulated systems (Pratapa et al., 2020), such that both tasks can be evaluated in unison, which emulate real systems. However, these simulated systems may be limited in their ability to fully reflect system behavior in real settings. [For example, our simulated knockout conditions may not exactly reflect realistic CRISPR effects, such as imperfect knockouts. We believe exploring settings with imperfect intervention is a fruitful direction for future work.](#)

In a similar vein, evaluating structure learning performance in real-data settings requires experimental validation ([which was accomplished by Tokumasu et al. \(2023\) to produce what we label as the Renge dataset.](#) Such experimental validation is costly to acquire and unrealistic in large systems. Thus, this experimental validation is limited to a small set of system variables, and is still at best an estimate of the ground truth network. Evaluating performance of structure learning methods in real-data settings remains an unsolved challenge. We also do not consider certain aspects of physical processes that are prevalent in natural systems, such as modeling population dynamics while considering interacting particles (Atanackovic et al., 2024) and/or [structure learning of non-stationary networks \(Lu et al., 2025\), but view these directions as natural extensions for STRUCTUREFLOW.](#) [Incorporating biological priors \(either through the choice of reference process \(Petrović et al., 2025; Zhang and Stumpf, 2025\) or by directly imposing knowledge of structural relationships on the network \(Roohani et al., 2024\)\) is another interesting direction for future work.](#) Lastly, using optimal couplings acquired using marginals across all timepoints is a directly applicable extension to STRUCTUREFLOW, and may help improve overall performance of both structure learning and dynamical inference (Chen et al., 2023; Theodoropoulos et al., 2025; Rohbeck et al., 2025).

We remark, STRUCTUREFLOW is designed to be built upon and extended to all aforementioned settings. In this work, we provide a methodological and experimental foundation amenable to further development.

540 REPRODUCIBILITY STATEMENT

541 To facilitate the reproducibility of our results and findings, all source code and datasets will be
542 released upon publication.

543 We also provide the necessary mathematical preliminaries, background, and introduction to our
544 methods in Section 2 and 3, respectively. We further outline experimental details in Appendix A,
545 where we expand on the mathematical details of NGM and RF, provide the experimental details we
546 used to train STRUCTUREFLOW, and clarify our data setups across synthetic (linear), BoolODE, and
547 single-cell CRISPR perturbation settings. In Appendix B, we provide details on our implementation,
548 including Algorithm 1, which outlines the core STRUCTUREFLOW training loop.

550 LLM USAGE STATEMENT

551 Large Language Models (LLMs) were not used during ideation or writing, and thus, are not regarded
552 as a contributor to this work.

555 REFERENCES

- 556 Atanackovic, L., Tong, A., Wang, B., Lee, L. J., Bengio, Y., and Hartford, J. S. (2023). Dyngfn:
557 Towards bayesian inference of gene regulatory networks with gflownets. *Advances in Neural*
558 *Information Processing Systems*, 36:74410–74428.
- 559 Atanackovic, L., Zhang, X., Amos, B., Blanchette, M., Lee, L. J., Bengio, Y., Tong, A., and
560 Neklyudov, K. (2024). Meta flow matching: Integrating vector fields on the wasserstein manifold.
561 *arXiv preprint arXiv:2408.14608*.
- 562 Bellot, A., Branson, K., and van der Schaar, M. (2022). Neural graphical modelling in continuous-
563 time: Consistency guarantees and algorithms. *International Conference on Learning Representa-*
564 *tions*.
- 565 Binnewies, M., Roberts, E. W., Kersten, K., Chan, V., Fearon, D. F., Merad, M., Coussens, L. M.,
566 Gabilovich, D. I., Ostrand-Rosenberg, S., Hedrick, C. C., Vonderheide, R. H., Pittet, M. J., Jain,
567 R. K., Zou, W., Howcroft, T. K., Woodhouse, E. C., Weinberg, R. A., and Krummel, M. F. (2018).
568 Understanding the tumor immune microenvironment (time) for effective therapy. *Nature Medicine*,
569 24(5):541–550.
- 570 Bortoli, V. D., Thornton, J., Heng, J., and Doucet, A. (2021). Diffusion schrödinger bridge with
571 applications to score-based generative modeling. In *Neural Information Processing Systems*
572 (*NeurIPS*).
- 573 Brouillard, P., Lachapelle, S., Lacoste, A., Lacoste-Julien, S., and Drouin, A. (2020). Differentiable
574 causal discovery from interventional data. *Advances in Neural Information Processing Systems*,
575 33:21865–21877.
- 576 Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2016). Discovering governing equations from data by
577 sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of*
578 *Sciences*, 113(15):3932–3937.
- 579 Bunne, C., Hsieh, Y.-P., Cuturi, M., and Krause, A. (2022). The Schrödinger bridge between gaussian
580 measures has a closed form. *arXiv preprint 2202.05722*.
- 581 Bunne, C., Stark, S. G., Gut, G., Del Castillo, J. S., Levesque, M., Lehmann, K.-V., Pelkmans,
582 L., Krause, A., and Rätsch, G. (2023). Learning single-cell perturbation responses using neural
583 optimal transport. *Nature Methods*, 20(11):1759–1768.
- 584 Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. (2018). Neural ordinary
585 differential equations. In *Advances in Neural Information Processing Systems*, volume 31.
- 586 Chen, T., Liu, G.-H., Tao, M., and Theodorou, E. A. (2023). Deep momentum multi-marginal
587 schrödinger bridge.
- 588 Chen, T., Liu, G.-H., and Theodorou, E. A. (2022). Likelihood training of Schrödinger bridge using
589 forward-backward SDEs theory. *International Conference on Learning Representations (ICLR)*.

- 594 Chizat, L., Peyré, G., Schmitzer, B., and Vialard, F.-X. (2019). Unbalanced optimal transport:
595 Dynamic and Kantorovich formulation.
- 596
- 597 Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. *Neural*
598 *Information Processing Systems (NIPS)*.
- 599 Dakhmouche, R., Lunati, I., and Gorji, H. (2025). Robust symbolic regression for dynamical system
600 identification. *Transactions on Machine Learning Research*.
- 601
- 602 Dinh, V. C. and Ho, L. S. (2020). Consistent feature selection for analytic deep neural networks.
603 *Advances in Neural Information Processing Systems*, 33:2420–2431.
- 604 Dixit, A., Parnas, O., Li, B., Chen, J., Fulco, C. P., Jerby-Aron, L., Marjanovic, N. D., Dionne,
605 D., Burks, T., Raychowdhury, R., et al. (2016). Perturb-seq: dissecting molecular circuits with
606 scalable single-cell RNA profiling of pooled genetic screens. *cell*, 167(7):1853–1866.
- 607
- 608 Elowitz, M. B., Levine, A. J., Siggia, E. D., and Swain, P. S. (2002). Stochastic gene expression in a
609 single cell. *Science*, 297(5584):1183–1186.
- 610 Erdős, P. and Rényi, A. (1959). On random graphs i. *Publicationes Mathematicae Debrecen*,
611 6:290–297.
- 612
- 613 Eyring, L., Klein, D., Uscidda, T., Palla, G., Kilbertus, N., Akata, Z., and Theis, F. (2024). Unbal-
614 ancedness in neural Monge maps improves unpaired domain translation.
- 615 Frishman, A. and Ronceray, P. (2020). Learning force fields from stochastic trajectories. *Physical*
616 *Review X*, 10(2).
- 617 Föllmer, H. (1988). Random fields and diffusion processes. *École d’Été de Probabilités de Saint-*
618 *Flour*.
- 619
- 620 Gao, N. P., Ud-Dean, S. M. M., Gandrillon, O., and Gunawan, R. (2018). Sincerities: Inferring gene
621 regulatory networks from time-stamped single cell transcriptional expression profiles. *Bioinfor-*
622 *matics*, 34(2):258–266.
- 623 Gulati, G. S., Sikandar, S. S., Wesche, D. J., Manjunath, A., Bharadwaj, A., Berger, M. J., Ilagan, F.,
624 Kuo, A. H., Hsieh, R. W., Cai, S., Zabala, M., Scheeren, F. A., Lobo, N. A., Qian, D., Yu, F. B.,
625 Dirbas, F. M., Clarke, M. F., and Newman, A. M. (2020). Single-cell transcriptional diversity is a
626 hallmark of developmental potential. *Science*, 367(6476):405–411.
- 627
- 628 Hashimoto, T. B., Gifford, D. K., and Jaakkola, T. S. (2016). Learning population-level diffusions
629 with generative recurrent networks. In *Proceedings of the 33rd International Conference on*
630 *Machine Learning*, pages 2417–2426.
- 631 Huang, K., Lopez, R., Hütter, J.-C., Kudo, T., Rios, A., and Regev, A. (2024). Sequential optimal
632 experimental design of perturbation screens guided by multi-modal priors. In *International*
633 *Conference on Research in Computational Molecular Biology*, pages 17–37. Springer.
- 634 Huynh-Thu, V. A. and Geurts, P. (2018). dyngenie3: Dynamical genie3 for the inference of gene
635 networks from time series expression data. *Scientific Reports*, 8(1):3384.
- 636
- 637 Huynh-Thu, V. A., Irrthum, A., Wehenkel, L., and Geurts, P. (2010). Inferring regulatory networks
638 from expression data using tree-based methods. *PLoS ONE*, 5(9):e12776.
- 639
- 640 Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal*
641 *of Machine Learning Research*, 6(24):695–709.
- 642 Ji, Y., Lotfollahi, M., Wolf, F. A., and Theis, F. J. (2021). Machine learning for perturbational
643 single-cell omics. *Cell Systems*, 12(6):522–537.
- 644 Kipf, T., Fetaya, E., Wang, K.-C., Welling, M., and Zemel, R. (2018). Neural relational inference for
645 interacting systems.
- 646
- 647 Lin, Z., Chang, S., Zweig, A., Kang, M., Azizi, E., and Knowles, D. A. (2025). Interpretable neural
odes for gene regulatory network discovery under perturbations. *arXiv preprint arXiv:2501.02409*.

- 648 Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. (2023a). Flow matching for
649 generative modeling.
650
- 651 Lipman, Y. et al. (2023b). Flow matching for generative modeling. *arXiv:2210.02747v2*.
- 652 Lu, Z., Zhang, W., Le, T., Wang, H., Sümbül, U., SheaBrown, E. T., and Mi, L. (2025). Netformer:
653 An interpretable model for recovering dynamical connectivity in neuronal population dynamics. In
654 *The Thirteenth International Conference on Learning Representations*.
655
- 656 Léonard, C. (2014). A survey of the schrödinger problem and some of its connections with optimal
657 transport. *Discrete and Continuous Dynamical Systems*, 34(4):1533–1574.
- 658 MacArthur, B. D. (2022). The geometry of cell fate. *Cell Systems*, 13(1):1–3.
659
- 660 Maoutsa, D., Reich, S., and Opper, M. (2020). Interacting particle solutions of fokker–planck
661 equations through gradient–log–density estimation. *Entropy*, 22(8):802.
- 662 Matsumoto, H., Kiryu, H., Furusawa, C., Ko, M. S. H., Ko, S. B. H., Gouda, N., Hayashi, T., and
663 Nikaido, I. (2017). Scode: an efficient regulatory network inference algorithm from single-cell
664 rna-seq during differentiation. *Bioinformatics*, 33(15):2314–2321.
665
- 666 Molè, M. A., Coorens, T. H. H., Shahbazi, M. N., Weberling, A., Weatherbee, B. A. T., Gantner,
667 C. W., Sancho-Serra, C., Richardson, L., Drinkwater, A., Syed, N., Engley, S., Snell, P., Christie,
668 L., Elder, K., Campbell, A., Fishel, S., Behjati, S., Vento-Tormo, R., and Zernicka-Goetz, M.
669 (2021). A single cell characterisation of human embryogenesis identifies pluripotency transitions
670 and putative anterior hypoblast centre. *Nature Communications*, 12(1).
- 671 Moscardó García, M., Aalto, A., Montanari, A. N., et al. (2025). Multi-omic network inference from
672 time-series data. *npj Systems Biology and Applications*, 11(114).
673
- 674 Neklyudov, K., Severo, D., and Makhzani, A. (2022). Action matching: A variational method for
675 learning stochastic dynamics from samples.
- 676 Norman, T. M., Horlbeck, M. A., Replogle, J. M., Ge, A. Y., Xu, A., Jost, M., Gilbert, L. A., and
677 Weissman, J. S. (2019). Exploring genetic interaction manifolds constructed from rich single-cell
678 phenotypes. *Science*, 365(6455):786–793.
- 679 Peidli, S., Green, T. D., Shen, C., Gross, T., Min, J., Garda, S., Yuan, B., Schumacher, L. J., Taylor-
680 King, J. P., Marks, D. S., et al. (2024). scperturb: harmonized single-cell perturbation data. *Nature*
681 *Methods*, pages 1–10.
682
- 683 Petrović, K., Atanackovic, L., Moro, V., Kapuśniak, K., Ceylan, I. I., Bronstein, M., Bose, A. J., and
684 Tong, A. (2025). Curly flow matching for learning non-gradient field dynamics. *arXiv preprint*
685 *arXiv:2510.26645*.
- 686 Pratapa, A., Jalihal, A. P., Law, J. N., Bharadwaj, A., and Murali, T. (2020). Benchmarking
687 algorithms for gene regulatory network inference from single-cell transcriptomic data. *Nature*
688 *Methods*, 17(2):147–154.
689
- 690 Qiu, X., Zhang, Y., Martin-Rufino, J. D., Weng, C., Hosseinzadeh, S., Yang, D., Pogson, A. N., Hein,
691 M. Y., Min, K. H. J., Wang, L., Grody, E. I., Shurtleff, M. J., Yuan, R., Xu, S., Ma, Y., Replogle,
692 J. M., Lander, E. S., Darmanis, S., Bahar, I., Sankaran, V. G., and Weissman, J. S. (2022). Mapping
693 transcriptomic vector fields of single cells. *Cell*, 185(4):690–711.e45.
- 694 Rizvi, A. H., Camara, P. G., Kandror, E. K., Roberts, T. J., Schieren, I., Maniatis, T., and Rabadan, R.
695 (2017). Single-cell topological rna-seq analysis reveals insights into cellular differentiation and
696 development. *Nature Biotechnology*, 35(6):551–560.
- 697 Rohbeck, M., Bunne, C., Brouwer, E. D., Huetter, J.-C., Biton, A., Chen, K. Y., Regev, A., and Lopez,
698 R. (2025). Modeling complex system dynamics with flow matching across time and conditions. In
699 *The Thirteenth International Conference on Learning Representations*.
700
- 701 Roohani, Y., Huang, K., and Leskovec, J. (2024). Predicting transcriptional outcomes of novel
multigene perturbations with gears. *Nature Biotechnology*, 42(6):927–935.

- 702 Schiebinger, G., Shu, J., Tabaka, M., et al. (2019). Optimal-transport analysis of single-cell gene
703 expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943.e22.
704
- 705 Sha, Y., Qiu, Y., Zhou, P., Nguyen, H. H., Xie, S., Zou, J., Lee, D., Gehring, J., Alizadeh, A.,
706 Tibshirani, R., Pritchard, J. K., and Zou, E. (2024). Reconstructing growth and dynamic trajectories
707 from single-cell transcriptomics data. *Nature Machine Intelligence*, 6:25–39.
- 708 Shi, Y., De Bortoli, V., Campbell, A., and Doucet, A. (2023). Diffusion Schrödinger bridge matching.
709 *arXiv preprint 2303.16852*.
- 710 Shrivastava, H. and Chajewska, U. (2023). Neural graphical models. In *European Conference on*
711 *Symbolic and Quantitative Approaches with Uncertainty*, pages 284–307. Springer.
712
- 713 Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2021). Score-based
714 generative modeling through stochastic differential equations.
- 715 Sun, F., Liu, Y., Wang, J.-X., and Sun, H. (2023). Symbolic physics learner: Discovering governing
716 equations via monte carlo tree search.
717
- 718 Theodoropoulos, P., Saravanos, A. D., Theodorou, E. A., and Liu, G.-H. (2025). Momentum
719 multi-marginal schrödinger bridge matching.
- 720 Tokumasu, T., Ishikawa, M., and Koyama, H. (2023). Renge: Inferring gene regulatory networks
721 using time-series single-cell rna-seq data with crispr perturbations. *bioRxiv*.
722
- 723 Tong, A. et al. (2024). Simulation-free schrödinger bridges via score and flow matching. *International*
724 *Conference on Artificial Intelligence and Statistics*.
- 725 Tong, A., Huang, J., Wolf, G., van Dijk, D., and Krishnaswamy, S. (2020). Trajectorynet: A dynamic
726 optimal transport network for modeling cellular dynamics.
727
- 728 Wang, B., Jennings, J., and Gong, W. (2023). Neural structure learning with stochastic differential
729 equations. *arXiv preprint arXiv:2311.03309*.
- 730 Wang, D., Jiang, Y., Zhang, Z., Gu, X., Zhou, P., and Sun, J. (2025). Joint velocity-growth flow
731 matching for single-cell dynamics modeling.
732
- 733 Wang, S., Al-Radhawi, M. A., Lauffenburger, D. A., and Sontag, E. D. (2024). Recovering biomolec-
734 ular network dynamics from single-cell omics data requires three time points. *NPJ Systems Biology*
735 *and Applications*, 10(1):97.
- 736 Weinreb, C., Wolock, S., Tusi, B. K., Socolovsky, M., and Klein, A. M. (2018). Fundamental limits
737 on dynamic inference from single-cell snapshots. 115(10):E2467–E2476.
738
- 739 Zhang, J., Cammarata, L., Squires, C., Sapsis, T. P., and Uhler, C. (2023). Active learning for optimal
740 intervention design in causal models. *Nature Machine Intelligence*, 5(10):1066–1075.
- 741 Zhang, S., Maddu, S., Qiu, X., and Chardès, V. (2025). Inferring stochastic dynamics with growth
742 from cross-sectional data.
743
- 744 Zhang, S. Y. (2024). Joint trajectory and network inference via reference fitting. *Proceedings of the*
745 *19th Machine Learning in Computational Biology meeting*.
- 746 Zhang, S. Y. and Stumpf, M. P. (2025). Learning non-equilibrium diffusions with schrödinger
747 bridges: from exactly solvable to simulation-free. *arXiv preprint arXiv:2505.16644*.
- 748 Zhao, W., Larschan, E., Sandstede, B., and Singh, R. (2024). Optimal transport reveals dynamic gene
749 regulatory networks via gene velocity estimation. *bioRxiv*. [Preprint].
750
- 751 Zheng, X., Aragam, B., Ravikumar, P. K., and Xing, E. P. (2018). Dags with no tears: Continuous
752 optimization for structure learning. *Advances in neural information processing systems*, 31.
753
754
755

756 A EXPERIMENTAL DETAILS

757 A.1 NEURAL GRAPHICAL MODELS FOR DYNAMICAL SYSTEMS

758 Unlike static graphical models, where variables are assumed to be independent and identically
759 distributed, dynamical systems require modeling dependencies in continuous time.

760 Neural graphical models (NGMs) parameterize the vector field $\mathbf{v}(\mathbf{x}(t))$ using deep neural networks
761 (Bellot et al., 2022). Specifically, a neural dynamic structural model (NDSM) is defined as:

$$762 dx_j(t) = v_j^\theta(\mathbf{x}(t)) dt + dw_j(t), \quad \mathbf{x}(t_0) = \mathbf{x}_0$$

763 where $v_j^\theta \in \mathcal{F}$ are analytic functions parameterized by neural networks with trainable weights $\theta \in \Theta$.
764 Each function v_j^θ can be expressed as:

$$765 v_j^\theta(\mathbf{x}) = \psi\left(\psi\left(\dots\psi\left(\mathbf{x}\theta_j^A\right)\theta_j^1\dots\right)\theta_j^K\right),$$

766 where $\theta_j^A \in \mathbb{R}^{d \times h}$ is the first (graph) layer, $\theta_j^k \in \mathbb{R}^{h \times h}$ for $k = 1, \dots, K - 1$ and $\theta_j^K \in \mathbb{R}^{h \times 1}$ are the
767 deeper layers, and ψ is an activation function. Directed edges in the inferred graph correspond to
768 non-zero partial derivatives of these networks, and indicate direct causal influences between variables.
769 The adjacency can be inferred if:

$$770 G_{ji} \neq 0 \iff \|\partial_i v_j^\theta\|_{L_2}$$

771 Now to recover G , structure learning can be formulated as a penalized optimization problem:

$$772 \arg \min_{\mathbf{v}^\theta} R_n(\theta) = \min_{\mathbf{v}^\theta} \frac{1}{n} \sum_{m=1}^n \|\mathbf{x}(t_m) - \hat{\mathbf{x}}(t_m)\|_2^2,$$

$$773 \text{subject to } d\mathbf{x}(t) = \mathbf{v}^\theta(\mathbf{x}(t)) dt \text{ and } \rho(\theta) \leq \eta,$$

774 where $\rho(\theta)$ is a regularization term. In (Bellot et al., 2022) this is often implemented using group
775 lasso (GL) or adaptive group lasso (AGL):

$$776 \rho_{\text{GL}}(\theta) = \lambda_{\text{GL}} \sum_{j=1}^d \sum_{i=1}^d \|(\theta_j^A)_{i,:}\|_2, \quad \rho_{\text{AGL}}(\theta) = \lambda_{\text{AGL}} \sum_{j=1}^d \sum_{i=1}^d \frac{\|(\theta_j^A)_{i,:}\|_2}{\|(\hat{\theta}_j^A)_{i,:}\|_2^\gamma}.$$

777 Here, $\lambda_{\text{AGL}}, \lambda_{\text{GL}}$ control the regularization strength, and $(\hat{\theta}_j^A)_{i,:}$ is an initial estimate of the parameters.

778 Given an estimate of \mathbf{v}^θ , NGMs can be extended to irregularly sampled and non-linear time series
779 data by numerically computing the forward trajectory with an ODE solver:

$$780 \hat{\mathbf{x}}(t_1), \hat{\mathbf{x}}(t_2), \dots, \hat{\mathbf{x}}(t_n) = \text{ODESolve}(\mathbf{v}^\theta, \mathbf{x}(t_0), t_1, t_2, \dots, t_n)$$

781 enforcing the constraint $d\mathbf{x}(t) = \mathbf{v}^\theta(\mathbf{x}(t)) dt$ while optimizing for $R_n(\theta)$ (Chen et al., 2018). While
782 the NGM was originally developed to identify causal dependencies in deterministic systems governed
783 by ODEs, our work extends their applicability to stochastic systems by using the probability flow
784 formulation of the underlying SDE. Differently from Bellot et al. (2022), we train an NGM which
785 parameterizes the probability-flow ODE, whose instantaneous drift is

$$786 \hat{\mathbf{v}}_t(\mathbf{x}; \Theta) = \mathbf{v}^\theta(\mathbf{x}) - \frac{\sigma^2}{2} \mathbf{s}_t^\phi(\mathbf{x}),$$

787 which is equivalent in marginal behavior to the original SDE but removes the stochasticity by
788 incorporating the score function $\nabla \log p_t(\mathbf{x})$. This allows us to retain the NGM for structure discovery,
789 while still modeling data generated from inherently noisy dynamics.

800 A.2 REFERENCE FITTING (RF)

801 In contrast to STRUCTUREFLOW, Reference Fitting (Zhang, 2024) (RF) starts from the principle
802 that the inferred couplings π should minimize entropy with respect to a reference process, with the
803 joint-optimization problem defined as:

$$804 \min_{A \in \mathbb{R}^{d \times d}, b \in \mathbb{R}^d} \min_{\pi \in \mathcal{C}(\mu, \mu')} \sigma^2 \text{KL}(\pi | K_{(A,b)}^\sigma) + \mathcal{R}(A, b)$$

805 where \mathcal{R} is a regularizer term, A and b are the linear and constant terms, respectively, of an Ornstein-
806 Uhlenbeck (OU) process of the form:

$$807 dX_t = (AX_t + b) dt + \sigma dB_t$$

810 and $K_{(A,b)}^\sigma$ is the transition kernel of the OU process.

811 Perturbations can then be modeled in this set-up by the linear interaction matrix $A^{(g)}$, where g
812 represents a knocked-out gene, and thus the g th is zeroed out. Thus, for any gene intervention
813 (knockout) g , the reference OU process takes the form:

$$814 \quad dX_t = \left(A \odot M^{(g)} \right) X_t^{(g)} dt + \sigma dB_t$$

815 where $M_{ij}^{(g)} = \mathbf{1}_{i \neq g}$ is the masking matrix for intervention g . This allows the RF process to learn
816 interactions which may have very low signal in the observational (wild-type) data, and thus cannot be
817 learned from the observed dynamics alone. The transition kernel, $K_{(A,b)}^\sigma$, is then approximated by
818 separating the drift and noise terms to model the kernel’s mean and covariance, respectively:

$$819 \quad \mu_t = e^{tA} x_0, \Sigma_t = \sigma^2 t I$$

820 This yields a transition kernel of the form:

$$821 \quad K_t(x, x') \propto \exp \left(- \frac{\|e^{tA} x - x'\|_2^2}{2\sigma^2 t} \right)$$

822 Running an alternating optimization over the couplings π and reference dynamics yields both the
823 reference process itself and couplings which correspond to this process.

824 **Trajectory fitting (ODE view).** Given v^θ , irregularly sampled trajectories can be obtained by
825 numerical integration,

$$826 \quad \hat{\mathbf{x}}(t_1), \dots, \hat{\mathbf{x}}(t_n) = \text{ODESolve}(v^\theta, \mathbf{x}(t_0), t_1, \dots, t_n),$$

827 and one may minimize a data misfit $R_n(\theta) = \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}(t_i) - \hat{\mathbf{x}}(t_i)\|_2^2$ with the above sparsity
828 penalty (Chen et al., 2018; Bellot et al., 2022).

829 **Stochastic dynamics via probability flow.** While NGMs were introduced for deterministic ODEs,
830 we extend them to stochastic systems by parameterizing the *probability-flow ODE* associated with an
831 SDE. Let the instantaneous probability-flow drift be

$$832 \quad \hat{v}_t(\mathbf{x}; \Theta) = v^\theta(\mathbf{x}) - \frac{\sigma^2}{2} \mathbf{s}_t^\phi(\mathbf{x}),$$

833 where $\Theta = \{\theta, \phi\}$, v^θ is the NGM (autonomous) drift defined above, and $\mathbf{s}_t^\phi(\mathbf{x}) = \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$
834 is a score network. This PF-ODE induces the same marginals p_t as the original SDE but removes
835 stochasticity by absorbing it into the score term. We therefore retain the NGM’s FC1 layer θ^A for
836 *structure discovery* (via GL/AGL on $A_{ji} = \|(\theta_j^A)_{i,:}\|_2$) while modeling data generated by noisy
837 dynamics.

838 A.3 ON THEORETICAL GUARANTEES FOR STRUCTURE RECOVERY FROM POPULATION 839 DYNAMICS VIA NGMS

840 **Edge Recovery in the NGM and STRUCTUREFLOW.** Although our vector field is parameterized
841 by a neural network, using a sparse weight matrix as the first layer of the network provides a principled
842 approach to encode, and likewise to recover, the underlying interaction structure of the non-linear
843 system (Bellot et al., 2022). For continuous-time dynamics Bellot et al. (2022) denote directed edge(s)
844 $x_k \rightarrow x_j$ whenever the drift f_j has a **functional dependence** on x_k , i.e. $\partial_k f_j \neq 0$. They show that
845 for any analytic neural network parameterization of f_j , every true functional dependency must pass
846 through the input layer and conversely, if $\partial_k f_j = 0$, there exists an equivalent parameterization where
847 the entire k -th column of the first-layer weight matrix is zero. This justifies the application group
848 lasso solely to the first-layer weights, because zeros in that layer correspond exactly to absent edges
849 in the underlying continuous system. Bellot et al. (2022) show that with enough data, adaptive group
850 lasso recovers the correct sparsity pattern. STRUCTUREFLOW relies on the same principle because
851 our autonomous vector field makes use of the same formulation of the NGM.

852 **On identifiability of underlying structure via the NGM and flow matching.** In general, the
853 question of identifiability for directed structure is challenging, even in the simplest possible case of
854 linear dynamics (see e.g. Wang et al. (2024)). Similarly, Zhang (2024) were not able to provide an
855 identifiability result but only show well-posedness of their optimization problem.

856 To our knowledge, identifiability in the joint trajectory inference and structure learning problem from
857 population snapshots is an open problem and is particularly challenging due to the interdependence
858 of the learned structure. We thus leave a general theoretical study of identifiability to future work. In
859 what follows, we reason that the NGM with Group Lasso recovers the true underlying structure in

the idealized setting where the drift function and score can be reconstructed exactly. We consider autonomous dynamics, where the drift field v is given by a NGM model with parameters θ_0

$$dX_t = v(X_t; \theta_0)dt + \sigma dB_t.$$

Equivalently, this can be written as a probability flow, with $u_t(x) = v(x) - D\nabla \log p_t$ and $D = \frac{1}{2}\sigma\sigma^\top$:

$$\partial_t p_t = -\nabla \cdot (p_t(x)u_t(x)).$$

Consider the case where samples of the flow $u_t(x)$ and score $\nabla \log p_t(x)$ can be accessed exactly, where $(t, x) \sim p_t(x)$. Writing $\hat{v}_t(x; \theta)$ to be the reconstructed vector field, parameterized as a NGM with parameter θ . The corresponding probability flow is

$$\hat{u}_t(x) = \hat{v}_t(x) - D\nabla \log p_t.$$

The flow matching objective with target $u_t(x)$ is thus

$$\min_{\theta} \mathbb{E}_{t \sim [0,1], x \sim p_t} \|\hat{u}_t(x) - u_t(x)\|_2^2 \implies \min_{\theta} \mathbb{E}_t \mathbb{E}_{x \sim p_t} \|\hat{v}_t(x; \theta) - v(x; \theta_0)\|_2^2.$$

In this idealized setting, flow matching amounts to least-square regression on the target vector field $v(x, \theta_0)$. Let $\{x_i\}_{i=1}^N$ a sample of size N drawn following $t \sim [0, 1]$ and $x \sim p_t$. Then, consider the empirical and population risks R_N, R

$$R_N(\theta) = \min_{\theta} \frac{1}{N} \sum_{i=1}^N \|\hat{v}_t(x_i; \theta) - v(x_i; \theta_0)\|_2^2,$$

$$R(\theta) = \mathbb{E}_{t,x} \|\hat{v}_t(x; \theta) - v(x; \theta_0)\|_2^2.$$

Our goal is to characterize identifiability of the structural graph from samples. Then this falls into the framework considered by Dinh and Ho (2020). Specifically, the results of [Dinh and Ho (2020), Theorem 3.6] guarantees local consistency of the learned graph from minimizing (1) in the limit where $n \rightarrow \infty$.

A.4 EXPERIMENT DETAILS

Model architectures. The learned flow model is a NGM as described above. The score and residual models are standard multilayer perceptrons with ReLU activations. These support time-varying and conditionally encoded inputs (knockout vectors). All networks are initialized using Gaussian weights: $\mathcal{N}(0, 0.1)$ for weights and $\mathcal{N}(0, 1 \times 10^{-2})$ for biases.

Trajectory simulation. Simulations are run between pairs of discrete timepoints (t_k, t_{k+1}) where $k = 0, \dots, T-1$ using 100 Euler steps. We consider two rollout strategies: (a) stochastic differential equation (SDE) sampling and (b) probability flow ODE (PF-ODE) sampling. Both start from p_{t_k} , the empirical distribution of cells at timepoint t_k , but differ in how the dynamics are integrated.

ODE SIMULATION The probability flow ODE (PF-ODE), given by

$$\dot{x}_t = v(x_t) - \frac{\sigma^2}{2} s_t^\phi(x_t), \quad x_0 \sim p_{t_k}, \quad (15)$$

describes the deterministic evolution of the density p_t using our learned score s_t^ϕ and flow $v(x_t)$. The score corrects the flow to match the evolving density, while the deterministic rollout removes sample-path noise, yielding smoother trajectories but potentially underestimating variability.

Evaluation metrics. We report both Wasserstein-2 (W2) distance and squared Maximum Mean Discrepancy (MMD²) between predicted and ground truth distributions. W2 is computed using the exact Earth Mover’s Distance from the POT library with a squared Euclidean cost matrix.

MMD² is computed using a radial basis function (RBF) kernel. For each evaluation, we average over five kernel bandwidths. Each bandwidth is defined via a kernel scale parameter $\sigma \in \{0.01, 0.1, 1, 10, 100\}$. The final MMD² score is the mean of the values computed under each of these five kernel settings.

For structure learning evaluation, we compute the Area Under the Receiver Operating Characteristic curve (AUROC) and Average Precision (AP) of the learned graph against the ground-truth. AUROC measures the trade-off between true positive and false positive rates across all classification thresholds. AP summarizes the precision-recall curve as the weighted mean of precisions achieved at each threshold.

Experiment protocols.

- **Leave-one-out (timepoint):** For each dataset, we iteratively exclude data from a single discrete timepoint t_k for $k \in \{1, \dots, T - 2\}$ from training (we only hold out the interior timepoints). The model is trained on the remaining timepoints $\{t_j : j \in \{0, \dots, T - 1\} \setminus \{k\}\}$ and then used to simulate the transition from timepoint t_{k-1} to t_k . We report the average performance across all held-out timepoints by training a separate model for each excluded t_k .
- **Leave-one-out (knockout):** For each dataset, we select the first three gene knockouts and exclude one per trained model. Each model is then used to simulate transitions from t_k to t_{k+1} conditioned on the held-out inter across discrete timepoints $k = 0, \dots, T - 2$. A separate model is trained for each held-out knockout, and results are averaged across the three excluded interventions.

Seed values. All experiments are repeated across fixed seed values. For structure learning experiments, we use 5 seeds: $\{1, 2, 3, 4, 5\}$. For computationally intensive experiments (leave-one-timepoint-out and leave-one-knockout-out) we use 3 seeds: $\{1, 2, 3\}$

Hardware. All experiments were implemented in PyTorch and run on a MacBook Pro with an Apple M1 Pro CPU and 32GB RAM. No GPU is needed. All models train in under 10 minutes.

A.5 HYPERPARAMETERS

We report all hyperparameters used in training our models. Unless otherwise specified, we use the AdamW optimizer with a fixed weight decay of 1×10^{-2} .

Synthetic datasets. Models are trained for 15,000 steps using a batch size of 64 and a learning rate of 3×10^{-3} . Regularization and model-specific settings are provided below:

Table 5: Hyperparameters for synthetic datasets.

Hyperparameter	Value
Training steps	15,000
Batch size	64
Learning rate	3×10^{-3}
Weight decay	1×10^{-2}
α (Score term influence)	0.1
ℓ_2 regularization	5×10^{-6}
Residual regularization	1×10^{-3}
Group Lasso regularization	0.04
Knockout hidden layers	[100]
Score model hidden layers	[100, 100]
Residual model hidden layers	[64, 64]
SDE noise scale σ	1.0

Real data. For the real data experiments, some hyperparameters are overridden:

Table 6: Overriden Hyperparameters for real data (RENGE).

Hyperparameter	Value
Training steps per fold	10,000
Learning rate	0.07
ℓ_2 regularization	1×10^{-9}
α (Score term influence)	0.73
Group Lasso regularization	0.0008
Knockout hidden layers	[128]
Score model hidden layers	[128, 128]
Residual model hidden layers	[128, 128]

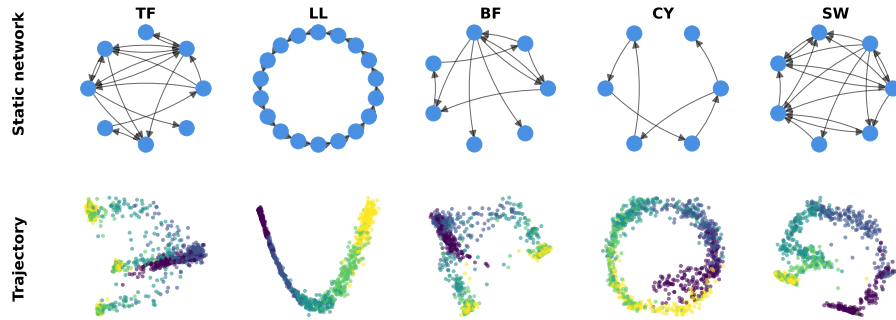


Figure 6: **BoolODE synthetic biological systems.** Network topologies and example trajectory visualizations for the six synthetic biological systems used in our experiments: trifurcating (TF), cyclical (CY), long linear (LL), swirling (SW), bifurcating (BF), and hematopoietic stem cell (HSC) systems. Figure is sourced from Zhang (2024).

We used Optuna to perform hyperparameter optimization, and selected configurations that maximized the area under the precision-recall curve (AUPR). The search space included key hyperparameters such as the learning rate, regularization strength, group lasso regularization, batch size, and model capacity (e.g., number of hidden units). For each trial, model performance was evaluated using cross-validation, and the trial achieving the highest mean AUPR across folds was selected as the optimal configuration. We note that the score term was of higher importance here. This is likely due to the fact that scRNA-sequencing data is much noisier than the synthetic systems.

A.6 SYNTHETIC LINEAR SYSTEM

We generate synthetic linear systems to evaluate structure learning performance across varying dimensionality and sparsity levels. The underlying network structure is generated using Erdos-Renyi random graphs (Erdős and Rényi, 1959), where edges are placed independently with probability p to achieve the desired sparsity level (5%, 20%, or 40% of possible edges).

For a system with d variables, we construct a random adjacency matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ where non-zero entries are sampled to avoid weak interactions. Specifically, edge weights are drawn from the union of two uniform distributions: $\mathcal{U}(-1.0, -0.5) \cup \mathcal{U}(0.5, 1.0)$, ensuring all edges have magnitude ≥ 0.5 . Positive and negative edges occur with equal probability, representing activating and inhibiting interactions respectively. The linear SDE governing the system dynamics is:

$$d\mathbf{x}_t = \mathbf{A}\mathbf{x}_t dt + \sigma d\mathbf{B}_t \quad (16)$$

where $\mathbf{x}_t \in \mathbb{R}^d$ represents the system state, \mathbf{A} encodes the linear interactions, and $\sigma = 0.1$ is the noise level. We simulate trajectories over $T = 5$ time points, generating $N = 1000$ samples per time point. The initial conditions are sampled from a multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, 0.5\mathbf{I})$.

A.7 BOOLODE

BoolODE was employed to simulate 1000 cells independently for each trajectory. To obtain time-resolved snapshots, we divided the simulation into $T=5$ discrete intervals. Expression outputs from simulations were log-transformed for use in subsequent analysis. Knockout trajectories were produced by altering Boolean rules so that the targeted gene retained only self-activation, while its initial expression level was set to zero.

To evaluate performance across different conditions, we considered gene expression values across the following conditions for each dataset type:

- HSC: Gata1, Fli1, Fog1, Eklf, Scl, Gfi1, EgrNab, cJun
- TF, BF, CY, SW, and LL: observational (wild-type), g2, g3, g4, g5, g6, g7, g8

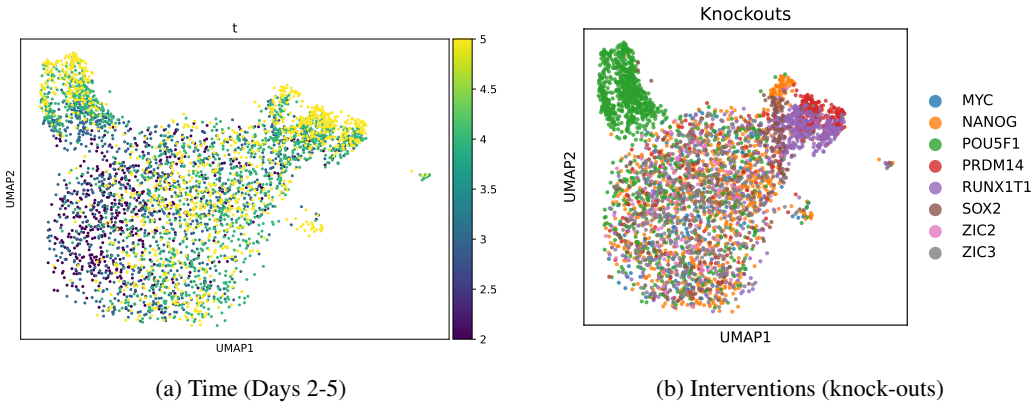
Since we considered 1000 cell expression trajectories, each measured across 800 continuous time and pseudo-time values, and divided across eight conditions, we were left with 100,000 cell measurements per condition. This was then further discretized into $T=5$ time bins.

A.8 SINGLE-CELL CRISPR PERTURBATION TIME-SERIES

Following Zhang (2024), we use raw count data from Tokumasu et al. (2023), retrieved from the Gene Expression Omnibus database (accession GSE213069). Zhang et al. use the following procedure:

- 1026 • Columns corresponding to gRNAs were removed.
- 1027 • Counts were normalized using the `scanpy.pp.normalize_total` function with default options.
- 1028 • The data was log-transformed.
- 1029 • Prior to dimensionality reduction, highly variable genes were selected using `scanpy.pp.highly_variable_genes`.

1032 We considered the set of 103 transcription factors that Zhang (2024) uses and included only cells that
 1033 received a single knockout. To construct the ChIP-seq reference, we obtained experimental binding
 1034 information from ChIP-Atlas (Zou et al., 2024) for the following TFs (within a 1kb window): Chd7,
 1035 Ctnnb1, Dnmt1, Foxh1, Jarid2, Kdm5b, Med1, Myc, Nanog, Nr5a2, Pou5f1, Prdm14, Sall4, Sox2,
 1036 Tcf3, Tcf7l1, Ubtf, Znf398 with a 1kb window. We then further reduced the knockouts to a group
 1037 of 8 most prominent ones as referenced in Zhang (2024). We show a visualization of the dataset in
 1038 UMAP space in Figure 7.



1051 (a) Time (Days 2-5) 1052 (b) Interventions (knock-outs)

1053 Figure 7: **Visualization of real (Renge) dataset.** (a) UMAP visualization showing temporal
 1054 progression of wildtype (unconditional) data across days 2-5, with colors indicating time points. (b)
 1055 UMAP visualization of intervention conditions showing trajectories under different gene knock-out
 1056 perturbations, with colors indicating the specific knocked-out gene.

1059 B IMPLEMENTATION DETAILS

1061 B.1 STRUCTUREFLOW ALGORITHM

1063 Algorithm 1 STRUCTUREFLOW Training Loop

1065 **Require:** Training data across conditions c , hyperparameters $\alpha, \sigma, \lambda_{GL}$, number of steps N , d
 1066 variables
 1067 1: Initialize parameters $\Theta = \{\theta, \phi\}$, Optimizer Opt
 1068 2: **for** step $n = 1$ to N **do**
 1069 3: $M^{(c)} \leftarrow M_{ji}^{(c)} = 0$ if $i = c, j \neq c, M_{ji}^{(c)} = 1$ otherwise **for all** $j, i \in 1, \dots, d \times 1, \dots, d$
 1070 4: $k^{(c)} \leftarrow k_j^{(c)} = 1$ if j perturbed under c , else $k_j^{(c)} = 0$ **for all** j
 1071 5: Sample $(\mathbf{x}_t, t, \mathbf{s}_t, \mathbf{v}_t^\circ)$ from SB for condition c
 1072 6: $\mathbf{s}_t^\phi \leftarrow \mathbf{s}_t^\phi(\mathbf{x}_t | \mathbf{k}^{(c)})$
 1073 7: $\hat{\mathbf{v}}_t \leftarrow \mathbf{v}^{\theta_j}(\mathbf{x}_t | M^{(c)}) - \frac{\sigma^2}{2} \mathbf{s}_t^\phi(\mathbf{x}_t | \mathbf{k}^{(c)})$ \triangleright Using $M^{(c)} \odot \theta_j^A$
 1074 8: $\mathcal{L} \leftarrow \sum_c \mathbb{E} \left[(1 - \alpha) \|\hat{\mathbf{v}}_t(\mathbf{x}_t; \Theta | M^{(c)}, \mathbf{k}^{(c)}) - \mathbf{v}_t^\circ\|^2 + \alpha \|\mathbf{s}_t^\phi(\mathbf{x}_t | \mathbf{k}^{(c)}) - \mathbf{s}_t\|^2 \right]$
 1075 9: $\mathbf{g} \leftarrow \nabla_{\Theta} \mathcal{L}$
 1076 10: Update parameters $\Theta \leftarrow Opt(\Theta, \mathbf{g})$
 1077 11: $\theta_j^A \leftarrow \text{Prox}_{\eta \lambda_{GL} \|\cdot\|_{\text{group}}}(\theta_j^A)$ \triangleright group-wise proximal update on first layer
 1078 12: **end for**

B.2 GENE REGULATORY NETWORK INFERENCE BASELINES

All baselines were trained with the same pre-processing steps, train-test splits, and evaluation metrics. In addition, we experimented with the hyper-parameters of the baseline methods. For methods that considered the same BoolODE dataset(s) (Zhang, 2024; Zhao et al., 2024), we used the respective hyper-parameters reported in the work. For other methods, we experimented with hyper-parameters settings, but found that the hyper-parameters set by the authors of the respective baseline methods performed generally well.

SCODE: Sparse CODing for Differential Equations (SCODE) from Matsumoto et al. (2017) is a method for reconstructing GRNs from single-cell time series by fitting a linear ODE system to gene expression trajectories. The method infers a sparse gene interaction matrix such that simulated expression profiles recapitulate observed temporal dynamics. In the synthetic case, we take the transposed expression matrix and pseudotime vector that is given to us in both the observational and interventional settings and write them to disk in a format compatible with SCODE. We then execute the model via the Ruby + R wrapper developed by Matsumoto et al. (2017), and read the mean interaction matrix from the resulting output files, which gives us an adjacency matrix. SCODE does not work with interventional data so we do not expect an increase in performance.

SINCERITIES: Single-Cell Network Inference by Covariance Regression for Time Series, from Gao et al. (2018) is an R-based algorithm for inferring directed GRNs from time-course single-cell data. SINCERITIES combines information from lagged covariances and local linear regression to estimate causal effects, and it outputs signed edge weights. We use the R script developed by the authors and the given expression matrix and pseudotime matrix to create a `GRNprediction.txt` file which gives a gene \times gene matrix.

GLASSO: Graphical Lasso is a classical method for inferring sparse undirected graphical models from high dimensional data. GLASSO estimates the precision (inverse covariance) matrix under an l_1 penalty to enforce sparsity, such that nonzero entries correspond to conditional dependencies between genes. This method does not work with time-series data. We first standardize the gene expression matrix, and choose the regularization parameter via cross-validation. The estimated precision matrix is negated (higher values correspond to stronger interactions), and the diagonal is set to zero. Our result is an adjacency matrix with symmetric edge weights that represent partial correlations between genes.

dynGENIE3: This method from Huynh-Thu and Geurts (2018) is an extension of the GENIE3 algorithm (Huynh-Thu et al., 2010) to reconstruct GRNs from time-series single-cell or bulk gene expression data. The method frames network inference as a feature selection problem, where for each gene, a random forest regression model is trained to predict its expression at each time point as a function of all other genes' expressions at earlier time points. The regulatory strength is aggregated over all of the trees and time lags, and yields a directed adjacency matrix. We aggregated our expression matrix `adata.X`, which is an `annData` object, by time points (mean expression per time) and store it in a list `X`, and keep the corresponding points in `t`. The dynGENIE3 model is then called with these two parameters as input and we save the resulting adjacency matrix. *By averaging across all cells at each time t we effectively create pseudo-bulk expressions, then stack these per-timepoint means into a matrix so as to be a single trajectory. This approach is corroborated by Moscardó García et al. (2025) who use a nearly identical procedure for dynGENIE3.*

RENGE: RENG Tokumasu et al. (2023) can use both temporal and interventional data. Renge uses pseudotime trajectories from single-cell data, and CRISPR knockout information, with a two-step inference strategy, that involves a regression model to predict expression as a function of genes, and network aggregation to aggregate the inferred effects across all perturbations and time points. This gives us a gene \times gene matrix based on both the endogenous dynamics and systematic perturbation present in the dataset. The main experiment from Tokumasu et al. (2023) was what was used in our paper and in Zhang (2024). We use the model that Tokumasu et al. (2023) provide in their codebase and do further experimentation, using that matrix as a baseline.

TIGON: Trajectory Inference with Growth via Optimal transport and Neural networks (TIGON) (Sha et al., 2024) models time-series scRNA-seq as a continuum of cell states whose density evolves under a reaction-transport dynamic. It represents a time-dependent cell-state density represented by two fields: a velocity field that models state transitions, and a growth field that accounts for birth/death and mass change. TIGON integrates the learned dynamics using ODE integration per time step, and learns to minimize a the Wasserstein-Fisher-Rao distance. At inference, TIGON learns to reconstruct the continuous trajectories and velocities for individual cells, estimates time-resolved

1134 growth maps, and can learn a gene-regulatory structure from the Jacobian of its flow. TIGON
 1135 uses either a PCA layer or an autoencoder to first encode the ambient data. This is a step for the
 1136 “dimensionless-solver” used within TIGON, which is optimized to operate in 10 dimensional space.
 1137 All of the other baselines, and STRUCTUREFLOW, are trained directly in the ambient space. Thus,
 1138 in order to provide a comparative setting, we train TIGON in the ambient space as well and forgo
 1139 the PCA/autoencoder step. We did not encounter challenges with applying TIGON to the high
 1140 dimensional real (Renge) biological system. TIGON uses Gaussian mixture models to represent p_t
 1141 (equivalent to our Eq. 2). We also train TIGON in PCA space, and then convert to the ambient space
 1142 using the following formula: $J_X = W^T J_z W$ where J_z is the latent Jacobian. We note no difference
 1143 in performance for structure learning.

1144 **OTVelo:** OTVelo Zhao et al. (2024) estimates gene velocity from cell trajectories reconstructed
 1145 via optimal transport. The method proceeds in two stages. The first stage is trajectory inference
 1146 where it aligns cell across consecutive time-points using Fused Gromov-Wasserstein (FGW) optimal
 1147 transport, which uses both expression state and local geometric structure. The second part uses OT
 1148 pairings to estimate the gene velocities via two different approaches. We consider two versions of this
 1149 model, OTVelo-Corr and OTVelo-Granger. OTVelo-Corr uses time-lagged correlation analysis that
 1150 generalizes the standard time-lagged correlation to distributions of cells matched by the OT pairings.
 1151 OTVelo-Granger, on the other hand, is a regularized linear regression model (Elastic Net) where the
 1152 velocity of a target gene at time $t + 1$ is predicted from the velocities of regulating gene at time t .
 1153 Given that we discretize all of our data into bins, and OTVelo tests their model on the same datasets,
 1154 simulated using BoolODE, we apply the model from Zhao et al. (2024) directly in our pipeline.

1155 B.3 UNBALANCED STRUCTUREFLOW

1156 In order to properly model cell proliferation and death, we formalize the extension of STRUCTURE-
 1157 FLOW to biological systems where probability mass is not conserved. We generalize the probability
 1158 flow formulation to include a growth term, derived from Unbalanced Optimal Transport (UOT).
 1159

1160 **Unbalanced Optimal Transport** In the standard STRUCTUREFLOW framework, couplings
 1161 $(x_0, x_1) \sim \pi(x_0, x_1)$ are acquired via a balanced Sinkhorn solver, which enforces strict mass
 1162 conservation ($\sum \pi_i = \frac{1}{N}$). To model cell proliferation and death, we replace this with UOT. The
 1163 universality of our method allows us to easily incorporate unbalancedness into STRUCTUREFLOW
 1164 by rescaling source and target measures Eyring et al. (2024). We compute the coupling π^* between
 1165 empirical snapshots $\hat{\rho}_{t_k}$ and $\hat{\rho}_{t_{k+1}}$ by minimizing the unbalanced Sinkhorn Divergence Wang et al.
 1166 (2025); Chizat et al. (2019):

$$1167 \pi^* = \arg \min_{\pi \in \mathbb{R}_+^{N \times M}} \sum_{i,j} C_{ij} \pi_{ij} + \epsilon H(\pi) + \rho [KL(\pi_{\#0} | \hat{\rho}_{t_k})] + KL(\pi_{\#1} | \hat{\rho}_{t_{k+1}})$$

1169 where $\epsilon H(x)$ is the entropic regularization coefficient, ρ is the marginal relaxation parameter, and
 1170 $\pi_{\#0}$ are row sums which control the amount of mass transported away from each starting cell. If
 1171 $\pi_{\#0} < \hat{\rho}_{t_k}$ then the transport plan used less mass than was available and if $\pi_{\#1} > \hat{\rho}_{t_{k+1}}$ then the
 1172 transport plan delivered more mass than was originally available relative to the source.

1173 **Stochastic Dynamics with Growth** To account for cellular growth, we adopt the formulation
 1174 proposed in Zhang et al. (2025), where the population density $\rho_t(x)$ evolves according to a Fokker-
 1175 Planck equation with a source term:

$$1176 \partial_t \rho_t(x) = -\nabla \cdot (\rho_t(x) \mathbf{v}_t(x)) + \frac{1}{2} \nabla \cdot (\sigma^2 \nabla \rho_t(x)) + g_t(x) \rho_t(x)$$

1177 where $\mathbf{v}_t(x)$ represents the drift term and $\frac{1}{2} \nabla \cdot (\sigma^2 \nabla \rho_t(x))$ is the diffusion coefficient. Unlike
 1178 eq. (2) where probability mass is conserved, the total mass $M_t = \int \rho_t(x) dx$ varies over time
 1179 according to the net growth rate of the system. We model the joint evolution where the local change
 1180 in log-density is governed by both the divergence of the flow and the local growth rate, defined as:

$$1181 \frac{d}{dt} \log \rho_t(x_t) = g_t(x_t) - \nabla \cdot u_t(x_t)$$

1182 where $u_t(x)$ is defined in eq. (3). This formulation then allows us to disentangle the vector field
 1183 \mathbf{v}_t from the growth field, so that the density increase is due to the growth field, and the vector
 1184 field only models transitions. Lastly to learn the growth field $g_\phi(x)$, we adopt the simulation-free
 1185 strategy proposed in Wang et al. (2025). Instead of solving a differential equation for mass evolution
 1186 during training, we interpret the row-sums of the unbalanced coupling matrix as static proxies for net
 1187 proliferation accumulated along a trajectory. For a source sample x_0^i , the target growth G_i is defined

1188 by the log-marginal of the transport plan π^* :

$$1189 \quad G_i = \log\left(\sum_j \pi_{ij}^*\right)$$

1191 This value represents the total probability mass at t_1 attributed to the lineage of x_0^i . If the sum is
 1192 greater than 1, the cell has proliferated, otherwise the cell has undergone apoptosis (died). To learn
 1193 the growth network, we regress the neural network $g_\phi(x, t)$ to the static target across the interpolated
 1194 trajectory:

$$1195 \quad \mathcal{L}_{\text{growth}}(\phi) = \mathbb{E}_{t, (x_0^i, x_1^j) \sim \pi^*} [\|g_\phi(x_t) - G_i\|^2]$$

1198 C NOVELTY STATEMENT

1200 This work introduces key methodological and empirical contributions to joint structure learning and
 1201 dynamics inference:

1202 **Decomposition of Stochastic Dynamics.** We decompose the probability flow into an autonomous
 1203 drift $v(x)$ capturing system structure and a time-dependent score function $s_t(x)$ modeling stochas-
 1204 ticity. This aligns with biophysical reality where regulatory mechanisms are time-invariant.

1205 **Conditional Dynamics via Interventions.** We incorporate interventional data through knockout
 1206 conditioning masks, enabling learning from perturbational experiments and interventional trajectory
 1207 inference on unseen perturbations.

1208 **Simulation-Free Training for High-Dimensional Systems.** Our framework extends simulation-free
 1209 training to conditional settings, enabling tractable learning on high-dimensional systems (up to 500
 1210 variables) while jointly inferring structure and dynamics.

1211 **Comprehensive Evaluation.** We establish extensive benchmarks including scalability analysis,
 1212 evaluation on six biologically plausible systems with interventional data, and validation on real
 1213 single-cell CRISPR time-series data with realistic constraints (4-5 time points).

1215 D BROADER IMPACT

1217 The genome is the fundamental identity of every living thing. The expression of DNA dictates life. It
 1218 is a multi-dimensional representation of the state of any living thing. The ultimate goal of our work
 1219 is to enable improve our understanding of how an organism’s genes interact, allowing us to predict
 1220 the trajectory of its life, and ultimately to computationally develop interventions that can positively
 1221 change that trajectory. The stochastic dynamical systems we consider here are used ubiquitously in
 1222 the biophysics and modelling literature. In a biological context, our vector field v can be related to
 1223 the metaphorical “Waddington’s landscape”.

1224 E ADDITIONAL RESULTS

1226 E.1 STRUCTURE LEARNING ON SYNTHETIC SYSTEMS WITH OBSERVATIONAL DATA

1227 In 8, we observe that STRUCTUREFLOW achieves competitive results in most systems across both
 1228 AP and AUROC metrics. Whereas other methods tend to fluctuate widely in performance across
 1229 these synthetic systems, STRUCTUREFLOW underperforms all other methods only on the SW system,
 1230 while beating most other benchmarks otherwise.

1232 E.2 DYNAMICAL INFERENCE OF CONDITIONAL POPULATION DYNAMICS ON SYNTHETIC 1233 SYSTEMS USING [SF]²M

1234 We note in Figure 11 that [SF]²M predictions cluster around the synthetic systems’ ground-truth
 1235 trajectories across both observational and gene knockout settings. This is a result of training across
 1236 both unconditional and conditional data.

1238 E.3 PREDICTION TO UNSEEN INTERVENTIONS (KNOCK-OUTS) ON SYNTHETIC SYSTEMS

1239 Table 7 shows performance of STRUCTUREFLOW and baselines on the leave-one-knockout-out
 1240 trajectory inference task. STRUCTUREFLOW outperforms RF on the TF and CY systems, and is
 1241 competitive on the Curated (HSC) system.

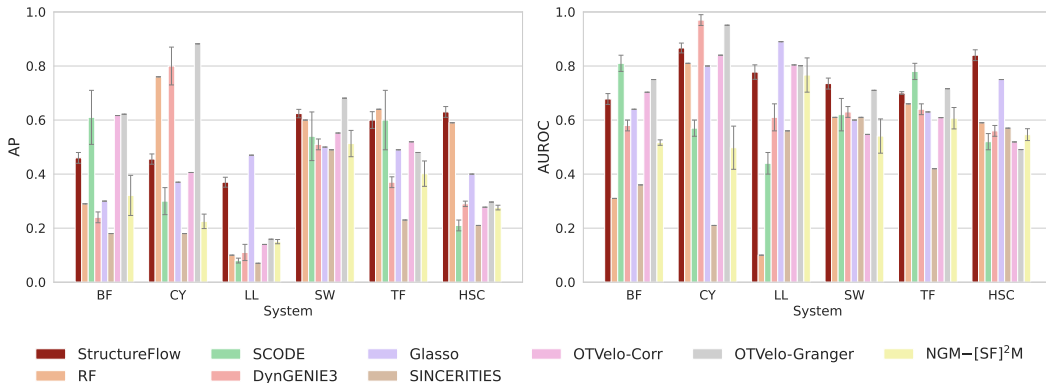


Figure 8: **Structure learning performance across synthetic systems using only observational data.** Shown are average precision (AP) and area under the ROC curve (AUROC) scores for models trained on only observational data

Table 7: **Comparison of left-out intervention (knock-out) prediction on simulated biological systems.** For each dataset, we report Wasserstein-2 (W_2) and MMD metrics. Lower is better.

	TF		CY		LL		HSC		BF		SW	
	$W_2 \downarrow$	MMD \downarrow	$W_2 \downarrow$	MMD \downarrow	$W_2 \downarrow$	MMD \downarrow	$W_2 \downarrow$	MMD \downarrow	$W_2 \downarrow$	MMD \downarrow	$W_2 \downarrow$	MMD \downarrow
[SF] ² M	0.964 ± 0.016	0.103 ± 0.003	0.967 ± 0.028	0.277 ± 0.017	1.082 ± 0.095	0.157 ± 0.010	0.886 ± 0.011	0.104 ± 0.003	1.138 ± 0.031	0.220 ± 0.005	1.404 ± 0.051	0.358 ± 0.015
RF	1.062 ± 0.000	0.098 ± 2.8e-8	1.889 ± 0.000	0.329 ± 1.5e-8	1.887 ± 0.005	0.232 ± 6.5e-4	0.984 ± 3.4e-10	0.065 ± 1.7e-8	1.176 ± 0.000	0.134 ± 1.1e-8	1.285 ± 0.000	0.284 ± 3.4e-8
STRUCTUREFLOW	1.012 ± 0.010	0.110 ± 4.1e-3	0.924 ± 0.036	0.250 ± 1.8e-2	1.123 ± 0.037	0.141 ± 6.1e-3	0.852 ± 0.011	0.120 ± 4.2e-3	1.079 ± 0.016	0.213 ± 7.1e-3	1.159 ± 0.032	0.330 ± 1.1e-2

E.4 OT-CFM ABLATION ON SYNTHETIC (TRIFURCATING) SYSTEM

In Table 15 we remove the stochastic noise term from the conditional flow matching objective, reducing the formulation to an ODE that models the drift component. This corresponds to the classical flow matching objective Lipman et al. (2023b), where the dynamics align deterministic trajectories. STRUCTUREFLOW is trained to model the process as a Schrödinger Bridge, where the forward-backward SDE explicitly models both drift and diffusion. From the results, we observe that discarding the stochastic component (OT-CFM) substantially degrades both classification (AUROC/AUPR) and distributional alignment metrics (W_2), highlighting that noise modeling contributes meaningfully to the structure learning task. For dynamical systems such as cellular trajectories observed across time points, modeling the process as a Schrödinger bridge rather than an ODE provides a more accurate representation of developmental variability and uncertainty.

Table 8: **Comparison of STRUCTUREFLOW with and without stochastic component on the TF synthetic system.** Results show AUROC, AUPR, and distributional distance metrics (W_2 , MMD).

	AUROC \uparrow	AUPR \uparrow	$W_2 \downarrow$	MMD \downarrow
STRUCTUREFLOW	0.968 ± 0.005	0.932 ± 0.011	0.795 ± 0.007	0.136 ± 0.002
STRUCTUREFLOW (OT-CFM)	0.791 ± 0.013	0.634 ± 0.037	0.846 ± 0.022	0.062 ± 0.004

E.5 SCALING AND RUNTIME OF VARIOUS BASELINES AND OUR METHOD

Here we provide an additional experiment comparing training time of all methods on the BoolODE TF system. For this experiment, all methods are run on the same device (GPU, Quadro RTX 6000), but we note STRUCTUREFLOW can be efficiently trained using a CPU. We only report paramter sizes for neural network-based methods. For STRUCTUREFLOW and the flow matching baseline, we use \cdot to denote the additional time used to compute EOT. We use the same number of timepoints (5), dimensionality (8), and number of cells (8000). We add flow matching as a reference point to STRUCTUREFLOW, where we remove regularization terms, the score model, and replace the NGM backbone with a traditional MLP backbone.

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344


1345

1346


1347

1348


1349




(a) Dataset BF (Observational Data)




(b) Dataset CY (Observational Data)




(c) Dataset SW (Observational Data)



(d) Dataset TF (Observational Data)



(e) Dataset LL (Observational Data)



(f) Dataset HSC (Observational Data)

Figure 9: Six of the synthetic datasets with only observational information. Refer to panels (a)–(f) when discussing individual datasets.

1350

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398


1399

1400


1401

1402


1403




(a) Dataset BF (Observational + Interventional Data)




(b) Dataset CY (Observational + Interventional Data)




(c) Dataset SW (Observational + Interventional Data)



(d) Dataset TF (Observational + Interventional Data)



(e) Dataset LL (Observational + Interventional Data)



(f) Dataset HSC (Observational + Interventional Data)

Figure 10: Six of the synthetic datasets with the addition of interventional information. Refer to panels (a)–(f) when discussing individual datasets.

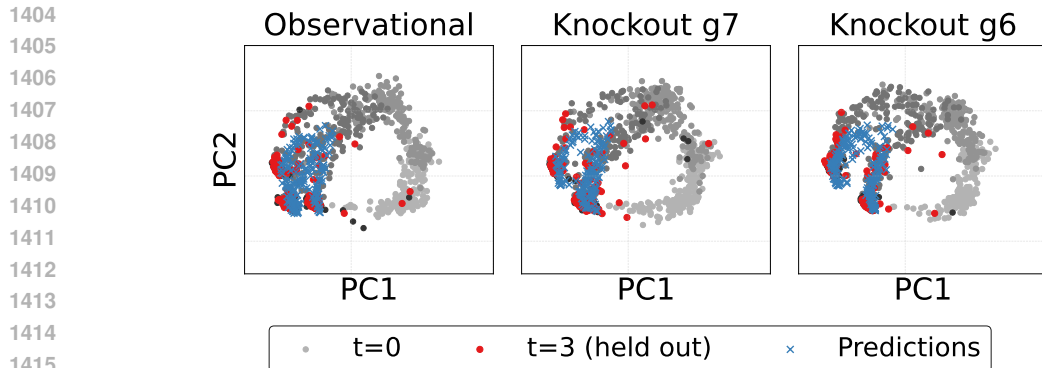


Figure 11: $[\text{SF}]^2\text{M}$ trajectory inference visualization. Performance on observational (wild-type) data and two interventions selected for their diverse trifurcating paths. These examples illustrate the trajectory inference problem, where timepoint 3 was excluded during training, and the model simulates with the ODE from timepoint 2 to timepoint 3. Ground truth data is shown in grey, with progressively darker shades as time increases, except for timepoint 3 which is highlighted in red. Blue represent the model’s predictions.

Table 9: Model wall-clock time on the Trifurcating (TF) BoolODE system. We include parameter counts for all neural network-based methods. We use (\cdot) to denote runtime of EOT.

Time [s]	Model	Parameter Size
263 + (3.3)	STRUCTUREFLOW	20716
1511.18	TIGON	43140
82.01	RF	-
3228.61	NGM-NODE	8008
116.17 + (3.3)	Flow matching	8008
208.58 + (3.3)	$[\text{SF}]^2\text{M}$	14892
131.22	OTVelo	-
5771.41	SDFL	-

E.6 ADDITIONAL ABLATIONS OF HYPERPARAMETERS

In order to assess the stability of our method we vary three hyperparameters that are crucial to the performance of our model: α , which varies the weight of the score and flow term; group lasso regularization which determines how sparse the adjacency matrix recovered by the autonomous flow term is; and the size of the hidden dimensions in our flow module (We use one layer in our case).

Table 10: Ablation over the group-lasso regularization term.

Group Lasso Reg.	0.0	0.0001	0.001	0.01	0.04	0.1	0.2	0.5
AUROC	0.943 ± 0.004	0.946 ± 0.004	0.944 ± 0.004	0.952 ± 0.002	0.971 ± 0.008	0.853 ± 0.032	0.851 ± 0.017	0.879 ± 0.016
AP	0.804 ± 0.014	0.816 ± 0.015	0.805 ± 0.014	0.833 ± 0.008	0.918 ± 0.026	0.677 ± 0.032	0.688 ± 0.015	0.708 ± 0.021
W_2	0.7623 ± 0.0047	0.7623 ± 0.0030	0.7608 ± 0.0050	0.7665 ± 0.0077	0.7886 ± 0.0048	0.8329 ± 0.0049	0.8398 ± 0.0059	0.8414 ± 0.0058

Table 11: Ablation over α . $\alpha = 1$ pertains to only using the score model. $\alpha = 0$ pertains to only using flow model.

α	0.1	0.2	0.5	0.8
AUROC	0.971 ± 0.008	0.965 ± 0.007	0.971 ± 0.009	0.937 ± 0.005
AP	0.917 ± 0.025	0.886 ± 0.026	0.918 ± 0.040	0.851 ± 0.011
W_2	0.7886 ± 0.0048	0.7814 ± 0.0007	0.7887 ± 0.0035	0.8085 ± 0.0042

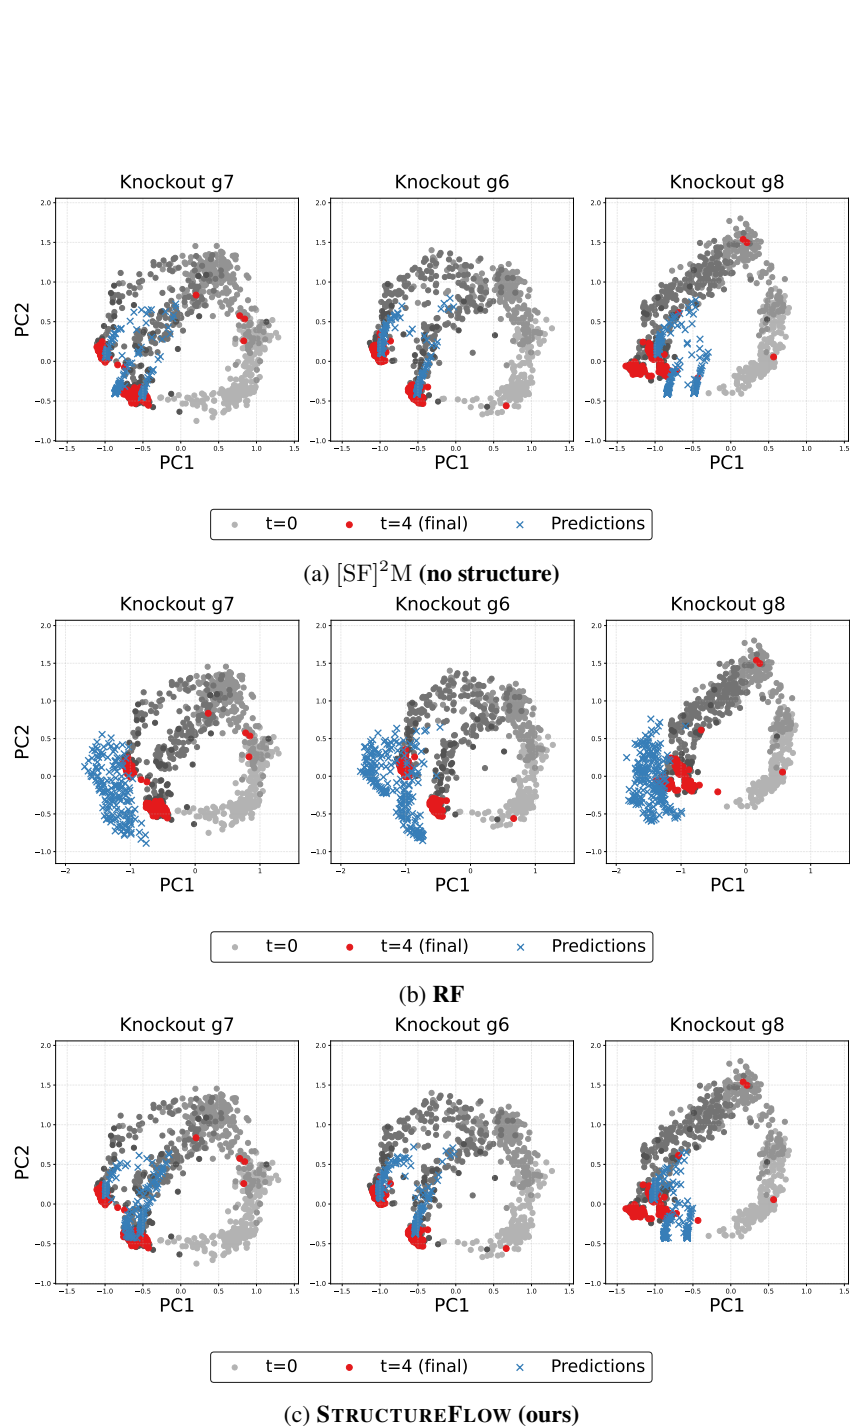


Figure 12: **Comparison of trajectory inference performance for left-out interventions (knock-outs).** We show PCA visualizations of model predictions for interventions that were excluded during training. Ground truth data is shown in grey, with model predictions overlaid. This demonstrates each method’s ability to generalize to unseen perturbational conditions.

Table 12: Ablation over size of hidden dimension (# of neurons) per layer for the NGM flow model.

Hidden dimension for NGM	10	50	100	200	256
AUROC	0.925 \pm 0.010	0.964 \pm 0.004	0.971 \pm 0.008	0.969 \pm 0.003	0.975 \pm 0.003
AP	0.832 \pm 0.020	0.910 \pm 0.013	0.917 \pm 0.025	0.916 \pm 0.014	0.932 \pm 0.012
W_2	0.7950 \pm 0.0024	0.7853 \pm 0.0026	0.7886 \pm 0.0048	0.7972 \pm 0.0055	0.7858 \pm 0.0208

E.7 ASSESSING THE STABILITY OF THE SCORE AND FLOW NETWORK IN STRUCTUREFLOW

Here we test two settings where we try freezing the weights of one of the modules. We jointly train s_ϕ and v_θ for $K \in \{10000, 15000\}$ iterations, then freeze the parameters of s_ϕ and continue training only v_θ for an additional $M = 5000$ iterations. We do the same experiment but freeze the flow term instead. In both settings our performance stays constant for structure learning, but is more sensitive for dynamics.

Table 13: Ablation over freezing flow and score model post-training when models pre-trained for 10000 iterations.

Freeze model	AP	AUROC	W_2
score	0.9011	0.9681	1.5462 \pm 0.0520
flow	0.9109	0.9700	0.9366 \pm 0.1992

Table 14: Ablation over freezing flow and score model post-training when models pre-trained for 15000 iterations.

Freeze model	AP	AUROC	W_2
score	0.9019	0.9668	1.5200 \pm 0.0569
flow	0.8843	0.9650	0.9337 \pm 0.2181

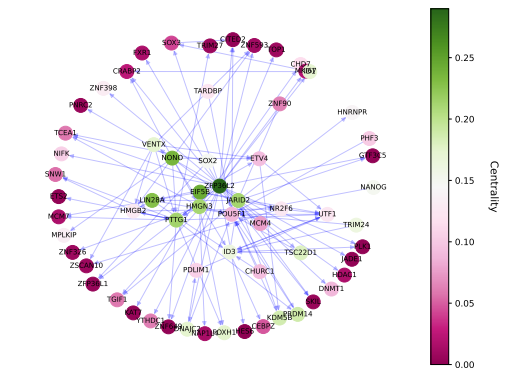
E.8 SCALING AND RUNTIME OF VARIOUS BASELINES AND OUR METHOD ON RENGE DATA

Table 15: Model wall-clock time on the Real (Renge) dataset. We include parameter counts for all neural network-based methods. We note that for STRUCTUREFLOW and $[SF]^2M$ the OT pairing was negligible and thus not included.

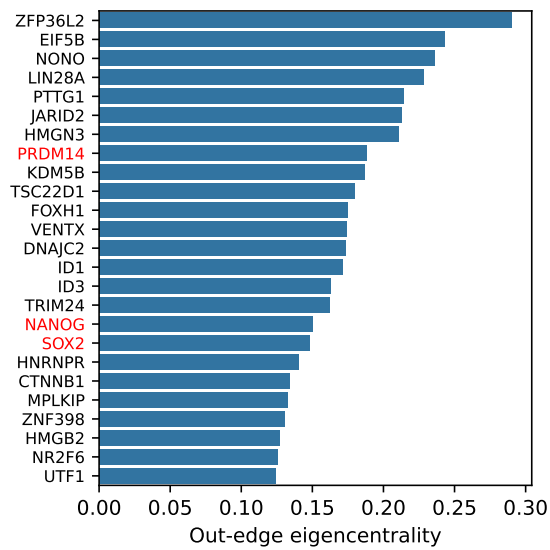
Time [s]	Model	Parameter Size
319.31 \pm 21.29	STRUCTUREFLOW	1440646
1542.88 \pm 26.67	TIGON (PCA)	43140
14112.49	TIGON (Ambient)	409448
338.12 \pm 0.78	RF	-
101.31 \pm 1.43	$[SF]^2M$	175310
12.21 \pm .125	OTVelo	-

E.9 STRUCTURE EXPERIMENTS FOR REAL DATA


1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619



(a) STRUCTUREFLOW inferred regulatory network graph.



(b) Out-edge eigencentralities scores for the STRUCTUREFLOW-inferred network.



(c) Ground truth regulatory network (heatmap). (d) StructFlow-inferred network (no threshold). (e) StructFlow-inferred network (99th quantile threshold).

Figure 13: **Analysis of gene regulatory network inference results.** (a) STRUCTUREFLOW graph representation of the inferred gene regulatory network. (b) The top 25 Centrality scores (out-edge eigencentralities) for nodes in the STRUCTUREFLOW-inferred network. (c) Ground truth regulatory network shown as a heatmap. This network is from the experimental binding information from the ChIP-atlas. (d) STRUCTUREFLOW-inferred network adjacency matrix (no thresholding). (e) STRUCTUREFLOW-inferred network thresholded at the 99th quantile to highlight the strongest edges.

E.10 EVALUATING STRUCTUREFLOW WITH IRREGULARLY SAMPLED TIME-POINTS

We evaluate STRUCTUREFLOW and Reference Fitting on the TF BoolODE system using population snapshots from non-uniform timepoints. To accomplish this, we bin the continuous timepoint data into randomly non-uniformly spaced time bins with a standard deviation of 30% of the time bin width (whereas our original experiments bin into discrete equally spaced timebins of $t = 1, 2, 3, \dots$). We run this experiment over 3 seeds, where the seeding randomizes the time bin spacing (making the data non-deterministic between runs, and thus RF is non-deterministic in this specific setting).

Table 16: Comparison of STRUCTUREFLOW and RF on irregularly sampled time-points for the TF system.

Metric	STRUCTUREFLOW	RF
GRN Graph AP	0.8501 ± 0.0348	0.9614 ± 0.0282
GRN Graph AUROC	0.9326 ± 0.0180	0.9812 ± 0.0099
Trajectory W_2	0.8491 ± 0.0168	1.4060 ± 0.0665
Trajectory MMD	0.1408 ± 0.0022	0.1977 ± 0.0152

E.11 COMPARISON WITH SYMBOLIC DISTRIBUTION FLOW LEARNING (SDFL)

We compare STRUCTUREFLOW to Symbolic Distribution Flow Learning (SDFL) (Dakhmouche et al., 2025), which tackles the population dynamics problem through learning explicit symbolic equations. While SDFL provides interpretable outputs through per-dimension equations, it requires significant computation time to run and cannot explicitly recover a learned matrix structure for the structure learning task. We evaluate SDFL on the trifurcating (TF) system for the trajectory inference task and report results below (we include STRUCTUREFLOW results from Table 1 for easy comparison). Note due to the methods training instability, we were unable to get 3 consecutive successful seeds, and therefore report only 1 seed result.

Table 17: Comparison of dynamical inference performance with SDFL on the TF system.

Method	$W_2 \downarrow$	MMD \downarrow
SDFL	3.4880	0.0785
STRUCTUREFLOW	0.789 ± 0.012	0.135 ± 0.004

APPENDIX REFERENCES

- Huynh-Thu, V. A., Irrthum, A., Wehenkel, L., and Geurts, P. (2010). Inferring regulatory networks from expression data using tree-based methods. *PLoS ONE*, 5(9):e12776.
- Matsumoto, H., Kiryu, H., Furusawa, C., Ko, M. S. H., Ko, S. B. H., Gouda, N., Hayashi, T., and Nikaido, I. (2017). Scode: an efficient regulatory network inference algorithm from single-cell rna-seq during differentiation. *Bioinformatics*, 33(15):2314–2321.
- Zou, Z., Ohta, T., and Oki, S. (2024). ChIP-Atlas 3.0: a data-mining suite to explore chromosome architecture together with large-scale regulome data. *Nucleic Acids Research*, page gkae358.